Oracle® Database

PL/SQL Packages and Types Reference

23ai
F46980-20
October 2024

Oracle Database PL/SQL Packages and Types Reference, 23ai

F46980-20

Copyright © 1996, 2024, Oracle and/or its affiliates.

Primary Author: Gunjan Jain

Contributing Authors: Louise Morin, Denis Raphaely, Lance Ashdown, Drue Baker, Donna Carver, Maitreyee Chaliha, Beethoven Cheng, Rhonda Day, Steve Fogel, Bryn Llewellyn, Paul Lane, David McDermid, Tony Morales, Anne Murphy, Chuck Murray, Sue Pelski, Kathy Rich, Antonio Romero, Vivian Schupmann, Cathy Shea, Margaret Taft, Kathy Taylor, Randy Urbano, Pat Huey, Drew Adams, Sarika Surumpudi, Rodney Ward, Rajesh Bhatiya, Tanmay Choudhury, Roopesh Ashok Kumar, Atif Chaudhry, Subha Vikram

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface

Audience cxlvii
Documentation Accessibility cxlvii
Related Documents cxlvii
Conventions cxlviii

Changes in This Release for Oracle Database PL/SQL Packages and Types Reference

Changes in this Release cl

1 Introduction to Oracle Supplied PL/SQL Packages & Types

1.1 Package Overview 1-1
1.1.1 Package Components 1-2
1.1.2 Using Oracle Supplied Packages 1-2
1.1.3 Creating New Packages 1-2
1.1.3.1 Separating the Specification and Body 1-3
1.1.4 Referencing Package Contents 1-5
1.2 Summary of Oracle Supplied PL/SQL Packages and Types 1-5

2 Oracle Application Express Packages APEX_APPLICATION Through APEX_ZIP

3 CTX_ADM

4 CTX_ANL

4.1 GET_COMPRESSION_RATIO Procedure 4-1

5 CTX_DDL

6 CTX_CLS

7 CTX_DOC

8 CTX_ENTITY

9 CTX_OUTPUT

10 CTX_QUERY

11 CTX_REPORT

12 CTX_THES

13 CTX_ULEXER

14 DBMS_ACTIVITY

14.1 Using DBMS_ACTIVITY 14-1
14.2 Summary of DBMS_ACTIVITY Subprograms 14-2
14.2.1 CONFIGURE Procedure 14-2
14.2.2 CREATE_SNAPSHOT Procedure 14-3
14.2.3 DELETE_SNAPSHOT Procedure 14-4
14.2.4 DELETE_SNAPSHOT Procedure 14-5

15 DBMS_ADDM

15.1 DBMS_ADDM Security Model 15-1
15.2 Summary of DBMS_ADDM Subprograms 15-1
15.2.1 ANALYZE_DB Procedure 15-3

15.2.2 ANALYZE_INST Procedure 15-4
15.2.3 ANALYZE_PARTIAL Procedure 15-5
15.2.4 COMPARE_CAPTURE_REPLAY_REPORT Function 15-7
15.2.5 COMPARE_DATABASES Function 15-8
15.2.6 COMPARE_INSTANCES Function 15-9
15.2.7 COMPARE_REPLAY_REPLAY_REPORT Function 15-10
15.2.8 DELETE Procedure 15-11
15.2.9 DELETE_FINDING_DIRECTIVE Procedure 15-12
15.2.10 DELETE_PARAMETER_DIRECTIVE Procedure 15-12
15.2.11 DELETE_SEGMENT_DIRECTIVE Procedure 15-13
15.2.12 DELETE_SQL_DIRECTIVE Procedure 15-13
15.2.13 FAILED_AUTO_TASKS_REPORT Function 15-14
15.2.14 GET_ASH_QUERY Function 15-14
15.2.15 GET_REPORT Function 15-15
15.2.16 INSERT_FINDING_DIRECTIVE Procedure 15-15
15.2.17 INSERT_PARAMETER_DIRECTIVE Procedure 15-16
15.2.18 INSERT_SEGMENT_DIRECTIVE Procedure 15-17
15.2.19 INSERT_SQL_DIRECTIVE Procedure 15-18
15.2.20 REAL_TIME_ADDM_REPORT Function 15-19
15.2.21 REEXECUTE_FAILED_AUTO_TASKS Procedure 15-20

16 DBMS_ADVANCED_REWRITE

16.1 DBMS_ADVANCED_REWRITE Security Model 16-1
16.2 Summary of DBMS_ADVANCED_REWRITE Subprograms 16-1
16.2.1 ALTER_REWRITE_EQUIVALENCE Procedure 16-2
16.2.2 BUILD_SAFE_REWRITE_EQUIVALENCE Procedure 16-3
16.2.3 DECLARE_REWRITE_EQUIVALENCE Procedures 16-3
16.2.4 DROP_REWRITE_EQUIVALENCE Procedure 16-5
16.2.5 VALIDATE_REWRITE_EQUIVALENCE Procedure 16-6

17 DBMS_ADVISOR

17.1 DBMS_ADVISOR Deprecated Subprograms 17-1
17.2 DBMS_ADVISOR Security Model 17-2
17.3 Summary of DBMS_ADVISOR Subprograms 17-2
17.3.1 ADD_SQLWKLD_REF Procedure 17-5
17.3.2 ADD_SQLWKLD_STATEMENT Procedure 17-6
17.3.3 ADD_STS_REF Procedure 17-8
17.3.4 CANCEL_TASK Procedure 17-9
17.3.5 COPY_SQLWKLD_TO_STS Procedure 17-10
17.3.6 CREATE_FILE Procedure 17-11

17.3.7 CREATE_OBJECT Procedure 17-12
17.3.8 CREATE_SQLWKLD Procedure 17-14
17.3.9 CREATE_TASK Procedures 17-15
17.3.10 DELETE_SQLWKLD Procedure 17-17
17.3.11 DELETE_SQLWKLD_REF Procedure 17-18
17.3.12 DELETE_SQLWKLD_STATEMENT Procedure 17-19
17.3.13 DELETE_STS_REF Procedure 17-20
17.3.14 DELETE_TASK Procedure 17-21
17.3.15 EXECUTE_TASK Procedure 17-21
17.3.16 GET_REC_ATTRIBUTES Procedure 17-23
17.3.17 GET_TASK_REPORT Function 17-24
17.3.18 GET_TASK_SCRIPT Function 17-25
17.3.19 IMPLEMENT_TASK Procedure 17-26
17.3.20 IMPORT_SQLWKLD_SCHEMA Procedure 17-27
17.3.21 IMPORT_SQLWKLD_SQLCACHE Procedure 17-28
17.3.22 IMPORT_SQLWKLD_STS Procedure 17-30
17.3.23 IMPORT_SQLWKLD_SUMADV Procedure 17-31
17.3.24 IMPORT_SQLWKLD_USER Procedure 17-33
17.3.25 INTERRUPT_TASK Procedure 17-34
17.3.26 MARK_RECOMMENDATION Procedure 17-35
17.3.27 QUICK_TUNE Procedure 17-36
17.3.28 RESET_SQLWKLD Procedure 17-38
17.3.29 RESET_TASK Procedure 17-39
17.3.30 SET_DEFAULT_SQLWKLD_PARAMETER Procedure 17-39
17.3.31 SET_DEFAULT_TASK_PARAMETER Procedure 17-40
17.3.32 SET_SQLWKLD_PARAMETER Procedure 17-41
17.3.33 SET_TASK_PARAMETER Procedure 17-43
17.3.34 TUNE_MVIEW Procedure 17-57
17.3.35 UPDATE_OBJECT Procedure 17-59
17.3.36 UPDATE_REC_ATTRIBUTES Procedure 17-61
17.3.37 UPDATE_SQLWKLD_ATTRIBUTES Procedure 17-62
17.3.38 UPDATE_SQLWKLD_STATEMENT Procedure 17-63
17.3.39 UPDATE_TASK_ATTRIBUTES Procedure 17-65

18 DBMS_AUTOIM

18.1 Using DBMS_AUTOIM 18-1
18.2 Summary of DBMS_AUTOIM Subprograms 18-1
18.2.1 SET_PARAMETER Procedure 18-1
18.2.2 ACTIVITY_REPORT Function 18-2

19 DBMS_ALERT

19.1 DBMS_ALERT Overview 19-1
19.2 DBMS_ALERT Security Model 19-1
19.3 DBMS_ALERT Constants 19-1
19.4 DBMS_ALERT Restrictions 19-2
19.5 DBMS_ALERT Exceptions 19-2
19.6 DBMS_ALERT Operational Notes 19-2
19.7 DBMS_ALERT Examples 19-3
19.8 Summary of DBMS_ALERT Subprograms 19-4
19.8.1 REGISTER Procedure 19-4
19.8.2 REMOVE Procedure 19-5
19.8.3 REMOVEALL Procedure 19-5
19.8.4 SET_DEFAULTS Procedure 19-5
19.8.5 SIGNAL Procedure 19-6
19.8.6 WAITANY Procedure 19-6
19.8.7 WAITONE Procedure 19-7

20 DBMS_APP_CONT

20.1 DBMS_APP_CONT Overview 20-1
20.2 DBMS_APP_CONT Security Model 20-2
20.3 Summary of DBMS_APP_CONT Subprograms 20-2
20.3.1 ADD_SQL_CONNECTION_TEST Procedure 20-2
20.3.2 APPLY_REPLAY_RULE Procedure 20-3
20.3.3 DELETE_SQL_CONNECTION_TEST Procedure 20-4
20.3.4 DISABLE_CONNECTION_TEST Procedure 20-4
20.3.5 ENABLE_CONNECTION_TEST Procedure 20-5
20.3.6 GET_LTXID_OUTCOME Procedure 20-6
20.3.7 GET_REPLAY_RULES Function 20-7
20.3.8 RESET_REPLAY_RULES Procedure 20-8

21 DBMS_APP_CONT_ADMIN

21.1 DBMS_APP_CONT_ADMIN Security Model 21-1
21.2 Summary of DBMS_APP_CONT_ADMIN Subprograms 21-1
21.2.1 ACCHK_CLEAR_FILTER Procedure 21-2
21.2.2 ACCHK_SET Procedure 21-3
21.2.3 ACCHK_SET_FILTER Procedure 21-5
21.2.4 ACCHK_SHOW_FILTERS Procedure 21-6
21.2.5 ADD_SQL_CONNECTION_TEST Procedure 21-6
21.2.6 DELETE_SQL_CONNECTION_TEST Procedure 21-7

21.2.7 DISABLE_CONNECTION_TEST Procedure 21-8
21.2.8 DISABLE_FAILOVER Procedure 21-8
21.2.9 ENABLE_AC Procedure 21-9
21.2.10 ENABLE_CONNECTION_TEST Procedure 21-10
21.2.11 ENABLE_RESET_STATE Procedure 21-11
21.2.12 ENABLE_TAC Procedure 21-11
21.2.13 MODIFY_SERVICE Procedure 21-12
21.2.14 SET_DRAINING Procedure 21-15
21.2.15 SET_LOAD_BALANCING Procedure 21-16

22 DBMS_APP_CONT_REPORT

22.1 DBMS_APP_CONT_REPORT Security Model 22-1
22.2 Summary of DBMS_APP_CONT_REPORT Subprograms 22-1
22.2.1 ACCHK_REPORT Procedure 22-1

23 DBMS_APPLICATION_INFO

23.1 DBMS_APPLICATION_INFO Overview 23-1
23.2 DBMS_APPLICATION_INFO Security Model 23-1
23.3 DBMS_APPLICATION_INFO Operational Notes 23-1
23.4 Summary of DBMS_APPLICATION_INFO Subprograms 23-2
23.4.1 READ_CLIENT_INFO Procedure 23-2
23.4.2 READ_MODULE Procedure 23-2
23.4.3 SET_ACTION Procedure 23-3
23.4.4 SET_CLIENT_INFO Procedure 23-4
23.4.5 SET_MODULE Procedure 23-5
23.4.6 SET_SESSION_LONGOPS Procedure 23-5

24 DBMS_APPLY_ADM

24.1 DBMS_APPLY_ADM Overview 24-1
24.2 DBMS_APPLY_ADM Security Model 24-2
24.3 DBMS_APPLY_ADM Deprecated Subprograms 24-3
24.4 Summary of DBMS_APPLY_ADM Subprograms 24-3
24.4.1 ALTER_APPLY Procedure 24-5
24.4.2 CLEAR_KEY_COLUMNS Procedure 24-10
24.4.3 COMPARE_OLD_VALUES Procedure 24-11
24.4.4 CREATE_APPLY Procedure 24-13
24.4.5 CREATE_OBJECT_DEPENDENCY Procedure 24-19
24.4.6 DELETE_ALL_ERRORS Procedure 24-20
24.4.7 DELETE_ERROR Procedure 24-20

24.4.8 DROP_APPLY Procedure 24-21
24.4.9 DROP_OBJECT_DEPENDENCY Procedure 24-22
24.4.10 EXECUTE_ALL_ERRORS Procedure 24-23
24.4.11 EXECUTE_ERROR Procedure 24-24
24.4.12 GET_ERROR_MESSAGE Function 24-26
24.4.13 HANDLE_COLLISIONS 24-28
24.4.14 SET_DML_CONFLICT_HANDLER Procedure 24-28
24.4.15 SET_DML_HANDLER Procedure 24-34
24.4.16 SET_ENQUEUE_DESTINATION Procedure 24-38
24.4.17 SET_EXECUTE Procedure 24-40
24.4.18 SET_GLOBAL_INSTANTIATION_SCN Procedure 24-41
24.4.19 SET_KEY_COLUMNS Procedures 24-45
24.4.20 SET_PARAMETER Procedure 24-48
24.4.21 SET_REPERROR_HANDLER Procedure 24-69
24.4.22 SET_SCHEMA_INSTANTIATION_SCN Procedure 24-70
24.4.23 SET_TABLE_INSTANTIATION_SCN Procedure 24-73
24.4.24 SET_UPDATE_CONFLICT_HANDLER Procedure 24-75
24.4.25 SET_VALUE_DEPENDENCY Procedure 24-79
24.4.26 START_APPLY Procedure 24-80
24.4.27 STOP_APPLY Procedure 24-81

25 DBMS_AQ

25.1 DBMS_AQ Security Model 25-1
25.2 DBMS_AQ Constants 25-2
25.3 DBMS_AQ Data Types 25-3
25.4 DBMS_AQ Operational Notes 25-5
25.5 Summary of DBMS_AQ Subprograms 25-5
25.5.1 BIND_AGENT Procedure 25-6
25.5.2 DEQUEUE Procedure 25-7
25.5.3 DEQUEUE_ARRAY Function 25-9
25.5.4 ENQUEUE Procedure 25-10
25.5.5 ENQUEUE_ARRAY Function 25-11
25.5.6 LISTEN Procedures 25-12
25.5.7 POST Procedure 25-13
25.5.8 REGISTER Procedure 25-14
25.5.9 SEEK Procedure 25-15
25.5.10 UNBIND_AGENT Procedure 25-16
25.5.11 UNREGISTER Procedure 25-16

26 DBMS_AQADM

26.1 DBMS_AQADM Security Model 26-1
26.2 DBMS_AQADM Constants 26-2
26.3 DBMS_AQADM Subprogram Groups 26-3
26.3.1 DBMS_AQADM Queue Table Subprograms 26-3
26.3.2 DBMS_AQADM Privilege Subprograms 26-4
26.3.3 DBMS_AQADM Queue Subprograms 26-4
26.3.4 DBMS_AQADM Subscriber Subprograms 26-5
26.3.5 DBMS_AQADM OKafka Subprograms 26-5
26.3.6 DBMS_AQADM Propagation Subprograms 26-5
26.3.7 DBMS_AQADM Miscellaneous Subprograms 26-6
26.3.8 DBMS_AQADM Agent Subprograms 26-6
26.3.9 DBMS_AQADM Alias Subprograms 26-7
26.4 Summary of DBMS_AQADM Subprograms 26-7
26.4.1 ADD_ALIAS_TO_LDAP Procedure 26-10
26.4.2 ADD_SUBSCRIBER Procedure 26-10
26.4.3 ALTER_AQ_AGENT Procedure 26-11
26.4.4 ALTER_PROPAGATION_SCHEDULE Procedure 26-12
26.4.5 ALTER_QUEUE Procedure 26-13
26.4.6 ALTER_QUEUE_TABLE Procedure 26-14
26.4.7 ALTER_SHARDED_QUEUE Procedure 26-15
26.4.8 ALTER_SUBSCRIBER Procedure 26-16
26.4.9 ALTER_TRANSACTIONAL_EVENT_QUEUE Procedure 26-16
26.4.10 CREATE_AQ_AGENT Procedure 26-17
26.4.11 CREATE_DATABASE_KAFKA_TOPIC Procedure 26-18
26.4.12 CREATE_NP_QUEUE Procedure 26-19
26.4.13 CREATE_QUEUE Procedure 26-20
26.4.14 CREATE_QUEUE_TABLE Procedure 26-21
26.4.15 CREATE_SHARDED_QUEUE Procedure 26-24
26.4.16 CREATE_EXCEPTION_QUEUE Procedure 26-25
26.4.17 CREATE_EQ EXCEPTION_QUEUE Procedure 26-26
26.4.18 CREATE_TRANSACTIONAL_EVENT_QUEUE Procedure 26-26
26.4.19 DEL_ALIAS_FROM_LDAP Procedure 26-27
26.4.20 DISABLE_DB_ACCESS Procedure 26-28
26.4.21 DISABLE_PROPAGATION_SCHEDULE Procedure 26-28
26.4.22 DROP_AQ_AGENT Procedure 26-29
26.4.23 DROP_DATABASE_KAFKA_TOPIC Procedure 26-29
26.4.24 DROP_QUEUE Procedure 26-30
26.4.25 DROP_QUEUE_TABLE Procedure 26-30
26.4.26 DROP_SHARDED_QUEUE Procedure 26-31
26.4.27 DROP_TRANSACTIONAL_EVENT_QUEUE Procedure 26-31

26.4.28 ENABLE_DB_ACCESS Procedure 26-32
26.4.29 ENABLE_JMS_TYPES Procedure 26-32
26.4.30 ENABLE_PROPAGATION_SCHEDULE Procedure 26-33
26.4.31 GET_MAX_STREAMS_POOL Procedure 26-33
26.4.32 GET_MIN_STREAMS_POOL Procedure 26-34
26.4.33 GET_QUEUE_PARAMETER Procedure 26-34
26.4.34 GET_WATERMARK Procedure 26-37
26.4.35 GRANT_QUEUE_PRIVILEGE Procedure 26-37
26.4.36 GRANT_SYSTEM_PRIVILEGE Procedure 26-37
26.4.37 MIGRATE_QUEUE_TABLE Procedure 26-38
26.4.38 PURGE_QUEUE_TABLE Procedure 26-39
26.4.39 QUEUE_SUBSCRIBERS Function 26-40
26.4.40 REMOVE_SUBSCRIBER Procedure 26-41
26.4.41 REVOKE_QUEUE_PRIVILEGE Procedure 26-41
26.4.42 REVOKE_SYSTEM_PRIVILEGE Procedure 26-42
26.4.43 SCHEDULE_PROPAGATION Procedure 26-42
26.4.44 SET_QUEUE_PARAMETER Procedure 26-44
26.4.45 SET_MAX_STREAMS_POOL Procedure 26-47
26.4.46 SET_MIN_STREAMS_POOL Procedure 26-47
26.4.47 SET_WATERMARK Procedure 26-47
26.4.48 START_QUEUE Procedure 26-48
26.4.49 STOP_QUEUE Procedure 26-48
26.4.50 UNSCHEDULE_PROPAGATION Procedure 26-49
26.4.51 UNSET_QUEUE_PARAMETER Procedure 26-50
26.4.52 VERIFY_QUEUE_TYPES Procedure 26-52

27 DBMS_AQELM

27.1 DBMS_AQELM Security Model 27-1
27.2 Summary of DBMS_AQELM Subprograms 27-1
27.2.1 SET_MAILHOST Procedure 27-1
27.2.2 SET_MAILPORT Procedure 27-2
27.2.3 SET_SENDFROM Procedure 27-2

28 DBMS_AQIN

28.1 DBMS_AQIN Security Model 28-1

29 DBMS_AQMIGTOOL

29.1 DBMS_AQMIGTOOL Security Model 29-1
29.2 DBMS_AQMIGTOOL Constants 29-1

29.3 Summary of DBMS_AQMIGTOOL Subprograms 29-2
29.3.1 CANCEL_MIGRATION Procedure 29-3
29.3.2 CHECK_MIGRATED_MESSAGES Procedure 29-4
29.3.3 CHECK_MIGRATION_TO_TXEVENTQ Procedure 29-4
29.3.4 CHECK_STATUS Procedure 29-5
29.3.5 CLEAR_UNSUPPORTED_FEATURE_TABLE Procedure 29-7
29.3.6 COMMIT_MIGRATION Procedure 29-8
29.3.7 DISABLE_MIGRATION_CHECK Procedure 29-8
29.3.8 INIT_MIGRATION Procedure 29-9
29.3.9 PURGE_QUEUE_MESSAGES Procedure 29-11
29.3.10 RECOVER_MIGRATION Procedure 29-12
29.3.11 RENAME_QUEUE Procedure 29-13

30 DBMS_ASSERT

30.1 DBMS_ASSERT Operational Notes 30-1
30.2 Summary of DBMS_ASSERT Subprograms 30-1
30.2.1 ENQUOTE_LITERAL Function 30-2
30.2.2 ENQUOTE_NAME Function 30-2
30.2.3 NOOP Functions 30-4
30.2.4 QUALIFIED_SQL_NAME Function 30-4
30.2.5 SCHEMA_NAME Function 30-5
30.2.6 SIMPLE_SQL_NAME Function 30-5
30.2.7 SQL_OBJECT_NAME Function 30-6

31 DBMS_AUDIT_MGMT

31.1 DBMS_AUDIT_MGMT Overview 31-1
31.2 DBMS_AUDIT_MGMT Deprecated and Desupported Subprograms 31-2
31.3 DBMS_AUDIT_MGMT Security Model 31-3
31.4 DBMS_AUDIT_MGMT Constants 31-3
31.5 DBMS_AUDIT_MGMT Views 31-5
31.6 DBMS_AUDIT_MGMT Subprogram Groups 31-5
31.6.1 DBMS_AUDIT_MGMT Audit Trail Management Subprograms 31-5
31.6.2 DBMS_AUDIT_MGMT Audit Trail Cleanup Subprograms 31-6
31.7 Summary of DBMS_AUDIT_MGMT Subprograms 31-7
31.7.1 ALTER_PARTITION_INTERVAL Procedure 31-8
31.7.2 CLEAN_AUDIT_TRAIL Procedure 31-8
31.7.3 CLEAR_AUDIT_TRAIL_PROPERTY Procedure 31-11
31.7.4 CLEAR_LAST_ARCHIVE_TIMESTAMP Procedure 31-12
31.7.5 CREATE_PURGE_JOB Procedure 31-14
31.7.6 DEINIT_CLEANUP Procedure 31-15

31.7.7 DROP_OLD_UNIFIED_AUDIT_TABLES Procedure 31-16
31.7.8 DROP_PURGE_JOB Procedure 31-17
31.7.9 GET_AUDIT_COMMIT_DELAY Function 31-17
31.7.10 GET_ AUDIT_TRAIL_PROPERTY_VALUE Function 31-18
31.7.11 GET_LAST_ARCHIVE_TIMESTAMP Function 31-19
31.7.12 INIT_CLEANUP Procedure 31-20
31.7.13 IS_CLEANUP_INITIALIZED Function 31-22
31.7.14 LOAD_UNIFIED_AUDIT_FILES Procedure 31-23
31.7.15 SET_AUDIT_TRAIL_LOCATION Procedure 31-24
31.7.16 SET_AUDIT_TRAIL_PROPERTY Procedure 31-26
31.7.17 SET_LAST_ARCHIVE_TIMESTAMP Procedure 31-28
31.7.18 SET_PURGE_JOB_INTERVAL Procedure 31-30
31.7.19 SET_PURGE_JOB_STATUS Procedure 31-31
31.7.20 TRANSFER_UNIFIED_AUDIT_RECORDS Procedure 31-32

32 DBMS_AUDIT_UTIL

32.1 DBMS_AUDIT_UTIL Overview 32-1
32.2 DBMS_AUDIT_UTIL Security Model 32-1
32.3 DBMS_AUDIT_UTIL Views 32-1
32.4 Summary of DBMS_AUDIT_UTIL Subprograms 32-2
32.4.1 DECODE_RLS_INFO_ATRAIL_FGA Function 32-2
32.4.2 DECODE_RLS_INFO_ATRAIL_STD Function 32-3
32.4.3 DECODE_RLS_INFO_ATRAIL_UNI Function 32-4
32.4.4 DECODE_RLS_INFO_ATRAIL_XML Function 32-5

33 DBMS_AUTO_CLUSTERING

33.1 DBMS_AUTO_CLUSTERING Overview 33-1
33.2 Summary of DBMS_AUTO_CLUSTERING Subprograms 33-1
33.2.1 CONFIGURE Procedure 33-2
33.2.2 RECOMMEND_CLUSTERING_METHOD Function 33-4
33.2.3 VERIFY_RECOMMENDATION Procedure 33-5
33.2.4 APPLY_RECOMMENDATION Procedure 33-6
33.2.5 GET_RECOMMENDATION Function 33-7
33.2.6 REPORT ACTIVITY Function 33-7
33.2.7 REPORT_LAST_ACTIVITY Function 33-8

34 DBMS_AUTO_SQLTUNE

34.1 DBMS_AUTO_SQLTUNE Overview 34-1
34.2 DBMS_AUTO_SQLTUNE Security Model 34-1

34.3 Summary of DBMS_AUTO_SQLTUNE Subprograms 34-1
34.3.1 EXECUTE_AUTO_TUNING_TASK Function and Procedure 34-2
34.3.2 REPORT_AUTO_TUNING_TASK Function 34-3
34.3.3 SET_AUTO_TUNING_TASK_PARAMETER Procedures 34-4

35 DBMS_AUTO_INDEX

35.1 DBMS_AUTO_INDEX Overview 35-1
35.2 Summary of DBMS_AUTO_INDEX Subprograms 35-1
35.2.1 CONFIGURE Procedure 35-1
35.2.2 DROP_AUTO_INDEXES Procedure 35-6
35.2.3 DROP_SECONDARY_INDEXES Procedure 35-7
35.2.4 REPORT_ACTIVITY Function 35-8
35.2.5 REPORT_LAST_ACTIVITY Function 35-10
35.2.6 RECOMMEND Function 35-12

36 DBMS_AUTO_MV

36.1 Using DBMS_AUTO_MV 36-1
36.2 Summary of DBMS_AUTO_MV Subprograms 36-1
36.2.1 CONFIGURE Procedure 36-2
36.2.2 DROP_AUTO_MVS Procedure 36-4
36.2.3 RECOMMEND Function 36-5
36.2.4 REFRESH Procedure 36-7
36.2.5 REPORT_ACTIVITY Function 36-7
36.2.6 REPORT_LAST_ACTIVITY Function 36-8

37 DBMS_AUTO_REPORT

37.1 DBMS_AUTO_REPORT Overview 37-1
37.2 DBMS_AUTO_REPORT Security Model 37-1
37.3 Summary of DBMS_AUTO_REPORT Subprograms 37-1
37.3.1 FINISH_REPORT_CAPTURE Procedure 37-2
37.3.2 REPORT_REPOSITORY_DETAIL Function 37-2
37.3.3 REPORT_REPOSITORY_DETAIL_XML Function 37-3
37.3.4 REPORT_REPOSITORY_LIST_XML Function 37-3
37.3.5 START_REPORT_CAPTURE Procedure 37-5
37.3.6 PURGE_ALL_REPORTS Function 37-5
37.3.7 PURGE_REPORTS_BY_RETENTION Function 37-5

38 DBMS_AUTO_TASK_ADMIN

38.1 DBMS_AUTO_TASK_ADMIN Deprecated Subprograms 38-1
38.2 DBMS_AUTO_TASK_ADMIN Security Model 38-1
38.3 DBMS_AUTO_TASK_ADMIN Constants 38-1
38.4 Summary of DBMS_AUTO_TASK_ADMIN Subprograms 38-2
38.4.1 DISABLE Procedures 38-2
38.4.2 ENABLE Procedures 38-3
38.4.3 GET_CLIENT_ATTRIBUTES Procedure 38-4
38.4.4 GET_P1_RESOURCES Procedure 38-4
38.4.5 OVERRIDE_PRIORITY Procedures 38-5
38.4.6 SET_CLIENT_SERVICE Procedure 38-6
38.4.7 SET_P1_RESOURCES Procedure 38-6

39 DBMS_AUTO_ZONEMAP

39.1 DBMS_AUTO_TASK_ADMIN Security Model 39-1
39.2 Summary of DBMS_AUTO_ZONEMAP Subprograms 39-1
39.2.1 ACTIVITY_REPORT Function 39-1
39.2.2 CONFIGURE Procedure 39-4

40 DBMS_AVTUNE

40.1 DBMS_AVTUNE Overview 40-1
40.2 DBMS_AVTUNE Security Model 40-1
40.3 Summary of DBMS_AVTUNE Subprograms 40-1
40.3.1 AUTO_CACHE_CREATE Procedure 40-2
40.3.2 AUTO_CACHE_DISABLE Procedure 40-3
40.3.3 AUTO_CACHE_ENABLE Procedure 40-3
40.3.4 AUTO_CACHE_MODIFY Procedure 40-4
40.3.5 AUTO_CACHE_REFRESH Procedure 40-4
40.3.6 AUTO_CACHE_REMOVE Procedure 40-5
40.3.7 AUTO_CACHE_TUNE Procedure 40-5
40.3.8 AUTO_CACHE_STAR_DISABLE Procedure 40-6
40.3.9 AUTO_CACHE_STAR_ENABLE Procedure 40-6
40.3.10 AUTO_CACHE_STAR_MODIFY Procedure 40-7
40.3.11 AUTO_CACHE_STAR_REFRESH Procedure 40-7

41 DBMS_AW_STATS

41.1 Using DBMS_AW_STATS 41-1
41.2 Summary of DBMS_AW_STATS Subprograms 41-1
41.2.1 ANALYZE Procedure 41-1

41.2.2 CLEAR Procedure 41-4

42 DBMS_BLOCKCHAIN_TABLE

42.1 DBMS_BLOCKCHAIN_TABLE Overview 42-1
42.2 DBMS_BLOCKCHAIN_TABLE Security Model 42-2
42.3 Summary of DBMS_BLOCKCHAIN_TABLE Subprograms 42-2
42.3.1 ADD_INTERVAL_PARTITIONING Procedure 42-4
42.3.2 COUNTERSIGN_ROW Procedure 42-5
42.3.3 COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure 42-6
42.3.4 DELETE_EXPIRED_ROWS Procedure 42-7
42.3.5 GET_BLOCKCHAIN_DIGEST Function 42-8
42.3.6 GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function 42-9
42.3.7 GET_BYTES_FOR_ROW_HASH Procedure 42-10
42.3.8 GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_COLUMNS Procedure 42-12
42.3.9 GET_BYTES_FOR_ROW_SIGNATURE Procedure 42-13
42.3.10 GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_KEY_COLUMNS Procedure 42-14
42.3.11 GET_SIGNED_BLOCKCHAIN_DIGEST Function 42-15
42.3.12 GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function 42-16
42.3.13 SIGN_ROW Procedure 42-18
42.3.14 SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure 42-20
42.3.15 SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE Procedure 42-21
42.3.16 SIGN_ROW_WITH_COUNTERSIGNATURE Procedure 42-22
42.3.17 VERIFY_ROWS Procedure 42-24
42.3.18 VERIFY_TABLE_BLOCKCHAIN Procedure 42-25
42.3.19 VERIFY_USER_BLOCKCHAIN_ROWS Procedure 42-26

43 DBMS_BLOCKER_RESOLVER

43.1 DBMS_BLOCKER_RESOLVER Overview 43-1
43.2 DBMS_BLOCKER_RESOLVER Security Model 43-1
43.3 DBMS_BLOCKER_RESOLVER Constants 43-2
43.4 DBMS_BLOCKER_RESOLVER Exceptions 43-3
43.5 Summary of DBMS_BLOCKER_RESOLVER Subprograms 43-3
43.5.1 SET Procedure 43-3

44 DBMS_CAPTURE_ADM

44.1 DBMS_CAPTURE_ADM Overview 44-1
44.2 DBMS_CAPTURE_ADM Security Model 44-1
44.3 Summary of DBMS_CAPTURE_ADM Subprograms 44-2

44.3.1 ABORT_GLOBAL_INSTANTIATION Procedure 44-3
44.3.2 ABORT_SCHEMA_INSTANTIATION Procedure 44-4
44.3.3 ABORT_SYNC_INSTANTIATION Procedure 44-5
44.3.4 ABORT_TABLE_INSTANTIATION Procedure 44-5
44.3.5 ALTER_CAPTURE Procedure 44-6
44.3.6 ALTER_SYNC_CAPTURE Procedure 44-11
44.3.7 BUILD Procedure 44-12
44.3.8 CREATE_CAPTURE Procedure 44-13
44.3.9 CREATE_SYNC_CAPTURE Procedure 44-23
44.3.10 DROP_CAPTURE Procedure 44-25
44.3.11 INCLUDE_EXTRA_ATTRIBUTE Procedure 44-26
44.3.12 PREPARE_GLOBAL_INSTANTIATION Procedure 44-27
44.3.13 PREPARE_SCHEMA_INSTANTIATION Procedure 44-28
44.3.14 PREPARE_SYNC_INSTANTIATION Function 44-29
44.3.15 PREPARE_TABLE_INSTANTIATION Procedure 44-30
44.3.16 SET_PARAMETER Procedure 44-31
44.3.17 START_CAPTURE Procedure 44-46
44.3.18 STOP_CAPTURE Procedure 44-46

45 DBMS_CACHEUTIL

45.1 DBMS_CACHEUTIL Overview 45-1
45.2 DBMS_CACHEUTIL Security Model 45-2
45.3 Summary of DBMS_CACHEUTIL Subprograms 45-2
45.3.1 DISSOLVE_AFFINITY Procedure 45-3
45.3.2 DISSOLVE_READMOSTLY Procedure 45-3
45.3.3 GRAB_AFFINITY Procedure 45-4
45.3.4 GRAB_READMOSTLY Procedure 45-5
45.3.5 LIST_READMOSTLY Procedure 45-6
45.3.6 OBJECT_DOWNCONVERT Procedure 45-6
45.3.7 TRUE_CACHE_KEEP Procedure 45-7
45.3.8 TRUE_CACHE_UNKEEP Procedure 45-7

46 DBMS_CLOUD

46.1 DBMS_CLOUD Overview 46-1
46.2 Installing the DBMS_CLOUD Package 46-1
46.3 DBMS_CLOUD Security Model 46-2
46.4 DBMS_CLOUD Exceptions 46-2
46.5 Summary of DBMS_CLOUD Subprograms 46-2
46.5.1 COPY_DATA Procedure 46-4
46.5.2 COPY_DATA Procedure for Avro or Parquet Files 46-6

46.5.3 CREATE_CREDENTIAL Procedure 46-8
46.5.4 CREATE_EXTERNAL_TABLE Procedure 46-11
46.5.5 CREATE_EXTERNAL_TABLE Procedure for Avro or Parquet Files 46-13
46.5.6 CREATE_EXTERNAL_PART_TABLE Procedure 46-15
46.5.7 CREATE_HYBRID_PART_TABLE Procedure 46-20
46.5.8 DELETE_ALL_OPERATIONS Procedure 46-21
46.5.9 DELETE_FILE Procedure 46-22
46.5.10 DELETE_OBJECT Procedure 46-22
46.5.11 DROP_CREDENTIAL Procedure 46-23
46.5.12 GET_OBJECT Procedure and Function 46-23
46.5.13 LIST_FILES Function 46-25
46.5.14 LIST_OBJECTS Function 46-26
46.5.15 PUT_OBJECT Procedure 46-27
46.5.16 UPDATE_CREDENTIAL Procedure 46-29
46.5.17 VALIDATE_EXTERNAL_TABLE Procedure 46-30
46.5.18 VALIDATE_EXTERNAL_PART_TABLE Procedure 46-31
46.5.19 VALIDATE_HYBRID_PART_TABLE Procedure 46-32
46.6 DBMS_CLOUD Package File URI Formats 46-33
46.6.1 Oracle Cloud Infrastructure Object Storage Native URI Format 46-33
46.6.2 Oracle Cloud Infrastructure Object Storage Swift URI Format 46-35
46.6.3 Oracle Cloud Infrastructure Object Storage URI Format Using Pre-Authenticated Request URL 46-35
46.6.4 URI Format Using Public URL 46-36
46.6.5 Oracle Cloud Infrastructure Object Storage Classic URI Format 46-37
46.6.6 Amazon S3 URI Format 46-37
46.6.7 Azure Blob Storage URI Format 46-38
46.7 DBMS_CLOUD Package Format Options 46-38
46.8 DBMS_CLOUD Avro and Parquet Support 46-44
46.8.1 DBMS_CLOUD Package Format Options for Avro or Parquet 46-44
46.8.2 DBMS_CLOUD Package Avro to Oracle Data Type Mapping 46-45
46.8.3 DBMS_CLOUD Package Parquet to Oracle Data Type Mapping 46-46
46.8.4 DBMS_CLOUD Package Avro and Parquet Complex Types 46-47
46.8.5 DBMS_CLOUD Package Avro and Parquet to Oracle Column Name Mapping 46-48

47 DBMS_COMPARISON

47.1 DBMS_COMPARISON Overview 47-1
47.2 DBMS_COMPARISON Security Model 47-2
47.3 DBMS_COMPARISON Constants 47-3
47.4 DBMS_COMPARISON Views 47-5
47.5 DBMS_COMPARISON Operational Notes 47-5
47.6 DBMS_COMPARISON Data Structures 47-9

47.6.1 COMPARISON_TYPE Record Type 47-10
47.7 Summary of DBMS_COMPARISON Subprograms 47-10
47.7.1 COMPARE Function 47-11
47.7.2 CONVERGE Procedure 47-12
47.7.3 CREATE_COMPARISON Procedure 47-13
47.7.4 DROP_COMPARISON Procedure 47-17
47.7.5 PURGE_COMPARISON Procedure 47-18
47.7.6 RECHECK Function 47-18

48 DBMS_COMPRESSION

48.1 DBMS_COMPRESSION Overview 48-1
48.2 DBMS_COMPRESSION Security Model 48-1
48.3 DBMS_COMPRESSION Constants 48-1
48.4 DBMS_COMPRESSION Data Structures 48-3
48.4.1 COMPREC Record Type 48-3
48.4.2 COMPRECLIST Table Type 48-4
48.5 Summary of DBMS_COMPRESSION Subprograms 48-4
48.5.1 GET_COMPRESSION_RATIO Procedure 48-4
48.5.2 GET_COMPRESSION_TYPE Function 48-12

49 DBMS_CONNECTION_POOL

49.1 Summary of DBMS_CONNECTION_POOL Subprograms 49-1
49.1.1 ADD_POOL Procedure 49-2
49.1.2 ALTER_PARAM Procedure 49-4
49.1.3 CONFIGURE_POOL Procedure 49-5
49.1.4 REMOVE_POOL Procedure 49-7
49.1.5 RESTORE_DEFAULTS Procedure 49-8
49.1.6 START_POOL Procedure 49-8
49.1.7 STOP_POOL Procedure 49-9

50 DBMS_CQ_NOTIFICATION

50.1 DBMS_CQ_NOTIFICATION Overview 50-1
50.2 DBMS_CQ_NOTIFICATION Security Model 50-2
50.3 DBMS_CQ_NOTIFICATION Constants 50-2
50.4 DBMS_CQ_NOTIFICATION Operational Notes 50-3
50.5 DBMS_CQ_NOTIFICATION Examples 50-4
50.6 DBMS_CQ_NOTIFICATION Data Structures 50-10
50.6.1 CQ_NOTIFICATION$_DESCRIPTOR Object Type 50-11
50.6.2 CQ_NOTIFICATION$_QUERY Object Type 50-11

50.6.3 CQ_NOTIFICATION$_QUERY_ARRAY Object (Array) Type 50-12
50.6.4 CQ_NOTIFICATION$_TABLE Object Type 50-12
50.6.5 CQ_NOTIFICATION$_TABLE_ARRAY Object (Array) Type 50-13
50.6.6 CQ_NOTIFICATION$_ROW Object Type 50-13
50.6.7 CQ_NOTIFICATION$_ROW_ARRAY Object (Array) Type 50-13
50.6.8 CQ_NOTIFICATION$_REG_INFO Object Type 50-14
50.7 Summary of DBMS_CQ_NOTIFICATION Subprograms 50-16
50.7.1 CQ_NOTIFICATION_QUERYID Function 50-17
50.7.2 DEREGISTER Procedure 50-17
50.7.3 ENABLE_REG Procedure 50-18
50.7.4 NEW_REG_START Function 50-18
50.7.5 REG_END Procedure 50-19
50.7.6 SET_ROWID_THRESHOLD Procedure 50-19

51 DBMS_CREDENTIAL

51.1 DBMS_CREDENTIAL Overview 51-1
51.2 DBMS_CREDENTIAL Security Model 51-1
51.3 DBMS_CREDENTIAL Operational Notes 51-2
51.4 Summary of DBMS_CREDENTIAL Subprograms 51-2
51.4.1 CREATE_CREDENTIAL Procedure 51-3
51.4.2 DISABLE_CREDENTIAL Procedure 51-4
51.4.3 DROP_CREDENTIAL Procedure 51-5
51.4.4 ENABLE_CREDENTIAL Procedure 51-6
51.4.5 UPDATE_CREDENTIAL Procedure 51-6

52 DBMS_CRYPTO

52.1 DBMS_CRYPTO Overview 52-2
52.2 DBMS_CRYPTO Security Model 52-3
52.3 DBMS_CRYPTO Constants 52-3
52.4 DBMS_CRYPTO Datatypes 52-5
52.5 DBMS_CRYPTO Algorithms 52-5
52.6 DBMS_CRYPTO Restrictions 52-8
52.7 DBMS_CRYPTO Exceptions 52-9
52.8 DBMS_CRYPTO Operational Notes 52-9
52.9 DBMS_CRYPTO Examples 52-11
52.10 Summary of DBMS_CRYPTO Subprograms 52-15
52.10.1 DECRYPT Functions 52-16
52.10.2 DECRYPT Procedures 52-17
52.10.3 ECDHDERIVE_SHAREDSECRET Function 52-18
52.10.4 ECDH_GENKEYPAIR Function 52-19

52.10.5 ENCRYPT Functions 52-19
52.10.6 ENCRYPT Procedures 52-21
52.10.7 HASH Function 52-22
52.10.8 HASH_LEN Function 52-23
52.10.9 MAC Function 52-24
52.10.10 KMACXOF Function 52-25
52.10.11 PKDECRYPT Function 52-25
52.10.12 PKENCRYPT Function 52-26
52.10.13 RANDOMBYTES Function 52-26
52.10.14 RANDOMINTEGER Function 52-27
52.10.15 RANDOMNUMBER Function 52-27
52.10.16 SIGN Function 52-27
52.10.17 VERIFY Function 52-29

53 DBMS_CSX_ADMIN

53.1 DBMS_CSX_ADMIN Overview 53-1
53.2 DBMS_CSX_ADMIN Security Model 53-1
53.3 DBMS_CSX_ADMIN Constants 53-1
53.4 Summary of DBMS_CSX_ADMIN Subprograms 53-2
53.4.1 GETTOKENTABLEINFO Procedure & Function 53-2
53.4.2 GETTOKENTABLEINFOBYTABLESPACE Procedure 53-3
53.4.3 NAMESPACEIDTABLE Function 53-4
53.4.4 PATHIDTABLE Function 53-4
53.4.5 QNAMEIDTABLE Function 53-4

54 DBMS_CUBE

54.1 Using DBMS_CUBE 54-1
54.1.1 DBMS_CUBE Security Model 54-2
54.2 Using SQL Aggregation Management 54-2
54.2.1 Subprograms in SQL Aggregation Management 54-3
54.2.2 Requirements for the Relational Materialized View 54-3
54.2.3 Permissions for Managing and Querying Cube Materialized Views 54-4
54.2.4 Example of SQL Aggregation Management 54-4
54.3 Upgrading 10g Analytic Workspaces 54-8
54.4 Summary of DBMS_CUBE Subprograms 54-13
54.4.1 BUILD Procedure 54-13
54.4.2 CREATE_EXPORT_OPTIONS Procedure 54-25
54.4.3 CREATE_IMPORT_OPTIONS Procedure 54-28
54.4.4 CREATE_MVIEW Function 54-29
54.4.5 DERIVE_FROM_MVIEW Function 54-34

54.4.6 DROP_MVIEW Procedure 54-35
54.4.7 EXPORT_XML Procedure 54-37
54.4.8 EXPORT_XML_TO_FILE Procedure 54-39
54.4.9 IMPORT_XML Procedure 54-40
54.4.10 INITIALIZE_CUBE_UPGRADE Procedure 54-43
54.4.11 REFRESH_MVIEW Procedure 54-44
54.4.12 UPGRADE_AW Procedure 54-46
54.4.13 VALIDATE_XML Procedure 54-47

55 DBMS_CUBE_ADVISE

55.1 DBMS_CUBE_ADVISE Security Model 55-1
55.2 Summary of DBMS_CUBE_ADVISE Subprograms 55-1
55.2.1 MV_CUBE_ADVICE Function 55-1
55.2.2 SET_CNS_EXCEPTION_LOG Procedure 55-4
55.2.3 TRACE Procedure 55-5

56 DBMS_CUBE_LOG

56.1 Using DBMS_CUBE_LOG 56-1
56.1.1 DBMS_CUBE_LOG — Logging Types 56-2
56.1.2 DBMS_CUBE_LOG — Logging Targets 56-2
56.1.3 DBMS_CUBE_LOG — Verbosity Levels 56-3
56.1.4 DBMS_CUBE_LOG Security Model 56-3
56.1.5 DBMS_CUBE_LOG — Creating Cube Logs 56-3
56.1.6 DBMS_CUBE_LOG — Cube Build Log 56-4
56.1.7 DBMS_CUBE_LOG — Cube Dimension Compile Log 56-5
56.1.8 DBMS_CUBE_LOG — Cube Operations Log 56-6
56.1.9 DBMS_CUBE_LOG — Cube Rejected Records Log 56-7
56.2 Summary of DBMS_CUBE_LOG Subprograms 56-7
56.2.1 DEFAULT_NAME Function 56-8
56.2.2 DISABLE Procedure 56-9
56.2.3 ENABLE Procedure 56-10
56.2.4 FLUSH Procedure 56-11
56.2.5 GET_LOG Procedure 56-12
56.2.6 GET_LOG_SPEC Function 56-13
56.2.7 GET_PARAMETER Function 56-14
56.2.8 LEVEL_HIGH Function 56-14
56.2.9 LEVEL_HIGHEST Function 56-15
56.2.10 LEVEL_LOW Function 56-15
56.2.11 LEVEL_LOWEST Function 56-16
56.2.12 LEVEL_MEDIUM Function 56-16

56.2.13 SET_LOG_SPEC Procedure 56-17
56.2.14 SET_PARAMETER Procedure 56-17
56.2.15 TABLE_CREATE Procedure 56-19
56.2.16 TARGET_FILE Function 56-20
56.2.17 TARGET_LOB Function 56-20
56.2.18 TARGET_TABLE Function 56-20
56.2.19 TARGET_TRACE Function 56-21
56.2.20 TYPE_BUILD Function 56-21
56.2.21 TYPE_DIMENSION_COMPILE Function 56-22
56.2.22 TYPE_OPERATIONS Function 56-22
56.2.23 TYPE_REJECTED_RECORDS Function 56-23
56.2.24 VERSION Function 56-23

57 DBMS_DATA_MINING

57.1 DBMS_DATA_MINING Overview 57-1
57.2 DBMS_DATA_MINING Security Model 57-3
57.3 DBMS_DATA_MINING — Machine Learning Functions 57-3
57.4 DBMS_DATA_MINING — Model Settings 57-5
57.4.1 DBMS_DATA_MINING — Algorithm Names 57-5
57.4.2 DBMS_DATA_MINING — Automatic Data Preparation 57-6
57.4.3 DBMS_DATA_MINING — Machine Learning Function Settings 57-7
57.4.4 DBMS_DATA_MINING — Global Settings 57-13
57.5 DBMS_DATA_MINING — Algorithm Specific Model Settings 57-18
57.5.1 DBMS_DATA_MINING — Algorithm Settings: ALGO_EXTENSIBLE_LANG 57-18
57.5.2 DBMS_DATA_MINING — Algorithm Settings: CUR Matrix Decomposition 57-20
57.5.3 DBMS_DATA_MINING — Algorithm Settings: Decision Tree 57-20
57.5.4 DBMS_DATA_MINING — Algorithm Settings: Expectation Maximization 57-22
57.5.5 DBMS_DATA_MINING — Algorithm Settings: Explicit Semantic Analysis 57-25
57.5.6 DBMS_DATA_MINING — Algorithm Settings: Exponential Smoothing 57-27
57.5.7 DBMS_DATA_MINING — Algorithm Settings: Generalized Linear Model 57-32
57.5.8 DBMS_DATA_MINING — Algorithm Settings: k-Means 57-35
57.5.9 DBMS_DATA_MINING - Algorithm Settings: Multivariate State Estimation Technique - Sequential Probability Ratio Test 57-37
57.5.10 DBMS_DATA_MINING — Algorithm Settings: Naive Bayes 57-38
57.5.11 DBMS_DATA_MINING — Algorithm Settings: Neural Network 57-39
57.5.12 DBMS_DATA_MINING — Algorithm Settings: Non-Negative Matrix Factorization 57-42
57.5.13 DBMS_DATA_MINING — Algorithm Settings: O-Cluster 57-42
57.5.14 DBMS_DATA_MINING — Algorithm Settings: Random Forest 57-43
57.5.15 DBMS_DATA_MINING — Algorithm Constants and Settings: Singular Value Decomposition 57-44
57.5.16 DBMS_DATA_MINING — Algorithm Settings: Support Vector Machine 57-45
57.5.17 DBMS_DATA_MINING — Algorithm Settings: XGBoost 57-47

57.6 DBMS_DATA_MINING — Solver Settings 57-57
57.6.1 DBMS_DATA_MINING - Solver Settings: Adam 57-57
57.6.2 DBMS_DATA_MINING — Solver Settings: ADMM 57-58
57.6.3 DBMS_DATA_MINING — Solver Settings: LBFGS 57-59
57.7 DBMS_DATA_MINING Datatypes 57-59
57.7.1 Deprecated Types 57-60
57.8 Summary of DBMS_DATA_MINING Subprograms 57-65
57.8.1 ADD_COST_MATRIX Procedure 57-67
57.8.2 ADD_PARTITION Procedure 57-70
57.8.3 ALTER_REVERSE_EXPRESSION Procedure 57-71
57.8.4 APPLY Procedure 57-74
57.8.5 COMPUTE_CONFUSION_MATRIX Procedure 57-77
57.8.6 COMPUTE_CONFUSION_MATRIX_PART Procedure 57-83
57.8.7 COMPUTE_LIFT Procedure 57-89
57.8.8 COMPUTE_LIFT_PART Procedure 57-94
57.8.9 COMPUTE_ROC Procedure 57-99
57.8.10 COMPUTE_ROC_PART Procedure 57-103
57.8.11 CREATE_MODEL Procedure 57-108
57.8.12 CREATE_MODEL2 Procedure 57-112
57.8.13 Create Model Using Registration Information 57-113
57.8.14 DROP_ALGORITHM Procedure 57-114
57.8.15 DROP_PARTITION Procedure 57-115
57.8.16 DROP_MODEL Procedure 57-115
57.8.17 EXPORT_MODEL Procedure 57-116
57.8.18 EXPORT_SERMODEL Procedure 57-119
57.8.19 FETCH_JSON_SCHEMA Procedure 57-120
57.8.20 GET_ASSOCIATION_RULES Function 57-120
57.8.21 GET_FREQUENT_ITEMSETS Function 57-125
57.8.22 GET_MODEL_COST_MATRIX Function 57-127
57.8.23 GET_MODEL_DETAILS_AI Function 57-128
57.8.24 GET_MODEL_DETAILS_EM Function 57-130
57.8.25 GET_MODEL_DETAILS_EM_COMP Function 57-131
57.8.26 GET_MODEL_DETAILS_EM_PROJ Function 57-133
57.8.27 GET_MODEL_DETAILS_GLM Function 57-135
57.8.28 GET_MODEL_DETAILS_GLOBAL Function 57-138
57.8.29 GET_MODEL_DETAILS_KM Function 57-140
57.8.30 GET_MODEL_DETAILS_NB Function 57-142
57.8.31 GET_MODEL_DETAILS_NMF Function 57-144
57.8.32 GET_MODEL_DETAILS_OC Function 57-145
57.8.33 GET_MODEL_SETTINGS Function 57-147
57.8.34 GET_MODEL_SIGNATURE Function 57-148
57.8.35 GET_MODEL_DETAILS_SVD Function 57-150

57.8.36 GET_MODEL_DETAILS_SVM Function 57-152
57.8.37 GET_MODEL_DETAILS_XML Function 57-154
57.8.38 GET_MODEL_TRANSFORMATIONS Function 57-156
57.8.39 GET_TRANSFORM_LIST Procedure 57-159
57.8.40 IMPORT_MODEL Procedure 57-162
57.8.41 IMPORT_ONNX_MODEL Procedure 57-166
57.8.42 IMPORT_SERMODEL Procedure 57-167
57.8.43 JSON Schema for R Extensible Algorithm 57-168
57.8.44 REGISTER_ALGORITHM Procedure 57-172
57.8.45 RANK_APPLY Procedure 57-173
57.8.46 REMOVE_COST_MATRIX Procedure 57-176
57.8.47 RENAME_MODEL Procedure 57-177

58 DBMS_DATA_MINING_TRANSFORM

58.1 DBMS_DATA_MINING_TRANSFORM Overview 58-1
58.2 DBMS_DATA_MINING_TRANSFORM Operational Notes 58-4
58.2.1 DBMS_DATA_MINING_TRANSFORM — About Transformation Lists 58-6
58.2.2 DBMS_DATA_MINING_TRANSFORM — About Stacking and Stack Procedures 58-8
58.2.3 DBMS_DATA_MINING_TRANSFORM — Nested Data Transformations 58-9
58.3 DBMS_DATA_MINING_TRANSFORM Security Model 58-13
58.4 DBMS_DATA_MINING_TRANSFORM Datatypes 58-13
58.5 DBMS_DATA_MINING_TRANSFORM Constants 58-14
58.6 Summary of DBMS_DATA_MINING_TRANSFORM Subprograms 58-15
58.6.1 CREATE_BIN_CAT Procedure 58-17
58.6.2 CREATE_BIN_NUM Procedure 58-18
58.6.3 CREATE_CLIP Procedure 58-20
58.6.4 CREATE_COL_REM Procedure 58-21
58.6.5 CREATE_MISS_CAT Procedure 58-23
58.6.6 CREATE_MISS_NUM Procedure 58-24
58.6.7 CREATE_NORM_LIN Procedure 58-25
58.6.8 DESCRIBE_STACK Procedure 58-27
58.6.9 GET_EXPRESSION Function 58-28
58.6.10 INSERT_AUTOBIN_NUM_EQWIDTH Procedure 58-29
58.6.11 INSERT_BIN_CAT_FREQ Procedure 58-33
58.6.12 INSERT_BIN_NUM_EQWIDTH Procedure 58-37
58.6.13 INSERT_BIN_NUM_QTILE Procedure 58-41
58.6.14 INSERT_BIN_SUPER Procedure 58-43
58.6.15 INSERT_CLIP_TRIM_TAIL Procedure 58-47
58.6.16 INSERT_CLIP_WINSOR_TAIL Procedure 58-49
58.6.17 INSERT_MISS_CAT_MODE Procedure 58-52
58.6.18 INSERT_MISS_NUM_MEAN Procedure 58-54

58.6.19 INSERT_NORM_LIN_MINMAX Procedure 58-57
58.6.20 INSERT_NORM_LIN_SCALE Procedure 58-59
58.6.21 INSERT_NORM_LIN_ZSCORE Procedure 58-61
58.6.22 SET_EXPRESSION Procedure 58-63
58.6.23 SET_TRANSFORM Procedure 58-65
58.6.24 STACK_BIN_CAT Procedure 58-66
58.6.25 STACK_BIN_NUM Procedure 58-68
58.6.26 STACK_CLIP Procedure 58-70
58.6.27 STACK_COL_REM Procedure 58-72
58.6.28 STACK_MISS_CAT Procedure 58-74
58.6.29 STACK_MISS_NUM Procedure 58-76
58.6.30 STACK_NORM_LIN Procedure 58-78
58.6.31 XFORM_BIN_CAT Procedure 58-80
58.6.32 XFORM_BIN_NUM Procedure 58-82
58.6.33 XFORM_CLIP Procedure 58-84
58.6.34 XFORM_COL_REM Procedure 58-86
58.6.35 XFORM_EXPR_NUM Procedure 58-87
58.6.36 XFORM_EXPR_STR Procedure 58-89
58.6.37 XFORM_MISS_CAT Procedure 58-92
58.6.38 XFORM_MISS_NUM Procedure 58-94
58.6.39 XFORM_NORM_LIN Procedure 58-95
58.6.40 XFORM_STACK Procedure 58-97

59 DBMS_DATAPUMP

59.1 DBMS_DATAPUMP Overview 59-1
59.2 DBMS_DATAPUMP Security Model 59-1
59.3 DBMS_DATAPUMP Constants 59-2
59.4 DBMS_DATAPUMP Data Structures 59-3
59.5 Summary of DBMS_DATAPUMP Subprograms 59-12
59.5.1 ADD_FILE Procedure 59-13
59.5.2 ATTACH Function 59-21
59.5.3 DATA_FILTER Procedures 59-22
59.5.4 DATA_REMAP Procedure 59-25
59.5.5 DETACH Procedure 59-26
59.5.6 GET_DUMPFILE_INFO Procedure 59-27
59.5.7 GET_STATUS Procedure 59-30
59.5.8 LOG_ENTRY Procedure 59-32
59.5.9 METADATA_FILTER Procedure 59-33
59.5.10 METADATA_REMAP Procedure 59-37
59.5.11 METADATA_TRANSFORM Procedure 59-39
59.5.12 OPEN Function 59-44

59.5.13 SET_PARALLEL Procedure 59-47
59.5.14 SET_PARAMETER Procedures 59-48
59.5.15 START_JOB Procedure 59-57
59.5.16 STOP_JOB Procedure 59-58
59.5.17 WAIT_FOR_JOB Procedure 59-60

60 DBMS_DB_VERSION

60.1 DBMS_DB_VERSION Overview 60-1
60.2 DBMS_DB_VERSION Constants 60-2
60.3 DBMS_DB_VERSION Examples 60-3

61 DBMS_DBCOMP

61.1 Using DBMS_DBCOMP 61-1
61.2 DBMS_DBCOMP Security Model 61-2
61.3 Summary of the DBMS_DBCOMP Subprogram 61-2
61.3.1 DBCOMP Procedure 61-2

62 DBMS_DBFS_CONTENT

62.1 DBMS_DBFS_CONTENT Overview 62-1
62.2 DBMS_DBFS_CONTENT Security Model 62-1
62.3 DBMS_DBFS_CONTENT Constants 62-2
62.4 Exceptions 62-13
62.5 DBMS_DBFS_CONTENT Operational Notes 62-13
62.6 DBMS_DBFS_CONTENT Data Structures 62-14
62.6.1 DBMS_DBFS_CONTENT FEATURE_T Record Type 62-15
62.6.2 DBMS_DBFS_CONTENT MOUNT_T Record Type 62-16
62.6.3 DBMS_DBFS_CONTENT PATH_ITEM_T Record Type 62-16
62.6.4 DBMS_DBFS_CONTENT PROP_ITEM_T Record Type 62-18
62.6.5 PROPERTY_T Record Type 62-19
62.6.6 DBMS_DBFS_CONTENT STORE_T Record Type 62-19
62.6.7 DBMS_DBFS_CONTENT FEATURES_T Table Type 62-20
62.6.8 MOUNTS_T Table Type 62-20
62.6.9 DBMS_DBFS_CONTENT PATH_ITEMS_T Table Type 62-20
62.6.10 DBMS_DBFS_CONTENT PROP_ITEMS_T Table Type 62-21
62.6.11 DBMS_DBFS_CONTENT PROPERTIES_T Table Type 62-21
62.6.12 STORES_T Table Type 62-21
62.7 Summary of DBMS_DBFS_CONTENT Subprograms 62-21
62.7.1 CHECKACCESS Function 62-24
62.7.2 CHECKSPI Functions and Procedures 62-24

62.7.3 CREATEDIRECTORY Procedures 62-25
62.7.4 CREATEFILE Procedures 62-26
62.7.5 CREATELINK Procedures 62-27
62.7.6 CREATEREFERENCE Procedures 62-28
62.7.7 DECODEFEATURES Function 62-28
62.7.8 DELETECONTENT Procedure 62-29
62.7.9 DELETEDIRECTORY Procedure 62-29
62.7.10 DELETEFILE Procedure 62-30
62.7.11 FEATURENAME Function 62-30
62.7.12 FLUSHSTATS Function 62-31
62.7.13 GETDEFAULTACL Procedure 62-31
62.7.14 GETDEFAULTASOF Procedure 62-31
62.7.15 GETDEFAULTCONTEXT Procedure 62-32
62.7.16 GETDEFAULTOWNER Procedure 62-32
62.7.17 GETDEFAULTPRINCIPAL Procedure 62-33
62.7.18 GETFEATURESBYMOUNT Function 62-33
62.7.19 GETFEATURESBYNAME Function 62-34
62.7.20 GETFEATURESBYPATH Function 62-34
62.7.21 GETPATH Procedures 62-34
62.7.22 GETPATHBYMOUNTID Function 62-37
62.7.23 GETPATHBYSTOREID Function 62-37
62.7.24 GETPATHNOWAIT Procedures 62-38
62.7.25 GETSTOREBYMOUNT Function 62-39
62.7.26 GETSTOREBYNAME Function 62-39
62.7.27 GETSTOREBYPATH Function 62-40
62.7.28 GETSTATS Procedure 62-40
62.7.29 GETTRACE Function 62-40
62.7.30 GETVERSION Function 62-41
62.7.31 ISPATHLOCKED Procedure 62-41
62.7.32 LIST Function 62-41
62.7.33 LISTALLPROPERTIES Function 62-42
62.7.34 LISTALLCONTENT Function 62-42
62.7.35 LISTMOUNTS Function 62-42
62.7.36 LISTSTORES Function 62-43
62.7.37 LOCKPATH Procedure 62-43
62.7.38 MOUNTSTORE Procedure 62-44
62.7.39 NORMALIZEPATH Functions 62-44
62.7.40 PROPANY Functions 62-45
62.7.41 PROPERTIESH2T Function 62-46
62.7.42 PROPERTIEST2H Function 62-46
62.7.43 PROPNUMBER Function 62-47
62.7.44 PROPRAW Function 62-47

62.7.45 PROPTIMESTAMP Function 62-47
62.7.46 PROPVARCHAR2 Function 62-48
62.7.47 PURGEALL Procedure 62-48
62.7.48 PURGEPATH Procedure 62-49
62.7.49 PUTPATH Procedures 62-49
62.7.50 REGISTERSTORE Procedure 62-51
62.7.51 RENAMEPATH Procedures 62-51
62.7.52 RESTOREALL Procedure 62-52
62.7.53 RESTOREPATH Procedure 62-53
62.7.54 SETDEFAULTACL Procedure 62-53
62.7.55 SETDEFAULTASOF Procedure 62-54
62.7.56 SETDEFAULTCONTEXT Procedure 62-54
62.7.57 SETDEFAULTOWNER Procedure 62-55
62.7.58 SETDEFAULTPRINCIPAL Procedure 62-55
62.7.59 SETPATH Procedures 62-56
62.7.60 SETSTATS Procedure 62-56
62.7.61 SETTRACE Procedure 62-57
62.7.62 SPACEUSAGE Procedure 62-57
62.7.63 TRACE Procedure 62-58
62.7.64 TRACEENABLED Function 62-59
62.7.65 UNLOCKPATH Procedure 62-60
62.7.66 UNMOUNTSTORE Procedure 62-60
62.7.67 UNREGISTERSTORE Procedure 62-61

63 DBMS_DBFS_CONTENT_SPI

63.1 ODBMS_DBFS_CONTENT_SPI Overview 63-1
63.2 DBMS_DBFS_CONTENT_SPI Security Model 63-2
63.3 DBMS_DBFS_CONTENT_SPI Operational Notes 63-2
63.4 Summary of DBMS_DBFS_CONTENT_SPI Subprograms 63-3
63.4.1 CHECKACCESS Function 63-4
63.4.2 CREATEDIRECTORY Procedure 63-5
63.4.3 CREATEFILE Procedure 63-6
63.4.4 CREATELINK Procedure 63-6
63.4.5 CREATEREFERENCE Procedure 63-7
63.4.6 DELETECONTENT Procedure 63-8
63.4.7 DELETEDIRECTORY Procedure 63-8
63.4.8 DELETEFILE Procedure 63-9
63.4.9 GETFEATURES Function 63-10
63.4.10 GETPATH Procedures 63-10
63.4.11 GETPATHBYSTOREID Function 63-11
63.4.12 GETPATHNOWAIT Procedure 63-12

63.4.13 GETSTOREID Function 63-13
63.4.14 GETVERSION Function 63-13
63.4.15 ISPATHLOCKED Procedure 63-14
63.4.16 LIST Function 63-14
63.4.17 LOCKPATH Procedure 63-15
63.4.18 PURGEALL Procedure 63-16
63.4.19 PURGEPATH Procedure 63-16
63.4.20 PUTPATH Procedures 63-17
63.4.21 RENAMEPATH Procedure 63-18
63.4.22 RESTOREALL Procedure 63-19
63.4.23 RESTOREPATH Procedure 63-19
63.4.24 SEARCH Function 63-20
63.4.25 SETPATH Procedure 63-20
63.4.26 SPACEUSAGE Procedure 63-21
63.4.27 UNLOCKPATH Procedure 63-22

64 DBMS_DBFS_HS

64.1 DBMS_DBFS_HS Overview 64-1
64.2 DBMS_DBFS_HS Security Model 64-1
64.3 DBMS_DBFS_HS Constants 64-1
64.4 DBMS_DBFS_HS Operational Notes 64-6
64.5 Summary of DBMS_DBFS_HS Subprograms 64-6
64.5.1 CLEANUPUNUSEDBACKUPFILES Procedure 64-7
64.5.2 CREATEBUCKET Procedure 64-7
64.5.3 CREATESTORE Procedure 64-8
64.5.4 DEREGSTORECOMMAND Function 64-9
64.5.5 DROPSTORE Procedure 64-9
64.5.6 FLUSHCACHE Procedure 64-10
64.5.7 GETSTOREPROPERTY Function 64-10
64.5.8 RECONFIGCACHE Procedure 64-11
64.5.9 REGISTERSTORECOMMAND Procedure 64-12
64.5.10 SENDCOMMAND Procedure 64-13
64.5.11 SETSTOREPROPERTY Procedure 64-13
64.5.12 STOREPUSH Procedure 64-17

65 DBMS_DBFS_SFS

65.1 DBMS_DBFS_SFS Overview 65-1
65.2 DBMS_DBFS_SFS Security Model 65-1
65.3 DBMS_DBFS_SFS Constants 65-1
65.4 Summary of DBMS_DBFS_SFS Subprograms 65-2

65.4.1 CREATEFILESYSTEM Procedure 65-3
65.4.2 CREATESTORE Procedure 65-4
65.4.3 DROPFILESYSTEM Procedures 65-5
65.4.4 INITFS Procedure 65-6

66 DBMS_DDL

66.1 DBMS_DDL Deprecated Subprograms 66-1
66.2 DBMS_DDL Security Model 66-1
66.3 DBMS_DDL Operational Notes 66-1
66.4 Summary of DBMS_DDL Subprograms 66-1
66.4.1 ALTER_COMPILE Procedure 66-2
66.4.2 ALTER_TABLE_NOT_REFERENCEABLE Procedure 66-3
66.4.3 ALTER_TABLE_REFERENCEABLE Procedure 66-4
66.4.4 CREATE_WRAPPED Procedures 66-4
66.4.5 IS_TRIGGER_FIRE_ONCE Function 66-6
66.4.6 SET_TRIGGER_FIRING_PROPERTY Procedures 66-7
66.4.7 WRAP Functions 66-8

67 DBMS_DEBUG

67.1 DBMS_DEBUG Overview 67-1
67.2 DBMS_DEBUG Constants 67-2
67.3 DBMS_DEBUG Variables 67-2
67.4 DBMS_DEBUG Exceptions 67-2
67.5 DBMS_DEBUG Operational Notes 67-4
67.6 DBMS_DEBUG Data Structures 67-11
67.6.1 BREAKPOINT_INFO Record Type 67-11
67.6.2 PROGRAM_INFO Record Type 67-12
67.6.3 RUNTIME_INFO Record Type 67-13
67.6.4 BACKTRACE_TABLE Table Type 67-13
67.6.5 BREAKPOINT_TABLE Table Type 67-13
67.6.6 INDEX_TABLE Table Type 67-14
67.6.7 VC2_TABLE Table Type 67-14
67.7 Summary of DBMS_DEBUG Subprograms 67-14
67.7.1 ATTACH_SESSION Procedure 67-15
67.7.2 CONTINUE Function 67-16
67.7.3 DEBUG_OFF Procedure 67-16
67.7.4 DEBUG_ON Procedure 67-17
67.7.5 DELETE_BREAKPOINT Function 67-17
67.7.6 DELETE_OER_BREAKPOINT Function 67-18
67.7.7 DETACH_SESSION Procedure 67-18

67.7.8 DISABLE_BREAKPOINT Function 67-18
67.7.9 ENABLE_BREAKPOINT Function 67-19
67.7.10 EXECUTE Procedure 67-20
67.7.11 GET_INDEXES Function 67-21
67.7.12 GET_MORE_SOURCE Procedure 67-22
67.7.13 GET_LINE_MAP Function 67-23
67.7.14 GET_RUNTIME_INFO Function 67-24
67.7.15 GET_TIMEOUT_BEHAVIOUR Function 67-24
67.7.16 GET_VALUE Function 67-25
67.7.17 INITIALIZE Function 67-27
67.7.18 PING Procedure 67-28
67.7.19 PRINT_BACKTRACE Procedure 67-28
67.7.20 PRINT_INSTANTIATIONS Procedure 67-29
67.7.21 PROBE_VERSION Procedure 67-30
67.7.22 SELF_CHECK Procedure 67-30
67.7.23 SET_BREAKPOINT Function 67-31
67.7.24 SET_OER_BREAKPOINT Function 67-32
67.7.25 SET_TIMEOUT Function 67-32
67.7.26 SET_TIMEOUT_BEHAVIOUR Procedure 67-33
67.7.27 SET_VALUE Function 67-33
67.7.28 SHOW_BREAKPOINTS Procedures 67-35
67.7.29 SHOW_FRAME_SOURCE Procedure 67-35
67.7.30 SHOW_SOURCE Procedures 67-36
67.7.31 SYNCHRONIZE Function 67-37
67.7.32 TARGET_PROGRAM_RUNNING Procedure 67-38

68 DBMS_DEBUG_JDWP

68.1 DBMS_DEBUG_JDWP Overview 68-1
68.2 DBMS_DEBUG_JDWP Security Model 68-1
68.3 Summary DBMS_DEBUG_JDWP Subprograms 68-2
68.3.1 CONNECT_TCP Procedure 68-2
68.3.2 CURRENT_SESSION_ID Function 68-4
68.3.3 CURRENT_SESSION_SERIAL Function 68-4
68.3.4 DISCONNECT Procedure 68-4
68.3.5 GET_NLS_PARAMETER Function 68-5
68.3.6 PROCESS_CONNECT_STRING Procedure 68-5
68.3.7 SET_NLS_PARAMETER Procedure 68-6

69 DBMS_DEBUG_JDWP_CUSTOM

69.1 DBMS_DEBUG_JDWP_CUSTOM Overview 69-1
69.2 DBMS_DEBUG_JDWP_CUSTOM Security Model 69-1
69.3 Summary DBMS_DEBUG_JDWP_CUSTOM Subprograms 69-1
69.3.1 CONNECT_DEBUGGER Procedure 69-2

70 DBMS_DESCRIBE

70.1 DBMS_DESCRIBE Overview 70-1
70.2 DBMS_DESCRIBE Security Model 70-1
70.3 DBMS_DESCRIBE Types 70-1
70.4 DBMS_DESCRIBE Exceptions 70-1
70.5 DBMS_DESCRIBE Examples 70-2
70.6 Summary of DBMS_DESCRIBE Subprograms 70-5
70.6.1 DESCRIBE_PROCEDURE Procedure 70-5

71 DBMS_DG

71.1 Using DBMS_DG 71-1
71.2 DBMS_DG Security Model 71-2
71.3 Summary of the DBMS_DG Subprogram 71-2
71.3.1 INITIATE_FS_FAILOVER Procedure 71-2

72 DBMS_DICTIONARY_CHECK

72.1 Overview of Oracle Database Dictionary Check 72-1
72.2 Using DBMS_DICTIONARY_CHECK 72-1
72.3 Summary of DBMS_DICTIONARY_CHECK Subprograms 72-8

73 DBMS_DIMENSION

73.1 DBMS_DIMENSION Security Model 73-1
73.2 Summary of DBMS_DIMENSION Subprograms 73-1
73.2.1 DESCRIBE_DIMENSION Procedure 73-1
73.2.2 VALIDATE_DIMENSION Procedure 73-2

74 DBMS_DISTRIBUTED_TRUST_ADMIN

74.1 DBMS_DISTRIBUTED_TRUST_ADMIN Overview 74-1
74.2 DBMS_DISTRIBUTED_TRUST_ADMIN Security Model 74-1
74.3 DBMS_DISTRIBUTED_TRUST_ADMIN Examples 74-2
74.4 Summary of DBMS_DISTRIBUTED_TRUST_ADMIN Subprograms 74-2

74.4.1 ALLOW_ALL Procedure 74-3
74.4.2 ALLOW_SERVER Procedure 74-3
74.4.3 DENY_ALL Procedure 74-4
74.4.4 DENY_SERVER Procedure 74-4

75 DBMS_DNFS

75.1 DBMS_DNFS Security Model 75-1
75.2 Summary of DBMS_DNFS Subprograms 75-1
75.2.1 CLONEDB_RENAMEFILE Procedure 75-1
75.2.2 RESTORE_DATAFILE_PERMISSIONS Procedure 75-2
75.2.3 UNMOUNTVOLUME Procedure 75-2

76 DBMS_DST

76.1 DBMS_DST Overview 76-1
76.2 DBMS_DST Security Model 76-1
76.3 DBMS_DST Views 76-2
76.4 Summary of DBMS_DST Subprograms 76-2
76.4.1 BEGIN_PREPARE Procedure 76-3
76.4.2 BEGIN_UPGRADE Procedure 76-3
76.4.3 CREATE_AFFECTED_TABLE Procedure 76-4
76.4.4 CREATE_ERROR_TABLE Procedure 76-4
76.4.5 CREATE_TRIGGER_TABLE Procedure 76-5
76.4.6 END_PREPARE Procedure 76-6
76.4.7 END_UPGRADE Procedure 76-6
76.4.8 FIND_AFFECTED_TABLES Procedure 76-6
76.4.9 UPGRADE_DATABASE Procedure 76-7
76.4.10 UPGRADE_SCHEMA Procedure 76-8
76.4.11 UPGRADE_TABLE Procedure 76-10

77 DBMS_EDITIONS_UTILITIES

77.1 DBMS_EDITIONS_UTILITIES Overview 77-1
77.2 DBMS_EDITIONS_UTILITIES Security Model 77-1
77.3 DBMS_EDITIONS_UTILITIES Exceptions 77-1
77.4 Summary of DBMS_EDITIONS_UTILITIES Subprograms 77-1
77.4.1 CLEAN_UNUSABLE_EDITIONS Procedure 77-2
77.4.2 SET_EDITIONING_VIEWS_READ_ONLY Procedure 77-2
77.4.3 SET_NULL_COLUMN_VALUES_TO_EXPR Procedure 77-3

78 DBMS_EPG

78.1 DBMS_EPG Overview 78-1
78.2 DBMS_EPG Security Model 78-1
78.3 DBMS_EPG Exceptions 78-1
78.4 DBMS_EPG Data Structures 78-2
78.5 DBMS_EPG Subprogram Groups 78-2
78.5.1 DBMS_EPG Configuration Subprograms 78-2
78.5.2 DBMS_EPG Authorization Subprograms 78-3
78.6 Summary of DBMS_EPG Subprograms 78-3
78.6.1 AUTHORIZE_DAD Procedure 78-4
78.6.2 CREATE_DAD Procedure 78-5
78.6.3 DEAUTHORIZE_DAD Procedure 78-5
78.6.4 DELETE_DAD_ATTRIBUTE Procedure 78-6
78.6.5 DELETE_GLOBAL_ATTRIBUTE Procedure 78-6
78.6.6 DROP_DAD Procedure 78-7
78.6.7 GET_ALL_DAD_ATTRIBUTES Procedure 78-7
78.6.8 GET_ALL_DAD_MAPPINGS Procedure 78-8
78.6.9 GET_ALL_GLOBAL_ATTRIBUTES Procedure 78-9
78.6.10 GET_DAD_ATTRIBUTE Function 78-9
78.6.11 GET_DAD_LIST Procedure 78-10
78.6.12 GET_GLOBAL_ATTRIBUTE Function 78-10
78.6.13 MAP_DAD Procedure 78-11
78.6.14 SET_DAD_ATTRIBUTE Procedure 78-12
78.6.15 SET_GLOBAL_ATTRIBUTE Procedure 78-14
78.6.16 UNMAP_DAD Procedure 78-15

79 DBMS_ERRLOG

79.1 DBMS_ERRLOG Security Model 79-1
79.2 Summary of DBMS_ERRLOG Subprograms 79-1
79.2.1 CREATE_ERROR_LOG Procedure 79-1

80 DBMS_FGA

80.1 DBMS_FGA Security Model 80-1
80.2 DBMS_FGA Operational Notes 80-1
80.3 Summary of DBMS_FGA Subprograms 80-1
80.3.1 ADD_POLICY Procedure 80-2
80.3.2 DISABLE_POLICY Procedure 80-5
80.3.3 DROP_POLICY Procedure 80-6

80.3.4 ENABLE_POLICY Procedure 80-7

81 DBMS_FILE_TRANSFER

81.1 DBMS_FILE_TRANSFER Overview 81-1
81.2 DBMS_FILE_TRANSFER Security Model 81-1
81.3 DBMS_FILE_TRANSFER Operational Notes 81-1
81.4 Summary of DBMS_FILE_TRANSFER Subprograms 81-2
81.4.1 COPY_FILE Procedure 81-2
81.4.2 GET_FILE Procedure 81-4
81.4.3 PUT_FILE Procedure 81-6

82 DBMS_FLASHBACK

82.1 DBMS_FLASHBACK Overview 82-1
82.2 DBMS_FLASHBACK Security Model 82-1
82.3 DBMS_FLASHBACK Types 82-2
82.4 DBMS_FLASHBACK Exceptions 82-2
82.5 DBMS_FLASHBACK Operational Notes 82-2
82.6 DBMS_FLASHBACK Examples 82-3
82.7 Summary of DBMS_FLASHBACK Subprograms 82-5
82.7.1 DISABLE Procedure 82-5
82.7.2 ENABLE_AT_SYSTEM_CHANGE_NUMBER Procedure 82-6
82.7.3 ENABLE_AT_TIME Procedure 82-6
82.7.4 GET_SYSTEM_CHANGE_NUMBER Function 82-7
82.7.5 TRANSACTION_BACKOUT Procedures 82-7

83 DBMS_FLASHBACK_ARCHIVE

83.1 DBMS_FLASHBACK_ARCHIVE Overview 83-2
83.2 DBMS_FLASHBACK_ARCHIVE Security Model 83-2
83.3 DBMS_FLASHBACK_ARCHIVE Constants 83-2
83.4 Summary of DBMS_FLASHBACK_ARCHIVE Subprograms 83-3
83.4.1 ADD_TABLE_TO_APPLICATION Procedure 83-4
83.4.2 CREATE_TEMP_HISTORY_TABLE Procedure 83-4
83.4.3 DISABLE_APPLICATION Procedure 83-5
83.4.4 DISABLE_ASOF_VALID_TIME Procedure 83-5
83.4.5 DISASSOCIATE_FBA Procedure 83-5
83.4.6 DROP_APPLICATION Procedure 83-6
83.4.7 ENABLE_APPLICATION Procedure 83-6
83.4.8 ENABLE_AT_VALID_TIME Procedure 83-6
83.4.9 EXTEND_MAPPINGS Procedure 83-7

83.4.10 GET_CURRENT_LIFESPAN_DIGEST Function 83-7
83.4.11 GET_SYS_CONTEXT Function 83-8
83.4.12 IMPORT_HISTORY Procedure 83-9
83.4.13 LOCK_DOWN_APPLICATION Procedure 83-10
83.4.14 PURGE_CONTEXT Procedure 83-10
83.4.15 REASSOCIATE_FBA Procedure 83-10
83.4.16 REGISTER_APPLICATION Procedure 83-11
83.4.17 REMOVE_TABLE_FROM_APPLICATION Procedure 83-12
83.4.18 SET_CONTEXT_LEVEL Procedure 83-12
83.4.19 VERIFY_BLOCKCHAIN_LIFESPAN Procedure 83-12

84 DBMS_FLASHBACK_ARCHIVE_MIGRATE

84.1 DBMS_FLASHBACK_ARCHIVE_MIGRATE Overview 84-1
84.2 DBMS_FLASHBACK_ARCHIVE_MIGRATE Security Model 84-1
84.3 DBMS_FLASHBACK_ARCHIVE_MIGRATE Operational Notes 84-2
84.4 DBMS_FLASHBACK_ARCHIVE_MIGRATE Examples 84-3
84.5 Summary of DBMS_FLASHBACK Subprograms 84-4
84.5.1 EXPORT Procedure 84-5
84.5.2 EXPORT_ANALYZE Procedure 84-5
84.5.3 IMPORT Procedure 84-6

85 DBMS_FREQUENT_ITEMSET

85.1 Summary of DBMS_FREQUENT_ITEMSET Subprograms 85-1
85.1.1 FI_HORIZONTAL Function 85-1
85.1.2 FI_TRANSACTIONAL Function 85-3

86 DBMS_FS

86.1 DBMS_FS Overview 86-1
86.2 DBMS_FS Security Model 86-1
86.3 Summary of DBMS_FS Subprograms 86-1
86.3.1 DESTROY_ORACLE_FS Procedure 86-2
86.3.2 FS_EXISTS Procedure 86-3
86.3.3 LIST_FILES Procedure 86-3
86.3.4 MAKE_ORACLE_FS Procedure 86-4
86.3.5 MOUNT_ORACLE_FS Procedure 86-5
86.3.6 UNMOUNT_ORACLE_FS Procedure 86-9

87 DBMS_GOLDENGATE_ADM

87.1 Using DBMS_GOLDENGATE_ADM 87-1
87.1.1 DBMS_GOLDENGATE_ADM Overview 87-1
87.1.2 DBMS_GOLDENGATE_ADM Security Model 87-2
87.2 Summary of DBMS_GOLDENGATE_ADM Subprograms 87-2
87.2.1 ADD_AUTO_CDR Procedure 87-3
87.2.2 ADD_AUTO_CDR_COLUMN_GROUP Procedure 87-5
87.2.3 ADD_AUTO_CDR_DELTA_RES Procedure 87-6
87.2.4 ALTER_AUTO_CDR Procedure 87-7
87.2.5 ALTER_AUTO_CDR_COLUMN_GROUP Procedure 87-10
87.2.6 DELETE_PROCREP_EXCLUSION_OBJ Procedure 87-10
87.2.7 GG_PROCEDURE_REPLICATION_ON Function 87-11
87.2.8 INSERT_PROCREP_EXCLUSION_OBJ Procedure 87-11
87.2.9 PURGE_TOMBSTONES Procedure 87-12
87.2.10 REMOVE_AUTO_CDR Procedure 87-12
87.2.11 REMOVE_AUTO_CDR_COLUMN_GROUP Procedure 87-13
87.2.12 REMOVE_AUTO_CDR_DELTA_RES Procedure 87-13

88 DBMS_GOLDENGATE_AUTH

88.1 DBMS_GOLDENGATE_AUTH Overview 88-1
88.2 DBMS_GOLDENGATE_AUTH Security Model 88-1
88.3 Summary of DBMS_GOLDENGATE_AUTH Subprograms 88-2
88.3.1 GRANT_ADMIN_PRIVILEGE Procedure 88-2
88.3.2 REVOKE_ADMIN_PRIVILEGE Procedure 88-5

89 DBMS_HADOOP

89.1 DBMS_HADOOP Overview 89-1
89.2 DBMS_HADOOP Security Model 89-1
89.3 Summary of DBMS_HADOOP Subprograms 89-1
89.3.1 CREATE_EXTDDL_FOR_HIVE Procedure 89-2
89.3.2 CREATE_HYBRID_PARTNED_TABLE Procedure 89-3
89.3.3 SYNCHRONIZE_PARTITIONS_FOR_HIVE Procedure 89-3

90 DBMS_HEAT_MAP

90.1 DBMS_HEAT_MAP Overview 90-1
90.2 DBMS_HEAT_MAP Security Model 90-1
90.3 Summary of DBMS_HEAT_MAP Subprograms 90-1
90.3.1 BLOCK_HEAT_MAP Function 90-2
90.3.2 EXTENT_HEAT_MAP Function 90-3

90.3.3 OBJECT_HEAT_MAP Function 90-4
90.3.4 SEGMENT_HEAT_MAP Procedure 90-5
90.3.5 TABLESPACE_HEAT_MAP Function 90-6

91 DBMS_HIERARCHY

91.1 DBMS_HIERARCHY Overview 91-1
91.2 DBMS_HIERARCHY Security Model 91-1
91.3 Summary of DBMS_HIERARCHY Subprograms 91-2
91.3.1 CREATE_VALIDATE_LOG_TABLE Procedure 91-3
91.3.2 CREATE_VIEW_FOR_FACT_ROWS Procedure 91-4
91.3.3 CREATE_VIEW_FOR_STAR_ROWS Procedure 91-5
91.3.4 GET_MV_SQL_FOR_AV_CACHE Function 91-6
91.3.5 GET_MV_SQL_FOR_STAR_CACHE Function 91-6
91.3.6 UPGRADE_VALIDATE_LOG_TABLE Procedure 91-7
91.3.7 VALIDATE_ANALYTIC_VIEW Function 91-7
91.3.8 VALIDATE_CHECK_SUCCESS Function 91-9
91.3.9 VALIDATE_HIERARCHY Function 91-10

92 DBMS_HM

92.1 DBMS_HM Security Model 92-1
92.2 Summary of DBMS_HM Subprograms 92-1
92.2.1 GET_RUN_REPORT Function 92-1
92.2.2 RUN_CHECK Procedure 92-2

93 DBMS_HPROF

93.1 DBMS_HPROF Security Model 93-1
93.2 Summary of DBMS_HPROF Subprograms 93-1
93.2.1 ANALYZE Function 93-2
93.2.2 CREATE_TABLES Procedure 93-4
93.2.3 START_PROFILING Procedure 93-4
93.2.4 STOP_PROFILING Procedure 93-5

94 DBMS_HS_PARALLEL

94.1 Using the DBMS_HS_PARALLEL Subprogram 94-1
94.2 Summary of DBMS_HS_PARALLEL Subprograms 94-1
94.2.1 CREATE_OR_REPLACE_VIEW Procedure 94-1
94.2.2 CREATE_TABLE_TEMPLATE Procedure 94-3
94.2.3 DROP_VIEW Procedure 94-3

94.2.4 LOAD_TABLE Procedure 94-4

95 DBMS_HS_PASSTHROUGH

95.1 DBMS_HS_PASSTHROUGH Overview 95-1
95.2 DBMS_HS_PASSTHROUGH Operational Notes 95-1
95.3 Summary of DBMS_HS_PASSTHROUGH Subprograms 95-1
95.3.1 BIND_INOUT_VARIABLE Procedure 95-2
95.3.2 BIND_INOUT_VARIABLE_RAW Procedure 95-3
95.3.3 BIND_OUT_VARIABLE Procedure 95-4
95.3.4 BIND_OUT_VARIABLE_RAW Procedure 95-5
95.3.5 BIND_VARIABLE Procedure 95-6
95.3.6 BIND_VARIABLE_RAW Procedure 95-7
95.3.7 CLOSE_CURSOR Procedure 95-8
95.3.8 EXECUTE_IMMEDIATE Procedure 95-9
95.3.9 EXECUTE_NON_QUERY Function 95-10
95.3.10 FETCH_ROW Function 95-10
95.3.11 GET_VALUE Procedure 95-11
95.3.12 GET_VALUE_RAW Procedure 95-12
95.3.13 OPEN_CURSOR Function 95-13
95.3.14 PARSE Procedure 95-14

96 DBMS_HYBRID_VECTOR

96.1 SEARCH 96-1

97 DBMS_ILM

97.1 DBMS_ILM Overview 97-1
97.2 DBMS_ILM Security Model 97-2
97.3 DBMS_ILM Constants 97-2
97.4 DBMS_ILM Exceptions 97-2
97.5 Summary of DBMS_ILM Subprograms 97-3
97.5.1 ADD_TO_ILM Procedure 97-3
97.5.2 ARCHIVESTATENAME Function 97-4
97.5.3 EXECUTE_ILM Procedure 97-4
97.5.4 EXECUTE_ILM_TASK Procedure 97-5
97.5.5 PREVIEW_ILM Procedure 97-6
97.5.6 REMOVE_FROM_ILM Procedure 97-6
97.5.7 STOP_ILM Procedure 97-7

98 DBMS_ILM_ADMIN

98.1 DBMS_ILM_ADMIN Overview 98-1
98.2 DBMS_ILM_ADMIN Security Model 98-1
98.3 DBMS_ILM_ADMIN Constants 98-1
98.4 Summary of DBMS_ILM_ADMIN Subprograms 98-4
98.4.1 CLEAR_HEAT_MAP_ALL Procedure 98-5
98.4.2 CLEAR_HEAT_MAP_TABLE Procedure 98-5
98.4.3 CUSTOMIZE_ILM Procedure 98-6
98.4.4 DISABLE_AUTO_OPTIMIZE Procedure 98-6
98.4.5 DISABLE_ILM Procedure 98-6
98.4.6 ENABLE_AUTO_OPTIMIZE Procedure 98-7
98.4.7 ENABLE_ILM Procedure 98-7
98.4.8 IGNORE_AUTO_OPTIMIZE_CRITERIA Procedure 98-7
98.4.9 SET_HEAT_MAP_ALL Procedure 98-7
98.4.10 SET_HEAT_MAP_START Procedure 98-8
98.4.11 SET_HEAT_MAP_TABLE Procedure 98-8

99 DBMS_IMMUTABLE_TABLE

99.1 DBMS_IMMUTABLE_TABLE Overview 99-1
99.2 DBMS_IMMUTABLE_TABLE Security Model 99-1
99.3 Summary of DBMS_IMMUTABLE_TABLE Subprograms 99-1
99.3.1 ADD_INTERVAL_PARTITIONING Procedure 99-2
99.3.2 DELETE_EXPIRED_ROWS Procedure 99-2

100 DBMS_INMEMORY

100.1 DBMS_INMEMORY Overview 100-1
100.2 DBMS_INMEMORY Security Model 100-2
100.3 Summary of DBMS_INMEMORY Subprograms 100-2
100.3.1 IME_DROP_EXPRESSIONS Procedure 100-2
100.3.2 POPULATE Procedure 100-3
100.3.3 REPOPULATE Procedure 100-4
100.3.4 SEGMENT_DEALLOCATE_VERSIONS Procedure 100-4

101 DBMS_INMEMORY_ADMIN

101.1 DBMS_INMEMORY_ADMIN Overview 101-1
101.2 DBMS_INMEMORY_ADMIN Security Model 101-3
101.3 DBMS_INMEMORY_ADMIN Operational Notes 101-3
101.4 Summary of DBMS_INMEMORY_ADMIN Subprograms 101-3
101.4.1 AIM_GET_PARAMETER Procedure 101-4

101.4.2 AIM_SET_PARAMETER Procedure 101-5
101.4.3 FASTSTART_DISABLE Procedure 101-6
101.4.4 FASTSTART_ENABLE Procedure 101-7
101.4.5 FASTSTART_MIGRATE_STORAGE Procedure 101-9
101.4.6 GET_FASTSTART_TABLESPACE Function 101-10
101.4.7 IME_CAPTURE_EXPRESSIONS Procedure 101-11
101.4.8 IME_CLOSE_CAPTURE_WINDOW Procedure 101-13
101.4.9 IME_DROP_ALL_EXPRESSIONS Procedure 101-14
101.4.10 IME_GET_CAPTURE_STATE Procedure 101-14
101.4.11 IME_OPEN_CAPTURE_WINDOW Procedure 101-15
101.4.12 IME_POPULATE_EXPRESSIONS Procedure 101-16
101.4.13 POPULATE_WAIT Function 101-16

102 DBMS_INMEMORY_ADVISE

102.1 DBMS_INMEMORY_ADVISE Overview 102-1
102.2 Summary of DBMS_INMEMORY_ADVISE Subprograms 102-1
102.3 IS_INMEMORY_ELIGIBLE Procedure 102-1

103 DBMS_IOT

103.1 Summary of DBMS_IOT Subprograms 103-1
103.1.1 BUILD_CHAIN_ROWS_TABLE Procedure 103-1
103.1.2 BUILD_EXCEPTIONS_TABLE Procedure 103-2

104 DBMS_JAVA

105 DBMS_JOB

105.1 DBMS_JOB Security Model 105-1
105.2 DBMS_JOB Operational Notes 105-1
105.3 Summary of DBMS_JOB Subprograms 105-3
105.3.1 BROKEN Procedure 105-3
105.3.2 CHANGE Procedure 105-4
105.3.3 INSTANCE Procedure 105-5
105.3.4 INTERVAL Procedure 105-6
105.3.5 NEXT_DATE Procedure 105-6
105.3.6 REMOVE Procedure 105-7
105.3.7 RUN Procedure 105-7
105.3.8 SUBMIT Procedure 105-8
105.3.9 USER_EXPORT Procedures 105-10

105.3.10 WHAT Procedure 105-10

106 DBMS_JSON_DUALITY

106.1 DBMS_JSON_DUALITY Overview 106-1
106.2 DBMS_JSON_DUALITY Security 106-2
106.3 Summary of DBMS_JSON_Duality Subprograms 106-2
106.3.1 INFER_SCHEMA Function 106-3
106.3.2 GENERATE_SCHEMA Function 106-5
106.3.3 INFER_AND_GENERATE_SCHEMA Function 106-6
106.3.4 IMPORT Procedure 106-7
106.3.5 BEGIN_TRANSACTION Procedure 106-8
106.3.6 COMMIT_TRANSACTION Procedure 106-9
106.3.7 REGISTER Procedure 106-9

107 DBMS_JSON

107.1 DBMS_JSON Overview 107-1
107.2 DBMS_JSON Security Model 107-1
107.3 DBMS_JSON Constants 107-1
107.4 Summary of DBMS_JSON Subprograms 107-4
107.4.1 ADD_VIRTUAL_COLUMNS Procedure 107-4
107.4.2 CREATE_VIEW Procedure 107-6
107.4.3 CREATE_VIEW_ON_PATH Procedure 107-7
107.4.4 DROP_VIRTUAL_COLUMNS Procedure 107-8
107.4.5 GET_INDEX_DATAGUIDE Function 107-9
107.4.6 GET_VIEW_SQL Function 107-10
107.4.7 JSON_TYPE_CONVERTIBLE_CHECK Procedure 107-11
107.4.8 RENAME_COLUMN Procedure 107-12

108 DBMS_JSON_SCHEMA

108.1 DBMS_JSON_SCHEMA Security Model 108-1
108.2 DBMS_JSON_SCHEMA Constants 108-1
108.3 Summary of DBMS_JSON_SCHEMA Subprograms 108-2
108.3.1 DESCRIBE Function 108-2
108.3.2 IS_SCHEMA_VALID Function 108-3
108.3.3 IS_VALID Function 108-3
108.3.4 IS_VALID Procedure 108-4
108.3.5 VALIDATE_REPORT Function 108-5

109 DBMS_KAFKA

109.1 DBMS_KAFKA Overview 109-1
109.2 DBMS_KAFKA LOADING Mode 109-3
109.2.1 CREATE_LOAD_APP 109-3
109.2.2 DROP_LOAD_APP 109-4
109.2.3 EXECUTE_LOAD_APP 109-5
109.3 DBMS_KAFKA Global Temporary Tables 109-7
109.3.1 LOAD_TEMP_TABLE 109-7
109.4 DBMS_KAFKA Streaming Mode 109-9
109.4.1 CREATE_STREAMING_APP 109-9
109.4.2 DROP_STREAMING_APP 109-11
109.4.3 INIT_OFFSET 109-12
109.4.4 INIT_OFFSET_TS (Milliseconds Since Epoch) 109-13
109.4.5 INIT_OFFSET_TS (Timestamp with Separate Timezone Parameter) 109-14
109.4.6 INIT_OFFSET_TS (Timestamp with Time Zone) 109-15
109.4.7 UPDATE_OFFSET 109-17
109.5 DBMS_KAFKA Seekable Mode 109-19
109.5.1 CREATE_SEEKABLE_APP 109-19
109.5.2 DROP_SEEKABLE_APP 109-20
109.5.3 SEEK_OFFSET_TS (Timestamp with Separate Timezone Parameter) 109-21
109.5.4 SEEK_OFFSET_TS (Timestamp in Milliseconds) 109-23
109.5.5 SEEK_OFFSET_TS (Timestamp with Timezone) 109-25
109.6 ADD_PARTITIONS 109-27
109.7 DROP_ALL_APPS 109-29
109.8 ENABLE_VIEW_QUERY 109-29
109.9 SET_TRACING 109-31
109.10 DBMS_KAFKA OPTIONS Passed to CREATE_xxx_APP 109-32

110 DBMS_KAFKA_ADM

110.1 DBMS_KAFKA_ADM Overview 110-1
110.2 DBMS_KAFKA_ADM Security Model 110-1
110.3 DBMS_KAFKA_ADM Constants 110-2
110.4 Summary of DBMS_KAFKA_ADM Procedures 110-3
110.5 CHECK_CLUSTER 110-4
110.6 DEREGISTER_CLUSTER 110-5
110.7 DISABLE_CLUSTER 110-6
110.8 ENABLE_CLUSTER 110-7
110.9 REGISTER_CLUSTER 110-8
110.10 UPDATE_CLUSTER_INFO 110-11

111 DBMS_LDAP

112 DBMS_LDAP_UTL

113 DBMS_LIBCACHE

113.1 DBMS_LIBCACHE Overview 113-1
113.2 DBMS_LIBCACHE Security Model 113-1
113.3 Summary of DBMS_LIBCACHE Subprograms 113-1
113.3.1 COMPILE_FROM_REMOTE Procedure 113-2

114 DBMS_LOB

114.1 DBMS_LOB Overview 114-1
114.2 DBMS_LOB Security Model 114-1
114.3 DBMS_LOB Constants 114-2
114.4 DBMS_LOB Datatypes 114-4
114.5 DBMS_LOB Operational Notes 114-4
114.6 DBMS_LOB Rules and Limits 114-7
114.7 DBMS_LOB Exceptions 114-10
114.8 Summary of DBMS_LOB Subprograms 114-11
114.8.1 APPEND Procedures 114-13
114.8.2 CLOB2FILE Procedure 114-14
114.8.3 CLOSE Procedure 114-15
114.8.4 COMPARE Functions 114-16
114.8.5 CONVERTTOBLOB Procedure 114-18
114.8.6 CONVERTTOCLOB Procedure 114-20
114.8.7 COPY Procedures 114-23
114.8.8 COPY_DBFS_LINK Procedures 114-24
114.8.9 COPY_FROM_DBFS_LINK 114-25
114.8.10 CREATETEMPORARY Procedures 114-26
114.8.11 DBFS_LINK_GENERATE_PATH Functions 114-27
114.8.12 ERASE Procedures 114-27
114.8.13 FILECLOSE Procedure 114-29
114.8.14 FILECLOSEALL Procedure 114-30
114.8.15 FILEEXISTS Function 114-30
114.8.16 FILEGETNAME Procedure 114-31
114.8.17 FILEISOPEN Function 114-32
114.8.18 FILEOPEN Procedure 114-33

114.8.19 FRAGMENT_DELETE Procedure 114-34
114.8.20 FRAGMENT_INSERT Procedures 114-35
114.8.21 FRAGMENT_MOVE Procedure 114-36
114.8.22 FRAGMENT_REPLACE Procedures 114-37
114.8.23 FREETEMPORARY Procedures 114-38
114.8.24 GET_DBFS_LINK Functions 114-39
114.8.25 GET_DBFS_LINK_STATE Procedures 114-40
114.8.26 GET_LOB_DEDUPLICATION_RATIO Function 114-41
114.8.27 GETCONTENTTYPE Functions 114-42
114.8.28 GET_STORAGE_LIMIT Function 114-43
114.8.29 GETCHUNKSIZE Functions 114-44
114.8.30 GETLENGTH Functions 114-45
114.8.31 GETOPTIONS Functions 114-46
114.8.32 INSTR Functions 114-47
114.8.33 ISOPEN Functions 114-49
114.8.34 ISREMOTE Function 114-50
114.8.35 ISSECUREFILE Function 114-50
114.8.36 ISTEMPORARY Functions 114-51
114.8.37 LOADBLOBFROMFILE Procedure 114-52
114.8.38 LOADCLOBFROMFILE Procedure 114-53
114.8.39 MOVE_TO_DBFS_LINK Procedures 114-56
114.8.40 OPEN Procedures 114-57
114.8.41 READ Procedures 114-58
114.8.42 SET_DBFS_LINK Procedures 114-60
114.8.43 SETCONTENTTYPE Procedure 114-61
114.8.44 SETOPTIONS Procedures 114-62
114.8.45 SUBSTR Functions 114-63
114.8.46 TRIM Procedures 114-65
114.8.47 WRITE Procedures 114-66
114.8.48 WRITEAPPEND Procedures 114-68

115 DBMS_LOCK

115.1 DBMS_LOCK Overview 115-1
115.2 DBMS_LOCK Security Model 115-1
115.3 DBMS_LOCK Constants 115-2
115.4 DBMS_LOCK. Rules and Limits 115-2
115.5 DBMS_LOCK Operational Notes 115-3
115.6 Summary of DBMS_LOCK Subprograms 115-3
115.6.1 ALLOCATE_UNIQUE Procedure 115-3
115.6.2 ALLOCATE_UNIQUE_AUTONOMOUS Procedure 115-5
115.6.3 CONVERT Function 115-6

115.6.4 RELEASE Function 115-7
115.6.5 REQUEST Function 115-8

116 DBMS_LOGMNR

116.1 DBMS_LOGMNR Overview 116-1
116.2 DBMS_LOGMNR Security Model 116-2
116.3 DBMS_LOGMNR Constants 116-2
116.4 DBMS_LOGMNR Views 116-4
116.5 DBMS_LOGMNR Operational Notes 116-4
116.6 Summary of DBMS_LOGMNR Subprograms 116-4
116.6.1 ADD_LOGFILE Procedure 116-5
116.6.2 COLUMN_PRESENT Function 116-7
116.6.3 END_LOGMNR Procedure 116-8
116.6.4 MINE_VALUE Function 116-9
116.6.5 REMOVE_LOGFILE Procedure 116-10
116.6.6 START_LOGMNR Procedure 116-11

117 DBMS_LOGSTDBY

117.1 DBMS_LOGSTDBY Overview 117-1
117.2 DBMS_LOGSTDBY Security Model 117-1
117.3 DBMS_LOGSTDBY Constants 117-2
117.4 Summary of DBMS_LOGSTDBY Subprograms 117-2
117.4.1 APPLY_SET Procedure 117-4
117.4.2 APPLY_UNSET Procedure 117-7
117.4.3 BUILD Procedure 117-7
117.4.4 INSTANTIATE_TABLE Procedure 117-8
117.4.5 IS_APPLY_SERVER Function 117-10
117.4.6 MAP_PRIMARY_SCN Function 117-10
117.4.7 PREPARE_FOR_NEW_PRIMARY Procedure 117-11
117.4.8 PURGE_SESSION Procedure 117-12
117.4.9 REBUILD Procedure 117-13
117.4.10 SET_TABLESPACE Procedure 117-13
117.4.11 SKIP Procedure 117-14
117.4.12 SKIP_ERROR Procedure 117-23
117.4.13 SKIP_TRANSACTION Procedure 117-26
117.4.14 UNSKIP Procedure 117-28
117.4.15 UNSKIP_ERROR Procedure 117-29
117.4.16 UNSKIP_TRANSACTION Procedure 117-30

118 DBMS_LOGMNR_D

118.1 DBMS_LOGMNR_D Overview 118-1
118.2 DBMS_LOGMNR_D Security Model 118-2
118.3 Summary of DBMS_LOGMNR_D Subprograms 118-2
118.3.1 BUILD Procedure 118-2
118.3.2 SET_TABLESPACE Procedure 118-5

119 DBMS_LOGSTDBY_CONTEXT

119.1 DBMS_LOGSTDBY_CONTEXT Overview 119-1
119.2 DBMS_LOGSTDBY_CONTEXT Security Model 119-2
119.3 Summary of DBMS_LOGSTDBY_CONTEXT Subprograms 119-2
119.3.1 CLEAR_ALL_CONTEXT Procedure 119-3
119.3.2 CLEAR_CONTEXT Procedure 119-3
119.3.3 GET_CONTEXT Procedure 119-3
119.3.4 SET_CONTEXT Procedure 119-4

120 DBMS_MEMOPTIMIZE

120.1 DBMS_MEMOPTIMIZE Overview 120-1
120.2 Summary of DBMS_MEMOPTIMIZE Subprograms 120-1
120.2.1 DROP_OBJECT Procedure 120-2
120.2.2 GET_APPLY_HWM_SEQID Function 120-3
120.2.3 GET_WRITE_HWM_SEQID Function 120-3
120.2.4 POPULATE Procedure 120-3
120.2.5 WRITE_END Procedure 120-4

121 DBMS_MEMOPTIMIZE_ADMIN

121.1 DBMS_MEMOPTIMIZE_ADMIN Overview 121-1
121.2 Summary of DBMS_MEMOPTIMIZE_ADMIN Subprograms 121-1
121.2.1 WRITES_FLUSH Procedure 121-2

122 DBMS_METADATA

122.1 DBMS_METADATA Overview 122-1
122.2 DBMS_METADATA Security Model 122-2
122.3 Rules and Limits 122-3
122.4 DBMS_METADATA Data Structures - Object and Table Types 122-3
122.5 DBMS_METADATA Subprogram Groupings 122-4
122.5.1 DBMS_METADATA Subprograms for Retrieving Multiple Objects From the Database 122-4

122.5.2 DBMS_METADATA Subprograms for Submitting XML to the Database 122-5
122.6 Summary of All DBMS_METADATA Subprograms 122-5
122.6.1 ADD_TRANSFORM Function 122-6
122.6.2 CLOSE Procedure 122-10
122.6.3 CONVERT Functions and Procedures 122-10
122.6.4 FETCH_xxx Functions and Procedures 122-12
122.6.5 GET_xxx Functions 122-15
122.6.6 GET_QUERY Function 122-19
122.6.7 OPEN Function 122-20
122.6.8 OPENW Function 122-27
122.6.9 PUT Function 122-28
122.6.10 SET_COUNT Procedure 122-29
122.6.11 SET_FILTER Procedure 122-30
122.6.12 SET_PARSE_ITEM Procedure 122-41
122.6.13 SET_TRANSFORM_PARAM and SET_REMAP_PARAM Procedures 122-44

123 DBMS_METADATA_DIFF

123.1 DBMS_METADATA_DIFF Overview 123-1
123.2 DBMS_METADATA_DIFF Security Model 123-1
123.3 DBMS_METADATA_DIFF — Browsing APIs for Fetching and Comparing Objects 123-2
123.4 Summary of DBMS_METADATA_DIFF Subprograms 123-3
123.4.1 OPENC Function 123-4
123.4.2 ADD_DOCUMENT Procedure 123-4
123.4.3 FETCH_CLOB Functions and Procedures 123-5
123.4.4 CLOSE Procedure 123-6

124 DBMS_MGD_ID_UTL

124.1 DBMS_MGD_ID_UTL Security Model 124-1
124.2 DBMS_MGD_ID_UTL Constants 124-2
124.3 DBMS_MGD_ID_UTL Exceptions 124-2
124.4 Summary of DBMS_MGD_ID_UTL Subprograms 124-3
124.4.1 ADD_SCHEME Procedure 124-4
124.4.2 CREATE_CATEGORY Function 124-7
124.4.3 EPC_TO_ORACLE_SCHEME Function 124-8
124.4.4 GET_CATEGORY_ID Function 124-11
124.4.5 GET_COMPONENTS Function 124-12
124.4.6 GET_ENCODINGS Function 124-13
124.4.7 GET_JAVA_LOGGING_LEVEL Function 124-13
124.4.8 GET_PLSQL_LOGGING_LEVEL Function 124-14
124.4.9 GET_SCHEME_NAMES Function 124-15

124.4.10 GET_TDT_XML Function 124-15
124.4.11 GET_VALIDATOR Function 124-17
124.4.12 REFRESH_CATEGORY Function 124-20
124.4.13 REMOVE_CATEGORY Procedure 124-21
124.4.14 REMOVE_PROXY Procedure 124-22
124.4.15 REMOVE_SCHEME Procedure 124-22
124.4.16 SET_JAVA_LOGGING_LEVEL Procedure 124-23
124.4.17 SET_PLSQL_LOGGING_LEVEL Procedure 124-23
124.4.18 SET_PROXY Procedure 124-24
124.4.19 VALIDATE_SCHEME Function 124-25

125 DBMS_MGWADM

125.1 DBMS_MGWADM Security Model 125-1
125.2 DBMS_MGWADM Deprecated Subprograms 125-1
125.3 DBMS_MGWADM Constants 125-2
125.4 DBMS_MGWADM Data Structures 125-5
125.4.1 DBMS_MGWADM SYS.MGW_MQSERIES_PROPERTIES Object Type 125-5
125.4.2 DBMS_MGWADM SYS.MGW_PROPERTIES Object Type 125-7
125.4.3 DBMS_MGWADM SYS.MGW_PROPERTY Object Type 125-8
125.4.4 SYS.MGW_TIBRV_PROPERTIES Object Type 125-9
125.5 Summary of DBMS_MGWADM Subprograms 125-10
125.5.1 ADD_SUBSCRIBER Procedure 125-12
125.5.2 ALTER_AGENT Procedures 125-15
125.5.3 ALTER_JOB Procedure 125-16
125.5.4 ALTER_MSGSYSTEM_LINK Procedure for TIB/Rendezvous 125-18
125.5.5 ALTER_MSGSYSTEM_LINK Procedure for WebSphere MQ 125-19
125.5.6 ALTER_PROPAGATION_SCHEDULE Procedure 125-20
125.5.7 ALTER_SUBSCRIBER Procedure 125-21
125.5.8 CLEANUP_GATEWAY Procedures 125-23
125.5.9 CREATE_AGENT Procedure 125-25
125.5.10 CREATE_JOB Procedure 125-27
125.5.11 CREATE_MSGSYSTEM_LINK Procedures for TIB/Rendezvous 125-30
125.5.12 CREATE_MSGSYSTEM_LINK Procedures for WebSphere MQ 125-31
125.5.13 DB_CONNECT_INFO Procedure 125-32
125.5.14 DISABLE_JOB Procedure 125-32
125.5.15 DISABLE_PROPAGATION_SCHEDULE Procedure 125-33
125.5.16 ENABLE_JOB Procedure 125-33
125.5.17 ENABLE_PROPAGATION_SCHEDULE Procedure 125-34
125.5.18 REGISTER_FOREIGN_QUEUE Procedure 125-34
125.5.19 REMOVE_AGENT Procedure 125-35
125.5.20 REMOVE_JOB Procedure 125-36

125.5.21 REMOVE_MSGSYSTEM_LINK Procedure 125-36
125.5.22 REMOVE_OPTION Procedure 125-37
125.5.23 REMOVE_SUBSCRIBER Procedure 125-38
125.5.24 RESET_JOB Procedure 125-39
125.5.25 RESET_SUBSCRIBER Procedure 125-40
125.5.26 SCHEDULE_PROPAGATION Procedure 125-40
125.5.27 SET_LOG_LEVEL Procedures 125-42
125.5.28 SET_OPTION Procedure 125-42
125.5.29 SHUTDOWN Procedures 125-44
125.5.30 STARTUP Procedures 125-44
125.5.31 UNREGISTER_FOREIGN_QUEUE Procedure 125-45
125.5.32 UNSCHEDULE_PROPAGATION Procedure 125-46

126 DBMS_MGWMSG

126.1 DBMS_MGWMSG Security Model 126-1
126.2 DBMS_MGWMSG Constants 126-1
126.3 DBMS_MGWMSG Types 126-3
126.3.1 DBMS_MGWMSG SYS.MGW_NAME_VALUE_T Type 126-3
126.3.2 DBMS_MGWMSG SYS.MGW_NAME_TYPE_ARRAY_T Type 126-6
126.3.3 DBMS_MGWMSG SYS.MGW_TEXT_VALUE_T Type 126-6
126.3.4 DBMS_MGWMSG SYS.MGW_RAW_VALUE_T Type 126-6
126.3.5 DBMS_MGWMSG SYS.MGW_BASIC_MSG_T Type 126-7
126.3.6 DBMS_MGWMSG SYS.MGW_NUMBER_ARRAY_T Type 126-8
126.3.7 DBMS_MGWMSG SYS.MGW_TIBRV_FIELD_T Type 126-8
126.3.8 DBMS_MGWMSG SYS.MGW_TIBRV_MSG_T Type 126-9
126.4 Summary of DBMS_MGWMSG Subprograms 126-17
126.4.1 LCR_TO_XML Function 126-18
126.4.2 NVARRAY_ADD Procedure 126-19
126.4.3 NVARRAY_FIND_NAME Function 126-19
126.4.4 NVARRAY_FIND_NAME_TYPE Function 126-20
126.4.5 NVARRAY_GET Function 126-21
126.4.6 NVARRAY_GET_BOOLEAN Function 126-21
126.4.7 NVARRAY_GET_BYTE Function 126-22
126.4.8 NVARRAY_GET_DATE Function 126-22
126.4.9 NVARRAY_GET_DOUBLE Function 126-23
126.4.10 NVARRAY_GET_FLOAT Function 126-23
126.4.11 NVARRAY_GET_INTEGER Function 126-24
126.4.12 NVARRAY_GET_LONG Function 126-24
126.4.13 NVARRAY_GET_RAW Function 126-25
126.4.14 NVARRAY_GET_SHORT Function 126-25
126.4.15 NVARRAY_GET_TEXT Function 126-26

126.4.16 XML_TO_LCR Function 126-27

127 DBMS_MLE

127.1 DBMS_MLE Overview 127-1
127.2 DBMS_MLE Security Model 127-2
127.3 Summary of DBMS_MLE Subprograms 127-2
127.3.1 CREATE_CONTEXT Function 127-3
127.3.2 DISABLE_CTX_STDERR Procedure 127-4
127.3.3 DISABLE_CTX_STDOUT Procedure 127-5
127.3.4 DISABLE_DEBUGGING Procedure 127-5
127.3.5 DISABLE_ICS_STDERR Procedure 127-6
127.3.6 DISABLE_ICS_STDOUT Procedure 127-6
127.3.7 DISABLE_STDERR Procedure 127-7
127.3.8 DISABLE_STDOUT Procedure 127-7
127.3.9 DROP_CONTEXT Procedure 127-8
127.3.10 ENABLE_DEBUGGING Procedure 127-9
127.3.11 EVAL Procedure 127-9
127.3.12 EXPORT_TO_MLE Procedure 127-11
127.3.13 GET_AVAILABLE_LANGUAGES Function 127-13
127.3.14 GET_CTX_ERROR_STACK Function 127-13
127.3.15 GET_ERROR_STACK Function 127-14
127.3.16 IMPORT_FROM_MLE Procedure 127-14
127.3.17 PARSE_DEBUG_OUTPUT Function 127-16
127.3.18 SET_CTX_STDERR Procedure 127-17
127.3.19 SET_CTX_STDERR_TO_DBMS_OUTPUT Procedure 127-18
127.3.20 SET_CTX_STDOUT Procedure 127-18
127.3.21 SET_CTX_STDOUT_TO_DBMS_OUTPUT Procedure 127-19
127.3.22 SET_ICS_STDERR Procedure 127-19
127.3.23 SET_ICS_STDERR_TO_DBMS_OUTPUT Procedure 127-20
127.3.24 SET_ICS_STDOUT Procedure 127-21
127.3.25 SET_ICS_STDOUT_TO_DBMS_OUTPUT Procedure 127-21
127.3.26 SET_STDERR Procedure 127-22
127.3.27 SET_STDERR_TO_DBMS_OUTPUT Procedure 127-23
127.3.28 SET_STDOUT Procedure 127-23
127.3.29 SET_STDOUT_TO_DBMS_OUTPUT Procedure 127-24

128 DBMS_MONITOR

128.1 Summary of DBMS_MONITOR Subprograms 128-1
128.1.1 CLIENT_ID_STAT_DISABLE Procedure 128-2
128.1.2 CLIENT_ID_STAT_ENABLE Procedure 128-2

128.1.3 CLIENT_ID_TRACE_DISABLE Procedure 128-3
128.1.4 CLIENT_ID_TRACE_ENABLE Procedure 128-3
128.1.5 DATABASE_TRACE_DISABLE Procedure 128-4
128.1.6 DATABASE_TRACE_ENABLE Procedure 128-4
128.1.7 SERV_MOD_ACT_STAT_DISABLE Procedure 128-5
128.1.8 SERV_MOD_ACT_STAT_ENABLE Procedure 128-5
128.1.9 SERV_MOD_ACT_TRACE_DISABLE Procedure 128-7
128.1.10 SERV_MOD_ACT_TRACE_ENABLE Procedure 128-8
128.1.11 SESSION_TRACE_DISABLE Procedure 128-9
128.1.12 SESSION_TRACE_ENABLE Procedure 128-10

129 DBMS_MVIEW

129.1 DBMS_MVIEW Operational Notes 129-1
129.2 DBMS_MVIEW Security Model 129-1
129.3 DBMS_MVIEW Rules and Limits 129-1
129.4 Summary of DBMS_MVIEW Subprograms 129-2
129.4.1 BEGIN_TABLE_REORGANIZATION Procedure 129-2
129.4.2 END_TABLE_REORGANIZATION Procedure 129-3
129.4.3 ESTIMATE_MVIEW_SIZE Procedure 129-3
129.4.4 EXPLAIN_MVIEW Procedure 129-4
129.4.5 EXPLAIN_REWRITE Procedure 129-5
129.4.6 I_AM_A_REFRESH Function 129-6
129.4.7 PMARKER Function 129-6
129.4.8 PURGE_DIRECT_LOAD_LOG Procedure 129-7
129.4.9 PURGE_LOG Procedure 129-7
129.4.10 PURGE_MVIEW_FROM_LOG Procedure 129-8
129.4.11 REFRESH Procedures 129-9
129.4.12 REFRESH_ALL_MVIEWS Procedure 129-11
129.4.13 REFRESH_DEPENDENT Procedures 129-13
129.4.14 REGISTER_MVIEW Procedure 129-15
129.4.15 UNREGISTER_MVIEW Procedure 129-16

130 DBMS_MVIEW_STATS

130.1 DBMS_MVIEW_STATS Overview 130-1
130.2 DBMS_MVIEW_STATS Security Model 130-1
130.3 Summary of DBMS_MVIEW_STATS Subprograms 130-1
130.3.1 PURGE_REFRESH_STATS Procedure 130-2
130.3.2 SET_MVREF_STATS_PARAMS Procedure 130-3
130.3.3 SET_SYSTEM_DEFAULT Procedure 130-4

131 DBMS_NETWORK_ACL_ADMIN

131.1 DBMS_NETWORK_ACL_ADMIN Overview 131-1
131.2 DBMS_NETWORK_ACL_ADMIN Deprecated Subprograms 131-1
131.3 DBMS_NETWORK_ACL_ADMIN Security Model 131-2
131.4 DBMS_NETWORK_ACL_ADMIN Constants 131-2
131.5 DBMS_NETWORK_ACL_ADMIN Exceptions 131-2
131.6 DBMS_NETWORK_ACL_ADMIN Examples 131-3
131.7 Summary of DBMS_NETWORK_ACL_ADMIN Subprograms 131-4
131.7.1 ADD_PRIVILEGE Procedure 131-6
131.7.2 APPEND_HOST_ACE Procedure 131-7
131.7.3 APPEND_HOST_ACL Procedure 131-8
131.7.4 APPEND_WALLET_ACE Procedure 131-9
131.7.5 APPEND_WALLET_ACL Procedure 131-10
131.7.6 ASSIGN_ACL Procedure 131-10
131.7.7 ASSIGN_WALLET_ACL Procedure 131-12
131.7.8 CHECK_PRIVILEGE Function 131-13
131.7.9 CHECK_PRIVILEGE_ACLID Function 131-14
131.7.10 CREATE_ACL Procedure 131-15
131.7.11 DELETE_PRIVILEGE Procedure 131-16
131.7.12 DROP_ACL Procedure 131-17
131.7.13 REMOVE_HOST_ACE Procedure 131-17
131.7.14 REMOVE_WALLET_ACE Procedure 131-18
131.7.15 SET_HOST_ACL Procedure 131-18
131.7.16 SET_WALLET_ACL Procedure 131-19
131.7.17 UNASSIGN_ACL Procedure 131-20
131.7.18 UNASSIGN_WALLET_ACL Procedure 131-21

132 DBMS_NETWORK_ACL_UTILITY

132.1 DBMS_NETWORK_ACL_UTILITY Security Model 132-1
132.2 DBMS_NETWORK_ACL_UTILITY Examples 132-1
132.3 Summary of DBMS_NETWORK_ACL_UTILITY Subprograms 132-2
132.3.1 CONTAINS_HOST Function 132-3
132.3.2 DOMAIN_LEVEL Function 132-4
132.3.3 DOMAINS Function 132-5
132.3.4 EQUALS_HOST Function 132-5

133 DBMS_ODCI

133.1 Summary of DBMS_ODCI Subprograms 133-1

133.1.1 ESTIMATE_CPU_UNITS Function 133-1

134 DBMS_OPTIM_BUNDLE

134.1 DBMS_OPTIM_BUNDLE Security Model 134-1
134.2 Summary of DBMS_OPTIM_BUNDLE Subprograms 134-1
134.2.1 ENABLE_OPTIM_FIXES Procedure 134-1
134.2.2 GETBUGSFORBUNDLE Procedure 134-3
134.2.3 LISTBUNDLESWITHFCFIXES Procedure 134-3
134.2.4 SET_FIX_CONTROLS Procedure 134-4

135 DBMS_OUTLN

135.1 DBMS_OUTLN Overview 135-1
135.2 DBMS_OUTLN Security Model 135-2
135.3 Summary of DBMS_OUTLN Subprograms 135-2
135.3.1 CLEAR_USED Procedure 135-2
135.3.2 CREATE_OUTLINE Procedure 135-3
135.3.3 DROP_BY_CAT Procedure 135-3
135.3.4 DROP_UNUSED Procedure 135-4
135.3.5 EXACT_TEXT_SIGNATURES Procedure 135-4
135.3.6 UPDATE_BY_CAT Procedure 135-4
135.3.7 UPDATE_SIGNATURES Procedure 135-4

136 DBMS_OUTPUT

136.1 DBMS_OUTPUT Overview 136-1
136.2 DBMS_OUTPUT Security Model 136-1
136.3 DBMS_OUTPUT Operational Notes 136-1
136.4 DBMS_OUTPUT Exceptions 136-2
136.5 DBMS_OUTPUT Rules and Limits 136-2
136.6 DBMS_OUTPUT Examples 136-2
136.7 DBMS_OUTPUT Data Structures 136-5
136.7.1 CHARARR Table Type 136-5
136.7.2 DBMS_OUTPUT DBMSOUTPUT_LINESARRAY Object Type 136-5
136.8 Summary of DBMS_OUTPUT Subprograms 136-6
136.8.1 DISABLE Procedure 136-6
136.8.2 ENABLE Procedure 136-6
136.8.3 GET_LINE Procedure 136-7
136.8.4 GET_LINES Procedure 136-8
136.8.5 NEW_LINE Procedure 136-8
136.8.6 PUT Procedure 136-9

136.8.7 PUT_LINE Procedure 136-10

137 DBMS_PARALLEL_EXECUTE

137.1 DBMS_PARALLEL_EXECUTE Overview 137-1
137.2 DBMS_PARALLEL_EXECUTE Security Model 137-1
137.3 DBMS_PARALLEL_EXECUTE Constants 137-2
137.4 DBMS_PARALLEL_EXECUTE Views 137-3
137.5 DBMS_PARALLEL_EXECUTE Exceptions 137-3
137.6 DBMS_PARALLEL_EXECUTE Examples 137-3
137.7 Summary of DBMS_PARALLEL_EXECUTE Subprograms 137-6
137.7.1 ADM_DROP_CHUNKS Procedure 137-7
137.7.2 ADM_DROP_TASK Procedure 137-7
137.7.3 ADM_TASK_STATUS Function 137-8
137.7.4 ADM_STOP_TASK Procedure 137-8
137.7.5 CREATE_TASK Procedure 137-8
137.7.6 CREATE_CHUNKS_BY_NUMBER_COL Procedure 137-9
137.7.7 CREATE_CHUNKS_BY_ROWID Procedure 137-9
137.7.8 CREATE_CHUNKS_BY_SQL Procedure 137-10
137.7.9 DROP_TASK Procedure 137-11
137.7.10 DROP_CHUNKS Procedure 137-11
137.7.11 GENERATE_TASK_NAME Function 137-11
137.7.12 GET_NUMBER_COL_CHUNK Procedure 137-12
137.7.13 GET_ROWID_CHUNK Procedure 137-12
137.7.14 PURGE_PROCESSED_CHUNKS Procedure 137-13
137.7.15 RESUME_TASK Procedures 137-14
137.7.16 RUN_TASK Procedure 137-16
137.7.17 SET_CHUNK_STATUS Procedure 137-17
137.7.18 STOP_TASK Procedure 137-18
137.7.19 TASK_STATUS Procedure 137-18

138 DBMS_PART

138.1 DBMS_PART Security Model 138-1
138.2 DBMS_PART Operational Notes 138-1
138.3 Summary of DBMS_PART Subprograms 138-1
138.3.1 CLEANUP_GIDX Procedure 138-2
138.3.2 CLEANUP_GIDX_JOB Procedure 138-2
138.3.3 CLEANUP_ONLINE_OP Procedure 138-3

139 DBMS_PCLXUTIL

139.1 DBMS_PCLXUTIL Overview 139-1
139.2 DBMS_PCLXUTIL Security Model 139-2
139.3 DBMS_PCLXUTIL Operational Notes 139-2
139.4 DBMS_PCLXUTIL Rules and Limits 139-2
139.5 Summary of DBMS_PCLXUTIL Subprograms 139-3
139.5.1 BUILD_PART_INDEX Procedure 139-3

140 DBMS_PDB

140.1 DBMS_PDB Overview 140-1
140.2 DBMS_PDB Security Model 140-2
140.3 Summary of DBMS_PDB Subprograms 140-2
140.3.1 CHECK_PLUG_COMPATIBILITY Function 140-3
140.3.2 DESCRIBE Procedure 140-4
140.3.3 EXPORTRMANBACKUP Procedure 140-4
140.3.4 RECOVER Procedure 140-5
140.3.5 REMOVE_LINK Procedure 140-6
140.3.6 SET_DATA_LINKED Procedure 140-6
140.3.7 SET_EXT_DATA_LINKED Procedure 140-7
140.3.8 SET_METADATA_LINKED Procedure 140-8
140.3.9 SET_PROFILE_EXPLICIT Procedure 140-9
140.3.10 SET_ROLE_EXPLICIT Procedure 140-9
140.3.11 SET_USER_EXPLICIT Procedure 140-10

141 DBMS_PDB_ALTER_SHARING

141.1 DBMS_PDB_ALTER_SHARING Overview 141-1
141.2 DBMS_PDB_ALTER_SHARING Security Model 141-2
141.3 Summary of DBMS_PDB_ALTER_SHARING Subprograms 141-2
141.3.1 REMOVE_LINK Procedure 141-3
141.3.2 SET_DATA_LINKED Procedure 141-3
141.3.3 SET_EXT_DATA_LINKED Procedure 141-4
141.3.4 SET_METADATA_LINKED Procedure 141-5
141.3.5 SET_PROFILE_EXPLICIT Procedure 141-6
141.3.6 SET_ROLE_EXPLICIT Procedure 141-6
141.3.7 SET_USER_EXPLICIT Procedure 141-7

142 DBMS_PERF

142.1 DBMS_PERF Overview 142-1
142.2 DBMS_PERF Security Model 142-1

142.3 Summary of DBMS_PERF Subprograms 142-1
142.3.1 REPORT_PERFHUB Function 142-2
142.3.2 REPORT_SESSION Function 142-3
142.3.3 REPORT_SQL Function 142-5

143 DBMS_PIPE

143.1 DBMS_PIPE Overview 143-1
143.2 DBMS_PIPE Security Model 143-2
143.3 DBMS_PIPE Constants 143-2
143.4 DBMS_PIPE Operational Notes 143-2
143.5 DBMS_PIPE Exceptions 143-5
143.6 DBMS_PIPE Examples 143-5
143.7 Summary of DBMS_PIPE Subprograms 143-14
143.7.1 CREATE_PIPE Function 143-15
143.7.2 GET_CREDENTIAL_NAME Function 143-17
143.7.3 GET_LOCATION_URI Function 143-17
143.7.4 NEXT_ITEM_TYPE Function 143-17
143.7.5 PACK_MESSAGE Procedures 143-18
143.7.6 PURGE Procedure 143-19
143.7.7 RECEIVE_MESSAGE Function 143-20
143.7.8 RESET_BUFFER Procedure 143-24
143.7.9 REMOVE_PIPE Function 143-24
143.7.10 SEND_MESSAGE Function 143-25
143.7.11 SET_CREDENTIAL_NAME Procedure 143-28
143.7.12 SET_LOCATION_URI Procedure 143-28
143.7.13 UNIQUE_SESSION_NAME Function 143-29
143.7.14 UNPACK_MESSAGE Procedures 143-29

144 DBMS_PLSQL_CODE_COVERAGE

144.1 DBMS_PLSQL_CODE_COVERAGE Overview 144-1
144.2 DBMS_PLSQL_CODE_COVERAGE Security Model 144-1
144.3 DBMS_PLSQL_CODE_COVERAGE Constants 144-2
144.4 DBMS_PLSQL_CODE_COVERAGE Data Structures 144-2
144.4.1 MAP_REC Record Type 144-2
144.4.2 T_MAP_REC Table Type 144-3
144.5 Summary of DBMS_PLSQL_CODE_COVERAGE Subprograms 144-3
144.5.1 CREATE_COVERAGE_TABLES Procedure 144-3
144.5.2 GET_BLOCK_MAP Function 144-4
144.5.3 START_COVERAGE Function 144-4

144.5.4 STOP_COVERAGE Procedure 144-5

145 DBMS_PREDICTIVE_ANALYTICS

145.1 DBMS_PREDICTIVE_ANALYTICS Overview 145-1
145.2 DBMS_PREDICTIVE_ANALYTICS Security Model 145-1
145.3 Summary of DBMS_PREDICTIVE_ANALYTICS Subprograms 145-2
145.3.1 EXPLAIN Procedure 145-2
145.3.2 PREDICT Procedure 145-4
145.3.3 PROFILE Procedure 145-6

146 DBMS_PREPROCESSOR

146.1 DBMS_PREPROCESSOR Overview 146-1
146.2 DBMS_PREPROCESSOR Operating Notes 146-1
146.3 DBMS_PREPROCESSOR Data Structures 146-2
146.3.1 DBMS_PREPROCESSOR SOURCE_LINES_T Table Type 146-2
146.4 Summary of DBMS_PREPROCESSOR Subprograms 146-2
146.4.1 GET_POST_PROCESSED_SOURCE Functions 146-3
146.4.2 PRINT_POST_PROCESSED_SOURCE Procedures 146-4

147 DBMS_PRIVILEGE_CAPTURE

147.1 DBMS_PRIVILEGE_CAPTURE Overview 147-1
147.2 DBMS_PRIVILEGE_CAPTURE Security Model 147-1
147.3 DBMS_PRIVILEGE_CAPTURE Constants 147-1
147.4 DBMS_PRIVILEGE_CAPTURE Examples 147-2
147.5 Summary of DBMS_PRIVILEGE_CAPTURE Subprograms 147-3
147.5.1 CAPTURE_DEPENDENCY_PRIVS Procedure 147-3
147.5.2 CREATE_CAPTURE Procedure 147-3
147.5.3 DELETE_RUN Procedure 147-5
147.5.4 DISABLE_CAPTURE Procedure 147-5
147.5.5 DROP_CAPTURE Procedure 147-6
147.5.6 ENABLE_CAPTURE Procedure 147-6
147.5.7 GENERATE_RESULT Procedure 147-7

148 DBMS_PROCESS

148.1 DBMS_PROCESS Overview 148-1
148.2 DBMS_PROCESS Security Model 148-1
148.3 Summary of DBMS_PROCESS Subprograms 148-2
148.3.1 CONFIGURE_POOL Procedure 148-2

148.3.2 START_POOL Procedure 148-3
148.3.3 STOP_POOL Procedure 148-4

149 DBMS_PROFILER

149.1 DBMS_PROFILER Overview 149-1
149.2 DBMS_PROFILER Security Model 149-3
149.3 DBMS_PROFILER Operational Notes 149-3
149.4 DBMS_PROFILER Exceptions 149-5
149.5 Summary of DBMS_PROFILER Subprograms 149-6
149.5.1 FLUSH_DATA Function and Procedure 149-6
149.5.2 GET_VERSION Procedure 149-6
149.5.3 INTERNAL_VERSION_CHECK Function 149-7
149.5.4 PAUSE_PROFILER Function and Procedure 149-7
149.5.5 RESUME_PROFILER Function and Procedure 149-7
149.5.6 START_PROFILER Functions and Procedures 149-7
149.5.7 STOP_PROFILER Function and Procedure 149-8

150 DBMS_PROPAGATION_ADM

150.1 DBMS_PROPAGATION_ADM Overview 150-1
150.2 DBMS_PROPAGATION_ADM Security Model 150-1
150.3 Summary of DBMS_PROPAGATION_ADM Subprograms 150-1
150.3.1 ALTER_PROPAGATION Procedure 150-2
150.3.2 CREATE_PROPAGATION Procedure 150-3
150.3.3 DROP_PROPAGATION Procedure 150-6
150.3.4 START_PROPAGATION Procedure 150-7
150.3.5 STOP_PROPAGATION Procedure 150-7

151 DBMS_QOPATCH

151.1 DBMS_QOPATCH Overview 151-1
151.2 DBMS_QOPATCH Security Model 151-1
151.3 DBMS_QOPATCH Operational Notes 151-1
151.4 DBMS_QOPATCH Exceptions 151-1
151.5 Summary of DBMS_QOPATCH Subprograms 151-2
151.5.1 CHECK_PATCH_INSTALLED Function 151-3
151.5.2 GET_OPATCH_BUGS Function 151-3
151.5.3 GET_OPATCH_COUNT Function 151-4
151.5.4 GET_OPATCH_DATA Function 151-4
151.5.5 GET_OPATCH_FILES Function 151-4
151.5.6 GET_OPATCH_INSTALL_INFO Function 151-5

151.5.7 GET_OPATCH_LIST Function 151-5
151.5.8 GET_OPATCH_LSINVENTORY 151-5
151.5.9 GET_OPATCH_OLAYS Function 151-5
151.5.10 GET_OPATCH_PREQS Function 151-6
151.5.11 GET_OPATCH_XSLT 151-6
151.5.12 GET_PATCH_DETAILS Function 151-6
151.5.13 GET_SQLPATCH_STATUS Procedure 151-7
151.5.14 IS_PATCH_INSTALLED Function 151-7
151.5.15 OPATCH_COMPARE_CURRENT Function 151-7
151.5.16 OPATCH_COMPARE_NODES Function 151-8
151.5.17 SET_CURRENT_OPINST Procedure 151-8

152 DBMS_RANDOM

152.1 DBMS_RANDOM Deprecated Subprograms 152-1
152.2 DBMS_RANDOM Security Model 152-1
152.3 DBMS_RANDOM Operational Notes 152-1
152.4 Summary of DBMS_RANDOM Subprograms 152-2
152.4.1 INITIALIZE Procedure 152-2
152.4.2 NORMAL Function 152-3
152.4.3 RANDOM Function 152-3
152.4.4 SEED Procedures 152-4
152.4.5 STRING Function 152-4
152.4.6 TERMINATE Procedure 152-5
152.4.7 VALUE Functions 152-5

153 DBMS_REDACT

153.1 DBMS_REDACT Overview 153-1
153.2 DBMS_REDACT Security Model 153-1
153.3 DBMS_REDACT Constants 153-2
153.4 DBMS_REDACT Operating Procedures 153-2
153.5 Summary of DBMS_REDACT Subprograms 153-7
153.5.1 ADD_POLICY Procedure 153-7
153.5.2 ALTER_POLICY Procedure 153-13
153.5.3 APPLY_POLICY_EXPR_TO_COL Procedure 153-20
153.5.4 CREATE_POLICY_EXPRESSION Procedure 153-21
153.5.5 DISABLE_POLICY Procedure 153-22
153.5.6 DROP_POLICY Procedure 153-23
153.5.7 DROP_POLICY_EXPRESSION Procedure 153-23
153.5.8 ENABLE_POLICY Procedure 153-24
153.5.9 UPDATE_FULL_REDACTION_VALUES Procedure 153-25

153.5.10 UPDATE_POLICY_EXPRESSION Procedure 153-26

154 DBMS_REDEFINITION

154.1 DBMS_REDEFINITION Overview 154-1
154.2 DBMS_REDEFINITION Security Model 154-1
154.3 DBMS_REDEFINITION Constants 154-1
154.4 DBMS_REDEFINITION Operational Notes 154-2
154.5 DBMS_REDEFINITION Rules and Limits 154-3
154.6 DBMS_REDEFINITION Examples 154-3
154.7 Summary of DBMS_REDEFINITION Subprograms 154-4
154.7.1 ABORT_REDEF_TABLE Procedure 154-5
154.7.2 ABORT_ROLLBACK Procedure 154-5
154.7.3 ABORT_UPDATE Procedure 154-6
154.7.4 CAN_REDEF_TABLE Procedure 154-7
154.7.5 COPY_TABLE_DEPENDENTS Procedure 154-7
154.7.6 EXECUTE_UPDATE Procedure 154-9
154.7.7 FINISH_REDEF_TABLE Procedure 154-9
154.7.8 REDEF_TABLE Procedure 154-10
154.7.9 REGISTER_DEPENDENT_OBJECT Procedure 154-12
154.7.10 ROLLBACK Procedure 154-13
154.7.11 SET_PARAM Procedure 154-13
154.7.12 START_REDEF_TABLE Procedure 154-14
154.7.13 SYNC_INTERIM_TABLE Procedure 154-16
154.7.14 UNREGISTER_DEPENDENT_OBJECT Procedure 154-16

155 DBMS_REFRESH

155.1 DBMS_REFRESH Overview 155-1
155.2 DBMS_REFRESH Security Model 155-1
155.3 Summary of DBMS_REFRESH Subprograms 155-1
155.3.1 ADD Procedure 155-2
155.3.2 CHANGE Procedure 155-2
155.3.3 DESTROY Procedure 155-4
155.3.4 MAKE Procedure 155-4
155.3.5 REFRESH Procedure 155-6
155.3.6 SUBTRACT Procedure 155-7

156 DBMS_REPAIR

156.1 DBMS_REPAIR Overview 156-1
156.2 DBMS_REPAIR Security Model 156-1

156.3 DBMS_REPAIR Constants 156-1
156.4 DBMS_REPAIR Operating Notes 156-2
156.5 DBMS_REPAIR Exceptions 156-3
156.6 DBMS_REPAIR Examples 156-4
156.7 Summary of DBMS_REPAIR Subprograms 156-4
156.7.1 ADMIN_TABLES Procedure 156-4
156.7.2 CHECK_OBJECT Procedure 156-5
156.7.3 DUMP_ORPHAN_KEYS Procedure 156-6
156.7.4 FIX_CORRUPT_BLOCKS Procedure 156-7
156.7.5 ONLINE_INDEX_CLEAN Function 156-8
156.7.6 REBUILD_FREELISTS Procedure 156-9
156.7.7 SEGMENT_FIX_STATUS Procedure 156-9
156.7.8 SKIP_CORRUPT_BLOCKS Procedure 156-10

157 DBMS_RESCONFIG

157.1 DBMS_RESCONFIG Overview 157-1
157.2 Summary of DBMS_RESCONFIG Subprograms 157-1
157.2.1 ADDREPOSITORYRESCONFIG Procedure 157-2
157.2.2 ADDRESCONFIG Procedure 157-3
157.2.3 APPENDRESCONFIG Procedure 157-3
157.2.4 DELETEREPOSITORYRESCONFIG Procedure 157-4
157.2.5 DELETERESCONFIG Procedures 157-5
157.2.6 GETLISTENERS Function 157-5
157.2.7 GETREPOSITORYRESCONFIG Function 157-6
157.2.8 GETREPOSITORYRESCONFIGPATHS Function 157-6
157.2.9 GETRESCONFIG Function 157-7
157.2.10 GETRESCONFIGPATHS Function 157-7
157.2.11 PATCHREPOSITORYRESCONFIGLIST Procedure 157-8

158 DBMS_RESOURCE_MANAGER

158.1 DBMS_RESOURCE_MANAGER Deprecated Subprograms 158-1
158.2 DBMS_RESOURCE_MANAGER Security Model 158-1
158.3 DBMS_RESOURCE_MANAGER Constants 158-2
158.4 Summary of DBMS_RESOURCE_MANAGER Subprograms 158-3
158.4.1 BEGIN_SQL_BLOCK Procedure 158-5
158.4.2 CALIBRATE_IO Procedure 158-5
158.4.3 CLEAR_PENDING_AREA Procedure 158-7
158.4.4 CREATE_CATEGORY Procedure 158-7
158.4.5 CREATE_CDB_PLAN Procedure 158-8
158.4.6 CREATE_CDB_PLAN_DIRECTIVE Procedure 158-8

158.4.7 CREATE_CDB_PROFILE_DIRECTIVE Procedure 158-10
158.4.8 CREATE_CONSUMER_GROUP Procedure 158-11
158.4.9 CREATE_PENDING_AREA Procedure 158-11
158.4.10 CREATE_PLAN Procedure 158-13
158.4.11 CREATE_PLAN_DIRECTIVE Procedure 158-14
158.4.12 CREATE_SIMPLE_PLAN Procedure 158-18
158.4.13 DELETE_CATEGORY Procedure 158-20
158.4.14 DELETE_CDB_PLAN Procedure 158-20
158.4.15 DELETE_CDB_PLAN_DIRECTIVE Procedure 158-21
158.4.16 DELETE_CDB_PROFILE_DIRECTIVE Procedure 158-21
158.4.17 DELETE_CONSUMER_GROUP Procedure 158-22
158.4.18 DELETE_PLAN Procedure 158-22
158.4.19 DELETE_PLAN_CASCADE Procedure 158-22
158.4.20 DELETE_PLAN_DIRECTIVE Procedure 158-23
158.4.21 DEQUEUE_PARALLEL_STATEMENT Procedure 158-23
158.4.22 END_SQL_BLOCK Procedure 158-24
158.4.23 SET_CONSUMER_GROUP_MAPPING Procedure 158-24
158.4.24 SET_CONSUMER_GROUP_MAPPING_PRI Procedure 158-25
158.4.25 SET_INITIAL_CONSUMER_GROUP Procedure 158-26
158.4.26 SUBMIT_PENDING_AREA Procedure 158-27
158.4.27 SWITCH_CONSUMER_GROUP_FOR_SESS Procedure 158-27
158.4.28 SWITCH_CONSUMER_GROUP_FOR_USER Procedure 158-28
158.4.29 SWITCH_PLAN Procedure 158-28
158.4.30 UPDATE_CATEGORY Procedure 158-29
158.4.31 UPDATE_CDB_AUTOTASK_DIRECTIVE Procedure 158-29
158.4.32 UPDATE_CDB_DEFAULT_DIRECTIVE Procedure 158-30
158.4.33 UPDATE_CDB_PLAN Procedure 158-31
158.4.34 UPDATE_CDB_PLAN_DIRECTIVE Procedure 158-32
158.4.35 UPDATE_CDB_PROFILE_DIRECTIVE Procedure 158-33
158.4.36 UPDATE_CONSUMER_GROUP Procedure 158-34
158.4.37 UPDATE_PLAN Procedure 158-35
158.4.38 UPDATE_PLAN_DIRECTIVE Procedure 158-36
158.4.39 VALIDATE_PENDING_AREA Procedure 158-40

159 DBMS_RESOURCE_MANAGER_PRIVS

159.1 Summary of DBMS_RESOURCE_MANAGER_PRIVS Subprograms 159-1
159.1.1 GRANT_SWITCH_CONSUMER_GROUP Procedure 159-1
159.1.2 GRANT_SYSTEM_PRIVILEGE Procedure 159-2
159.1.3 REVOKE_SWITCH_CONSUMER_GROUP Procedure 159-3
159.1.4 REVOKE_SYSTEM_PRIVILEGE Procedure 159-4

160 DBMS_RESULT_CACHE

160.1 DBMS_RESULT_CACHE Security Model 160-1
160.2 DBMS_RESULT_CACHE Constants 160-1
160.3 Summary of DBMS_RESULT_CACHE Subprograms 160-2
160.3.1 BLACK_LIST Function 160-3
160.3.2 BLACK_LIST_ADD Procedure 160-3
160.3.3 BLACK_LIST_CLEAR Procedure 160-3
160.3.4 BLACK_LIST_REMOVE Procedure 160-3
160.3.5 BYPASS Procedure 160-3
160.3.6 FLUSH Function & Procedure 160-4
160.3.7 IGNORE_LIST Function 160-5
160.3.8 IGNORE_LIST_ADD Procedure 160-5
160.3.9 IGNORE_LIST_CLEAR Procedure 160-6
160.3.10 IGNORE_LIST_REMOVE Procedure 160-6
160.3.11 INVALIDATE Functions & Procedures 160-7
160.3.12 INVALIDATE_OBJECT Functions & Procedures 160-7
160.3.13 MEMORY_REPORT Procedure 160-8
160.3.14 OBJECT_BLACK_LIST Function 160-8
160.3.15 OBJECT_BLACK_LIST_ADD Procedure 160-8
160.3.16 OBJECT_BLACK_LIST_CLEAR Procedure 160-9
160.3.17 OBJECT_BLACK_LIST_REMOVE Procedure 160-9
160.3.18 STATUS Function 160-9
160.3.19 BLOCK_LIST Function 160-9
160.3.20 BLOCK_LIST_ADD Procedure 160-9
160.3.21 BLOCK_LIST_CLEAR Procedure 160-10
160.3.22 BLOCK_LIST_REMOVE Procedure 160-10
160.3.23 OBJECT_BLOCK_LIST Function 160-10
160.3.24 OBJECT_BLOCK_LIST_ADD Procedure 160-10
160.3.25 OBJECT_BLOCK_LIST_CLEAR Procedure 160-11
160.3.26 OBJECT_BLOCK_LIST_REMOVE Procedure 160-11

161 DBMS_RESUMABLE

161.1 DBMS_RESUMABLE Operational Notes 161-1
161.2 Summary of DBMS_RESUMABLE Subprograms 161-1
161.2.1 ABORT Procedure 161-2
161.2.2 GET_SESSION_TIMEOUT Function 161-2
161.2.3 GET_TIMEOUT Function 161-3
161.2.4 SET_SESSION_TIMEOUT Procedure 161-3
161.2.5 SET_TIMEOUT Procedure 161-4

161.2.6 SPACE_ERROR_INFO Function 161-4

162 DBMS_RLS

162.1 DBMS_RLS Overview 162-1
162.2 DBMS_RLS Security Model 162-2
162.3 DBMS_RLS Constants 162-2
162.4 DBMS_RLS Operational Notes 162-3
162.5 DBMS_RLS Rules and Limits 162-3
162.6 Summary of DBMS_RLS Subprograms 162-3
162.6.1 ADD_GROUPED_POLICY Procedure 162-4
162.6.2 ADD_POLICY Procedure 162-5
162.6.3 ADD_POLICY_CONTEXT Procedure 162-10
162.6.4 ALTER_POLICY Procedure 162-11
162.6.5 ALTER_GROUPED_POLICY Procedure 162-12
162.6.6 CREATE_POLICY_GROUP Procedure 162-13
162.6.7 DELETE_POLICY_GROUP Procedure 162-13
162.6.8 DISABLE_GROUPED_POLICY Procedure 162-14
162.6.9 DROP_GROUPED_POLICY Procedure 162-15
162.6.10 DROP_POLICY Procedure 162-15
162.6.11 DROP_POLICY_CONTEXT Procedure 162-16
162.6.12 ENABLE_GROUPED_POLICY Procedure 162-16
162.6.13 ENABLE_POLICY Procedure 162-17
162.6.14 REFRESH_GROUPED_POLICY Procedure 162-18
162.6.15 REFRESH_POLICY Procedure 162-18

163 DBMS_ROLLING

163.1 DBMS_ROLLING Overview 163-1
163.2 DBMS_ROLLING Security Model 163-2
163.3 Summary of DBMS_ROLLING Subprograms 163-2
163.3.1 INIT_PLAN Procedure 163-3
163.3.2 DESTROY_PLAN Procedure 163-4
163.3.3 BUILD_PLAN Procedure 163-4
163.3.4 SET_PARAMETER Procedure 163-5
163.3.5 START_PLAN Procedure 163-8
163.3.6 SWITCHOVER Procedure 163-9
163.3.7 FINISH_PLAN Procedure 163-10
163.3.8 ROLLBACK_PLAN Procedure 163-11

164 DBMS_ROWID

164.1 DBMS_ROWID Security Model 164-1
164.2 DBMS_ROWID Types 164-1
164.3 DBMS_ROWID Exceptions 164-2
164.4 DBMS_ROWID Operational Notes 164-3
164.5 DBMS_ROWID Examples 164-4
164.6 Summary of DBMS_ROWID Subprograms 164-4
164.6.1 ROWID_BLOCK_NUMBER Function 164-4
164.6.2 ROWID_CREATE Function 164-5
164.6.3 ROWID_INFO Procedure 164-6
164.6.4 ROWID_OBJECT Function 164-7
164.6.5 ROWID_RELATIVE_FNO Function 164-8
164.6.6 ROWID_ROW_NUMBER Function 164-8
164.6.7 ROWID_TO_ABSOLUTE_FNO Function 164-9
164.6.8 ROWID_TO_EXTENDED Function 164-10
164.6.9 ROWID_TO_RESTRICTED Function 164-11
164.6.10 ROWID_TYPE Function 164-12
164.6.11 ROWID_VERIFY Function 164-12

165 DBMS_RULE

165.1 DBMS_RULE Overview 165-1
165.2 DBMS_RULE Security Model 165-1
165.3 Summary of DBMS_RULE Subprograms 165-1
165.3.1 CLOSE_ITERATOR Procedure 165-2
165.3.2 EVALUATE Procedure 165-3
165.3.3 EVALUATE_EXPRESSION Procedure 165-6
165.3.4 EVALUATE_EXPRESSION_ITERATOR Procedure 165-7
165.3.5 EVALUATE_RULE Procedure 165-8
165.3.6 EVALUATE_RULE_ITERATOR Procedure 165-9
165.3.7 GET_NEXT_HIT Function 165-9
165.3.8 GET_NEXT_RESULT Function 165-10
165.3.9 IS_FAST Procedure 165-10

166 DBMS_RULE_ADM

166.1 DBMS_RULE_ADM Overview 166-1
166.2 DBMS_RULE_ADM Security Model 166-1
166.3 Summary of DBMS_RULE_ADM Subprograms 166-1
166.3.1 ADD_RULE Procedure 166-2
166.3.2 ALTER_EVALUATION_CONTEXT Procedure 166-4

166.3.3 ALTER_RULE Procedure 166-6
166.3.4 CREATE_EVALUATION_CONTEXT Procedure 166-8
166.3.5 CREATE_RULE Procedure 166-9
166.3.6 CREATE_RULE_SET Procedure 166-11
166.3.7 DROP_EVALUATION_CONTEXT Procedure 166-11
166.3.8 DROP_RULE Procedure 166-12
166.3.9 DROP_RULE_SET Procedure 166-13
166.3.10 GRANT_OBJECT_PRIVILEGE Procedure 166-14
166.3.11 GRANT_SCHEMA_PRIVILEGE Procedure 166-15
166.3.12 GRANT_SYSTEM_PRIVILEGE Procedure 166-16
166.3.13 REMOVE_RULE Procedure 166-18
166.3.14 REVOKE_OBJECT_PRIVILEGE Procedure 166-19
166.3.15 REVOKE_SCHEMA_PRIVILEGE Procedure 166-19
166.3.16 REVOKE_SYSTEM_PRIVILEGE Procedure 166-20

167 DBMS_SAGA_ADM

167.1 DBMS_SAGA_ADM Overview 167-1
167.2 DBMS_SAGA_ADM Security Model 167-1
167.3 Summary of DBMS_SAGA_ADM Subprograms 167-1
167.3.1 ADD_BROKER Procedure 167-2
167.3.2 DROP_BROKER Procedure 167-2
167.3.3 ADD_COORDINATOR Procedure 167-3
167.3.4 DROP_COORDINATOR Procedure 167-4
167.3.5 ADD_PARTICIPANT Procedure 167-4
167.3.6 DROP_PARTICIPANT Procedure 167-6
167.3.7 REGISTER_SAGA_CALLBACK Procedure 167-6

168 DBMS_SAGA

168.1 DBMS_SAGA Overview 168-1
168.2 DBMS_SAGA Security Model 168-1
168.3 Summary of DBMS_SAGA Subprograms 168-1
168.3.1 BEGIN_SAGA Function 168-2
168.3.2 GET_SAGA_ID Function 168-2
168.3.3 SET_SAGA_ID Procedure 168-2
168.3.4 COMMIT_SAGA Procedure 168-3
168.3.5 ROLLBACK_SAGA Procedure 168-3
168.3.6 IS_INCOMPLETE Function 168-4
168.3.7 SET_INCOMPLETE Procedure 168-4
168.3.8 SEND_REQUEST Procedure 168-5

169 DBMS_SCHEDULER

169.1 DBMS_SCHEDULER Deprecated Subprograms 169-1
169.2 DBMS_SCHEDULER Security Model 169-1
169.3 DBMS_SCHEDULER Rules and Limits 169-1
169.4 DBMS_SCHEDULER Operational Notes 169-2
169.5 DBMS_SCHEDULER Data Structures 169-11
169.5.1 DBMS_SCHEDULER JOBARG Object Type 169-12
169.5.2 JOBARG_ARRAY Table Type 169-13
169.5.3 DBMS_SCHEDULER JOB_DEFINITION Object Type 169-13
169.5.4 JOB_DEFINITION_ARRAY Table Type 169-17
169.5.5 JOBATTR Object Type 169-17
169.5.6 JOBATTR_ARRAY Table Type 169-18
169.5.7 SCHEDULER$_STEP_TYPE Object Type 169-18
169.5.8 SCHEDULER$_STEP_TYPE_LIST Table Type 169-19
169.5.9 SCHEDULER$_EVENT_INFO Object Type 169-19
169.5.10 SCHEDULER_FILEWATCHER_RESULT Object Type 169-20
169.5.11 SCHEDULER_FILEWATCHER_REQUEST Object Type 169-21
169.6 Summary of DBMS_SCHEDULER Subprograms 169-22
169.6.1 ADD_EVENT_QUEUE_SUBSCRIBER Procedure 169-25
169.6.2 ADD_GROUP_MEMBER Procedure 169-26
169.6.3 ADD_JOB_EMAIL_NOTIFICATION Procedure 169-27
169.6.4 ADD_TO_INCOMPATIBILITY Procedure 169-29
169.6.5 ALTER_CHAIN Procedure 169-30
169.6.6 ALTER_RUNNING_CHAIN Procedure 169-32
169.6.7 CLOSE_WINDOW Procedure 169-34
169.6.8 COPY_JOB Procedure 169-34
169.6.9 CREATE_CHAIN Procedure 169-35
169.6.10 CREATE_CREDENTIAL Procedure 169-36
169.6.11 CREATE_DATABASE_DESTINATION Procedure 169-37
169.6.12 CREATE_EVENT_SCHEDULE Procedure 169-38
169.6.13 CREATE_FILE_WATCHER Procedure 169-39
169.6.14 CREATE_GROUP Procedure 169-40
169.6.15 CREATE_INCOMPATIBILITY Procedure 169-42
169.6.16 CREATE_JOB Procedure 169-43
169.6.17 CREATE_JOB_CLASS Procedure 169-51
169.6.18 CREATE_JOBS Procedure 169-53
169.6.19 CREATE_PROGRAM Procedure 169-54
169.6.20 CREATE_RESOURCE Procedure 169-57
169.6.21 CREATE_SCHEDULE Procedure 169-58
169.6.22 CREATE_WINDOW Procedure 169-59
169.6.23 DEFINE_ANYDATA_ARGUMENT Procedure 169-61

169.6.24 DEFINE_CHAIN_EVENT_STEP Procedure 169-62
169.6.25 DEFINE_CHAIN_RULE Procedure 169-63
169.6.26 DEFINE_CHAIN_STEP Procedure 169-66
169.6.27 DEFINE_METADATA_ARGUMENT Procedure 169-67
169.6.28 DEFINE_PROGRAM_ARGUMENT Procedure 169-68
169.6.29 DISABLE Procedure 169-70
169.6.30 DROP_AGENT_DESTINATION Procedure 169-72
169.6.31 DROP_CHAIN Procedure 169-73
169.6.32 DROP_CHAIN_RULE Procedure 169-74
169.6.33 DROP_CHAIN_STEP Procedure 169-75
169.6.34 DROP_CREDENTIAL Procedure 169-75
169.6.35 DROP_DATABASE_DESTINATION Procedure 169-76
169.6.36 DROP_FILE_WATCHER Procedure 169-77
169.6.37 DROP_GROUP Procedure 169-78
169.6.38 DROP_INCOMPATIBILITY Procedure 169-78
169.6.39 DROP_JOB Procedure 169-79
169.6.40 DROP_JOB_CLASS Procedure 169-80
169.6.41 DROP_PROGRAM Procedure 169-81
169.6.42 DROP_PROGRAM_ARGUMENT Procedure 169-81
169.6.43 DROP_RESOURCE Procedure 169-82
169.6.44 DROP_SCHEDULE Procedure 169-83
169.6.45 DROP_WINDOW Procedure 169-83
169.6.46 DUMP_IN_MEMORY_TRACE Procedure 169-84
169.6.47 ENABLE Procedure 169-84
169.6.48 END_DETACHED_JOB_RUN Procedure 169-85
169.6.49 EVALUATE_CALENDAR_STRING Procedure 169-86
169.6.50 EVALUATE_RUNNING_CHAIN Procedure 169-88
169.6.51 GENERATE_JOB_NAME Function 169-89
169.6.52 GET_AGENT_INFO Function 169-89
169.6.53 GET_AGENT_VERSION Function 169-90
169.6.54 GET_ATTRIBUTE Procedure 169-91
169.6.55 GET_FILE Procedure 169-91
169.6.56 GET_SCHEDULER_ATTRIBUTE Procedure 169-93
169.6.57 OPEN_WINDOW Procedure 169-94
169.6.58 PURGE_LOG Procedure 169-95
169.6.59 PUT_FILE Procedure 169-96
169.6.60 REMOVE_EVENT_QUEUE_SUBSCRIBER Procedure 169-97
169.6.61 REMOVE_FROM_INCOMPATIBILITY Procedure 169-97
169.6.62 REMOVE_GROUP_MEMBER Procedure 169-98
169.6.63 REMOVE_JOB_EMAIL_NOTIFICATION Procedure 169-99
169.6.64 RESET_JOB_ARGUMENT_VALUE Procedure 169-99
169.6.65 RUN_CHAIN Procedure 169-100

169.6.66 RUN_JOB Procedure 169-102
169.6.67 SET_AGENT_REGISTRATION_PASS Procedure 169-103
169.6.68 SET_ATTRIBUTE Procedure 169-104
169.6.69 SET_ATTRIBUTE_NULL Procedure 169-119
169.6.70 SET_JOB_ANYDATA_VALUE Procedure 169-120
169.6.71 SET_JOB_ARGUMENT_VALUE Procedure 169-121
169.6.72 SET_JOB_ATTRIBUTES Procedure 169-122
169.6.73 SET_RESOURCE_CONSTRAINT Procedure 169-123
169.6.74 SET_SCHEDULER_ATTRIBUTE Procedure 169-124
169.6.75 STOP_JOB Procedure 169-125

170 DBMS_SEARCH

171 DBMS_SERVER_ALERT

171.1 DBMS_SERVER_ALERT Security Model 171-1
171.2 DBMS_SERVER_ALERT Object Types 171-1
171.3 DBMS_SERVER_ALERT Relational Operators 171-2
171.4 DBMS_SERVER_ALERT Supported Metrics 171-2
171.5 Summary of DBMS_SERVER_ALERT Subprograms 171-6
171.5.1 EXPAND_MESSAGE Function 171-6
171.5.2 GET_THRESHOLD Procedure 171-7
171.5.3 SET_THRESHOLD Procedure 171-8

172 DBMS_SERVICE

172.1 DBMS_SERVICE Overview 172-1
172.2 DBMS_SERVICE Security Model 172-2
172.3 DBMS_SERVICE Constants 172-2
172.4 DBMS_SERVICE Operating Procedures 172-4
172.5 DBMS_SERVICE Exceptions 172-5
172.6 Summary of DBMS_SERVICE Subprograms 172-6
172.6.1 CREATE_SERVICE Procedure 172-7
172.6.2 DELETE_SERVICE Procedure 172-11
172.6.3 DISCONNECT_SESSION Procedure 172-11
172.6.4 MODIFY_SERVICE Procedure 172-12
172.6.5 START_SERVICE Procedure 172-16
172.6.6 STOP_SERVICE Procedure 172-17

173 DBMS_SESSION

173.1 DBMS_SESSION Security Model 173-1
173.2 DBMS_SESSION Operational Notes 173-1
173.3 DBMS_SESSION Data Structures 173-1
173.3.1 DBMS_SERVICE INTEGER_ARRAY Table Type 173-1
173.3.2 DBMS_SERVICE LNAME_ARRAY Table Type 173-1
173.4 Summary of DBMS_SESSION Subprograms 173-2
173.4.1 CLEAR_ALL_CONTEXT Procedure 173-3
173.4.2 CLEAR_ALL_LOCAL_CONTEXTS Procedure 173-3
173.4.3 CLEAR_CONTEXT Procedure 173-3
173.4.4 CLEAR_IDENTIFIER Procedure 173-4
173.4.5 CLOSE_DATABASE_LINK Procedure 173-4
173.4.6 CURRENT_IS_ROLE_ENABLED Function 173-5
173.4.7 FREE_UNUSED_USER_MEMORY Procedure 173-5
173.4.8 GET_PACKAGE_MEMORY_UTILIZATION Procedure 173-7
173.4.9 IS_ROLE_ENABLED Function 173-8
173.4.10 IS_SESSION_ALIVE Function 173-8
173.4.11 LIST_CONTEXT Procedures 173-9
173.4.12 MODIFY_PACKAGE_STATE Procedure 173-9
173.4.13 RESET_PACKAGE Procedure 173-13
173.4.14 SESSION_IS_ROLE_ENABLED Function 173-14
173.4.15 SESSION_TRACE_DISABLE Procedure 173-14
173.4.16 SESSION _TRACE_ENABLE Procedure 173-15
173.4.17 SET_CONTEXT Procedure 173-15
173.4.18 SET_EDITION_DEFERRED Procedure 173-17
173.4.19 SET_IDENTIFIER Procedure 173-17
173.4.20 SET_NLS Procedure 173-18
173.4.21 SET_ROLE Procedure 173-18
173.4.22 SET_SQL_TRACE Procedure 173-18
173.4.23 SLEEP Procedure 173-19
173.4.24 SWITCH_CURRENT_CONSUMER_GROUP Procedure 173-19
173.4.25 UNIQUE_SESSION_ID Function 173-21
173.4.26 USE_DEFAULT_EDITION_ALWAYS procedure 173-21
173.4.27 USE_DEFAULT_EDITION_DEFERRED procedure 173-22

174 DBMS_SFW_ACL_ADMIN

174.1 DBMS_SFW_ACL_ADMIN Security Model 174-1
174.2 DBMS_SFW_ACL_ADMIN Operational Notes 174-1
174.3 DBMS_SFW_ACL_ADMIN Examples 174-2
174.4 Summary of DBMS_SFW_ACL_ADMIN Subprograms 174-3

174.4.1 COMMIT_ACL Procedure 174-3
174.4.2 IP_ADD_ACE Procedure 174-3
174.4.3 IP_ADD_PDB_ACE Procedure 174-4
174.4.4 IP_REMOVE_ACE Procedure 174-4
174.4.5 IP_REMOVE_ACL Procedure 174-5
174.4.6 IP_REMOVE_PDB_ACE Procedure 174-5
174.4.7 IP_REMOVE_PDB_ACL Procedure 174-6

175 DBMS_SHARED_POOL

175.1 DBMS_SHARED_POOL Overview 175-1
175.2 DBMS_SHARED_POOL Operational Notes 175-1
175.3 Summary of DBMS_SHARED_POOL Subprograms 175-1
175.3.1 ABORTED_REQUEST_THRESHOLD Procedure 175-2
175.3.2 KEEP Procedure 175-2
175.3.3 MARKHOT Procedure 175-4
175.3.4 PURGE Procedure 175-5
175.3.5 PURGE_ROW_CACHE Procedure 175-6
175.3.6 SIZES Procedure 175-7
175.3.7 UNKEEP Procedure 175-7
175.3.8 UNMARKHOT Procedure 175-8

176 DBMS_SHARDING_DIRECTORY

176.1 DBMS_SHARDING_DIRECTORY Overview 176-1
176.2 DBMS_SHARDING_DIRECTORY Security Model 176-1
176.3 DBMS_SHARDING_DIRECTORY Public Constants 176-2
176.4 Summary of DBMS_SHARDING_DIRECTORY Subprograms 176-2
176.4.1 addKeyToPartition Procedure 176-2
176.4.2 addKeyToPartitionCommit Procedure 176-3
176.4.3 flagKeyForSplit Procedure 176-4
176.4.4 removeKey Procedure 176-4
176.4.5 removeKeyCommit Procedure 176-5
176.4.6 setAssignmentRule 176-5

177 DBMS_SODA

177.1 DBMS_SODA Security Model 177-1
177.2 Summary of DBMS_SODA Subprograms 177-1
177.2.1 CREATE_COLLECTION Function 177-2
177.2.2 DROP_COLLECTION Function 177-3
177.2.3 GET_AS_OF_SCN Function 177-4

177.2.4 GET_AS_OF_TIMESTAMP Function 177-4
177.2.5 GET_DEFAULT_METADATA_AS_CLOB Function 177-4
177.2.6 GET_DEFAULT_METADATA_AS_VARCHAR2 Function 177-4
177.2.7 LIST_COLLECTION_NAMES Function 177-4
177.2.8 OPEN_COLLECTION Function 177-5
177.3 Summary of SODA Online Redefinition Subprograms 177-6
177.3.1 ABORT_REDEF_COLLECTION Procedure 177-7
177.3.2 CAN_REDEF_COLLECTION Procedure 177-7
177.3.3 COPY_COLLECTION_DEPENDENTS Procedure 177-7
177.3.4 CREATE_INTERIM_COLLECTION Procedure 177-8
177.3.5 FINISH_REDEF_COLLECTION Procedure 177-9
177.3.6 START_REDEF_COLLECTION Procedure 177-9
177.3.7 SYNC_INTERIM_COLLETION Procedure 177-10
177.3.8 TO_UUID Function 177-10

178 DBMS_SPACE

178.1 DBMS_SPACE Security Model 178-1
178.2 DBMS_SPACE Data Structures 178-1
178.2.1 DBMS_SPACE CREATE_TABLE_COST_COLINFO Object Type 178-1
178.2.2 DBMS_SPACE ASA_RECO_ROW Record Type 178-2
178.2.3 DBMS_SPACE ASA_RECO_ROW_TB Table Type 178-3
178.3 Summary of DBMS_SPACE Subprograms 178-3
178.3.1 DBMS_SPACE ASA_RECOMMENDATIONS Function 178-3
178.3.2 CREATE_INDEX_COST Procedure 178-5
178.3.3 CREATE_TABLE_COST Procedures 178-5
178.3.4 FREE_BLOCKS Procedure 178-7
178.3.5 ISDATAFILEDROPPABLE_NAME Procedure 178-8
178.3.6 OBJECT_DEPENDENT_SEGMENTS Function 178-9
178.3.7 OBJECT_GROWTH_TREND Function 178-10
178.3.8 SHRINK_TABLESPACE Procedure 178-11
178.3.9 SPACE_USAGE Procedures 178-13
178.3.10 UNUSED_SPACE Procedure 178-16

179 DBMS_SPACE_ADMIN

179.1 DBMS_SPACE_ADMIN Security Model 179-1
179.2 DBMS_SPACE_ADMIN Constants 179-1
179.3 DBMS_SPACE_ADMIN Operational Notes 179-3
179.4 Summary of DBMS_SPACE_ADMIN Subprograms 179-3
179.4.1 ASSM_SEGMENT_VERIFY Procedure 179-4
179.4.2 ASSM_TABLESPACE_VERIFY Procedure 179-5

179.4.3 DROP_EMPTY_SEGMENTS Procedure 179-6
179.4.4 GET_SEGADV_ATTRIB Procedure 179-7
179.4.5 MATERIALIZE_DEFERRED_SEGMENTS Procedure 179-8
179.4.6 SEGMENT_CORRUPT Procedure 179-8
179.4.7 SEGMENT_DROP_CORRUPT Procedure 179-9
179.4.8 SEGMENT_DUMP Procedure 179-10
179.4.9 SEGMENT_VERIFY Procedure 179-11
179.4.10 SET_SEGADV_ATTRIB Procedure 179-11
179.4.11 TABLESPACE_FIX_BITMAPS Procedure 179-12
179.4.12 TABLESPACE_FIX_SEGMENT_STATES Procedure 179-13
179.4.13 TABLESPACE_MIGRATE_FROM_LOCAL Procedure 179-13
179.4.14 TABLESPACE_MIGRATE_TO_LOCAL Procedure 179-14
179.4.15 TABLESPACE_REBUILD_BITMAPS Procedure 179-15
179.4.16 TABLESPACE_REBUILD_QUOTAS Procedure 179-15
179.4.17 TABLESPACE_RELOCATE_BITMAPS Procedure 179-16
179.4.18 TABLESPACE_VERIFY Procedure 179-17

180 DBMS_SPD

180.1 DBMS_SPD Overview 180-1
180.2 DBMS_SPD Security Model 180-1
180.3 DBMS_SPD Views 180-1
180.4 Summary of DBMS_SPD Subprograms 180-1
180.4.1 ALTER_SQL_PLAN_DIRECTIVE Procedure 180-2
180.4.2 CREATE_STGTAB_DIRECTIVE Procedure 180-3
180.4.3 DROP_SQL_PLAN_DIRECTIVE Procedure 180-4
180.4.4 FLUSH_SQL_PLAN_DIRECTIVE Procedure 180-4
180.4.5 GET_PREFS Function 180-5
180.4.6 PACK_STGTAB_DIRECTIVE Function 180-5
180.4.7 SET_PREFS Procedure 180-7
180.4.8 UNPACK_STGTAB_DIRECTIVE Function 180-8

181 DBMS_SPM

181.1 DBMS_SPM Overview 181-1
181.2 DBMS_SPM Security Model 181-2
181.3 DBMS_SPM Constants 181-2
181.4 DBMS_SPM Examples 181-2
181.5 DBMS_SPM Data Structures 181-2
181.5.1 DBMS_SPM NAMELIST Table Type 181-3
181.6 Summary of DBMS_SPM Subprograms 181-3
181.6.1 ACCEPT_SQL_PLAN_BASELINE Procedure 181-4

181.6.2 ADD_VERIFIED_SQL_PLAN_BASELINE Function 181-5
181.6.3 ALTER_SQL_PLAN_BASELINE Function 181-8
181.6.4 CANCEL_EVOLVE_TASK Procedure 181-10
181.6.5 CONFIGURE Procedure 181-10
181.6.6 CREATE_EVOLVE_TASK Function 181-13
181.6.7 CREATE_STGTAB_BASELINE Procedure 181-14
181.6.8 DROP_EVOLVE_TASK Procedure 181-15
181.6.9 DROP_SQL_PLAN_BASELINE Function 181-15
181.6.10 EVOLVE_SQL_PLAN_BASELINE Function 181-16
181.6.11 EXECUTE_EVOLVE_TASK Function 181-17
181.6.12 IMPLEMENT_EVOLVE_TASK Function 181-18
181.6.13 INTERRUPT_EVOLVE_TASK Procedure 181-18
181.6.14 LOAD_PLANS_FROM_AWR Function 181-19
181.6.15 LOAD_PLANS_FROM_CURSOR_CACHE Functions 181-20
181.6.16 LOAD_PLANS_FROM_SQLSET Function 181-21
181.6.17 MIGRATE_STORED_OUTLINE Functions 181-22
181.6.18 PACK_STGTAB_BASELINE Function 181-24
181.6.19 RESET_EVOLVE_TASK Procedure 181-25
181.6.20 RESUME_EVOLVE_TASK Procedure 181-25
181.6.21 REPORT_AUTO_EVOLVE_TASK Function 181-25
181.6.22 REPORT_EVOLVE_TASK Function 181-26
181.6.23 SET_EVOLVE_TASK_PARAMETER Procedure 181-27
181.6.24 UNPACK_STGTAB_BASELINE Function 181-29

182 DBMS_SQL

182.1 DBMS_SQL Overview 182-1
182.2 DBMS_SQL Security Model 182-2
182.3 DBMS_SQL Constants 182-3
182.4 DBMS_SQL Operational Notes 182-4
182.5 DBMS_SQL Execution Flow 182-4
182.5.1 OPEN_CURSOR 182-5
182.5.2 PARSE 182-5
182.5.3 BIND_VARIABLE, BIND_VARIABLE_PKG or BIND_ARRAY 182-7
182.5.4 DEFINE_COLUMN, DEFINE_COLUMN_LONG, or DEFINE_ARRAY 182-7
182.5.5 EXECUTE 182-7
182.5.6 FETCH_ROWS or EXECUTE_AND_FETCH 182-7
182.5.7 VARIABLE_VALUE, VARIABLE_VALUE_PKG, COLUMN_VALUE, or COLUMN_VALUE_LONG 182-8
182.5.8 CLOSE_CURSOR 182-8
182.6 DBMS_SQL Exceptions 182-8
182.7 DBMS_SQL Examples 182-9

182.8 DBMS_SQL Data Structures 182-15
182.8.1 DBMS_SQL DESC_REC Record Type 182-16
182.8.2 DBMS_SQL DESC_REC2 Record Type 182-17
182.8.3 DBMS_SQL DESC_REC3 Record Type 182-18
182.8.4 DBMS_SQL DESC_REC4 Record Type 182-19
182.8.5 DBMS_SQL BFILE_TABLE Table Type 182-20
182.8.6 DBMS_SQL BINARY_DOUBLE_TABLE Table Type 182-20
182.8.7 DBMS_SQL BINARY_FLOAT_TABLE Table Type 182-20
182.8.8 DBMS_SQL BLOB_TABLE Table Type 182-21
182.8.9 DBMS_SQL BOOLEAN_TABLE Table Type 182-21
182.8.10 DBMS_SQL CLOB_TABLE Table Type 182-21
182.8.11 DBMS_SQL DATE_TABLE Table Type 182-21
182.8.12 DBMS_SQL DESC_TAB Table Type 182-21
182.8.13 DBMS_SQL DESC_TAB2 Table Type 182-21
182.8.14 DBMS_SQL DESC_TAB3 Table Type 182-22
182.8.15 DBMS_SQL DESC_TAB4 Table Type 182-22
182.8.16 DBMS_SQL INTERVAL_DAY_TO_SECOND_TABLE Table Type 182-22
182.8.17 DBMS_SQL INTERVAL_YEAR_TO_MONTH_TABLE Table Type 182-22
182.8.18 DBMS_SQL JSON_TABLE Table Type 182-23
182.8.19 DBMS_SQL NUMBER_TABLE Table Type 182-23
182.8.20 DBMS_SQL TIME_TABLE Table Type 182-23
182.8.21 DBMS_SQL TIME_WITH_TIME_ZONE_TABLE Table Type 182-23
182.8.22 DBMS_SQL TIMESTAMP_TABLE Table Type 182-23
182.8.23 DBMS_SQL TIMESTAMP_WITH_LTZ_TABLE Table Type 182-24
182.8.24 DBMS_SQL TIMESTAMP_WITH_TIME_ZONE_TABLE Table Type 182-24
182.8.25 DBMS_SQL UROWID_TABLE Table Type 182-24
182.8.26 DBMS_SQL VARCHAR2_TABLE Table Type 182-24
182.8.27 DBMS_SQL VARCHAR2A Table Type 182-24
182.8.28 DBMS_SQL VARCHAR2S Table Type 182-25
182.8.29 DBMS_SQL VECTOR Table Type 182-25
182.9 Summary of DBMS_SQL Subprograms 182-25
182.9.1 BIND_ARRAY Procedures 182-27
182.9.2 BIND_VARIABLE Procedures 182-32
182.9.3 BIND_VARIABLE_PKG Procedure 182-34
182.9.4 CLOSE_CURSOR Procedure 182-35
182.9.5 COLUMN_VALUE Procedure 182-36
182.9.6 COLUMN_VALUE_LONG Procedure 182-39
182.9.7 DEFINE_ARRAY Procedure 182-40
182.9.8 DEFINE_COLUMN Procedures 182-43
182.9.9 DEFINE_COLUMN_CHAR Procedure 182-45
182.9.10 DEFINE_COLUMN_LONG Procedure 182-46
182.9.11 DEFINE_COLUMN_RAW Procedure 182-46

182.9.12 DEFINE_COLUMN_ROWID Procedure 182-47
182.9.13 DESCRIBE_COLUMNS Procedure 182-48
182.9.14 DESCRIBE_COLUMNS2 Procedure 182-49
182.9.15 DESCRIBE_COLUMNS3 Procedure 182-50
182.9.16 EXECUTE Function 182-51
182.9.17 EXECUTE_AND_FETCH Function 182-52
182.9.18 FETCH_ROWS Function 182-52
182.9.19 GET_NEXT_RESULT Procedures 182-53
182.9.20 IS_OPEN Function 182-54
182.9.21 LAST_ERROR_POSITION Function 182-55
182.9.22 LAST_ROW_COUNT Function 182-55
182.9.23 LAST_ROW_ID Function 182-55
182.9.24 LAST_SQL_FUNCTION_CODE Function 182-56
182.9.25 OPEN_CURSOR Functions 182-56
182.9.26 PARSE Procedures 182-57
182.9.27 RETURN_RESULT Procedures 182-61
182.9.28 TO_CURSOR_NUMBER Function 182-62
182.9.29 TO_REFCURSOR Function 182-63
182.9.30 VARIABLE_VALUE Procedures 182-65
182.9.31 VARIABLE_VALUE_PKG Procedure 182-66

183 DBMS_SQL_FIREWALL

183.1 DBMS_SQL_FIREWALL Overview 183-1
183.2 DBMS_SQL_FIREWALL Security Model 183-1
183.3 DBMS_SQL_FIREWALL Constants 183-2
183.4 Summary of DBMS_SQL_FIREWALL Subprograms 183-3
183.4.1 ADD_ALLOWED_CONTEXT Procedure 183-4
183.4.2 APPEND_ALLOW_LIST Procedure 183-5
183.4.3 APPEND_ALLOW_LIST_SINGLE_SQL Procedure 183-6
183.4.4 CREATE_CAPTURE Procedure 183-8
183.4.5 DELETE_ALLOWED_CONTEXT Procedure 183-9
183.4.6 DELETE_ALLOWED_SQL Procedure 183-10
183.4.7 DISABLE Procedure 183-10
183.4.8 DISABLE_ALLOW_LIST Procedure 183-11
183.4.9 DROP_ALLOW_LIST Procedure 183-11
183.4.10 DROP_CAPTURE Procedure 183-12
183.4.11 ENABLE Procedure 183-12
183.4.12 ENABLE_ALLOW_LIST Procedure 183-13
183.4.13 EXCLUDE Procedure 183-14
183.4.14 EXPORT_ALLOW_LIST Procedure 183-14
183.4.15 FLUSH_LOGS Procedure 183-15

183.4.16 GENERATE_ALLOW_LIST Procedure 183-16
183.4.17 IMPORT_ALLOW_LIST Procedure 183-17
183.4.18 INCLUDE Procedure 183-18
183.4.19 PURGE_LOG Procedure 183-18
183.4.20 START_CAPTURE Procedure 183-19
183.4.21 STOP_CAPTURE Procedure 183-20
183.4.22 UPDATE_ALLOW_LIST_ENFORCEMENT Procedure 183-20

184 DBMS_SQL_MONITOR

184.1 DBMS_SQL_MONITOR Overview 184-1
184.2 DBMS_SQL_MONITOR Security Model 184-1
184.3 DBMS_SQL_MONITOR Constants 184-2
184.4 Summary of DBMS_SQL_MONITOR Subprograms 184-2
184.4.1 BEGIN_OPERATION Function 184-2
184.4.2 END_OPERATION Procedure 184-3
184.4.3 REPORT_SQL_MONITOR Function 184-4
184.4.4 REPORT_SQL_MONITOR_XML Function 184-8
184.4.5 REPORT_SQL_MONITOR_LIST Function 184-8
184.4.6 REPORT_SQL_MONITOR_LIST_XML Function 184-10

185 DBMS_SQL_TRANSLATOR

185.1 DBMS_SQL_TRANSLATOR Security Model 185-1
185.2 DBMS_SQL_TRANSLATOR Constants 185-1
185.3 DBMS_SQL_TRANSLATOR Operational Notes 185-2
185.4 DBMS_SQL_TRANSLATOR Exceptions 185-4
185.5 DBMS_SQL_TRANSLATOR Examples 185-4
185.6 Summary of DBMS_SQL_TRANSLATOR Subprograms 185-4
185.6.1 CREATE_PROFILE Procedure 185-5
185.6.2 DEREGISTER_SQL_TRANSLATION Procedure 185-6
185.6.3 DEREGISTER_ERROR_TRANSLATION Procedure 185-7
185.6.4 DROP_PROFILE Procedure 185-8
185.6.5 ENABLE_ERROR_TRANSLATION Procedure 185-8
185.6.6 ENABLE_SQL_TRANSLATION Procedure 185-9
185.6.7 EXPORT_PROFILE Procedure 185-10
185.6.8 IMPORT_PROFILE Procedure 185-11
185.6.9 REGISTER_ERROR_TRANSLATION Procedure 185-12
185.6.10 REGISTER_SQL_TRANSLATION Procedure 185-13
185.6.11 SET_ATTRIBUTE Procedure 185-14
185.6.12 SQL_HASH Function 185-15
185.6.13 SQL_ID Function 185-16

185.6.14 TRANSLATE_ERROR Procedure 185-17
185.6.15 TRANSLATE_SQL Procedure 185-18

186 DBMS_SQLDIAG

186.1 DBMS_SQLDIAG Overview 186-1
186.2 DBMS_SQLDIAG Security Model 186-1
186.3 DBMS_SQLDIAG Constants 186-1
186.4 Summary of DBMS_SQLDIAG Subprograms 186-4
186.4.1 ACCEPT_SQL_PATCH Function & Procedure 186-5
186.4.2 ALTER_SQL_PATCH Procedure 186-7
186.4.3 CANCEL_DIAGNOSIS_TASK Procedure 186-7
186.4.4 CREATE_DIAGNOSIS_TASK Functions 186-8
186.4.5 CREATE_SQL_PATCH Function 186-9
186.4.6 CREATE_STGTAB_SQLPATCH Procedure 186-10
186.4.7 DROP_DIAGNOSIS_TASK Procedure 186-11
186.4.8 DROP_SQL_PATCH Procedure 186-11
186.4.9 EXECUTE_DIAGNOSIS_TASK Procedure 186-11
186.4.10 EXPLAIN_SQL_TESTCASE Function 186-12
186.4.11 EXPORT_SQL_TESTCASE Procedures 186-12
186.4.12 EXPORT_SQL_TESTCASE_DIR_BY_INC Function 186-16
186.4.13 EXPORT_SQL_TESTCASE_DIR_BY_TXT Function 186-18
186.4.14 GET_FIX_CONTROL Function 186-19
186.4.15 GET_SQL Function 186-20
186.4.16 IMPORT_SQL_TESTCASE Procedures 186-20
186.4.17 INCIDENTID_2_SQL Procedure 186-22
186.4.18 INTERRUPT_DIAGNOSIS_TASK Procedure 186-23
186.4.19 LOAD_SQLSET_FROM_TCB Function 186-23
186.4.20 PACK_STGTAB_SQLPATCH Procedure 186-23
186.4.21 REPLAY_SQL_TESTCASE Function 186-24
186.4.22 REPORT_SQL Function 186-26
186.4.23 REPORT_DIAGNOSIS_TASK Function 186-27
186.4.24 RESET_DIAGNOSIS_TASK Procedure 186-28
186.4.25 RESUME_DIAGNOSIS_TASK Procedure 186-28
186.4.26 SET_DIAGNOSIS_TASK_PARAMETER Procedure 186-28
186.4.27 SQL_DIAGNOSE_AND_REPAIR Function 186-29
186.4.28 UNPACK_STGTAB_SQLPATCH Procedure 186-30

187 DBMS_SQLPA

187.1 DBMS_SQLPA Overview 187-1
187.2 DBMS_SQLPA Security Model 187-1

187.3 Summary of DBMS_SQLPA Subprograms 187-1
187.3.1 CANCEL_ANALYSIS_TASK Procedure 187-2
187.3.2 CREATE_ANALYSIS_TASK Functions 187-3
187.3.3 DROP_ANALYSIS_TASK Procedure 187-5
187.3.4 EXECUTE_ANALYSIS_TASK Function & Procedure 187-6
187.3.5 INTERRUPT_ANALYSIS_TASK Procedure 187-8
187.3.6 REPORT_ANALYSIS_TASK Function 187-8
187.3.7 RESET_ANALYSIS_TASK Procedure 187-10
187.3.8 RESUME_ANALYSIS_TASK Procedure 187-10
187.3.9 SET_ANALYSIS_TASK_PARAMETER Procedures 187-11
187.3.10 SET_ANALYSIS_DEFAULT_PARAMETER Procedures 187-15

188 DBMS_SQLQ

188.1 DBMS_SQLQ Overview 188-1
188.2 Summary of DBMS_SQLQ Subprograms 188-1
188.2.1 ALTER_QUARANTINE Procedure 188-2
188.2.2 CREATE_QUARANTINE_BY_SQL_ID Function 188-3
188.2.3 CREATE_QUARANTINE_BY_SQL_TEXT Function 188-4
188.2.4 CREATE_STGTAB_QUARANTINE Procedure 188-5
188.2.5 DROP_QUARANTINE Procedure 188-6
188.2.6 GET_PARAM_VALUE_QUARANTINE Function 188-6
188.2.7 PACK_STGTAB_QUARANTINE Function 188-7
188.2.8 UNPACK_STGTAB_QUARANTINE Function 188-8

189 DBMS_SQLSET

189.1 DBMS_SQLSET Overview 189-1
189.2 DBMS_SQLSET Security Model 189-2
189.3 DBMS_SQLSET Data Structures 189-2
189.3.1 SQLSET_ROW Object Type 189-2
189.4 Summary of DBMS_SQLSET Subprograms 189-4
189.4.1 ADD_REFERENCE Function 189-5
189.4.2 CAPTURE_CURSOR_CACHE Procedure 189-6
189.4.3 CREATE_SQLSET Procedure and Function 189-8
189.4.4 CREATE_STGTAB Procedure 189-9
189.4.5 DELETE_SQLSET Procedure 189-10
189.4.6 DROP_SQLSET Procedure 189-10
189.4.7 LOAD_SQLSET Procedure 189-11
189.4.8 PACK_STGTAB Procedure 189-15
189.4.9 REMAP_STGTAB Procedure 189-17
189.4.10 REMOVE_REFERENCE Procedure 189-18

189.4.11 SELECT_CURSOR_CACHE Function 189-19
189.4.12 SELECT_SQL_TRACE Function 189-24
189.4.13 SELECT_SQLPA_TASK Function 189-27
189.4.14 SELECT_SQLSET Function 189-28
189.4.15 SELECT_WORKLOAD_REPOSITORY Function 189-30
189.4.16 UNPACK_STGTAB Procedure 189-32
189.4.17 UPDATE_SQLSET Procedures 189-33

190 DBMS_SQLTUNE

190.1 DBMS_SQLTUNE Overview 190-1
190.2 DBMS_SQLTUNE Security Model 190-4
190.3 DBMS_SQLTUNE Data Structures 190-5
190.3.1 SQLSET_ROW Object Type 190-5
190.4 DBMS_SQLTUNE Subprogram Groups 190-7
190.4.1 DBMS_SQLTUNE SQL Tuning Advisor Subprograms 190-7
190.4.2 DBMS_SQLTUNE SQL Profile Subprograms 190-8
190.4.3 DBMS_SQLTUNE SQL Tuning Set Subprograms 190-8
190.4.4 DBMS_SQLTUNE Real-Time SQL Monitoring Subprograms 190-9
190.4.5 DBMS_SQLTUNE SQL Performance Reporting Subprograms 190-10
190.5 Summary of DBMS_SQLTUNE Subprograms 190-10
190.5.1 ACCEPT_ALL_SQL_PROFILES Procedure 190-13
190.5.2 ACCEPT_SQL_PROFILE Procedure and Function 190-15
190.5.3 ADD_SQLSET_REFERENCE Function 190-19
190.5.4 ALTER_SQL_PROFILE Procedure 190-20
190.5.5 CANCEL_TUNING_TASK Procedure 190-21
190.5.6 CAPTURE_CURSOR_CACHE_SQLSET Procedure 190-22
190.5.7 CREATE_SQL_PLAN_BASELINE Procedure 190-24
190.5.8 CREATE_SQLSET Procedure and Function 190-25
190.5.9 CREATE_STGTAB_SQLPROF Procedure 190-25
190.5.10 CREATE_STGTAB_SQLSET Procedure 190-26
190.5.11 CREATE_TUNING_TASK Functions 190-28
190.5.12 DELETE_SQLSET Procedure 190-34
190.5.13 DROP_SQL_PROFILE Procedure 190-35
190.5.14 DROP_SQLSET Procedure 190-35
190.5.15 DROP_TUNING_TASK Procedure 190-36
190.5.16 EXECUTE_TUNING_TASK Function and Procedure 190-37
190.5.17 IMPLEMENT_TUNING_TASK Procedure 190-38
190.5.18 INTERRUPT_TUNING_TASK Procedure 190-39
190.5.19 LOAD_SQLSET Procedure 190-40
190.5.20 PACK_STGTAB_SQLPROF Procedure 190-44
190.5.21 PACK_STGTAB_SQLSET Procedure 190-45

190.5.22 REMAP_STGTAB_SQLPROF Procedure 190-46
190.5.23 REMAP_STGTAB_SQLSET Procedure 190-47
190.5.24 REMOVE_SQLSET_REFERENCE Procedure 190-49
190.5.25 REPORT_AUTO_TUNING_TASK Function 190-50
190.5.26 REPORT_SQL_DETAIL Function 190-51
190.5.27 REPORT_SQL_HISTORY Function 190-55
190.5.28 REPORT_SQL_HISTORY_LIST Function 190-56
190.5.29 REPORT_SQL_MONITOR Function 190-56
190.5.30 REPORT_SQL_MONITOR_LIST Function 190-61
190.5.31 REPORT_TUNING_TASK Function 190-64
190.5.32 REPORT_TUNING_TASK_XML Function 190-66
190.5.33 RESET_TUNING_TASK Procedure 190-67
190.5.34 RESUME_TUNING_TASK Procedure 190-68
190.5.35 SCHEDULE_TUNING_TASK Function 190-69
190.5.36 SCRIPT_TUNING_TASK Function 190-71
190.5.37 SELECT_CURSOR_CACHE Function 190-73
190.5.38 SELECT_SQL_TRACE Function 190-77
190.5.39 SELECT_SQLPA_TASK Function 190-80
190.5.40 SELECT_SQLSET Function 190-81
190.5.41 SELECT_WORKLOAD_REPOSITORY Function 190-84
190.5.42 SET_TUNING_TASK_PARAMETER Procedures 190-86
190.5.43 SQLTEXT_TO_SIGNATURE Function 190-88
190.5.44 UNPACK_STGTAB_SQLPROF Procedure 190-89
190.5.45 UNPACK_STGTAB_SQLSET Procedure 190-90
190.5.46 UPDATE_SQLSET Procedures 190-91

191 DBMS_STAT_FUNCS

191.1 Summary of DBMS_STAT_FUNCS Subprograms 191-1
191.1.1 EXPONENTIAL_DIST_FIT Procedure 191-1
191.1.2 NORMAL_DIST_FIT Procedure 191-2
191.1.3 POISSON_DIST_FIT Procedure 191-2
191.1.4 SUMMARY Procedure 191-3
191.1.5 UNIFORM_DIST_FIT Procedure 191-4
191.1.6 WEIBULL_DIST_FIT Procedure 191-5

192 DBMS_STATS

192.1 DBMS_STATS Overview 192-1
192.2 DBMS_STATS Deprecated Subprograms 192-2
192.3 DBMS_STATS Types 192-2
192.4 DBMS_STATS Constants 192-4

192.5 DBMS_STATS Operational Notes 192-4
192.6 DBMS_STATS Data Structures 192-11
192.6.1 DBMS_STATS STAT_REC Record Type 192-11
192.7 Summary of DBMS_STATS Subprograms 192-12
192.7.1 ALTER_STATS_HISTORY_RETENTION Procedure 192-18
192.7.2 CANCEL_ADVISOR_TASK Procedure 192-18
192.7.3 CONFIGURE_ADVISOR_FILTER Function 192-20
192.7.4 CONFIGURE_ADVISOR_OBJ_FILTER Function 192-25
192.7.5 CONFIGURE_ADVISOR_OPR_FILTER Functions 192-26
192.7.6 CONFIGURE_ADVISOR_RULE_FILTER Function 192-29
192.7.7 CONVERT_RAW_VALUE Procedures 192-30
192.7.8 CONVERT_RAW_VALUE_NVARCHAR Procedure 192-31
192.7.9 CONVERT_RAW_VALUE_ROWID Procedure 192-32
192.7.10 COPY_TABLE_STATS Procedure 192-32
192.7.11 CREATE_ADVISOR_TASK Function 192-34
192.7.12 CREATE_EXTENDED_STATS Function 192-35
192.7.13 CREATE_STAT_TABLE Procedure 192-37
192.7.14 DELETE_COLUMN_STATS Procedure 192-37
192.7.15 DELETE_DATABASE_PREFS Procedure 192-39
192.7.16 DELETE_DATABASE_STATS Procedure 192-43
192.7.17 DELETE_DICTIONARY_STATS Procedure 192-44
192.7.18 DELETE_FIXED_OBJECTS_STATS Procedure 192-46
192.7.19 DELETE_INDEX_STATS Procedure 192-47
192.7.20 DELETE_PENDING_STATS Procedure 192-48
192.7.21 DELETE_PLSQL_PREFS Procedure 192-49
192.7.22 DELETE_PROCESSING_RATE Procedure 192-50
192.7.23 DELETE_SCHEMA_PREFS Procedure 192-50
192.7.24 DELETE_SCHEMA_STATS Procedure 192-56
192.7.25 DELETE_SYSTEM_STATS Procedure 192-57
192.7.26 DELETE_TABLE_PREFS Procedure 192-58
192.7.27 DELETE_TABLE_STATS Procedure 192-64
192.7.28 DIFF_TABLE_STATS_IN_HISTORY Function 192-66
192.7.29 DIFF_TABLE_STATS_IN_PENDING Function 192-66
192.7.30 DIFF_TABLE_STATS_IN_STATTAB Function 192-67
192.7.31 DROP_ADVISOR_TASK Procedure 192-68
192.7.32 DROP_EXTENDED_STATS Procedure 192-69
192.7.33 DROP_STAT_TABLE Procedure 192-70
192.7.34 EXECUTE_ADVISOR_TASK Function 192-70
192.7.35 EXPORT_COLUMN_STATS Procedure 192-72
192.7.36 EXPORT_DATABASE_PREFS Procedure 192-73
192.7.37 EXPORT_DATABASE_STATS Procedure 192-74
192.7.38 EXPORT_DICTIONARY_STATS Procedure 192-75

192.7.39 EXPORT_FIXED_OBJECTS_STATS Procedure 192-76
192.7.40 EXPORT_INDEX_STATS Procedure 192-76
192.7.41 EXPORT_PENDING_STATS Procedure 192-77
192.7.42 EXPORT_SCHEMA_PREFS Procedure 192-78
192.7.43 EXPORT_SCHEMA_STATS Procedure 192-79
192.7.44 EXPORT_SYSTEM_STATS Procedure 192-80
192.7.45 EXPORT_TABLE_PREFS Procedure 192-80
192.7.46 EXPORT_TABLE_STATS Procedure 192-81
192.7.47 FLUSH_DATABASE_MONITORING_INFO Procedure 192-83
192.7.48 GATHER_DATABASE_STATS Procedures 192-83
192.7.49 GATHER_DICTIONARY_STATS Procedure 192-87
192.7.50 GATHER_FIXED_OBJECTS_STATS Procedure 192-91
192.7.51 GATHER_INDEX_STATS Procedure 192-92
192.7.52 GATHER_PROCESSING_RATE Procedure 192-94
192.7.53 GATHER_SCHEMA_STATS Procedures 192-94
192.7.54 GATHER_SYSTEM_STATS Procedure 192-99
192.7.55 GATHER_TABLE_STATS Procedure 192-101
192.7.56 GENERATE_STATS Procedure 192-105
192.7.57 GET_ADVISOR_OPR_FILTER Procedure 192-106
192.7.58 GET_ADVISOR_RECS Function 192-107
192.7.59 GET_COLUMN_STATS Procedures 192-108
192.7.60 GET_INDEX_STATS Procedures 192-110
192.7.61 GET_PARAM Function 192-113
192.7.62 GET_PLSQL_PREFS Procedure 192-113
192.7.63 GET_PREFS Function 192-114
192.7.64 GET_STATS_HISTORY_AVAILABILITY Function 192-124
192.7.65 GET_STATS_HISTORY_RETENTION Function 192-124
192.7.66 GET_SYSTEM_STATS Procedure 192-124
192.7.67 GET_TABLE_STATS Procedure 192-126
192.7.68 IMPLEMENT_ADVISOR_TASK Function 192-128
192.7.69 IMPORT_COLUMN_STATS Procedure 192-130
192.7.70 IMPORT_DATABASE_PREFS Procedure 192-132
192.7.71 IMPORT_DATABASE_STATS Procedure 192-132
192.7.72 IMPORT_DICTIONARY_STATS Procedure 192-134
192.7.73 IMPORT_FIXED_OBJECTS_STATS Procedure 192-135
192.7.74 IMPORT_INDEX_STATS Procedure 192-136
192.7.75 IMPORT_SCHEMA_PREFS Procedure 192-137
192.7.76 IMPORT_SCHEMA_STATS Procedure 192-138
192.7.77 IMPORT_SYSTEM_STATS Procedure 192-139
192.7.78 IMPORT_TABLE_PREFS Procedure 192-140
192.7.79 IMPORT_TABLE_STATS Procedure 192-141
192.7.80 INTERRUPT_ADVISOR_TASK Procedure 192-142

192.7.81 LOCK_PARTITION_STATS Procedure 192-144
192.7.82 LOCK_SCHEMA_STATS Procedure 192-144
192.7.83 LOCK_TABLE_STATS Procedure 192-145
192.7.84 MERGE_COL_USAGE Procedure 192-146
192.7.85 PREPARE_COLUMN_VALUES Procedures 192-147
192.7.86 PREPARE_COLUMN_VALUES_NVARCHAR Procedure 192-149
192.7.87 PREPARE_COLUMN_VALUES_ROWID Procedure 192-151
192.7.88 PUBLISH_PENDING_STATS Procedure 192-153
192.7.89 PURGE_STATS Procedure 192-154
192.7.90 REMAP_STAT_TABLE Procedure 192-154
192.7.91 REPORT_ADVISOR_TASK Function 192-155
192.7.92 REPORT_COL_USAGE Function 192-159
192.7.93 REPORT_GATHER_AUTO_STATS Function 192-160
192.7.94 REPORT_GATHER_DATABASE_STATS Functions 192-162
192.7.95 REPORT_GATHER_DICTIONARY_STATS Functions 192-167
192.7.96 REPORT_GATHER_FIXED_OBJ_STATS Function 192-170
192.7.97 REPORT_GATHER_SCHEMA_STATS Functions 192-173
192.7.98 REPORT_GATHER_TABLE_STATS Function 192-178
192.7.99 REPORT_SINGLE_STATS_OPERATION Function 192-183
192.7.100 REPORT_STATS_OPERATIONS Function 192-185
192.7.101 RESET_ADVISOR_TASK Procedure 192-187
192.7.102 RESET_COL_USAGE Procedure 192-188
192.7.103 RESET_GLOBAL_PLSQL_PREFS_DEF Procedure 192-188
192.7.104 RESET_PARAM_DEFAULTS Procedure 192-189
192.7.105 RESTORE_DATABASE_STATS Procedure 192-189
192.7.106 RESTORE_DICTIONARY_STATS Procedure 192-190
192.7.107 RESTORE_FIXED_OBJECTS_STATS Procedure 192-191
192.7.108 RESTORE_SCHEMA_STATS Procedure 192-191
192.7.109 RESTORE_SYSTEM_STATS Procedure 192-192
192.7.110 RESTORE_TABLE_STATS Procedure 192-193
192.7.111 RESUME_ADVISOR_TASK Procedure 192-194
192.7.112 SCRIPT_ADVISOR_TASK Function 192-195
192.7.113 SEED_COL_USAGE Procedure 192-198
192.7.114 SET_ADVISOR_TASK_PARAMETER Procedure 192-199
192.7.115 SET_COLUMN_STATS Procedures 192-200
192.7.116 SET_DATABASE_PREFS Procedure 192-202
192.7.117 SET_GLOBAL_PREFS Procedure 192-209
192.7.118 SET_INDEX_STATS Procedures 192-220
192.7.119 SET_PARAM Procedure 192-223
192.7.120 SET_PLSQL_PREFS Procedure 192-224
192.7.121 SET_GLOBAL_PLSQL_PREFS Procedure 192-226
192.7.122 SET_PROCESSING_RATE Procedure 192-227

192.7.123 SET_SCHEMA_PREFS Procedure 192-229
192.7.124 SET_SYSTEM_STATS Procedure 192-236
192.7.125 SET_TABLE_PREFS Procedure 192-237
192.7.126 SET_TABLE_STATS Procedure 192-246
192.7.127 SHOW_EXTENDED_STATS_NAME Function 192-249
192.7.128 TRANSFER_STATS Procedure 192-249
192.7.129 UNLOCK_PARTITION_STATS Procedure 192-250
192.7.130 UNLOCK_SCHEMA_STATS Procedure 192-251
192.7.131 UNLOCK_TABLE_STATS Procedure 192-251
192.7.132 UPGRADE_STAT_TABLE Procedure 192-252

193 DBMS_STORAGE_MAP

193.1 DBMS_STORAGE_MAP Overview 193-1
193.2 DBMS_STORAGE_MAP Operational Notes 193-2
193.3 Summary of DBMS_STORAGE_MAP Subprograms 193-2
193.3.1 DROP_ALL Function 193-2
193.3.2 DROP_ELEMENT Function 193-3
193.3.3 DROP_FILE Function 193-3
193.3.4 LOCK_MAP Procedure 193-4
193.3.5 MAP_ALL Function 193-4
193.3.6 MAP_ELEMENT Function 193-4
193.3.7 MAP_FILE Function 193-5
193.3.8 MAP_OBJECT Function 193-6
193.3.9 RESTORE Function 193-6
193.3.10 SAVE Function 193-7
193.3.11 UNLOCK_MAP Procedure 193-7

194 DBMS_SYNC_REFRESH

194.1 DBMS_SYNC_REFRESH Overview 194-1
194.2 DBMS_SYNC_REFRESH Security Model 194-1
194.3 Summary of DBMS_SYNC_REFRESH Subprograms 194-2
194.3.1 ABORT_REFRESH Procedure 194-2
194.3.2 ALTER_REFRESH_STATS_RETENTION Procedure 194-3
194.3.3 CAN_SYNCREF_TABLE Procedure 194-4
194.3.4 EXECUTE_REFRESH Procedure 194-5
194.3.5 GET_ALL_GROUP_IDS Function 194-6
194.3.6 GET_GROUP_ID Function 194-6
194.3.7 GET_GROUP_ID_LIST Function 194-7
194.3.8 PREPARE_REFRESH Procedure 194-7
194.3.9 PREPARE_STAGING_LOG Procedure 194-8

194.3.10 PURGE_REFRESH_STATS Procedure 194-10
194.3.11 REGISTER_MVIEWS 194-10
194.3.12 REGISTER_PARTITION_OPERATION Procedure 194-11
194.3.13 UNREGISTER_MVIEWS 194-12
194.3.14 UNREGISTER_PARTITION_OPERATION Procedure 194-12

195 DBMS_TABLE_DATA

195.1 DBMS_TABLE_DATA Overview 195-1
195.2 DBMS_TABLE_DATA Security Model 195-1
195.3 Summary of DBMS_TABLE_DATA Subprograms 195-2
195.3.1 GET_BYTES_FOR_COLUMN Procedure 195-2
195.3.2 GET_BYTES_FOR_COLUMNS Procedure 195-3
195.3.3 GET_BYTES_FOR_ROW Procedure 195-3

196 DBMS_TDB

196.1 DBMS_TDB Overview 196-1
196.2 DBMS_TDB Security Model 196-2
196.3 DBMS_TDB Constants 196-2
196.4 DBMS_TDB Views 196-2
196.5 DBMS_TDB Operational Notes 196-2
196.6 Summary of DBMS_TDB Subprograms 196-2
196.6.1 CHECK_DB Function 196-3
196.6.2 CHECK_EXTERNAL Function 196-4

197 DBMS_TF

197.1 DBMS_TF Overview 197-1
197.2 DBMS_TF Security Model 197-1
197.3 DBMS_TF Constants 197-2
197.4 DBMS_TF Operational Notes 197-2
197.4.1 PTF Client Interface 197-3
197.4.1.1 DESCRIBE Only Polymorphic Table Function 197-5
197.4.2 PTF Server Side Interface 197-5
197.4.3 Read Columns 197-6
197.4.4 Pass-Through Columns 197-6
197.4.5 State Management 197-7
197.4.5.1 CSTORE Compilation State Management 197-7
197.4.5.2 XSTORE Execution State Management 197-9
197.4.6 Method Name Overrides 197-9
197.4.7 Using the COLUMNS Pseudo-Operator 197-11

197.4.8 Query Transformations 197-11
197.4.9 Parallel Execution 197-12
197.5 DBMS_TF Execution Flow 197-12
197.6 DBMS_TF Restrictions 197-14
197.7 DBMS_TF Examples 197-15
197.7.1 Echo Polymorphic Table Function Example 197-16
197.8 DBMS_TF Data Structures 197-19
197.8.1 Supported Types Collections 197-20
197.8.2 COLUMN_DATA_T Record Type 197-21
197.8.3 COLUMN_METADATA_T Record Type 197-22
197.8.4 COLUMN_T Record Type 197-23
197.8.5 DESCRIBE_T Record Type 197-23
197.8.6 ENV_T Record Type 197-23
197.8.7 PARALLEL_ENV_T Record Type 197-24
197.8.8 TABLE_T Record Type 197-25
197.8.9 COLUMNS_NEW_T Table Type 197-25
197.8.10 COLUMNS_T Table Type 197-25
197.8.11 COLUMNS_WITH_TYPE_T Table Type 197-25
197.8.12 TABLE_COLUMNS_T Table Type 197-26
197.8.13 ROW_SET_T Table Type 197-26
197.8.14 XID_T Subtype 197-26
197.9 Summary of DBMS_TF Subprograms 197-26
197.9.1 COLUMN_TYPE_NAME Function 197-27
197.9.2 COL_TO_CHAR Function 197-28
197.9.3 CSTORE_EXISTS Function 197-29
197.9.4 CSTORE_GET Procedure 197-29
197.9.5 GET_COL Procedure 197-30
197.9.6 GET_ENV Function 197-32
197.9.7 GET_ROW_SET Procedure 197-32
197.9.7.1 Stack Polymorphic Table Function Example 197-33
197.9.8 GET_XID Function 197-36
197.9.9 PUT_COL Procedure 197-36
197.9.9.1 Rand_col Polymorphic Table Function Example 197-37
197.9.10 PUT_ROW_SET Procedure 197-40
197.9.10.1 Split Polymorphic Table Function Example 197-40
197.9.11 ROW_REPLICATION Procedure 197-43
197.9.12 ROW_TO_CHAR Function 197-45
197.9.13 SUPPORTED_TYPE Function 197-46
197.9.14 TRACE Procedure 197-46
197.9.15 XSTORE_CLEAR Procedure 197-47
197.9.16 XSTORE_EXISTS Function 197-48
197.9.17 XSTORE_GET Procedure 197-48

197.9.17.1 Row_num Polymorphic Table Function Example 197-49
197.9.18 XSTORE_REMOVE Procedure 197-51
197.9.19 XSTORE_SET Procedure 197-52

198 DBMS_TNS

198.1 DBMS_TNS Overview 198-1
198.2 DBMS_TNS Security Model 198-1
198.3 Summary of DBMS_TNS Subprograms 198-1
198.3.1 RESOLVE_TNSNAME Function 198-2

199 DBMS_TRACE

199.1 DBMS_TRACE Overview 199-1
199.2 DBMS_TRACE Security Model 199-1
199.3 DBMS_TRACE Constants 199-1
199.4 DBMS_TRACE Restrictions 199-2
199.5 DBMS_TRACE Operational Notes 199-3
199.6 Summary of DBMS_TRACE Subprograms 199-5
199.6.1 CLEAR_PLSQL_TRACE Procedure 199-5
199.6.2 GET_PLSQL_TRACE_LEVEL Function 199-6
199.6.3 PLSQL_TRACE_VERSION Procedure 199-6
199.6.4 SET_PLSQL_TRACE Procedure 199-6

200 DBMS_TRANSACTION

200.1 DBMS_TRANSACTION Security Model 200-1
200.2 Summary of DBMS_TRANSACTION Subprograms 200-1
200.2.1 ADVISE_COMMIT Procedure 200-2
200.2.2 ADVISE_NOTHING Procedure 200-3
200.2.3 ADVISE_ROLLBACK Procedure 200-3
200.2.4 COMMIT Procedure 200-3
200.2.5 COMMIT_COMMENT Procedure 200-3
200.2.6 COMMIT_FORCE Procedure 200-3
200.2.7 GET_TRANSACTION_ID Function 200-4
200.2.8 GET_TRANSACTION_TIMEOUT Function 200-4
200.2.9 GET_TRANSACTION_TYPE Function 200-4
200.2.10 LOCAL_TRANSACTION_ID Function 200-5
200.2.11 PURGE_LOST_DB_ENTRY Procedure 200-5
200.2.12 PURGE_MIXED Procedure 200-7
200.2.13 READ_ONLY Procedure 200-7
200.2.14 READ_WRITE Procedure 200-8

200.2.15 ROLLBACK Procedure 200-8
200.2.16 ROLLBACK_FORCE Procedure 200-8
200.2.17 ROLLBACK_SAVEPOINT Procedure 200-8
200.2.18 SAVEPOINT Procedure 200-9
200.2.19 START_TRANSACTION Function 200-9
200.2.20 STEP_ID Function 200-10
200.2.21 SUSPEND_TRANSACTION Procedure 200-10
200.2.22 USE_ROLLBACK_SEGMENT Procedure 200-11

201 DBMS_TRANSFORM

201.1 Summary of DBMS_TRANSFORM Subprograms 201-1
201.1.1 CREATE_TRANSFORMATION Procedure 201-1
201.1.2 DROP_TRANSFORMATION Procedure 201-2
201.1.3 MODIFY_TRANSFORMATION Procedure 201-3

202 DBMS_TSDP_MANAGE

202.1 DBMS_TSDP_MANAGE Overview 202-1
202.2 DBMS_TSDP_MANAGE Security Model 202-1
202.3 Summary of DBMS_TSDP_MANAGE Subprograms 202-2
202.3.1 ADD_SENSITIVE_COLUMN Procedure 202-2
202.3.2 ALTER_SENSITIVE_COLUMN Procedure 202-3
202.3.3 ADD_SENSITIVE_TYPE Procedure 202-4
202.3.4 DROP_SENSITIVE_COLUMN Procedure 202-4
202.3.5 DROP_SENSITIVE_TYPE Procedure 202-5
202.3.6 DROP_SENSITIVE_TYPE_SOURCE Procedure 202-5
202.3.7 IMPORT_DISCOVERY_RESULT Procedure 202-6
202.3.8 IMPORT_SENSITIVE_TYPES Procedure 202-6
202.3.9 REMOVE_DISCOVERY_RESULT Procedure 202-7

203 DBMS_TSDP_PROTECT

203.1 DBMS_TSDP_PROTECT Overview 203-1
203.2 DBMS_TSDP_PROTECT Security Model 203-1
203.3 DBMS_TSDP_PROTECT Constants 203-2
203.4 DBMS_TSDP_PROTECT Data Structures 203-2
203.4.1 FEATURE_OPTIONS Table Type 203-2
203.4.2 POLICY_CONDITIONS Table Type 203-2
203.5 Summary of DBMS_TSDP_PROTECT Subprograms 203-2
203.5.1 ADD_POLICY Procedure 203-3
203.5.2 ALTER_POLICY Procedure 203-4

203.5.3 ASSOCIATE_POLICY Procedure 203-5
203.5.4 DISABLE_PROTECTION_COLUMN Procedure 203-6
203.5.5 DISABLE_PROTECTION_SOURCE Procedure 203-7
203.5.6 DISABLE_PROTECTION_TYPE Procedure 203-7
203.5.7 DROP_POLICY Procedure 203-7
203.5.8 ENABLE_PROTECTION_COLUMN Procedure 203-9
203.5.9 ENABLE_PROTECTION_SOURCE Procedure 203-9
203.5.10 ENABLE_PROTECTION_TYPE Procedure 203-10

204 DBMS_TTS

204.1 DBMS_TTS Security Model 204-1
204.2 DBMS_TTS Exceptions 204-1
204.3 DBMS_TTS Operational Notes 204-1
204.4 Summary of DBMS_TTS Subprograms 204-2
204.4.1 DOWNGRADE Procedure 204-2
204.4.2 TRANSPORT_SET_CHECK Procedure 204-2

205 DBMS_TYPES

205.1 DBMS_TYPES Constants 205-1
205.2 DBMS_TYPES Exceptions 205-2

206 DBMS_UMF

206.1 Summary of DBMS_UMF Subprograms 206-1
206.1.1 CONFIGURE_NODE Procedure 206-2
206.1.2 CREATE_LINK Procedure 206-2
206.1.3 CREATE_TOPOLOGY Procedure 206-3
206.1.4 DROP_LINK Procedure 206-3
206.1.5 DROP_TOPOLOGY Procedure 206-4
206.1.6 ENABLE_SERVICE Procedure 206-4
206.1.7 GET_NODE_ID_LOCAL Function 206-5
206.1.8 GET_NODE_NAME_LOCAL Function 206-6
206.1.9 GET_TARGET_ID Function 206-6
206.1.10 GET_TOPOLOGY_NAME_LOCAL Function 206-6
206.1.11 QUERY_LINK_INFO Procedure 206-7
206.1.12 QUERY_NODE_INFO Procedures 206-7
206.1.13 REGISTER_NODE Function and Procedure 206-8
206.1.14 SWITCH_DESTINATION Procedure 206-9
206.1.15 UNCONFIGURE_NODE Procedure 206-10

206.1.16 UNREGISTER_NODE Procedure 206-10

207 DBMS_UNDO_ADV

207.1 Summary of DBMS_UNDO_ADV Subprograms 207-1
207.1.1 BEST_POSSIBLE_RETENTION Function 207-2
207.1.2 RBU_MIGRATION Function 207-3
207.1.3 LONGEST_QUERY Function 207-3
207.1.4 REQUIRED_RETENTION Function 207-4
207.1.5 REQUIRED_UNDO_SIZE Function 207-5
207.1.6 UNDO_ADVISOR Function 207-6
207.1.7 UNDO_AUTOTUNE Function 207-7
207.1.8 UNDO_HEALTH Function 207-7
207.1.9 UNDO_INFO Function 207-9

208 DBMS_USER_CERTS

208.1 DBMS_USER_CERTS Overview 208-1
208.2 DBMS_USER_CERTS Security Model 208-1
208.3 Summary of DBMS_USER_CERTS Subprograms 208-2
208.3.1 ADD_CERTIFICATE Procedure 208-2
208.3.2 ADD_COPY Procedure 208-2
208.3.3 DROP_CERTIFICATE Procedure 208-3

209 DBMS_USERDIAG

209.1 DBMS_USERDIAG Overview 209-1
209.2 Summary of DBMS_USERDIAG Subprograms 209-1
209.2.1 CHECK_SQL_TRACE_EVENT Procedure 209-2
209.2.2 ENABLE_SQL_TRACE_EVENT Procedure 209-2
209.2.3 GET_CALL_ERROR_MSG Function 209-3
209.2.4 GET_CALL_STATUS Function 209-3
209.2.5 SET_EXCEPTION_MODE Procedure 209-4
209.2.6 SET_TRACEFILE_IDENTIFIER Procedure 209-4
209.2.7 TRACE Procedure 209-4

210 DBMS_UTILITY

210.1 DBMS_UTILITY Deprecated Subprograms 210-1
210.2 DBMS_UTILITY Security Model 210-1
210.3 DBMS_UTILITY Constants 210-2
210.4 DBMS_UTILITY Exceptions 210-2

210.5 DBMS_UTILITY Data Structures 210-2
210.5.1 DBMS_UTILITY INSTANCE_RECORD Record Type 210-3
210.5.2 DBMS_UTILITY DBLINK_ARRAY TABLE Type 210-3
210.5.3 DBMS_UTILITY INDEX_TABLE_TYPE Table Type 210-3
210.5.4 DBMS_UTILITY INSTANCE_TABLE Table Type 210-3
210.5.5 DBMS_UTILITY LNAME_ARRAY Table Type 210-4
210.5.6 DBMS_UTILITY NAME_ARRAY Table Type 210-4
210.5.7 DBMS_UTILITY NUMBER_ARRAY Table Type 210-4
210.5.8 DBMS_UTILITY UNCL_ARRAY Table Type 210-4
210.6 Summary of DBMS_UTILITY Subprograms 210-4
210.6.1 ACTIVE_INSTANCES Procedure 210-6
210.6.2 ANALYZE_DATABASE Procedure 210-6
210.6.3 ANALYZE_PART_OBJECT Procedure 210-7
210.6.4 ANALYZE_SCHEMA Procedure 210-8
210.6.5 CANONICALIZE Procedure 210-9
210.6.6 COMMA_TO_TABLE Procedures 210-10
210.6.7 COMPILE_SCHEMA Procedure 210-10
210.6.8 CREATE_ALTER_TYPE_ERROR_TABLE Procedure 210-11
210.6.9 CURRENT_INSTANCE Function 210-12
210.6.10 DATA_BLOCK_ADDRESS_BLOCK Function 210-12
210.6.11 DATA_BLOCK_ADDRESS_FILE Function 210-13
210.6.12 DB_VERSION Procedure 210-13
210.6.13 EXEC_DDL_STATEMENT Procedure 210-14
210.6.14 EXPAND_SQL_TEXT Procedure 210-14
210.6.15 FORMAT_CALL_STACK Function 210-15
210.6.16 FORMAT_ERROR_BACKTRACE Function 210-15
210.6.17 FORMAT_ERROR_STACK Function 210-18
210.6.18 GET_CPU_TIME Function 210-18
210.6.19 GET_DEPENDENCY Procedure 210-18
210.6.20 GET_ENDIANNESS Function 210-19
210.6.21 GET_HASH_VALUE Function 210-19
210.6.22 GET_PARAMETER_VALUE Function 210-20
210.6.23 GET_SQL_HASH Function 210-21
210.6.24 GET_TIME Function 210-22
210.6.25 GET_TZ_TRANSITIONS Procedure 210-22
210.6.26 INVALIDATE Procedure 210-23
210.6.27 IS_BIT_SET Function 210-25
210.6.28 IS_CLUSTER_DATABASE Function 210-26
210.6.29 MAKE_DATA_BLOCK_ADDRESS Function 210-26
210.6.30 NAME_RESOLVE Procedure 210-27
210.6.31 NAME_TOKENIZE Procedure 210-28
210.6.32 OLD_CURRENT_SCHEMA Function 210-29

210.6.33 OLD_CURRENT_USER Function 210-29
210.6.34 PORT_STRING Function 210-29
210.6.35 SQLID_TO_SQLHASH Function 210-29
210.6.36 TABLE_TO_COMMA Procedures 210-30
210.6.37 VALIDATE Procedure 210-30
210.6.38 WAIT_ON_PENDING_DML Function 210-31

211 DBMS_VECTOR

211.1 Summary of DBMS_VECTOR Subprograms 211-1
211.1.1 CREATE_CREDENTIAL 211-2
211.1.2 CREATE_INDEX 211-5
211.1.3 DISABLE_CHECKPOINT 211-6
211.1.4 DROP_CREDENTIAL 211-6
211.1.5 ENABLE_CHECKPOINT 211-7
211.1.6 DROP_ONNX_MODEL Procedure 211-8
211.1.7 GET_INDEX_STATUS 211-8
211.1.8 INDEX_ACCURACY_QUERY 211-10
211.1.9 INDEX_ACCURACY_REPORT 211-11
211.1.10 INDEX_VECTOR_MEMORY_ADVISOR 211-11
211.1.11 LOAD_ONNX_MODEL 211-13
211.1.11.1 JSON Metadata Parameters for ONNX Models 211-15
211.1.12 LOAD_ONNX_MODEL_CLOUD 211-21
211.1.13 QUERY 211-22
211.1.14 REBUILD_INDEX 211-23
211.1.15 RERANK 211-24
211.1.16 UTL_TO_CHUNKS 211-28
211.1.17 UTL_TO_EMBEDDING and UTL_TO_EMBEDDINGS 211-34
211.1.18 UTL_TO_GENERATE_TEXT 211-41

212 DBMS_VECTOR_CHAIN

212.1 Summary of DBMS_VECTOR_CHAIN Subprograms 212-1
212.1.1 CREATE_CREDENTIAL 212-2
212.1.2 CREATE_LANG_DATA 212-5
212.1.3 CREATE_PREFERENCE 212-6
212.1.4 CREATE_VOCABULARY 212-14
212.1.5 DROP_CREDENTIAL 212-16
212.1.6 DROP_LANG_DATA 212-17
212.1.7 DROP_PREFERENCE 212-17
212.1.8 DROP_VOCABULARY 212-18
212.1.9 RERANK 212-18

212.1.10 UTL_TO_CHUNKS 212-22
212.1.11 UTL_TO_EMBEDDING and UTL_TO_EMBEDDINGS 212-28
212.1.12 UTL_TO_GENERATE_TEXT 212-35
212.1.13 UTL_TO_TEXT 212-43
212.1.14 UTL_TO_SUMMARY 212-44
212.2 Supported Languages and Data File Locations 212-51

213 DBMS_WARNING

213.1 DBMS_WARNING Security Model 213-1
213.2 Summary of DBMS_WARNING Subprograms 213-1
213.2.1 ADD_WARNING_SETTING_CAT Procedure 213-2
213.2.2 ADD_WARNING_SETTING_NUM Procedure 213-2
213.2.3 GET_CATEGORY Function 213-3
213.2.4 GET_WARNING_SETTING_CAT Function 213-4
213.2.5 GET_WARNING_SETTING_NUM Function 213-4
213.2.6 GET_WARNING_SETTING_STRING Function 213-4
213.2.7 SET_WARNING_SETTING_STRING Procedure 213-5

214 DBMS_WM

215 DBMS_WORKLOAD_CAPTURE

215.1 DBMS_WORKLOAD_CAPTURE Overview 215-1
215.2 DBMS_WORKLOAD_CAPTURE Security Model 215-1
215.3 Summary of DBMS_WORKLOAD_CAPTURE Subprograms 215-2
215.3.1 ADD_FILTER Procedures 215-2
215.3.2 DECRYPT_CAPTURE Procedure 215-4
215.3.3 DELETE_CAPTURE_INFO Procedure 215-5
215.3.4 DELETE_FILTER Procedure 215-5
215.3.5 ENCRYPT_CAPTURE Procedure 215-6
215.3.6 EXPORT_AWR Procedure 215-6
215.3.7 FINISH_CAPTURE Procedure 215-7
215.3.8 GET_CAPTURE_INFO Function 215-8
215.3.9 IMPORT_AWR Function 215-8
215.3.10 REPORT Function 215-9
215.3.11 START_CAPTURE Procedure 215-10

216 DBMS_WORKLOAD_REPLAY

216.1 DBMS_WORKLOAD_REPLAY Overview 216-1
216.2 DBMS_WORKLOAD_REPLAY Security Model 216-1
216.3 Summary of DBMS_WORKLOAD_REPLAY Subprograms 216-2
216.3.1 ADD_CAPTURE Function 216-4
216.3.2 ADD_FILTER Procedure 216-5
216.3.3 ADD_SCHEDULE_ORDERING Function 216-5
216.3.4 ASSIGN_GROUP_TO_INSTANCE Procedure 216-6
216.3.5 BEGIN_REPLAY_SCHEDULE Procedure 216-7
216.3.6 CALIBRATE Function 216-7
216.3.7 CANCEL_REPLAY Procedure 216-8
216.3.8 COMPARE_PERIOD_REPORT Procedure 216-9
216.3.9 COMPARE_SQLSET_REPORT Function 216-9
216.3.10 CREATE_FILTER_SET Procedure 216-10
216.3.11 DELETE_FILTER Procedure 216-11
216.3.12 DELETE_REPLAY_INFO Procedure 216-11
216.3.13 END_REPLAY_SCHEDULE Procedure 216-12
216.3.14 EXPORT_AWR Procedure 216-12
216.3.15 GENERATE_CAPTURE_SUBSET Procedure 216-12
216.3.16 GET_DIVERGING_STATEMENT Function 216-13
216.3.17 GET_REPLAY_DIRECTORY Function 216-14
216.3.18 GET_REPLAY_INFO Function 216-14
216.3.19 GET_REPLAY_TIMEOUT Procedure 216-15
216.3.20 IMPORT_AWR Function 216-15
216.3.21 INITIALIZE_CONSOLIDATED_REPLAY Procedure 216-16
216.3.22 INITIALIZE_REPLAY Procedure 216-17
216.3.23 IS_REPLAY_PAUSED Function 216-19
216.3.24 LOAD_LONG_SQLTEXT Procedure 216-19
216.3.25 PAUSE_REPLAY Procedure 216-19
216.3.26 POPULATE_DIVERGENCE Procedure 216-20
216.3.27 PREPARE_CONSOLIDATED_REPLAY Procedure 216-20
216.3.28 PREPARE_REPLAY Procedure 216-22
216.3.29 PROCESS_CAPTURE Procedure 216-26
216.3.30 REMAP_CONNECTION Procedure 216-27
216.3.31 REMOVE_CAPTURE Procedure 216-28
216.3.32 REMOVE_SCHEDULE_ORDERING Procedure 216-29
216.3.33 REPORT Function 216-29
216.3.34 RESUME_REPLAY Procedure 216-30
216.3.35 REUSE_REPLAY_FILTER_SET Procedure 216-30
216.3.36 SET_ADVANCED_PARAMETER Procedure 216-31
216.3.37 SET_REPLAY_DIRECTORY Procedure 216-32

216.3.38 SET_REPLAY_TIMEOUT Procedure 216-32
216.3.39 SET_SQL_MAPPING Procedure 216-33
216.3.40 SET_USER_MAPPING Procedure 216-34
216.3.41 START_CONSOLIDATED_REPLAY Procedure 216-34
216.3.42 START_REPLAY Procedure 216-35
216.3.43 USE_FILTER_SET Procedure 216-35

217 DBMS_WORKLOAD_REPOSITORY

217.1 DBMS_WORKLOAD_REPOSITORY Examples 217-1
217.2 DBMS_WORKLOAD_REPOSITORY Data Structures 217-1
217.2.1 DBMS_WORKLOAD_REPOSITORY AWR_BASELINE_METRIC_TYPE Object Type 217-2
217.2.2 DBMS_WORKLOAD_REPOSITORY AWR_BASELINE_METRIC_TYPE_TABLE Table Type 217-2
217.2.3 DBMS_WORKLOAD_REPOSITORY AWRRPT_INSTANCE_LIST_TYPE Table Type 217-3
217.3 Summary of DBMS_WORKLOAD_REPOSITORY Subprograms 217-3
217.3.1 ADD_COLORED_SQL Procedure 217-5
217.3.2 ASH_GLOBAL_REPORT_HTML Function 217-5
217.3.3 ASH_GLOBAL_REPORT_TEXT Function 217-7
217.3.4 ASH_REPORT_ANALYTICS Function 217-9
217.3.5 ASH_REPORT_HTML Function 217-10
217.3.6 ASH_REPORT_TEXT Function 217-12
217.3.7 AWR_DIFF_REPORT_HTML Function 217-15
217.3.8 AWR_DIFF_REPORT_TEXT Function 217-15
217.3.9 AWR_EXP Procedure 217-16
217.3.10 AWR_GLOBAL_DIFF_REPORT_HTML Functions 217-17
217.3.11 AWR_GLOBAL_DIFF_REPORT_TEXT Functions 217-18
217.3.12 AWR_GLOBAL_REPORT_HTML Functions 217-19
217.3.13 AWR_GLOBAL_REPORT_TEXT Functions 217-20
217.3.14 AWR_IMP Procedure 217-21
217.3.15 AWR_REPORT_HTML Function 217-21
217.3.16 AWR_REPORT_TEXT Function 217-22
217.3.17 AWR_SET_REPORT_THRESHOLDS Procedure 217-23
217.3.18 AWR_SQL_REPORT_HTML Function 217-24
217.3.19 AWR_SQL_REPORT_TEXT Function 217-25
217.3.20 CONTROL_RESTRICTED_SNAPSHOT Procedure 217-26
217.3.21 CREATE_BASELINE Functions and Procedures 217-26
217.3.22 CREATE_BASELINE_TEMPLATE Procedures 217-28
217.3.23 CREATE_REMOTE_SNAPSHOT Function and Procedure 217-29
217.3.24 CREATE_SNAPSHOT Function and Procedure 217-30
217.3.25 DROP_BASELINE Procedure 217-31

217.3.26 DROP_BASELINE_TEMPLATE Procedure 217-32
217.3.27 DROP_SNAPSHOT_RANGE Procedure 217-32
217.3.28 LOCAL_AWR_DBID Function 217-33
217.3.29 MODIFY_BASELINE_WINDOW_SIZE Procedure 217-33
217.3.30 MODIFY_SNAPSHOT_SETTINGS Procedures 217-34
217.3.31 PURGE_SQL_DETAILS Procedure 217-35
217.3.32 REGISTER_REMOTE_DATABASE Procedures 217-36
217.3.33 REMOVE_COLORED_SQL Procedure 217-36
217.3.34 RENAME_BASELINE Procedure 217-37
217.3.35 SELECT_BASELINE_DETAILS Function 217-37
217.3.36 SELECT_BASELINE_METRIC Function 217-38
217.3.37 UNREGISTER_REMOTE_DATABASE Procedures 217-38
217.3.38 UPDATE_DATAFILE_INFO Procedure 217-39
217.3.39 UPDATE_OBJECT_INFO Procedure 217-39

218 DBMS_XDB

218.1 DBMS_XDB Overview 218-1
218.2 DBMS_XDB Security Model 218-2
218.3 DBMS_XDB Constants 218-2
218.4 Summary of DBMS_XDB Subprograms 218-3
218.4.1 ACLCHECKPRIVILEGES Function 218-4
218.4.2 APPENDRESOURCEMETADATA Procedure 218-5
218.4.3 CHANGEOWNER Procedure 218-6
218.4.4 CHANGEPRIVILEGES Function 218-7
218.4.5 CHECKPRIVILEGES Function 218-7
218.4.6 CREATEFOLDER Function 218-8
218.4.7 CREATEOIDPATH Function 218-9
218.4.8 CREATERESOURCE Functions 218-9
218.4.9 DELETERESOURCE Procedure 218-11
218.4.10 DELETERESOURCEMETADATA Procedures 218-12
218.4.11 ENABLEDIGESTAUTHENTICATION Procedure 218-13
218.4.12 EXISTSRESOURCE Function 218-13
218.4.13 GETACLDOCUMENT Function 218-14
218.4.14 GETCONTENTBLOB Function 218-14
218.4.15 GETCONTENTCLOB Function 218-15
218.4.16 GETCONTENTVARCHAR2 Function 218-16
218.4.17 GETCONTENTXMLREF Function 218-16
218.4.18 GETCONTENTXMLTYPE Function 218-17
218.4.19 GETHTTPREQUESTHEADER Function 218-17
218.4.20 GETLOCKTOKEN Procedure 218-18
218.4.21 GETPRIVILEGES Function 218-19

218.4.22 GETRESOID Function 218-19
218.4.23 GETXDB_TABLESPACE Function 218-20
218.4.24 HASBLOBCONTENT Function 218-20
218.4.25 HASCHARCONTENT Function 218-21
218.4.26 HASXMLCONTENT Function 218-21
218.4.27 HASXMLREFERENCE Function 218-22
218.4.28 ISFOLDER Function 218-22
218.4.29 LINK Procedures 218-23
218.4.30 LOCKRESOURCE Function 218-24
218.4.31 PROCESSLINKS Procedure 218-25
218.4.32 PURGERESOURCEMETADATA Procedure 218-25
218.4.33 RENAMERESOURCE Procedure 218-26
218.4.34 SETACL Procedure 218-26
218.4.35 SPLITPATH Procedure 218-27
218.4.36 TOUCHRESOURCE Procedure 218-28
218.4.37 UNLOCKRESOURCE Function 218-28
218.4.38 UPDATERESOURCEMETADATA Procedures 218-29

219 DBMS_XA

219.1 DBMS_XA Overview 219-1
219.2 DBMS_XA Security Model 219-1
219.3 DBMS_XA Constants 219-2
219.4 DBMS_XA Operational Notes 219-3
219.5 DBMS_XA Data Structures 219-4
219.5.1 DBMS_XA DBMS_XA_XID Object Type 219-4
219.5.2 DBMS_XA DBMS_XA_XID_ARRAY Table Type 219-5
219.6 Summary of DBMS_XA Subprograms 219-5
219.6.1 DIST_TXN_SYNC Procedure 219-6
219.6.2 XA_COMMIT Function 219-6
219.6.3 XA_END Function 219-7
219.6.4 XA_FORGET Function 219-8
219.6.5 XA_GETLASTOER Function 219-8
219.6.6 XA_PREPARE Function 219-8
219.6.7 XA_RECOVER Function 219-9
219.6.8 XA_ROLLBACK Function 219-9
219.6.9 XA_SETTIMEOUT Function 219-10
219.6.10 XA_START Function 219-10

220 DBMS_XDB_ADMIN

220.1 DBMS_XDB_ADMIN Security Model 220-1
220.2 Summary of DBMS_XDB_ADMIN Subprograms 220-1
220.2.1 CREATENONCEKEY Procedure 220-1
220.2.2 INSTALLDEFAULTWALLET Procedure 220-2
220.2.3 MOVEXDB_TABLESPACE Procedure 220-2
220.2.4 REBUILDHIERARCHICALINDEX Procedure 220-2

221 DBMS_XDB_CONFIG

221.1 DBMS_XDB_CONFIG Overview 221-1
221.2 DBMS_XDB_CONFIG Security Model 221-2
221.3 DBMS_XDB_CONFIG Constants 221-2
221.4 Summary of DBMS_XDB_CONFIG Subprograms 221-2
221.4.1 ADDHTTPEXPIREMAPPING Procedure 221-4
221.4.2 ADDMIMEMAPPING Procedure 221-4
221.4.3 ADDSCHEMALOCMAPPING Procedure 221-5
221.4.4 ADDSERVLET Procedure 221-5
221.4.5 ADDSERVLETMAPPING Procedure 221-6
221.4.6 ADDSERVLETSECROLE Procedure 221-7
221.4.7 ADDXMLEXTENSION Procedure 221-7
221.4.8 CFG_GET Function 221-8
221.4.9 CFG_REFRESH Procedure 221-8
221.4.10 CFG_UPDATE Procedure 221-8
221.4.11 DELETEHTTPEXPIREMAPPING Procedure 221-8
221.4.12 DELETEMIMEMAPPING Procedure 221-9
221.4.13 DELETESCHEMALOCMAPPING Procedure 221-9
221.4.14 DELETESERVLET Procedure 221-9
221.4.15 DELETESERVLETMAPPING Procedure 221-10
221.4.16 DELETESERVLETSECROLE Procedure 221-10
221.4.17 DELETEXMLEXTENSION Procedure 221-10
221.4.18 ENABLEDIGESTAUTHENTICATION Procedure 221-11
221.4.19 GETFTPPORT Function 221-11
221.4.20 GETHTTPCONFIGREALM Function 221-11
221.4.21 GETHTTPPORT Function 221-11
221.4.22 GETHTTPSPORT Function 221-11
221.4.23 GETLISTENERENDPOINT Procedure 221-12
221.4.24 GETREMOTEHTTPPORT Function 221-12
221.4.25 GETREMOTEHTTPSPORT Function 221-12
221.4.26 ISGLOBALPORTENABLED Function 221-12
221.4.27 SETFTPPORT Procedure 221-13

221.4.28 SETHTTPPORT Procedure 221-13
221.4.29 SETHTTPCONFIGREALM Procedure 221-13
221.4.30 SETHTTPSPORT Procedure 221-14
221.4.31 SETLISTENERENDPOINT Procedure 221-14
221.4.32 SETLISTENERLOCALACCESS Procedure 221-15
221.4.33 SETREMOTEHTTPPORT Procedure 221-15
221.4.34 SETREMOTEHTTPSPORT Procedure 221-15
221.4.35 USEDPORT Function 221-16

222 DBMS_XDB_CONSTANTS

222.1 DBMS_XDB_CONSTANTS Security Model 222-1
222.2 Summary of DBMS_XDB_CONSTANTS Subprograms 222-1
222.2.1 ENCODING_DEFAULT Function 222-2
222.2.2 ENCODING_ISOLATIN1 Function 222-3
222.2.3 ENCODING_UTF8 Function 222-3
222.2.4 ENCODING_WIN1252 Function 222-3
222.2.5 NAMESPACE_ACL Function 222-3
222.2.6 NAMESPACE_METADATA Function 222-4
222.2.7 NAMESPACE_ORACLE Function 222-4
222.2.8 NAMESPACE_ORACLE_XDB Function 222-4
222.2.9 NAMESPACE_RESOURCE Function 222-4
222.2.10 NAMESPACE_RESOURCE_EVENT Function 222-5
222.2.11 NAMESPACE_RESOURCE_CONFIG Function 222-5
222.2.12 NAMESPACE_XDBSCHEMA Function 222-5
222.2.13 NAMESPACE_XMLDIFF Function 222-5
222.2.14 NAMESPACE_XMLINSTANCE Function 222-6
222.2.15 NAMESPACE_XMLSCHEMA Function 222-6
222.2.16 NSPREFIX_ACL_ACL Function 222-6
222.2.17 NSPREFIX_RESCONFIG_RC Function 222-6
222.2.18 NSPREFIX_RESOURCE_R Function 222-7
222.2.19 NSPREFIX_XDB_XDB Function 222-7
222.2.20 NSPREFIX_XMLINSTANCE_XSI Function 222-7
222.2.21 NSPREFIX_XMLDIFF_XD Function 222-7
222.2.22 NSPREFIX_XMLSCHEMA_XSD Function 222-8
222.2.23 SCHEMAURL_ACL Function 222-8
222.2.24 SCHEMAELEM_RES_ACL Function 222-8
222.2.25 SCHEMAELEM_RESCONTENT_BINARY Function 222-8
222.2.26 SCHEMAELEM_RESCONTENT_TEXT Function 222-9
222.2.27 SCHEMAURL_RESOURCE Function 222-9
222.2.28 SCHEMAURL_XDBSCHEMA Function 222-9
222.2.29 XDBSCHEMA_PREFIXES Function 222-9

222.2.30 XSD_ATTRIBUTE Function 222-10
222.2.31 XSD_COMPLEX_TYPE Function 222-10
222.2.32 XSD_ELEMENT Function 222-10
222.2.33 XSD_GROUP Function 222-10

223 DBMS_XDB_REPOS

223.1 DBMS_XDB_REPOS Overview 223-1
223.2 DBMS_XDB_REPOS Security Model 223-1
223.3 DBMS_XDB_REPOS Constants 223-2
223.4 Summary of DBMS_XDB_REPOS Subprograms 223-2
223.4.1 ACLCHECKPRIVILEGES Function 223-4
223.4.2 APPENDRESOURCEMETADATA Procedure 223-4
223.4.3 CHANGEOWNER Procedure 223-5
223.4.4 CHANGEPRIVILEGES Function 223-5
223.4.5 CHECKPRIVILEGES Function 223-6
223.4.6 CREATEFOLDER Function 223-6
223.4.7 CREATEOIDPATH Function 223-7
223.4.8 CREATERESOURCE Functions 223-7
223.4.9 DELETERESOURCE Procedure 223-9
223.4.10 DELETERESOURCEMETADATA Procedures 223-9
223.4.11 EXISTSRESOURCE Function 223-10
223.4.12 GETACLDOCUMENT Function 223-10
223.4.13 GETCONTENTBLOB Function 223-11
223.4.14 GETCONTENTCLOB Function 223-11
223.4.15 GETCONTENTVARCHAR2 Function 223-12
223.4.16 GETCONTENTXMLREF Function 223-12
223.4.17 GETCONTENTXMLTYPE Function 223-13
223.4.18 GETLOCKTOKEN Procedure 223-13
223.4.19 GETPRIVILEGES Function 223-13
223.4.20 GETRESOID Function 223-14
223.4.21 GETXDB_TABLESPACE Function 223-14
223.4.22 HASBLOBCONTENT Function 223-14
223.4.23 HASCHARCONTENT Function 223-15
223.4.24 HASXMLCONTENT Function 223-15
223.4.25 HASXMLREFERENCE Function 223-16
223.4.26 ISFOLDER Function 223-16
223.4.27 LINK Procedures 223-16
223.4.28 LOCKRESOURCE Function 223-17
223.4.29 PROCESSLINKS Procedure 223-18
223.4.30 PURGERESOURCEMETADATA Procedure 223-18
223.4.31 RENAMERESOURCE Procedure 223-18

223.4.32 SETACL Procedure 223-19
223.4.33 SPLITPATH Procedure 223-19
223.4.34 TOUCHRESOURCE Procedure 223-20
223.4.35 UNLOCKRESOURCE Function 223-20
223.4.36 UPDATERESOURCEMETADATA Procedures 223-20

224 DBMS_XDB_VERSION

224.1 DBMS_XDB_VERSION Security Model 224-1
224.2 Summary of DBMS_XDB_VERSION Subprograms 224-1
224.2.1 CHECKIN Function 224-2
224.2.2 CHECKOUT Procedure 224-3
224.2.3 GETCONTENTSBLOBBYRESID Function 224-3
224.2.4 GETCONTENTSCLOBBYRESID Function 224-3
224.2.5 GETCONTENTSXMLBYRESID Function 224-4
224.2.6 GETPREDECESSORS Function 224-4
224.2.7 GETPREDSBYRESID Function 224-5
224.2.8 GETRESOURCEBYRESID Function 224-5
224.2.9 GETSUCCESSORS Function 224-5
224.2.10 GETSUCCSBYRESID Function 224-6
224.2.11 MAKEVERSIONED Function 224-6
224.2.12 UNCHECKOUT Function 224-7

225 DBMS_XDBRESOURCE

225.1 DBMS_XDBRESOURCE Overview 225-1
225.2 DBMS_XDBRESOURCE Security Model 225-1
225.3 Summary of DBMS_XDBRESOURCE Subprograms 225-1
225.3.1 FREERESOURCE Procedure 225-3
225.3.2 GETACL Function 225-4
225.3.3 GETACLDOCFROMRES Function 225-4
225.3.4 GETAUTHOR Function 225-4
225.3.5 GETCHARACTERSET Function 225-5
225.3.6 GETCOMMENT Function 225-5
225.3.7 GETCONTENTBLOB Function 225-5
225.3.8 GETCONTENTCLOB Function 225-6
225.3.9 GETCONTENTREF Function 225-6
225.3.10 GETCONTENTTYPE Function 225-6
225.3.11 GETCONTENTXML Function 225-7
225.3.12 GETCONTENTVARCHAR2 Function 225-7
225.3.13 GETCREATIONDATE Function 225-7
225.3.14 GETCREATOR Function 225-8

225.3.15 GETCUSTOMMETADATA Function 225-8
225.3.16 GETDISPLAYNAME Function 225-9
225.3.17 GETLANGUAGE Function 225-9
225.3.18 GETLASTMODIFIER Function 225-9
225.3.19 GETMODIFICATIONDATE Function 225-10
225.3.20 GETOWNER Function 225-10
225.3.21 GETREFCOUNT Function 225-10
225.3.22 GETVERSIONID Function 225-11
225.3.23 HASACLCHANGED Function 225-11
225.3.24 HASAUTHORCHANGED Function 225-11
225.3.25 HASCHANGED Function 225-12
225.3.26 HASCHARACTERSETCHANGED Function 225-12
225.3.27 HASCOMMENTCHANGED Function 225-12
225.3.28 HASCONTENTCHANGED Function 225-13
225.3.29 HASCONTENTTYPECHANGED Function 225-13
225.3.30 HASCREATIONDATECHANGED Function 225-14
225.3.31 HASCREATORCHANGED Function 225-14
225.3.32 HASCUSTOMMETADATACHANGED Function 225-14
225.3.33 HASDISPLAYNAMECHANGED Function 225-15
225.3.34 HASLANGUAGECHANGED Function 225-15
225.3.35 HASLASTMODIFIERCHANGED Function 225-15
225.3.36 HASMODIFICATIONDATECHANGED Function 225-16
225.3.37 HASOWNERCHANGED Function 225-16
225.3.38 HASREFCOUNTCHANGED Function 225-17
225.3.39 HASVERSIONIDCHANGED Function 225-17
225.3.40 ISFOLDER Function 225-17
225.3.41 ISNULL Function 225-18
225.3.42 MAKEDOCUMENT Function 225-18
225.3.43 SAVE Procedure 225-19
225.3.44 SETACL Procedure 225-19
225.3.45 SETAUTHOR Procedure 225-19
225.3.46 SETCHARACTERSET Procedure 225-20
225.3.47 SETCOMMENT Procedure 225-20
225.3.48 SETCONTENT Procedures 225-20
225.3.49 SETCONTENTTYPE Procedure 225-21
225.3.50 SETCUSTOMMETADATA Procedure 225-21
225.3.51 SETDISPLAYNAME Procedure 225-22
225.3.52 SETLANGUAGE Procedure 225-22
225.3.53 SETOWNER Procedure 225-23

226 DBMS_XDBZ

226.1 DBMS_XDBZ Security Model 226-1
226.2 DBMS_XDBZ Constants 226-1
226.3 Summary of DBMS_XDBZ Subprograms 226-2
226.3.1 CREATENONCEKEY Procedure 226-2
226.3.2 DISABLE_HIERARCHY Procedure 226-3
226.3.3 ENABLE_HIERARCHY Procedure 226-3
226.3.4 GET_ACLOID Function 226-4
226.3.5 GET_USERID Function 226-4
226.3.6 IS_HIERARCHY_ENABLED Function 226-5
226.3.7 PURGELDAPCACHE Function 226-6

227 DBMS_XEVENT

227.1 DBMS_XEVENT Security Model 227-1
227.2 DBMS_XEVENT Constants 227-1
227.3 Subprogram Groups 227-2
227.3.1 DBMS_XEVENT XDBEvent Type Subprograms 227-2
227.3.2 DBMS_XEVENT XDBRepositoryEvent Type Subprograms 227-3
227.3.3 DBMS_XEVENT XDBHandlerList Type Subprograms 227-4
227.3.4 DBMS_XEVENT XDBHandler Type Subprograms 227-4
227.3.5 DBMS_XEVENT XDBPath Type Subprograms 227-5
227.3.6 DBMS_XEVENT XDBLink Type Subprograms 227-5
227.4 Summary of DBMS_XEVENT Subprograms 227-5
227.4.1 CLEAR Procedure 227-7
227.4.2 GETAPPLICATIONDATA Function 227-8
227.4.3 GETCHILDOID Function 227-8
227.4.4 GETCURRENTUSER Function 227-9
227.4.5 GETEVENT Function 227-9
227.4.6 GETFIRST Function 227-10
227.4.7 GETHANDLERLIST Function 227-10
227.4.8 GETINTERFACE Function 227-11
227.4.9 GETLANGUAGE Function 227-11
227.4.10 GETLINK Function 227-12
227.4.11 GETLINKNAME Function 227-12
227.4.12 GETLOCK Function 227-13
227.4.13 GETLANGUAGE Function 227-13
227.4.14 GETNAME Function 227-14
227.4.15 GETNEXT Function 227-14
227.4.16 GETOLDRESOURCE Function 227-15
227.4.17 GETOPENACCESSMODE Function 227-15

227.4.18 GETOPENDENYMODE Function 227-16
227.4.19 GETOUTPUTSTREAM Function 227-16
227.4.20 GETPARAMETER Function 227-17
227.4.21 GETPARENT Function 227-17
227.4.22 GETPARENTNAME Function 227-18
227.4.23 GETPARENTOID Function 227-18
227.4.24 GETPARENTPATH Function 227-19
227.4.25 GETPATH Function 227-19
227.4.26 GETRESOURCE Function 227-20
227.4.27 GETSCHEMA Function 227-20
227.4.28 GETSOURCE Function 227-21
227.4.29 GETUPDATEBYTECOUNT Function 227-21
227.4.30 GETUPDATEBYTEOFFSET Function 227-22
227.4.31 GETXDBEVENT Function 227-22
227.4.32 ISNULL Functions 227-23
227.4.33 REMOVE Procedure 227-24
227.4.34 SETRENDERPATH Procedure 227-24
227.4.35 SETRENDERSTREAM Procedure 227-25

228 DBMS_XMLDOM

228.1 DBMS_XMLDOM Overview 228-1
228.2 DBMS_XMLDOM Security Model 228-2
228.3 DBMS_XMLDOM Constants 228-3
228.4 DBMS_XMLDOM Types 228-3
228.5 DBMS_XMLDOM Exceptions 228-4
228.6 DBMS_XMLDOM Subprogram Groups 228-4
228.6.1 DBMS_XMLDOM DOMNode Subprograms 228-5
228.6.2 DBMS_XMLDOM DOMAttr Subprograms 228-6
228.6.3 DBMS_XMLDOM DOMCDataSection Subprograms 228-7
228.6.4 DBMS_XMLDOM DOMCharacterData Subprograms 228-7
228.6.5 DBMS_XMLDOM DOMComment Subprograms 228-8
228.6.6 DBMS_XMLDOM DOMDocument Subprograms 228-8
228.6.7 DBMS_XMLDOM DOMDocumentFragment Subprograms 228-9
228.6.8 DBMS_XMLDOM DOMDocumentType Subprograms 228-9
228.6.9 DBMS_XMLDOM DOMElement Subprograms 228-10
228.6.10 DBMS_XMLDOM DOMEntity Subprograms 228-10
228.6.11 DBMS_XMLDOM DOMEntityReference Subprograms 228-11
228.6.12 DBMS_XMLDOM DOMImplementation Subprograms 228-11
228.6.13 DBMS_XMLDOM DOMNamedNodeMap Subprograms 228-11
228.6.14 DBMS_XMLDOM DOMNodeList Subprograms 228-12
228.6.15 DBMS_XMLDOM DOMNotation Subprograms 228-12

228.6.16 DBMS_XMLDOM DOMProcessingInstruction Subprograms 228-12
228.6.17 DBMS_XMLDOM DOMText Subprograms 228-12
228.7 Summary of DBMS_XMLDOM Subprograms 228-13
228.7.1 ADOPTNODE Function 228-20
228.7.2 APPENDCHILD Function 228-20
228.7.3 APPENDDATA Procedure 228-21
228.7.4 CLONENODE Function 228-21
228.7.5 CREATEATTRIBUTE Functions 228-22
228.7.6 CREATECDATASECTION Function 228-23
228.7.7 CREATECOMMENT Function 228-23
228.7.8 CREATEDOCUMENT Function 228-24
228.7.9 CREATEDOCUMENTFRAGMENT Function 228-24
228.7.10 CREATEELEMENT Functions 228-25
228.7.11 CREATEENTITYREFERENCE Function 228-26
228.7.12 CREATEPROCESSINGINSTRUCTION Function 228-26
228.7.13 CREATETEXTNODE Function 228-27
228.7.14 DELETEDATA Procedure 228-27
228.7.15 FINDENTITY Function 228-28
228.7.16 FINDNOTATION Function 228-28
228.7.17 FREEDOCFRAG Procedure 228-29
228.7.18 FREEDOCUMENT Procedure 228-29
228.7.19 FREEELEMENT Procedure 228-30
228.7.20 FREENODE Procedure 228-30
228.7.21 FREENODELIST Procedure 228-31
228.7.22 GETATTRIBUTE Functions 228-31
228.7.23 GETATTRIBUTENODE Functions 228-32
228.7.24 GETATTRIBUTES Function 228-33
228.7.25 GETCHARSET Function 228-33
228.7.26 GETCHILDNODES Function 228-34
228.7.27 GETCHILDRENBYTAGNAME Functions 228-34
228.7.28 GETDATA Functions 228-35
228.7.29 GETDOCTYPE Function 228-35
228.7.30 GETDOCUMENTELEMENT Function 228-36
228.7.31 GETELEMENTSBYTAGNAME Functions 228-36
228.7.32 GETENTITIES Function 228-37
228.7.33 GETEXPANDEDNAME Procedure and Functions 228-38
228.7.34 GETFIRSTCHILD Function 228-38
228.7.35 GETIMPLEMENTATION Function 228-39
228.7.36 GETLASTCHILD Function 228-39
228.7.37 GETLENGTH Functions 228-40
228.7.38 GETLOCALNAME Procedure and Functions 228-40
228.7.39 GETNAME Functions 228-41

228.7.40 GETNAMEDITEM Function 228-42
228.7.41 GETNAMESPACE Procedure and Functions 228-42
228.7.42 GETNEXTSIBLING Function 228-43
228.7.43 GETNODETYPE Function 228-43
228.7.44 GETNODENAME Function 228-44
228.7.45 GETNODEVALUE Function 228-44
228.7.46 GETNODEVALUEASBINARYSTREAM Function & Procedure 228-45
228.7.47 GETNODEVALUEASCHARACTERSTREAM Function & Procedure 228-46
228.7.48 GETNOTATIONNAME Function 228-46
228.7.49 GETNOTATIONS Function 228-47
228.7.50 GETTARGET Function 228-47
228.7.51 GETOWNERDOCUMENT Function 228-48
228.7.52 GETOWNERELEMENT Function 228-48
228.7.53 GETPARENTNODE Function 228-49
228.7.54 GETPREFIX Function 228-49
228.7.55 GETPREVIOUSSIBLING Function 228-50
228.7.56 GETPUBLICID Functions 228-50
228.7.57 GETQUALIFIEDNAME Functions 228-51
228.7.58 GETSCHEMANODE Function 228-51
228.7.59 GETSPECIFIED Function 228-52
228.7.60 GETSTANDALONE Function 228-52
228.7.61 GETSYSTEMID Functions 228-53
228.7.62 GETTAGNAME Function 228-53
228.7.63 GETVALUE Function 228-54
228.7.64 GETVERSION Function 228-54
228.7.65 GETXMLTYPE Function 228-55
228.7.66 HASATTRIBUTE Functions 228-55
228.7.67 HASATTRIBUTES Function 228-56
228.7.68 HASCHILDNODES Function 228-56
228.7.69 HASFEATURE Function 228-57
228.7.70 IMPORTNODE Function 228-57
228.7.71 INSERTBEFORE Function 228-58
228.7.72 INSERTDATA Procedure 228-59
228.7.73 ISNULL Functions 228-59
228.7.74 ITEM Functions 228-62
228.7.75 MAKEATTR Function 228-63
228.7.76 MAKECDATASECTION Function 228-63
228.7.77 MAKECHARACTERDATA Function 228-64
228.7.78 MAKECOMMENT Function 228-64
228.7.79 MAKEDOCUMENT Function 228-65
228.7.80 MAKEDOCUMENTFRAGMENT Function 228-65
228.7.81 MAKEDOCUMENTTYPE Function 228-66

228.7.82 MAKEELEMENT Function 228-66
228.7.83 MAKEENTITY Function 228-67
228.7.84 MAKEENTITYREFERENCE Function 228-67
228.7.85 MAKENODE Functions 228-68
228.7.86 MAKENOTATION Function 228-70
228.7.87 MAKEPROCESSINGINSTRUCTION Function 228-70
228.7.88 MAKETEXT Function 228-71
228.7.89 NEWDOMDOCUMENT Functions 228-71
228.7.90 NORMALIZE Procedure 228-72
228.7.91 REMOVEATTRIBUTE Procedures 228-72
228.7.92 REMOVEATTRIBUTENODE Function 228-73
228.7.93 REMOVECHILD Function 228-74
228.7.94 REMOVENAMEDITEM Function 228-74
228.7.95 REPLACECHILD Function 228-75
228.7.96 REPLACEDATA Procedure 228-76
228.7.97 RESOLVENAMESPACEPREFIX Function 228-76
228.7.98 SETATTRIBUTE Procedures 228-77
228.7.99 SETATTRIBUTENODE Functions 228-78
228.7.100 SETCHARSET Procedure 228-78
228.7.101 SETDATA Procedures 228-79
228.7.102 SETDOCTYPE Procedure 228-79
228.7.103 SETNAMEDITEM Function 228-80
228.7.104 SETNODEVALUE Procedure 228-81
228.7.105 SETNODEVALUEASBINARYSTREAM Function & Procedure 228-81
228.7.106 SETNODEVALUEASCHARACTERSTREAM Function & Procedure 228-82
228.7.107 SETPREFIX Procedure 228-83
228.7.108 SETSTANDALONE Procedure 228-83
228.7.109 SETVALUE Procedure 228-84
228.7.110 SETVERSION Procedure 228-84
228.7.111 SPLITTEXT Function 228-85
228.7.112 SUBSTRINGDATA Function 228-85
228.7.113 USEBINARYSTREAM Function 228-86
228.7.114 WRITETOBUFFER Procedures 228-86
228.7.115 WRITETOCLOB Procedures 228-87
228.7.116 WRITETOFILE Procedures 228-88

229 DBMS_XMLGEN

229.1 DBMS_XMLGEN Security Model 229-1
229.2 Summary of DBMS_XMLGEN Subprograms 229-1
229.2.1 CLOSECONTEXT Procedure 229-2
229.2.2 CONVERT Functions 229-3

229.2.3 GETNUMROWSPROCESSED Function 229-3
229.2.4 GETXML Functions 229-4
229.2.5 GETXMLTYPE Functions 229-5
229.2.6 NEWCONTEXT Functions 229-6
229.2.7 NEWCONTEXTFROMHIERARCHY Function 229-6
229.2.8 RESTARTQUERY Procedure 229-7
229.2.9 SETCONVERTSPECIALCHARS Procedure 229-7
229.2.10 SETMAXROWS Procedure 229-8
229.2.11 SETNULLHANDLING Procedure 229-8
229.2.12 SETROWSETTAG Procedure 229-9
229.2.13 SETROWTAG Procedure 229-9
229.2.14 SETSKIPROWS Procedure 229-10
229.2.15 USEITEMTAGSFORCOLL Procedure 229-10
229.2.16 USENULLATTRIBUTEINDICATOR Procedure 229-11

230 DBMS_XMLINDEX

230.1 DBMS_XMLINDEX Overview 230-1
230.2 DBMS_XMLINDEX Security Model 230-1
230.3 Summary of DBMS_XMLINDEX Subprograms 230-1
230.3.1 CREATEDATEINDEX Procedure 230-2
230.3.2 CREATENUMBERINDEX Procedure 230-3
230.3.3 DROPPARAMETER Procedure 230-3
230.3.4 MODIFYPARAMETER Procedure 230-4
230.3.5 PROCESS_PENDING Procedure 230-4
230.3.6 REGISTERPARAMETER Procedure 230-5
230.3.7 SYNCINDEX Procedure 230-6

231 DBMS_XMLPARSER

231.1 DBMS_XMLPARSER Security Model 231-1
231.2 Summary of DBMS_XMLPARSER Subprograms 231-1
231.2.1 FREEPARSER 231-2
231.2.2 GETDOCTYPE 231-3
231.2.3 GETDOCUMENT 231-3
231.2.4 GETRELEASEVERSION 231-3
231.2.5 GETVALIDATIONMODE 231-4
231.2.6 NEWPARSER 231-4
231.2.7 PARSE 231-4
231.2.8 PARSEBUFFER 231-5
231.2.9 PARSECLOB 231-5
231.2.10 PARSEDTD 231-6

231.2.11 PARSEDTDBUFFER 231-6
231.2.12 PARSEDTDCLOB 231-7
231.2.13 SETBASEDIR 231-7
231.2.14 SETDOCTYPE Procedure 231-8
231.2.15 SETERRORLOG Procedure 231-8
231.2.16 SETPRESERVEWHITESPACE 231-8
231.2.17 SETVALIDATIONMODE 231-9
231.2.18 SHOWWARNINGS 231-9

232 DBMS_XMLSCHEMA

232.1 DBMS_XMLSCHEMA Overview 232-1
232.2 DBMS_XMLSCHEMA Security Model 232-1
232.3 DBMS_XMLSCHEMA Constants 232-1
232.4 Views 232-3
232.5 DBMS_XMLSCHEMA Operational Notes 232-4
232.6 Summary of DBMS_XMLSCHEMA Subprograms 232-4
232.6.1 COMPILESCHEMA Procedure 232-5
232.6.2 COPYEVOLVE Procedure 232-5
232.6.3 DELETESCHEMA Procedure 232-7
232.6.4 INPLACEEVOLVE Procedure 232-8
232.6.5 PURGESCHEMA Procedure 232-10
232.6.6 REGISTERSCHEMA Procedures 232-10
232.6.7 REGISTERURI Procedure 232-14

233 DBMS_XMLSCHEMA_ANNOTATE

233.1 DBMS_XMLSCHEMA_ANNOTATE Overview 233-1
233.2 DBMS_XMLSCHEMA_ANNOTATE Security Model 233-1
233.3 Summary of DBMS_XMLSCHEMA_ANNOTATE Subprograms 233-1
233.3.1 ADDXDBNAMESPACE Procedure 233-3
233.3.2 DISABLEDEFAULTTABLECREATION Procedure 233-3
233.3.3 DISABLEMAINTAINDOM Procedure 233-4
233.3.4 ENABLEDEFAULTTABLECREATION Procedure 233-5
233.3.5 ENABLEMAINTAINDOM Procedure 233-5
233.3.6 GETSCHEMAANNOTATIONS Function 233-6
233.3.7 GETSIDXDEFFROMVIEW Function 233-7
233.3.8 PRINTWARNINGS Procedure 233-7
233.3.9 REMOVEANYSTORAGE Procedure 233-8
233.3.10 REMOVEDEFAULTTABLE Procedure 233-8
233.3.11 REMOVEMAINTAINDOM Procedure 233-9
233.3.12 REMOVEOUTOFLINE Procedure 233-9

233.3.13 REMOVESQLCOLLTYPE Procedure 233-10
233.3.14 REMOVESQLNAME Procedure 233-11
233.3.15 REMOVESQLTYPE Procedure 233-12
233.3.16 REMOVESQLTYPEMAPPING Procedure 233-12
233.3.17 REMOVETABLEPROPS Procedure 233-13
233.3.18 REMOVETIMESTAMPWITHTIMEZONE Procedure 233-13
233.3.19 SETANYSTORAGE Procedure 233-14
233.3.20 SETDEFAULTTABLE Procedure 233-15
233.3.21 SETOUTOFLINE Procedure 233-15
233.3.22 SETSCHEMAANNOTATATIONS Procedure 233-17
233.3.23 SETSQLCOLLTYPE Procedure 233-18
233.3.24 SETSQLNAME Procedure 233-19
233.3.25 SETSQLTYPE Procedure 233-20
233.3.26 SETSQLTYPEMAPPING Procedure 233-21
233.3.27 SETTABLEPROPS Procedure 233-21
233.3.28 SETTIMESTAMPWITHTIMEZONE Procedure 233-22

234 DBMS_XMLSCHEMA_UTIL

234.1 DBMS_XMLSCHEMA_UTIL Security Model 234-1
234.2 Summary of DBMS_XMLSCHEMA_UTIL Subprograms 234-1
234.2.1 CONFORMING Function 234-1
234.2.2 VALIDATE Procedure 234-2

235 DBMS_XMLSTORAGE_MANAGE

235.1 DBMS_XMLSTORAGE_MANAGE Overview 235-1
235.2 DBMS_XMLSTORAGE_MANAGE Security Model 235-1
235.3 Summary of DBMS_XMLSTORAGE_MANAGE Subprograms 235-1
235.3.1 DISABLEINDEXESANDCONSTRAINTS Procedure 235-2
235.3.2 ENABLEINDEXESANDCONSTRAINTS Procedure 235-4
235.3.3 EXCHANGEPOSTPROC Procedure 235-5
235.3.4 EXCHANGEPREPROC Procedure 235-5
235.3.5 INDEXXMLREFERENCES Procedure 235-5
235.3.6 REFPARTITIONEXCHANGEIN Procedure 235-6
235.3.7 REFPARTITIONEXCHANGEOUT Procedure 235-7
235.3.8 RENAMECOLLECTIONTABLE Procedure 235-8
235.3.9 SCOPEXMLREFERENCES Procedure 235-10
235.3.10 XPATH2TABCOLMAPPING Function 235-10

236 DBMS_XMLSTORE

236.1 DBMS_XMLSTORE Security Model 236-1
236.2 Types 236-1
236.3 Summary of DBMS_XMLSTORE Subprograms 236-2
236.3.1 CLEARKEYCOLUMNLIST 236-2
236.3.2 CLEARUPDATECOLUMNLIST 236-2
236.3.3 CLOSECONTEXT 236-3
236.3.4 DELETEXML 236-3
236.3.5 INSERTXML 236-4
236.3.6 NEWCONTEXT 236-4
236.3.7 SETKEYCOLUMN 236-5
236.3.8 SETROWTAG 236-5
236.3.9 SETUPDATECOLUMN 236-5
236.3.10 UPDATEXML 236-6

237 DBMS_XPLAN

237.1 DBMS_XPLAN Overview 237-1
237.2 DBMS_XPLAN Security Model 237-2
237.3 DBMS_XPLAN Data Structures 237-2
237.3.1 DBMS_XPLAN PLAN_OBJECT_LIST Table Type 237-2
237.4 Examples 237-4
237.5 Summary of DBMS_XPLAN Subprograms 237-7
237.5.1 COMPARE_PLANS Function 237-7
237.5.2 DIFF_PLAN Function 237-9
237.5.3 DISPLAY Function 237-10
237.5.4 DISPLAY_AWR Function 237-12
237.5.5 DISPLAY_CURSOR Function 237-15
237.5.6 DISPLAY_PLAN Function 237-18
237.5.7 DISPLAY_SQL_PLAN_BASELINE Function 237-20
237.5.8 DISPLAY_SQLSET Function 237-21
237.5.9 DISPLAY_WORKLOAD_REPOSITORY Function 237-25

238 DBMS_XSLPROCESSOR

238.1 DBMS_XSLPROCESSOR Overview 238-1
238.2 DBMS_XSLPROCESSOR Security Model 238-1
238.3 Summary of DBMS_XSLPROCESSOR Subprograms 238-2
238.3.1 FREEPROCESSOR Procedure 238-2
238.3.2 FREESTYLESHEET Procedure 238-2
238.3.3 NEWPROCESSOR Function 238-3

238.3.4 NEWSTYLESHEET Functions 238-3
238.3.5 PROCESSXSL Functions and Procedures 238-4
238.3.6 REMOVEPARAM Procedure 238-6
238.3.7 RESETPARAMS Procedure 238-6
238.3.8 SELECTNODES Function 238-6
238.3.9 SELECTSINGLENODE Function 238-7
238.3.10 SETERRORLOG Procedure 238-7
238.3.11 SETPARAM Procedure 238-8
238.3.12 SHOWWARNINGS Procedure 238-8
238.3.13 TRANSFORMNODE Function 238-9
238.3.14 VALUEOF Function and Procedure 238-9

239 DBMS_XSTREAM_ADM

239.1 DBMS_XSTREAM_ADM Overview 239-1
239.2 DBMS_XSTREAM_ADM Security Model 239-2
239.3 DBMS_XSTREAM_ADM Operational Notes 239-2
239.4 Summary of DBMS_XSTREAM_ADM Subprograms 239-3
239.4.1 ADD_COLUMN Procedure 239-5
239.4.2 ADD_GLOBAL_PROPAGATION_RULES Procedure 239-8
239.4.3 ADD_GLOBAL_RULES Procedure 239-12
239.4.4 ADD_OUTBOUND Procedure 239-16
239.4.5 ADD_SCHEMA_PROPAGATION_RULES Procedure 239-20
239.4.6 ADD_SCHEMA_RULES Procedure 239-24
239.4.7 ADD_SUBSET_OUTBOUND_RULES Procedure 239-28
239.4.8 ADD_SUBSET_PROPAGATION_RULES Procedure 239-31
239.4.9 ADD_SUBSET_RULES Procedure 239-34
239.4.10 ADD_TABLE_PROPAGATION_RULES Procedure 239-39
239.4.11 ADD_TABLE_RULES Procedure 239-43
239.4.12 ALTER_INBOUND Procedure 239-47
239.4.13 ALTER_OUTBOUND Procedure 239-48
239.4.14 CREATE_INBOUND Procedure 239-53
239.4.15 CREATE_OUTBOUND Procedure 239-55
239.4.16 DELETE_COLUMN Procedure 239-60
239.4.17 DROP_INBOUND Procedure 239-62
239.4.18 DROP_OUTBOUND Procedure 239-62
239.4.19 ENABLE_GG_XSTREAM_FOR_STREAMS Procedure 239-63
239.4.20 GET_MESSAGE_TRACKING Function 239-64
239.4.21 GET_TAG Function 239-65
239.4.22 IS_GG_XSTREAM_FOR_STREAMS Function 239-65
239.4.23 KEEP_COLUMNS Procedure 239-66
239.4.24 MERGE_STREAMS Procedure 239-68

239.4.25 MERGE_STREAMS_JOB Procedure 239-71
239.4.26 PURGE_SOURCE_CATALOG Procedure 239-74
239.4.27 RECOVER_OPERATION Procedure 239-75
239.4.28 REMOVE_QUEUE Procedure 239-77
239.4.29 REMOVE_RULE Procedure 239-78
239.4.30 REMOVE_SUBSET_OUTBOUND_RULES Procedure 239-79
239.4.31 REMOVE_XSTREAM_CONFIGURATION Procedure 239-80
239.4.32 RENAME_COLUMN Procedure 239-82
239.4.33 RENAME_SCHEMA Procedure 239-84
239.4.34 RENAME_TABLE Procedure 239-86
239.4.35 SET_MESSAGE_TRACKING Procedure 239-88
239.4.36 SET_PARAMETER Procedure 239-89
239.4.37 SET_TAG Procedure 239-90
239.4.38 SET_UP_QUEUE Procedure 239-91
239.4.39 SPLIT_STREAMS Procedure 239-93
239.4.40 START_OUTBOUND Procedure 239-97
239.4.41 STOP_OUTBOUND Procedure 239-97

240 XSTREAM_CAPTURE

240.1 XSTREAM_CAPTURE Overview 240-1
240.2 XSTREAM_CAPTURE Security Model 240-1
240.3 Summary of XSTREAM_CAPTURE Subprograms 240-1
240.3.1 XSTREAM_APPLY Procedure 240-2
240.3.2 GRANT_REMOTE_ADMIN_ACCESS Procedure 240-5
240.3.3 REVOKE_ADMIN_PRIVILEGE Procedure 240-6
240.3.4 REVOKE_REMOTE_ADMIN_ACCESS Procedure 240-8

241 DEBUG_EXTPROC

241.1 DEBUG_EXTPROC Security Model 241-1
241.2 DEBUG_EXTPROC Operational Notes 241-1
241.3 Rules and Limits 241-2
241.4 Summary of DEBUG_EXTPROC Subprograms 241-2
241.4.1 STARTUP_EXTPROC_AGENT Procedure 241-2

242 HTF

242.1 HTF Deprecated Subprograms 242-1
242.2 HTF Operational Notes 242-1
242.3 Rules and Limits 242-2
242.4 HTF Examples 242-2

242.5 Summary of Tags 242-2
242.6 Summary of HTF Subprograms 242-5
242.6.1 ADDRESS Function 242-11
242.6.2 ANCHOR Function 242-12
242.6.3 ANCHOR2 Function 242-12
242.6.4 APPLETCLOSE Function 242-13
242.6.5 APPLETOPEN Function 242-13
242.6.6 AREA Function 242-14
242.6.7 BASE Function 242-15
242.6.8 BASEFONT Function 242-15
242.6.9 BGSOUND Function 242-16
242.6.10 BIG Function 242-16
242.6.11 BLOCKQUOTECLOSE Function 242-17
242.6.12 BLOCKQUOTEOPEN Function 242-17
242.6.13 BODYCLOSE Function 242-18
242.6.14 BODYOPEN Function 242-18
242.6.15 BOLD Function 242-19
242.6.16 BR Function 242-19
242.6.17 CENTER Function 242-20
242.6.18 CENTERCLOSE Function 242-20
242.6.19 CENTEROPEN Function 242-20
242.6.20 CITE Function 242-21
242.6.21 CODE Function 242-21
242.6.22 COMMENT Function 242-22
242.6.23 DFN Function 242-22
242.6.24 DIRLISTCLOSE Function 242-23
242.6.25 DIRLISTOPEN Function 242-23
242.6.26 DIV Function 242-24
242.6.27 DLISTCLOSE Function 242-24
242.6.28 DLISTDEF Function 242-25
242.6.29 DLISTOPEN Function 242-25
242.6.30 DLISTTERM Function 242-26
242.6.31 EM Function 242-26
242.6.32 EMPHASIS Function 242-27
242.6.33 ESCAPE_SC Function 242-27
242.6.34 ESCAPE_URL Function 242-28
242.6.35 FONTCLOSE Function 242-29
242.6.36 FONTOPEN Function 242-29
242.6.37 FORMAT_CELL Function 242-30
242.6.38 FORMCHECKBOX Function 242-30
242.6.39 FORMCLOSE Function 242-31
242.6.40 FORMFILE Function 242-31

242.6.41 FORMHIDDEN Function 242-32
242.6.42 FORMIMAGE Function 242-32
242.6.43 FORMOPEN Function 242-33
242.6.44 FORMPASSWORD Function 242-34
242.6.45 FORMRADIO Function 242-34
242.6.46 FORMRESET Function 242-35
242.6.47 FORMSELECTCLOSE Function 242-35
242.6.48 FORMSELECTOPEN Function 242-36
242.6.49 FORMSELECTOPTION Function 242-37
242.6.50 FORMSUBMIT Function 242-37
242.6.51 FORMTEXT Function 242-38
242.6.52 FORMTEXTAREA Function 242-38
242.6.53 FORMTEXTAREA2 Function 242-39
242.6.54 FORMTEXTAREACLOSE Function 242-40
242.6.55 FORMTEXTAREAOPEN Function 242-40
242.6.56 FORMTEXTAREAOPEN2 Function 242-41
242.6.57 FRAME Function 242-42
242.6.58 FRAMESETCLOSE Function 242-42
242.6.59 FRAMESETOPEN Function 242-43
242.6.60 HEADCLOSE Function 242-43
242.6.61 HEADER Function 242-44
242.6.62 HEADOPEN Function 242-44
242.6.63 HR Function 242-45
242.6.64 HTMLCLOSE Function 242-45
242.6.65 HTMLOPEN Function 242-45
242.6.66 IMG Function 242-46
242.6.67 IMG2 Function 242-46
242.6.68 ISINDEX Function 242-47
242.6.69 ITALIC Function 242-48
242.6.70 KBD Function 242-48
242.6.71 KEYBOARD Function 242-49
242.6.72 LINE Function 242-49
242.6.73 LINKREL Function 242-50
242.6.74 LINKREV Function 242-50
242.6.75 LISTHEADER Function 242-51
242.6.76 LISTINGCLOSE Function 242-52
242.6.77 LISTINGOPEN Function 242-52
242.6.78 LISTITEM Function 242-52
242.6.79 MAILTO Function 242-53
242.6.80 MAPCLOSE Function 242-54
242.6.81 MAPOPEN Function 242-54
242.6.82 MENULISTCLOSE Function 242-54

242.6.83 MENULISTOPEN Function 242-55
242.6.84 META Function 242-55
242.6.85 NL Function 242-56
242.6.86 NOBR Function 242-56
242.6.87 NOFRAMESCLOSE Function 242-57
242.6.88 NOFRAMESOPEN Function 242-57
242.6.89 OLISTCLOSE Function 242-57
242.6.90 OLISTOPEN Function 242-58
242.6.91 PARA Function 242-58
242.6.92 PARAGRAPH Function 242-59
242.6.93 PARAM Function 242-59
242.6.94 PLAINTEXT Function 242-60
242.6.95 PRECLOSE Function 242-60
242.6.96 PREOPEN Function 242-61
242.6.97 PRINT Functions 242-61
242.6.98 PRN Functions 242-62
242.6.99 S Function 242-63
242.6.100 SAMPLE Function 242-63
242.6.101 SCRIPT Function 242-64
242.6.102 SMALL Function 242-64
242.6.103 STRIKE Function 242-65
242.6.104 STRONG Function 242-65
242.6.105 STYLE Function 242-66
242.6.106 SUB Function 242-66
242.6.107 SUP Function 242-67
242.6.108 TABLECAPTION Function 242-67
242.6.109 TABLECLOSE Function 242-68
242.6.110 TABLEDATA Function 242-68
242.6.111 TABLEHEADER Function 242-69
242.6.112 TABLEOPEN Function 242-70
242.6.113 TABLEROWCLOSE Function 242-70
242.6.114 TABLEROWOPEN Function 242-71
242.6.115 TELETYPE Function 242-71
242.6.116 TITLE Function 242-72
242.6.117 ULISTCLOSE Function 242-72
242.6.118 ULISTOPEN Function 242-73
242.6.119 UNDERLINE Function 242-73
242.6.120 VARIABLE Function 242-74
242.6.121 WBR Function 242-74

243 HTP

243.1 HTP Operational Notes 243-1
243.2 HTP Rules and Limits 243-2
243.3 HTP Examples 243-2
243.4 HTP Summary of Tags 243-2
243.5 Summary of HTP Subprograms 243-5
243.5.1 ADDRESS Procedure 243-11
243.5.2 ANCHOR Procedure 243-11
243.5.3 ANCHOR2 Procedure 243-12
243.5.4 APPLETCLOSE Procedure 243-13
243.5.5 APPLETOPEN Procedure 243-13
243.5.6 AREA Procedure 243-14
243.5.7 BASE Procedure 243-15
243.5.8 BASEFONT Procedure 243-15
243.5.9 BGSOUND Procedure 243-16
243.5.10 BIG Procedure 243-16
243.5.11 BLOCKQUOTECLOSE Procedure 243-17
243.5.12 BLOCKQUOTEOPEN Procedure 243-17
243.5.13 BODYCLOSE Procedure 243-17
243.5.14 BODYOPEN Procedure 243-18
243.5.15 BOLD Procedure 243-18
243.5.16 BR Procedure 243-19
243.5.17 CENTER Procedure 243-19
243.5.18 CENTERCLOSE Procedure 243-20
243.5.19 CENTEROPEN Procedure 243-20
243.5.20 CITE Procedure 243-20
243.5.21 CODE Procedure 243-21
243.5.22 COMMENT Procedure 243-21
243.5.23 DFN Procedure 243-22
243.5.24 DIRLISTCLOSE Procedure 243-22
243.5.25 DIRLISTOPEN Procedure 243-23
243.5.26 DIV Procedure 243-23
243.5.27 DLISTCLOSE Procedure 243-24
243.5.28 DLISTDEF Procedure 243-24
243.5.29 DLISTOPEN Procedure 243-25
243.5.30 DLISTTERM Procedure 243-25
243.5.31 EM Procedure 243-26
243.5.32 EMPHASIS Procedure 243-26
243.5.33 ESCAPE_SC Procedure 243-27
243.5.34 FONTCLOSE Procedure 243-27
243.5.35 FONTOPEN Procedure 243-28

243.5.36 FORMCHECKBOX Procedure 243-28
243.5.37 FORMCLOSE Procedure 243-29
243.5.38 FORMOPEN Procedure 243-29
243.5.39 FORMFILE Procedure 243-30
243.5.40 FORMHIDDEN Procedure 243-30
243.5.41 FORMIMAGE Procedure 243-31
243.5.42 FORMPASSWORD Procedure 243-32
243.5.43 FORMRADIO Procedure 243-32
243.5.44 FORMRESET Procedure 243-33
243.5.45 FORMSELECTCLOSE Procedure 243-33
243.5.46 FORMSELECTOPEN Procedure 243-34
243.5.47 FORMSELECTOPTION Procedure 243-35
243.5.48 FORMSUBMIT Procedure 243-35
243.5.49 FORMTEXT Procedure 243-36
243.5.50 FORMTEXTAREA Procedure 243-36
243.5.51 FORMTEXTAREA2 Procedure 243-37
243.5.52 FORMTEXTAREACLOSE Procedure 243-38
243.5.53 FORMTEXTAREAOPEN Procedure 243-38
243.5.54 FORMTEXTAREAOPEN2 Procedure 243-39
243.5.55 FRAME Procedure 243-39
243.5.56 FRAMESETCLOSE Procedure 243-40
243.5.57 FRAMESETOPEN Procedure 243-40
243.5.58 HEADCLOSE Procedure 243-41
243.5.59 HEADER Procedure 243-41
243.5.60 HEADOPEN Procedure 243-42
243.5.61 HR Procedure 243-42
243.5.62 HTMLCLOSE Procedure 243-43
243.5.63 HTMLOPEN Procedure 243-43
243.5.64 IMG Procedure 243-43
243.5.65 IMG2 Procedure 243-44
243.5.66 ISINDEX Procedure 243-45
243.5.67 ITALIC Procedure 243-45
243.5.68 KBD Procedure 243-46
243.5.69 KEYBOARD Procedure 243-46
243.5.70 LINE Procedure 243-47
243.5.71 LINKREL Procedure 243-47
243.5.72 LINKREV Procedure 243-48
243.5.73 LISTHEADER Procedure 243-49
243.5.74 LISTINGCLOSE Procedure 243-49
243.5.75 LISTINGOPEN Procedure 243-49
243.5.76 LISTITEM Procedure 243-50
243.5.77 MAILTO Procedure 243-50

243.5.78 MAPCLOSE Procedure 243-51
243.5.79 MAPOPEN Procedure 243-51
243.5.80 MENULISTCLOSE Procedure 243-52
243.5.81 MENULISTOPEN Procedure 243-52
243.5.82 META Procedure 243-53
243.5.83 NL Procedure 243-53
243.5.84 NOBR Procedure 243-54
243.5.85 NOFRAMESCLOSE Procedure 243-54
243.5.86 NOFRAMESOPEN Procedure 243-55
243.5.87 OLISTCLOSE Procedure 243-55
243.5.88 OLISTOPEN Procedure 243-55
243.5.89 PARA Procedure 243-56
243.5.90 PARAGRAPH Procedure 243-56
243.5.91 PARAM Procedure 243-57
243.5.92 PLAINTEXT Procedure 243-57
243.5.93 PRECLOSE Procedure 243-58
243.5.94 PREOPEN Procedure 243-58
243.5.95 PRINT Procedures 243-59
243.5.96 PRINTS Procedure 243-59
243.5.97 PRN Procedures 243-60
243.5.98 PS Procedure 243-61
243.5.99 S Procedure 243-61
243.5.100 SAMPLE Procedure 243-62
243.5.101 SCRIPT Procedure 243-62
243.5.102 SMALL Procedure 243-63
243.5.103 STRIKE Procedure 243-64
243.5.104 STRONG Procedure 243-64
243.5.105 STYLE Procedure 243-65
243.5.106 SUB Procedure 243-65
243.5.107 SUP Procedure 243-66
243.5.108 TABLECAPTION Procedure 243-66
243.5.109 TABLECLOSE Procedure 243-67
243.5.110 TABLEDATA Procedure 243-67
243.5.111 TABLEHEADER Procedure 243-68
243.5.112 TABLEOPEN Procedure 243-68
243.5.113 TABLEROWCLOSE Procedure 243-69
243.5.114 TABLEROWOPEN Procedure 243-69
243.5.115 TELETYPE Procedure 243-70
243.5.116 TITLE Procedure 243-70
243.5.117 ULISTCLOSE Procedure 243-71
243.5.118 ULISTOPEN Procedure 243-71
243.5.119 UNDERLINE Procedure 243-72

243.5.120 VARIABLE Procedure 243-72
243.5.121 WBR Procedure 243-73

244 OWA_CACHE

244.1 OWA_CACHE Constants 244-1
244.2 Summary of OWA_CACHE Subprograms 244-1
244.2.1 DISABLE Procedure 244-2
244.2.2 GET_ETAG Function 244-2
244.2.3 GET_LEVEL Function 244-2
244.2.4 SET_CACHE Procedure 244-2
244.2.5 SET_EXPIRES Procedure 244-3
244.2.6 SET_NOT_MODIFIED Procedure 244-3
244.2.7 SET_SURROGATE_CONTROL Procedure 244-4

245 OWA_COOKIE

245.1 OWA_CUSTOM Overview 245-1
245.2 OWA_COOKIE Types 245-1
245.3 OWA_COOKIE Rules and Limits 245-1
245.4 Summary of OWA_COOKIE Subprograms 245-2
245.4.1 GET Function 245-2
245.4.2 GET_ALL Procedure 245-2
245.4.3 REMOVE Procedure 245-3
245.4.4 SEND procedure 245-3

246 OWA_CUSTOM

246.1 OWA_CUSTOM Constants 246-1
246.2 Summary of OWA_CUSTOM Subprograms 246-1
246.2.1 AUTHORIZE Function 246-1

247 OWA_IMAGE

247.1 OWA_IMAGE Overview 247-1
247.2 OWA_IMAGE Types 247-1
247.3 OWA_IMAGE Variables 247-1
247.4 OWA_IMAGE Examples 247-2
247.5 Summary of OWA_IMAGE Subprograms 247-2
247.5.1 GET_X Function 247-2
247.5.2 GET_Y Function 247-3

248 OWA_OPT_LOCK

248.1 OWA_OPT_LOCK Overview 248-1
248.2 OWA_TEXT Types 248-1
248.3 Summary of OWA_OPT_LOCK Subprograms 248-2
248.3.1 CHECKSUM Functions 248-2
248.3.2 GET_ROWID Function 248-3
248.3.3 STORE_VALUES Procedure 248-3
248.3.4 VERIFY_VALUES Function 248-4

249 OWA_PATTERN

249.1 OWA_PATTERN Types 249-1
249.2 OWA_PATTERN Operational Notes 249-1
249.2.1 OWA_PATTERN Wildcards 249-2
249.2.2 OWA_PATTERN Quantifiers 249-2
249.2.3 OWA_PATTERN Flags 249-3
249.3 Summary of OWA_PATTERN Subprograms 249-3
249.3.1 AMATCH Function 249-3
249.3.2 CHANGE Functions and Procedures 249-5
249.3.3 GETPAT Procedure 249-6
249.3.4 MATCH Function 249-6

250 OWA_SEC

250.1 OWA_SEC Operational Notes 250-1
250.2 Summary of OWA_SEC Subprograms 250-1
250.2.1 GET_CLIENT_HOSTNAME Function 250-1
250.2.2 GET_CLIENT_IP Function 250-2
250.2.3 GET_PASSWORD Function 250-2
250.2.4 GET_USER_ID Function 250-2
250.2.5 SET_AUTHORIZATION Procedure 250-3
250.2.6 SET_PROTECTION_REALM Procedure 250-4

251 OWA_TEXT

251.1 OWA_OPT_LOCK Types 251-1
251.1.1 MULTI_LINE DATA TYPE 251-1
251.1.2 ROW_LIST DATA TYPE 251-1
251.1.3 VC_ARR DATA TYPE 251-2
251.2 Summary of OWA_TEXT Subprograms 251-2
251.2.1 ADD2MULTI Procedure 251-2
251.2.2 NEW_ROW_LIST Function and Procedure 251-3

251.2.3 PRINT_MULTI Procedure 251-3
251.2.4 PRINT_ROW_LIST Procedure 251-4
251.2.5 STREAM2MULTI Procedure 251-4

252 OWA_UTIL

252.1 OWA_UTIL Overview 252-1
252.2 OWA_UTIL Types 252-1
252.2.1 DATETYPE Datatype 252-1
252.2.2 IDENT_ARR Datatype 252-2
252.2.3 IP_ADDRESS Datatype 252-2
252.3 Summary of OWA_UTIL Subprograms 252-2
252.3.1 BIND_VARIABLES Function 252-3
252.3.2 CALENDARPRINT Procedures 252-3
252.3.3 CELLSPRINT Procedures 252-4
252.3.4 CHOOSE_DATE Procedure 252-6
252.3.5 GET_CGI_ENV Function 252-7
252.3.6 GET_OWA_SERVICE_PATH Function 252-7
252.3.7 GET_PROCEDURE Function 252-8
252.3.8 HTTP_HEADER_CLOSE Procedure 252-8
252.3.9 LISTPRINT Procedure 252-8
252.3.10 MIME_HEADER Procedure 252-9
252.3.11 PRINT_CGI_ENV Procedure 252-10
252.3.12 REDIRECT_URL Procedure 252-10
252.3.13 SHOWPAGE Procedure 252-11
252.3.14 SHOWSOURCE Procedure 252-12
252.3.15 SIGNATURE procedure 252-12
252.3.16 STATUS_LINE Procedure 252-13
252.3.17 TABLEPRINT Function 252-13
252.3.18 TODATE Function 252-16
252.3.19 WHO_CALLED_ME Procedure 252-17

253 SDO_CS

254 SDO_CSW_PROCESS

255 SDO_GCDR

256 SDO_GEOM

257 SDO_GEOR

258 SDO_GEOR_ADMIN

259 SDO_GEOR_AGGR

260 SDO_GEOR_RA

261 SDO_GEOR_UTL

262 SDO_LRS

263 SDO_MIGRATE

264 SDO_NET

265 SDO_NFE

266 SDO_OLS

267 SDO_PC_PKG

268 SDO_SAM

269 SDO_TIN_PKG

270 SDO_TOPO

271 SDO_TOPO_MAP

272 SDO_TUNE

273 SDO_UTIL

274 SDO_WFS_LOCK

275 SDO_WFS_PROCESS

276 SEM_APIS

277 SEM_OLS

278 SEM_PERF

279 SEM_RDFCTX

280 SEM_RDFSA

281 UTL_CALL_STACK

281.1 UTL_CALL_STACK Overview 281-1
281.2 UTL_CALL_STACK Security Model 281-2
281.3 UTL_CALL_STACK Operational Notes 281-3
281.4 UTL_CALL_STACK Exceptions 281-3
281.5 UTL_CALL_STACK Data Structures 281-3
281.5.1 UNIT_QUALIFIED_NAME 281-3
281.6 Summary of UTL_CALL_STACK Subprograms 281-4
281.6.1 BACKTRACE_DEPTH Function 281-4
281.6.2 BACKTRACE_LINE Function 281-5
281.6.3 BACKTRACE_UNIT Function 281-5
281.6.4 CURRENT_EDITION Function 281-5
281.6.5 CONCATENATE_SUBPROGRAM Function 281-6
281.6.6 DYNAMIC_DEPTH Function 281-6
281.6.7 ERROR_DEPTH Function 281-6
281.6.8 ERROR_MSG Function 281-7
281.6.9 ERROR_NUMBER Function 281-7
281.6.10 LEXICAL_DEPTH Function 281-8
281.6.11 OWNER Function 281-8
281.6.12 UNIT_LINE Function 281-8
281.6.13 SUBPROGRAM Function 281-9

282 UTL_COLL

282.1 Summary of UTL_COLL Subprograms 282-1
282.1.1 IS_LOCATOR Function 282-1

283 UTL_COMPRESS

283.1 UTL_COMPRESS Constants 283-1
283.2 UTL_COMPRESS Exceptions 283-1
283.3 UTL_COMPRESS Operational Notes 283-1
283.4 Summary of UTL_COMPRESS Subprograms 283-2
283.4.1 ISOPEN Function 283-2
283.4.2 LZ_COMPRESS Functions and Procedures 283-3
283.4.3 LZ_COMPRESS_ADD Procedure 283-4
283.4.4 LZ_COMPRESS_CLOSE 283-4
283.4.5 LZ_COMPRESS_OPEN 283-5
283.4.6 LZ_UNCOMPRESS Functions and Procedures 283-6
283.4.7 LZ_UNCOMPRESS_EXTRACT Procedure 283-6
283.4.8 LZ_UNCOMPRESS_OPEN Function 283-7
283.4.9 LZ_UNCOMPRESS_CLOSE Procedure 283-8

284 UTL_ENCODE

284.1 Summary of UTL_ENCODE Subprograms 284-1
284.1.1 BASE64_DECODE Function 284-2
284.1.2 BASE64_ENCODE Function 284-2
284.1.3 MIMEHEADER_DECODE Function 284-3
284.1.4 MIMEHEADER_ENCODE Function 284-4
284.1.5 QUOTED_PRINTABLE_DECODE Function 284-5
284.1.6 QUOTED_PRINTABLE_ENCODE Function 284-5
284.1.7 TEXT_DECODE Function 284-6
284.1.8 TEXT_ENCODE Function 284-7
284.1.9 UUDECODE Function 284-8
284.1.10 UUENCODE Function 284-9

285 UTL_FILE

285.1 UTL_FILE Security Model 285-1
285.2 UTL_FILE Operational Notes 285-2
285.3 UTL_FILE Rules and Limits 285-3
285.4 UTL_FILE Exceptions 285-3
285.5 UTL_FILE Examples 285-4
285.6 UTL_FILE Data Structures 285-5
285.6.1 FILETYPE Record Type 285-5
285.7 Summary of UTL_FILE Subprograms 285-6
285.7.1 FCLOSE Procedure 285-7
285.7.2 FCLOSE_ALL Procedure 285-7

285.7.3 FCOPY Procedure 285-8
285.7.4 FFLUSH Procedure 285-8
285.7.5 FGETATTR Procedure 285-9
285.7.6 FGETPOS Function 285-10
285.7.7 FOPEN Function 285-10
285.7.8 FOPEN_NCHAR Function 285-12
285.7.9 FREMOVE Procedure 285-13
285.7.10 FRENAME Procedure 285-13
285.7.11 FSEEK Procedure 285-14
285.7.12 GET_LINE Procedure 285-15
285.7.13 GET_LINE_NCHAR Procedure 285-16
285.7.14 GET_RAW Procedure 285-17
285.7.15 IS_OPEN Function 285-18
285.7.16 NEW_LINE Procedure 285-18
285.7.17 PUT Procedure 285-19
285.7.18 PUT_LINE Procedure 285-19
285.7.19 PUT_LINE_NCHAR Procedure 285-20
285.7.20 PUT_NCHAR Procedure 285-21
285.7.21 PUTF Procedure 285-22
285.7.22 PUTF_NCHAR Procedure 285-23
285.7.23 PUT_RAW Procedure 285-24

286 UTL_HTTP

286.1 UTL_HTTP Overview 286-1
286.2 UTL_HTTP Security Model 286-2
286.3 UTL_HTTP Constants 286-2
286.4 UTL_HTTP Exceptions 286-5
286.5 UTL_HTTP Examples 286-7
286.5.1 UTL_HTTP General Usage 286-7
286.5.2 UTL_HTTP Retrieving HTTP Response Headers 286-7
286.5.3 UTL_HTTP Handling HTTP Authentication 286-8
286.5.4 UTL_HTTP Handling HTTP Digest Authentication 286-9
286.5.5 UTL_HTTP Retrieving and Restoring Cookies 286-9
286.5.6 UTL_HTTP Making HTTP Request with Private Wallet and Cookie Table 286-10
286.5.7 UTL_HTTP Using a Proxy Server 286-11
286.6 UTL_HTTP Data Structures 286-12
286.6.1 REQ Type 286-12
286.6.2 REQUEST_CONTEXT_KEY Type 286-12
286.6.3 RESP Type 286-13
286.6.4 COOKIE and COOKIE_TABLE Types 286-13
286.6.5 CONNECTION Type 286-14

286.7 UTL_HTTP Operations 286-15
286.7.1 UTL_HTTP Operational Flow 286-15
286.7.2 UTL_HTTP Simple HTTP Fetches 286-18
286.7.3 UTL_HTTP HTTP Requests 286-18
286.7.4 UTL_HTTP HTTP Responses 286-19
286.7.5 UTL_HTTP HTTP Cookies 286-19
286.7.6 UTL_HTTP HTTP Persistent Connections 286-19
286.7.7 UTL_HTTP Error Conditions 286-20
286.7.8 UTL_HTTP Session Settings 286-20
286.7.9 UTL_HTTP Request Context 286-20
286.7.10 UTL_HTTP External Password Store 286-21
286.8 UTL_HTTP Subprogram Groups 286-21
286.8.1 UTL_HTTP Simple HTTP Fetches in a Single Call Subprograms 286-21
286.8.2 UTL_HTTP Session Settings Subprograms 286-21
286.8.3 UTL_HTTP HTTP Requests Subprograms 286-22
286.8.4 UTL_HTTP HTTP Request Contexts Subprograms 286-23
286.8.5 UTL_HTTP HTTP Responses Subprograms 286-23
286.8.6 UTL_HTTP HTTP Cookies Subprograms 286-24
286.8.7 UTL_HTTP HTTP Persistent Connections Subprograms 286-24
286.8.8 UTL_HTTP Error Conditions Subprograms 286-25
286.9 Summary of UTL_HTTP Subprograms 286-25
286.9.1 ADD_COOKIES Procedure 286-29
286.9.2 BEGIN_REQUEST Function 286-30
286.9.3 CLEAR_COOKIES Procedure 286-31
286.9.4 CLOSE_PERSISTENT_CONN Procedure 286-32
286.9.5 CLOSE_PERSISTENT_CONNS Procedure 286-32
286.9.6 CREATE_REQUEST_CONTEXT Function 286-33
286.9.7 DESTROY_REQUEST_CONTEXT Procedure 286-35
286.9.8 END_REQUEST Procedure 286-35
286.9.9 END_RESPONSE Procedure 286-36
286.9.10 GET_AUTHENTICATION Procedure 286-36
286.9.11 GET_BODY_CHARSET Procedure 286-37
286.9.12 GET_COOKIE_COUNT Function 286-38
286.9.13 GET_COOKIE_SUPPORT Procedure 286-38
286.9.14 GET_COOKIES Function 286-39
286.9.15 GET_DETAILED_EXCP_SUPPORT Procedure 286-39
286.9.16 GET_DETAILED_SQLCODE Function 286-40
286.9.17 GET_DETAILED_SQLERRM Function 286-40
286.9.18 GET_FOLLOW_REDIRECT Procedure 286-40
286.9.19 GET_HEADER Procedure 286-41
286.9.20 GET_HEADER_BY_NAME Procedure 286-42
286.9.21 GET_HEADER_COUNT Function 286-42

286.9.22 GET_PERSISTENT_CONN_COUNT Function 286-43
286.9.23 GET_PERSISTENT_CONN_SUPPORT Procedure 286-43
286.9.24 GET_PERSISTENT_CONNS Procedure 286-44
286.9.25 GET_PROXY Procedure 286-45
286.9.26 GET_RESPONSE Function 286-45
286.9.27 GET_RESPONSE_ERROR_CHECK Procedure 286-47
286.9.28 GET_TRANSFER_TIMEOUT Procedure 286-47
286.9.29 READ_LINE Procedure 286-48
286.9.30 READ_RAW Procedure 286-49
286.9.31 READ_TEXT Procedure 286-50
286.9.32 REQUEST Function 286-51
286.9.33 REQUEST_PIECES Function 286-53
286.9.34 SET_AUTHENTICATION Procedure 286-56
286.9.35 SET_AUTHENTICATION_FROM_WALLET Procedure 286-57
286.9.36 SET_BODY_CHARSET Procedures 286-59
286.9.37 SET_COOKIE_SUPPORT Procedures 286-60
286.9.38 SET_DETAILED_EXCP_SUPPORT Procedure 286-61
286.9.39 SET_FOLLOW_REDIRECT Procedures 286-62
286.9.40 SET_HEADER Procedure 286-63
286.9.41 SET_PERSISTENT_CONN_SUPPORT Procedure 286-64
286.9.42 SET_PROXY Procedure 286-67
286.9.43 SET_RESPONSE_ERROR_CHECK Procedure 286-68
286.9.44 SET_TRANSFER_TIMEOUT Procedure 286-68
286.9.45 SET_WALLET Procedure 286-69
286.9.46 WRITE_LINE Procedure 286-70
286.9.47 WRITE_RAW Procedure 286-71
286.9.48 WRITE_TEXT Procedure 286-72

287 UTL_I18N

287.1 UTL_I18N Overview 287-1
287.2 UTL_I18N Security Model 287-2
287.3 UTL_I18N Constants 287-2
287.4 UTL_I18N Data Structures 287-3
287.4.1 CHARSET_RESULT_RECORD Record Type 287-4
287.4.2 LANGUAGE_CHARSET_RESULT_RECORD Record Type 287-4
287.4.3 LANGUAGE_RESULT_RECORD Record Type 287-5
287.4.4 CHARSET_RESULT_TABLE Table Type 287-5
287.4.5 LANGUAGE_CHARSET_RESULT_TABLE Table Type 287-5
287.4.6 LANGUAGE_RESULT_TABLE Table Type 287-6
287.5 Summary of UTL_I18N Subprograms 287-6
287.5.1 DETECT_CHARSET Procedure 287-7

287.5.2 DETECT_LANGUAGE Procedure 287-9
287.5.3 DETECT_LANGUAGE_CHARSET Procedure 287-11
287.5.4 ESCAPE_REFERENCE Function 287-14
287.5.5 GET_COMMON_TIME_ZONES Function 287-15
287.5.6 GET_DEFAULT_CHARSET Function 287-15
287.5.7 GET_DEFAULT_ISO_CURRENCY Function 287-16
287.5.8 GET_DEFAULT_LINGUISTIC_SORT Function 287-17
287.5.9 GET_LOCAL_LANGUAGES Function 287-17
287.5.10 GET_LOCAL_LINGUISTIC_SORTS Function 287-18
287.5.11 GET_LOCAL_TERRITORIES Function 287-19
287.5.12 GET_LOCAL_TIME_ZONES Function 287-20
287.5.13 GET_MAX_CHARACTER_SIZE Function 287-21
287.5.14 GET_TRANSLATION Function 287-22
287.5.15 MAP_CHARSET Function 287-23
287.5.16 MAP_FROM_SHORT_LANGUAGE Function 287-24
287.5.17 MAP_LANGUAGE_FROM_ISO Function 287-25
287.5.18 MAP_LOCALE_TO_ISO Function 287-26
287.5.19 MAP_TERRITORY_FROM_ISO Function 287-27
287.5.20 MAP_TO_SHORT_LANGUAGE Function 287-27
287.5.21 RAW_TO_CHAR Functions 287-28
287.5.22 RAW_TO_NCHAR Functions 287-30
287.5.23 STRING_TO_RAW Function 287-31
287.5.24 TRANSLITERATE Function 287-32
287.5.25 UNESCAPE_REFERENCE Function 287-34
287.5.26 VALIDATE_CHARACTER_ENCODING Functions 287-35

288 UTL_IDENT

288.1 UTL_IDENT Overview 288-1
288.2 UTL_IDENT Security Model 288-1
288.3 UTL_IDENT Constants 288-1

289 UTL_INADDR

289.1 UTL_INADDR Security Model 289-1
289.2 UTL_INADDR Exceptions 289-1
289.3 UTL_INADDR Examples 289-1
289.4 Summary of UTL_INADDR Subprograms 289-2
289.4.1 GET_HOST_ADDRESS Function 289-2
289.4.2 GET_HOST_NAME Function 289-3

290 UTL_LMS

290.1 UTL_LMS Security Model 290-1
290.2 Summary of UTL_LMS Subprograms 290-1
290.2.1 FORMAT_MESSAGE Function 290-1
290.2.2 GET_MESSAGE Function 290-2

291 UTL_MAIL

291.1 UTL_MAIL Security Model 291-1
291.2 UTL_MAIL Operational Notes 291-1
291.3 UTL_MAIL Rules and Limits 291-1
291.4 Summary of UTL_MAIL Subprograms 291-2
291.4.1 SEND Procedure 291-2
291.4.2 SEND_ATTACH_RAW Procedure 291-3
291.4.3 SEND_ATTACH_VARCHAR2 Procedure 291-4

292 UTL_MATCH

292.1 UTL_MATCH Overview 292-1
292.2 UTL_MATCH Security Model 292-1
292.3 Summary of UTL_MATCH Subprograms 292-2
292.3.1 EDIT_DISTANCE Function 292-2
292.3.2 EDIT_DISTANCE_SIMILARITY Function 292-2
292.3.3 JARO_WINKLER Function 292-3
292.3.4 JARO_WINKLER_SIMILARITY Function 292-4

293 UTL_NLA

293.1 UTL_NLA Overview 293-1
293.2 UTL_NLA Rules and Limits 293-2
293.3 UTL_NLA Security Model 293-2
293.4 Subprogram Groups 293-2
293.4.1 UTL_NLA BLAS Level 1 (Vector-Vector Operations) Subprograms 293-2
293.4.2 UTL_NLA BLAS Level 2 (Matrix-Vector Operations) Subprograms 293-3
293.4.3 UTL_NLA BLAS Level 3 (Matrix-Matrix Operations) Subprograms 293-4
293.4.4 UTL_NLA LAPACK Driver Routines (Linear Equations) Subprograms 293-5
293.4.5 UTL_NLA LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms 293-6
293.5 Summary of UTL_NLA Subprograms 293-7
293.5.1 BLAS_ASUM Functions 293-13
293.5.2 BLAS_AXPY Procedures 293-14
293.5.3 BLAS_COPY Procedures 293-15

293.5.4 BLAS_DOT Functions 293-16
293.5.5 BLAS_GBMV Procedures 293-17
293.5.6 BLAS_GEMM Procedures 293-19
293.5.7 BLAS_GEMV Procedures 293-21
293.5.8 BLAS_GER Procedures 293-23
293.5.9 BLAS_IAMAX Functions 293-24
293.5.10 BLAS_NRM2 Functions 293-25
293.5.11 BLAS_ROT Procedures 293-26
293.5.12 BLAS_ROTG Procedures 293-27
293.5.13 BLAS_SCAL Procedures 293-27
293.5.14 BLAS_SPMV Procedures 293-28
293.5.15 BLAS_SPR Procedures 293-30
293.5.16 BLAS_SPR2 Procedures 293-31
293.5.17 BLAS_SBMV Procedures 293-33
293.5.18 BLAS_SWAP Procedures 293-34
293.5.19 BLAS_SYMM Procedures 293-35
293.5.20 BLAS_SYMV Procedures 293-37
293.5.21 BLAS_SYR Procedures 293-39
293.5.22 BLAS_SYR2 Procedures 293-40
293.5.23 BLAS_SYR2K Procedures 293-42
293.5.24 BLAS_SYRK Procedures 293-44
293.5.25 BLAS_TBMV Procedures 293-46
293.5.26 BLAS_TBSV Procedures 293-48
293.5.27 BLAS_TPMV Procedures 293-50
293.5.28 BLAS_TPSV Procedures 293-51
293.5.29 BLAS_TRMM Procedures 293-53
293.5.30 BLAS_TRMV Procedures 293-55
293.5.31 BLAS_TRSM Procedures 293-56
293.5.32 BLAS_TRSV Procedures 293-58
293.5.33 LAPACK_GBSV Procedures 293-60
293.5.34 LAPACK_GEES Procedures 293-61
293.5.35 LAPACK_GELS Procedures 293-63
293.5.36 LAPACK_GESDD Procedures 293-65
293.5.37 LAPACK_GESV Procedure 293-67
293.5.38 LAPACK_GESVD Procedures 293-69
293.5.39 LAPACK_GEEV Procedures 293-71
293.5.40 LAPACK_GTSV Procedure 293-73
293.5.41 LAPACK_PBSV Procedures 293-75
293.5.42 LAPACK_POSV Procedures 293-76
293.5.43 LAPACK_PPSV Procedures 293-78
293.5.44 LAPACK_PTSV Procedures 293-80
293.5.45 LAPACK_SBEV Procedures 293-81

293.5.46 LAPACK_SBEVD Procedures 293-83
293.5.47 LAPACK_SPEV Procedures 293-85
293.5.48 LAPACK_SPEVD Procedures 293-86
293.5.49 LAPACK_SPSV Procedures 293-88
293.5.50 LAPACK_STEV Procedures 293-90
293.5.51 LAPACK_STEVD Procedures 293-91
293.5.52 LAPACK_SYEV Procedures 293-93
293.5.53 LAPACK_SYEVD Procedures 293-94
293.5.54 LAPACK_SYSV Procedures 293-95

294 UTL_RAW

294.1 UTL_RAW Overview 294-1
294.2 UTL_RAW Operational Notes 294-1
294.3 Summary of UTL_RAW Subprograms 294-1
294.3.1 BIT_AND Function 294-3
294.3.2 BIT_COMPLEMENT Function 294-4
294.3.3 BIT_OR Function 294-4
294.3.4 BIT_XOR Function 294-5
294.3.5 CAST_FROM_BINARY_DOUBLE Function 294-6
294.3.6 CAST_FROM_BINARY_FLOAT Function 294-7
294.3.7 CAST_FROM_BINARY_INTEGER Function 294-8
294.3.8 CAST_FROM_NUMBER Function 294-8
294.3.9 CAST_TO_BINARY_DOUBLE Function 294-9
294.3.10 CAST_TO_BINARY_FLOAT Function 294-10
294.3.11 CAST_TO_BINARY_INTEGER Function 294-11
294.3.12 CAST_TO_NUMBER Function 294-11
294.3.13 CAST_TO_NVARCHAR2 Function 294-12
294.3.14 CAST_TO_RAW Function 294-13
294.3.15 CAST_TO_VARCHAR2 Function 294-13
294.3.16 COMPARE Function 294-14
294.3.17 CONCAT Function 294-15
294.3.18 CONVERT Function 294-16
294.3.19 COPIES Function 294-17
294.3.20 LENGTH Function 294-18
294.3.21 OVERLAY Function 294-18
294.3.22 REVERSE Function 294-20
294.3.23 SUBSTR Function 294-20
294.3.24 TRANSLATE Function 294-21
294.3.25 TRANSLITERATE Function 294-23
294.3.26 XRANGE Function 294-25

295 UTL_RECOMP

295.1 UTL_RECOMP Overview 295-1
295.2 UTL_RECOMP Operational Notes 295-1
295.3 UTL_RECOMP Examples 295-2
295.4 Summary of UTL_RECOMP Subprograms 295-2
295.4.1 RECOMP_PARALLEL Procedure 295-2
295.4.2 RECOMP_SERIAL Procedure 295-4

296 UTL_REF

296.1 UTL_REF Overview 296-1
296.2 UTL_REF Security Model 296-1
296.3 UTL_REF Types 296-2
296.4 UTL_REF Exceptions 296-2
296.5 Summary of UTL_REF Subprograms 296-3
296.5.1 DELETE_OBJECT Procedure 296-3
296.5.2 LOCK_OBJECT Procedure 296-4
296.5.3 SELECT_OBJECT Procedure 296-5
296.5.4 UPDATE_OBJECT Procedure 296-6

297 UTL_RPADV

297.1 UTL_RPADV Overview 297-1
297.2 DBMS_COMPARISON Security Model 297-1
297.3 UTL_RPADV Operational Notes 297-1
297.4 Summary of UTL_RPADV Subprograms 297-10
297.4.1 ALTER_MONITORING Procedure 297-11
297.4.2 COLLECT_STATS Procedure 297-12
297.4.3 IS_MONITORING Function 297-13
297.4.4 SHOW_STATS Procedure 297-14
297.4.5 SHOW_STATS_HTML Procedure 297-16
297.4.6 START_MONITORING Procedure 297-18
297.4.7 STOP_MONITORING Procedure 297-20

298 UTL_SMTP

298.1 UTL_SMTP Overview 298-1
298.2 UTL_SMTP Security Model 298-2
298.3 UTL_SMTP Constants 298-2
298.4 UTL_SMTP Types 298-2
298.5 UTL_SMTP Reply Codes 298-4
298.6 UTL_SMTP Operational Notes 298-5

298.7 UTL_SMTP Exceptions 298-5
298.8 UTL_SMTP Rules and Limits 298-6
298.9 UTL_SMTP Examples 298-6
298.10 Summary of UTL_SMTP Subprograms 298-7
298.10.1 AUTH Function and Procedure 298-8
298.10.2 CLOSE_CONNECTION Procedure 298-9
298.10.3 CLOSE_DATA Function and Procedure 298-10
298.10.4 COMMAND Function and Procedure 298-10
298.10.5 COMMAND_REPLIES Function 298-11
298.10.6 DATA Function and Procedure 298-12
298.10.7 EHLO Function and Procedure 298-13
298.10.8 HELO Function and Procedure 298-14
298.10.9 HELP Function 298-15
298.10.10 MAIL Function and Procedure 298-15
298.10.11 NOOP Function and Procedure 298-16
298.10.12 OPEN_CONNECTION Functions 298-17
298.10.13 OPEN_DATA Function and Procedure 298-19
298.10.14 QUIT Function and Procedure 298-20
298.10.15 RCPT Function 298-20
298.10.16 RSET Function and Procedure 298-21
298.10.17 STARTTLS Function and Procedure 298-22
298.10.18 VRFY Function 298-23
298.10.19 WRITE_DATA Procedure 298-24
298.10.20 WRITE_RAW_DATA Procedure 298-25

299 UTL_TCP

299.1 UTL_TCP Overview 299-1
299.2 UTL_TCP Security Model 299-1
299.3 UTL_TCP Types 299-1
299.4 UTL_TCP Exceptions 299-3
299.5 UTL_TCP Rules and Limits 299-3
299.6 UTL_TCP Examples 299-3
299.7 Summary of UTL_TCP Subprograms 299-4
299.7.1 AVAILABLE Function 299-5
299.7.2 CLOSE_ALL_CONNECTIONS Procedure 299-6
299.7.3 CLOSE_CONNECTION Procedure 299-7
299.7.4 FLUSH Procedure 299-7
299.7.5 GET_LINE Function 299-8
299.7.6 GET_LINE_NCHAR Function 299-8
299.7.7 GET_RAW Function 299-9
299.7.8 GET_TEXT Function 299-10

299.7.9 GET_TEXT_NCHAR Function 299-11
299.7.10 OPEN_CONNECTION Function 299-12
299.7.11 READ_LINE Function 299-14
299.7.12 READ_RAW Function 299-15
299.7.13 READ_TEXT Function 299-16
299.7.14 SECURE_CONNECTION Procedure 299-17
299.7.15 WRITE_LINE Function 299-17
299.7.16 WRITE_RAW Function 299-18
299.7.17 WRITE_TEXT Function 299-18

300 UTL_URL

300.1 UTL_URL Overview 300-1
300.2 UTL_URL Exceptions 300-2
300.3 UTL_URL Examples 300-2
300.4 Summary of UTL_URL Subprograms 300-2
300.4.1 ESCAPE Function 300-3
300.4.2 UNESCAPE Function 300-4

301 WPG_DOCLOAD

301.1 WPG_DOCLOAD Constants 301-1
301.2 Summary of WPG_DOCLOAD Subprograms 301-1
301.2.1 DOWNLOAD_FILE Procedures 301-2

302 ANYDATA TYPE

302.1 ANYDATA TYPE Restrictions 302-1
302.2 ANYDATA TYPE Operational Notes 302-1
302.3 Summary of ANYDATA Subprograms 302-2
302.3.1 BEGINCREATE Static Procedure 302-3
302.3.2 ENDCREATE Member Procedure 302-3
302.3.3 GET* Member Functions 302-4
302.3.4 GETTYPE Member Function 302-6
302.3.5 GETTYPENAME Member Function 302-7
302.3.6 PIECEWISE Member Procedure 302-7
302.3.7 SET* Member Procedures 302-8

303 ANYDATASET TYPE

303.1 ANYDATASET TYPE Construction 303-1
303.2 Summary of ANYDATASET TYPE Subprograms 303-1

303.2.1 ADDINSTANCE Member Procedure 303-2
303.2.2 BEGINCREATE Static Procedure 303-3
303.2.3 ENDCREATE Member Procedure 303-3
303.2.4 GET* Member Functions 303-3
303.2.5 GETCOUNT Member Function 303-6
303.2.6 GETINSTANCE Member Function 303-6
303.2.7 GETTYPE Member Function 303-7
303.2.8 GETTYPENAME Member Function 303-7
303.2.9 PIECEWISE Member Procedure 303-8
303.2.10 SET* Member Procedures 303-9

304 ANYTYPE TYPE

304.1 Summary of ANYTYPE Subprograms 304-1
304.1.1 BEGINCREATE Static Procedure 304-2
304.1.2 SETINFO Member Procedure 304-2
304.1.3 ADDATTR Member Procedure 304-3
304.1.4 ENDCREATE Member Procedure 304-4
304.1.5 GETINFO Member Function 304-5
304.1.6 GETATTRELEMINFO Member Function 304-6
304.1.7 GETANYTYPEFROMPERSISTENT Function 304-7

305 Oracle Database Advanced Queuing (AQ) Types

305.1 Advanced Queuing (AQ)Types Security Model 305-1
305.2 Oracle Database Advanced Queuing (AQ) Types —Summary of Types 305-1
305.2.1 AQ$_AGENT Type 305-2
305.2.2 AQ$_AGENT_LIST_T Type 305-2
305.2.3 AQ$_DESCRIPTOR Type 305-3
305.2.4 AQ$_NTFN_DESCRIPTOR Type 305-3
305.2.5 AQ$_NTFN_MSGID_ARRAY Type 305-4
305.2.6 AQ$_POST_INFO Type 305-4
305.2.7 AQ$_POST_INFO_LIST Type 305-4
305.2.8 AQ$_PURGE_OPTIONS_T Type 305-5
305.2.9 AQ$_RECIPIENT_LIST_T Type 305-5
305.2.10 AQ$_REG_INFO Type 305-6
305.2.11 AQ$_REG_INFO_LIST Type 305-8
305.2.12 AQ$_SUBSCRIBER_LIST_T Type 305-8
305.2.13 DEQUEUE_OPTIONS_T Type 305-9
305.2.14 ENQUEUE_OPTIONS_T Type 305-11
305.2.15 QUEUE_PROPS_T Type 305-12
305.2.16 SEEK_INPUT_T Type 305-13

305.2.17 SEEK_OUTPUT_T Type 305-13
305.2.18 SYS.MSG_PROP_T Type 305-14
305.2.19 MESSAGE_PROPERTIES_T Type 305-16
305.2.20 MESSAGE_PROPERTIES_ARRAY_T Type 305-19
305.2.21 MSGID_ARRAY_T Type 305-20

306 DBFS Content Interface Types

306.1 DDBFS Content Interface Types Overview 306-1
306.2 DBFS Content Interface Types Security Model 306-1
306.3 DBFS Content Interface Types —Data Structures 306-1
306.3.1 DBMS_DBFS_CONTENT_CONTEXT_T Object Type 306-2
306.3.2 DBMS_DBFS_CONTENT_LIST_ITEM_T Object Type 306-2
306.3.3 DBMS_DBFS_CONTENT_PROPERTY_T Object Type 306-3
306.3.4 DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type 306-3
306.3.5 DBMS_DBFS_CONTENT_PROPERTIES_T Table Type 306-4
306.3.6 DBMS_DBFS_CONTENT_RAW_T Table Type 306-4

307 Database URI TYPEs

307.1 Summary of URITYPE Supertype Subprograms 307-1
307.1.1 UriType Supertype GETBLOB 307-2
307.1.2 UriType Supertype GETCLOB 307-3
307.1.3 UriType Supertype GETCONTENTTYPE 307-3
307.1.4 UriType Supertype GETEXTERNALURL 307-3
307.1.5 UriType Supertype GETURL 307-4
307.1.6 UriType Supertype GETXML 307-4
307.2 Summary of HTTPURITYPE Subtype Subprograms 307-4
307.2.1 HttpUriType Subtype CREATEURI 307-5
307.2.2 HttpUriType Subtype GETBLOB 307-5
307.2.3 HttpUriType Subtype GETCLOB 307-6
307.2.4 HttpUriType Subtype GETCONTENTTYPE 307-6
307.2.5 HttpUriType Subtype GETEXTERNALURL 307-7
307.2.6 HttpUriType Subtype GETURL 307-7
307.2.7 HttpUriType Subtype GETXML 307-7
307.2.8 HttpUriType Subtype HTTPURITYPE 307-8
307.3 Summary of DBURITYPE Subtype Subprogams 307-8
307.3.1 DBUriType Subtype CREATEURI 307-9
307.3.2 DBUriType Subtype DBURITYPE 307-9
307.3.3 DBUriType Subtype GETBLOB 307-9
307.3.4 DBUriType Subtype GETCLOB 307-10
307.3.5 DBUriType Subtype GETCONTENTTYPE 307-10

307.3.6 DBUriType Subtype GETEXTERNALURL 307-11
307.3.7 DBUriType Subtype GETURL 307-11
307.3.8 DBUriType Subtype GETXML 307-11
307.4 Summary of XDBURITYPE Subtype Subprograms 307-11
307.4.1 XDBUriType Subtype CREATEURI 307-12
307.4.2 XDBUriType Subtype GETBLOB 307-12
307.4.3 XDBUriType Subtype GETCLOB 307-13
307.4.4 XDBUriType Subtype GETCONTENTTYPE 307-14
307.4.5 XDBUriType Subtype GETEXTERNALURL 307-14
307.4.6 XDBUriType Subtype GETURL 307-14
307.4.7 XDBUriType Subtype GETXML 307-14
307.4.8 XDBUriType subtype XDBURITYPE 307-15
307.5 Summary of URIFACTORY Package Subprograms 307-15
307.5.1 UriFactory Package GETURI 307-16
307.5.2 UriFactory Package ESCAPEURI 307-16
307.5.3 UriFactory Package UNESCAPEURI 307-16
307.5.4 UriFactory Package REGISTERURLHANDLER 307-17
307.5.5 UriFactory Package UNREGISTERURLHANDLER 307-17

308 JMS Types

308.1 JMS Types Overview 308-1
308.2 JMS Types Security Model 308-2
308.3 Java Versus PL/SQL Datatypes 308-2
308.4 More on Bytes, Stream and Map Messages 308-3
308.5 Upcasting and Downcasting Between General and Specific Messages 308-6
308.6 JMS Types Error Reporting 308-6
308.7 Oracle JMS Type Constants 308-7
308.8 CONVERT_JMS_SELECTOR 308-8
308.9 Summary of JMS Types 308-10
308.9.1 SYS.AQ$_JMS_MESSAGE Type 308-10
308.9.2 SYS.AQ$_JMS_TEXT_MESSAGE Type 308-16
308.9.3 SYS.AQ$_JMS_BYTES_MESSAGE Type 308-19
308.9.4 SYS.AQ$_JMS_MAP_MESSAGE Type 308-28
308.9.5 SYS.AQ$_JMS_STREAM_MESSAGE Type 308-37
308.9.6 SYS.AQ$_JMS_OBJECT_MESSAGE Type 308-46
308.9.7 SYS.AQ$_JMS_NAMEARRAY Type 308-46
308.9.8 SYS.AQ$_JMS_VALUE Type 308-46
308.9.9 SYS.AQ$_JMS_EXCEPTION Type 308-47

309 JSON Data Structures

309.1 JSON_ELEMENT_T Object Type 309-1
309.2 JSON_OBJECT_T Object Type 309-4
309.3 JSON_ARRAY_T Object Type 309-8
309.4 JSON_SCALAR_T Object Type 309-10
309.5 JSON_KEY_LIST Type 309-10

310 Logical Change Record TYPEs

310.1 Logical Change Record TYPEs Overview 310-1
310.2 Logical Change Record Types Security Model 310-1
310.3 Summary of Logical Change Record Types 310-1
310.3.1 LCR$_DDL_RECORD Type 310-2
310.3.2 LCR$_ROW_RECORD Type 310-10
310.3.3 Common Subprograms for LCR$_DDL_RECORD and LCR$_ROW_RECORD 310-34
310.3.4 LCR$_ROW_LIST Type 310-44
310.3.5 LCR$_ROW_UNIT Type 310-45

311 MGD_ID Package Types

311.1 MGD_ID Package Types Security Model 311-1
311.2 Summary of Types 311-2
311.2.1 MGD_ID_COMPONENT Object Type 311-2
311.2.2 MGD_ID_COMPONENT_VARRAY Object Type 311-2
311.2.3 MGD_ID Object Type 311-3
311.3 Summary of MGD_ID Subprograms 311-3
311.3.1 MGD_ID Constructor Function 311-4
311.3.2 FORMAT Function 311-7
311.3.3 GET_COMPONENT Function 311-8
311.3.4 TO_STRING Function 311-9
311.3.5 TRANSLATE Function 311-10

312 Rule TYPEs

312.1 Rule TYPEs Overview 312-1
312.2 Rule TYPEs Security Model 312-1
312.3 Summary of Rule Types 312-1
312.3.1 RE$ATTRIBUTE_VALUE Type 312-3
312.3.2 RE$ATTRIBUTE_VALUE_LIST Type 312-3
312.3.3 RE$COLUMN_VALUE Type 312-3
312.3.4 RE$COLUMN_VALUE_LIST Type 312-4
312.3.5 RE$NAME_ARRAY Type 312-4

312.3.6 RE$NV_ARRAY Type 312-4
312.3.7 RE$NV_LIST Type 312-4
312.3.8 RE$NV_NODE Type 312-7
312.3.9 RE$RULE_HIT Type 312-7
312.3.10 RE$RULE_HIT_LIST Type 312-8
312.3.11 RE$RULE_NAME Type 312-8
312.3.12 RE$RULE_NAME_LIST Type 312-8
312.3.13 RE$TABLE_ALIAS Type 312-8
312.3.14 RE$TABLE_ALIAS_LIST Type 312-9
312.3.15 RE$TABLE_VALUE Type 312-10
312.3.16 RE$TABLE_VALUE_LIST Type 312-10
312.3.17 RE$VALUE Type 312-10
312.3.18 RE$VALUE_LIST Type 312-11
312.3.19 RE$VARIABLE_TYPE Type 312-11
312.3.20 RE$VARIABLE_TYPE_LIST Type 312-13
312.3.21 RE$VARIABLE_VALUE Type 312-13
312.3.22 RE$VARIABLE_VALUE_LIST Type 312-13

313 SODA Types

313.1 SODA Types Overview 313-1
313.2 SODA Types Security Model 313-1
313.3 Summary of SODA Types 313-2
313.3.1 SODA_Collection_T Type 313-2
313.3.1.1 CREATE_INDEX Function 313-3
313.3.1.2 CREATE_VIEW_FROM_DG Function 313-4
313.3.1.3 DROP_INDEX Function 313-5
313.3.1.4 FIND Function 313-6
313.3.1.5 FIND_ONE Function 313-6
313.3.1.6 GET_DATA_GUIDE Function 313-7
313.3.1.7 GET_INDEX Function 313-7
313.3.1.8 GET_METADATA Function 313-8
313.3.1.9 GET_NAME Function 313-8
313.3.1.10 INSERT_ONE Function 313-8
313.3.1.11 INSERT_ONE_AND_GET Function 313-9
313.3.1.12 LIST_INDEXES Function 313-9
313.3.1.13 REMOVE_ONE Function 313-10
313.3.1.14 REPLACE_ONE Function 313-10
313.3.1.15 REPLACE_ONE_AND_GET Function 313-11
313.3.1.16 SAVE Function 313-12
313.3.1.17 SAVE_AND_GET Function 313-12
313.3.1.18 TRUNCATE Function 313-13

313.3.2 SODA_Document_T Type 313-13
313.3.2.1 GET_BLOB Function 313-14
313.3.2.2 GET_CLOB Function 313-14
313.3.2.3 GET_CREATED_ON Function 313-15
313.3.2.4 GET_DATA_TYPE Function 313-15
313.3.2.5 GET_JSON Function 313-15
313.3.2.6 GET_KEY Function 313-16
313.3.2.7 GET_LAST_MODIFIED Function 313-16
313.3.2.8 GET_MEDIA_TYPE Function 313-16
313.3.2.9 GET_VARCHAR2 Function 313-17
313.3.2.10 GET_VERSION Function 313-17
313.3.2.11 SODA_Document_T Function 313-17
313.3.3 SODA_OPERATION_T Type 313-18
313.3.3.1 ACQUIRE_LOCK Function 313-20
313.3.3.2 AS_OF_SCN Function 313-21
313.3.3.3 AS_OF_TIMESTAMP Function 313-21
313.3.3.4 COUNT Function 313-22
313.3.3.5 FILTER Function 313-22
313.3.3.6 GET_CURSOR Function 313-23
313.3.3.7 GET_DATA_GUIDE Function 313-23
313.3.3.8 GET_ONE Function 313-24
313.3.3.9 HINT Function 313-24
313.3.3.10 KEY Function 313-25
313.3.3.11 KEYS Function 313-25
313.3.3.12 LIMIT Function 313-26
313.3.3.13 REMOVE Function 313-26
313.3.3.14 REPLACE_ONE Function 313-26
313.3.3.15 REPLACE_ONE_AND_GET Function 313-27
313.3.3.16 SAMPLE Function 313-28
313.3.3.17 SKIP Function 313-28
313.3.3.18 VERSION Function 313-29
313.3.4 SODA_CURSOR_T Type 313-29
313.3.4.1 CLOSE Function 313-30
313.3.4.2 HAS_NEXT Function 313-30
313.3.4.3 NEXT Function 313-30

314 UTL Streams Types

314.1 UTL Streams Types Security Model 314-1
314.2 Summary of UTL Binary Streams Types 314-1
314.2.1 UTL_BINARYINPUTSTREAM Type 314-1
314.2.2 UTL_BINARYOUTPUTSTREAM Type 314-2

314.2.3 UTL_CHARACTERINPUTSTREAM Type 314-3
314.2.4 UTL_CHARACTEROUTPUTSTREAM Type 314-4

315 XMLTYPE

315.1 Summary of XMLType Subprograms 315-1
315.1.1 CREATENONSCHEMABASEDXML 315-2
315.1.2 CREATESCHEMABASEDXML 315-2
315.1.3 CREATEXML 315-3
315.1.4 EXISTSNODE 315-5
315.1.5 EXTRACT 315-5
315.1.6 GETBLOBVAL 315-6
315.1.7 GETCLOBVAL 315-6
315.1.8 GETNAMESPACE 315-6
315.1.9 GETNUMBERVAL 315-6
315.1.10 GETROOTELEMENT 315-7
315.1.11 GETSCHEMAURL 315-7
315.1.12 GETSTRINGVAL 315-7
315.1.13 ISFRAGMENT 315-7
315.1.14 ISSCHEMABASED 315-7
315.1.15 ISSCHEMAVALID 315-8
315.1.16 ISSCHEMAVALIDATED 315-8
315.1.17 SCHEMAVALIDATE 315-8
315.1.18 SETSCHEMAVALIDATED 315-9
315.1.19 TOOBJECT 315-9
315.1.20 TRANSFORM 315-9
315.1.21 XMLTYPE 315-10

Index

Preface

This Preface contains these topics:

- Audience
- Related Documents
- Conventions

Audience

Oracle Database PL/SQL Packages and Types Reference is intended for programmers, systems analysts, project managers, and others interested in developing database applications. This manual assumes a working knowledge of application programming and familiarity with SQL to access information in relational database systems. Some sections also assume a knowledge of basic object-oriented programming.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Related Documents

For more information, see the following documents:

- *Oracle Database Development Guide*
- *Oracle Database PL/SQL Language Reference*

Many of the examples in this book use the sample schemas, which are installed by default when you select the Basic Installation option with an Oracle Database installation. Refer to *Oracle Database Sample Schemas* for information on how these schemas were created and how you can use them yourself.

Printed documentation is available for sale in the Oracle Store at

`http://shop.oracle.com/`

To download free release notes, installation documentation, technical briefs, or other collateral, please visit the Oracle Technology Network (OTN).

`http://docs.oracle.com/`

Conventions

This section describes the conventions used in the text and code examples of this documentation set. It describes:

- Conventions in Text
- Conventions in Code Examples

Conventions in Text

We use various conventions in text to help you more quickly identify special terms. The following table describes those conventions and provides examples of their use.

Convention	Meaning	Example
Bold	Bold typeface indicates terms that are defined in the text or terms that appear in a glossary, or both.	When you specify this clause, you create an **index-organized table**.
Italics	Italic typeface indicates book titles or emphasis.	*Oracle Database Concepts* Ensure that the recovery catalog and target database do *not* reside on the same disk.
`UPPERCASE monospace (fixed-width) font`	Uppercase monospace typeface indicates elements supplied by the system. Such elements include parameters, privileges, datatypes, RMAN keywords, SQL keywords, SQL*Plus or utility commands, packages and methods, as well as system-supplied column names, database objects and structures, usernames, and roles.	You can specify this clause only for a `NUMBER` column. You can back up the database by using the `BACKUP` command. Query the `TABLE_NAME` column in the `USER_TABLES` data dictionary view. Use the `DBMS_STATS.GENERATE_STATS` procedure.
`lowercase monospace (fixed-width) font`	Lowercase monospace typeface indicates executable programs, filenames, directory names, and sample user-supplied elements. Such elements include computer and database names, net service names and connect identifiers, user-supplied database objects and structures, column names, packages and classes, usernames and roles, program units, and parameter values. *Note:* Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.	Enter `sqlplus` to start SQL*Plus. The password is specified in the `orapwd` file. Back up the datafiles and control files in the `/disk1/oracle/dbs` directory. The `department_id`, `department_name`, and `location_id` columns are in the `hr.departments` table. Set the `QUERY_REWRITE_ENABLED` initialization parameter to `true`. The `JRepUtil` class implements these methods.
`lowercase italic monospace (fixed-width) font`	Lowercase italic monospace font represents placeholders or variables.	You can specify the *`parallel_clause`*. Run *`old_release`*`.SQL` where *`old_release`* refers to the release you installed prior to upgrading.

Conventions in Code Examples

Code examples illustrate SQL, PL/SQL, SQL*Plus, or other command-line statements. They are displayed in a monospace (fixed-width) font and separated from normal text as shown in this example:

```
SELECT username FROM dba_users WHERE username = 'MIGRATE';
```

The following table describes typographic conventions used in code examples and provides examples of their use.

Convention	Meaning	Example
`[ ]`	Anything enclosed in brackets is optional.	`DECIMAL (digits [ , precision ])`
`{ }`	Braces are used for grouping items.	`{ENABLE | DISABLE}`
`|`	A vertical bar represents a choice of two options.	`{ENABLE | DISABLE}` `[COMPRESS | NOCOMPRESS]`
`...`	Ellipsis points mean repetition in syntax descriptions. In addition, ellipsis points can mean an omission in code examples or text.	`CREATE TABLE ... AS subquery;` `SELECT col1, col2, ... , coln FROM employees;`
Other symbols	You must use symbols other than brackets ([ ]), braces ({ }), vertical bars (|), and ellipsis points (...) exactly as shown.	`acctbal NUMBER(11,2);` `acct CONSTANT NUMBER(4) := 3;`
Italics	Italicized text indicates placeholders or variables for which you must supply particular values.	`enter password` `DB_NAME = database_name`

`UPPERCASE`	Uppercase typeface indicates elements supplied by the system. We show these terms in uppercase in order to distinguish them from terms you define. Unless terms appear in brackets, enter them in the order and with the spelling shown. Because these terms are not case sensitive, you can use them in either UPPERCASE or lowercase.	`SELECT last_name, employee_id FROM employees;`
		`SELECT * FROM USER_TABLES;`
		`DROP TABLE hr.employees;`

`lowercase`	Lowercase typeface indicates user-defined programmatic elements, such as names of tables, columns, or files.
	Note: Some programmatic elements use a mixture of UPPERCASE and lowercase. Enter these elements as shown.
		`SELECT last_name, employee_id FROM employees;`
		`sqlplus user_id/password`
		`CREATE USER mjones IDENTIFIED BY ty3MU9;`

Changes in This Release for Oracle Database PL/SQL Packages and Types Reference

This preface contains:

- Changes in this Release

Changes in this Release

Review changes in this release of *Oracle Database PL/SQL Packages and Types Reference*.

- New Features
- Deprecated Features
- Desupported Features

New Features

This section lists the major new features in Oracle Database release 23ai *PL/SQL Packages and Types Reference*.

The following features are new in this release:

- DBMS_AUTOIM provides functions to manage the execution of the Automatic In-Memory feature.
- DBMS_SPM supports the SQL plan management feature by providing an interface for the DBA or other users to perform controlled manipulation of plan history and SQL plan baselines maintained for various SQL statements.
- DBMS_INMEMORY_ADVISE enables you to determine if a workload can benefit from Oracle's Database In-Memory feature.
- The `DBMS_VECTOR` package is a lightweight package that simplifies common operations with Oracle AI Vector Search, such as chunking and embedding data, generating text for prompts and images, or managing vector indexes.

  DBMS_VECTOR
- The `DBMS_VECTOR_CHAIN` package enables advanced operations with Oracle AI Vector Search, such as chunking and embedding data, generating text for prompts and images along with text processing and end-to-end similarity search.

  DBMS_VECTOR_CHAIN
- The `DBMS_HYBRID_VECTOR` package enables you to search by both vectors and keywords against hybrid vector indexes.

  DBMS_HYBRID_VECTOR
- `DBMS_DICTIONARY_CHECK` is a read-only and lightweight PL/SQL package procedure that helps you identify database dictionary inconsistencies.

DBMS_DICTIONARY_CHECK

- The `UTL_HTTP` package supports digest SHA-256.

  UTL_HTTP

- The `UTL_I18N` package supports three new procedures that detect the most likely Oracle character sets and Oracle languages for a given data sample: `DETECT_CHARSET`, `DETECT_LANGUAGE`, and `DETECT_LANGUAGE_CHARSET`.

  UTL_I18N

- The `DBMS_SAGA_ADM` package provides a collection of saga administration functions and procedures to define and manage saga participants, coordinators, and brokers.

  DBMS_SAGA_ADM

- The `DBMS_SAGA` package provides a collection of saga functions and procedures to initiate and finalize sagas.

  DBMS_SAGA

- The DBMS_FLASHBACK_ARCHIVE package includes the GET_CURRENT_LIFESPAN_DIGEST Function that generates the current lifespan digest for the specified row in a user table, and the VERIFY_BLOCKCHAIN_LIFESPAN Procedure that verifies the contents of the current lifespans, historical lifespans, or all lifespans of the user table rows that are protected using blockchain Flashback Archive.
- The DBMS_KAFKA package provides a PL/SQL interface for enabling Oracle SQL access to topics in Kafka clusters.
- The `DBMS_SEARCH` package enables you to create, maintain, and query ubiquitous search indexes.

  DBMS_SEARCH

- The `DBMS_SQL_FIREWALL` package enables users to administer SQL Firewall.

  DBMS_SQL_FIREWALL

- The `DBMS_CRYPTO` package has the following new features.

  DBMS_CRYPTO

  - New APIs for elliptic-curve Diffie–Hellman (ECDH) operations
    - `ECDH_GENKEYPAIR`: This function generates an EC public/private key pair
    - `ECDHDERIVE_SHAREDSECRET`: This function derives shared secret using private key of local application and public key from the remote application.
  - New `PKENCRYPT`/`PKDECRYPT` algorithm: `PKENCRYPT_RSA_PKCS1_OAEP_SHA2`
  - New chain modes `GCM CCM`
  - New `DBMS_CRYPTO` block cipher suites `AES_CCM_NONE` and `AES_GCM_NONE`
  - New signature and verification algorithms:
    - `SIGN_SHA224_ECDSA`
    - `SIGN_SHA256_ECDSA`
    - `SIGN_SHA384_ECDSA`
    - `SIGN_SHA512_ECDSA`
    - `SIGN_ECDSA`

- The DBMS_SCHEDULER package has a new DUMP_IN_MEMORY_TRACE procedure that dumps the scheduler in-memory trace buffer of the specified process state object address into the current trace file of the requester process.

  DBMS_SCHEDULER

- The DBMS_SQL and DBMS_TF packages are updated to support the ISO SQL standard-compliant BOOLEAN data type.

  DBMS_SQL

  DBMS_TF

- The DBMS_COMPRESSION package is updated to support the Advanced Low Index Compression for IOTs.

  DBMS_COMPRESSION

- The DBMS_SHARDING_DIRECTORY package is added to support the Oracle Globally Distributed Database directory-based data distribution method.

  DBMS_SHARDING_DIRECTORY

- The DBMS_AUTO_CLUSTERING package is added to support the automatic clustering and zone mapping.

  DBMS_AUTO_CLUSTERING

- The DBMS_AUTO_INDEX package has a new parameter AUTO_INDEX_INCLUDE_DML_COST for the DBMS_AUTO_INDEX.CONFIGURE procedure.

  DBMS_AUTO_INDEX

- The DBMS_BLOCKCHAIN_TABLE package has enhancements related to Blockchain Table User Chains, Blockchain Table Delegate Signer, and Blockchain Table Countersignature.

  DBMS_BLOCKCHAIN_TABLE

- The DBMS_DATA_MINING package is enhanced to support the following settings:

  – IMPORT_ONNX_MODEL procedure. To learn more, see IMPORT_ONNX_MODEL Procedure.

  Generalized Linear Model setting:

  – GLMS_LINK_FUNCTION. See DBMS_DATA_MINING — Algorithm Settings: Generalized Linear Model.

  XGBoost settings:

  – xgboost_interaction_constraints
  – xgboost_decrease_constraints
  – xgboost_increase_constraints
  – objective: survival:aft
  – xgboost_aft_loss_distribution
  – xgboost_aft_loss_distribution_scale
  – xgboost_aft_right_bound_column_name
    See DBMS_DATA_MINING — Algorithm Settings: XGBoost.

  Explicit Semantic Analysis settings:

  – ESAS_EMBEDDINGS

  - `ESAS_EMBEDDING_SIZE`
    See DBMS_DATA_MINING — Algorithm Settings: Explicit Semantic Analysis.

  Expectation Maximization setting:

  - `EMCS_OUTLIER_RATE`
    See DBMS_DATA_MINING — Algorithm Settings: Expectation Maximization.

  Exponential Smoothing settings:

  - `EXSM_SERIES_LIST`
  - `EXSM_INITVL_OPTIMIZE`
    See DBMS_DATA_MINING — Algorithm Settings: Exponential Smoothing.

  k-Means setting:

  - `KMNS_WINSORIZE`
    See DBMS_DATA_MINING — Algorithm Settings: *k*-Means.

  Global settings:

  - `ODMS_BOXCOX`
  - `ODMS_EXPLOSION_MIN_SUPP`
    See DBMS_DATA_MINING — Global Settings.

- The `DBMS_USERDIAG` package is a new package that allows you to perform a limited set of diagnosis operations on the PDB, such as establish trace operations. Most of the regular diagnostic mechanisms available in a CDB have been restricted out of security concerns.

  See DBMS_USERDIAG.

- The `DBMS_PIPE` package has been enhanced to support singleton pipes as well as persistent messaging using Cloud Object stores.

  DBMS_PIPE

- The `DBMS_MLE` package is updated to support the use of JavaScript modules.

  DBMS_MLE

- The `DBMS_AQMIGTOOL` package simplifies migration from Oracle Database Advanced Queuing (AQ) to Transactional Event Queue (TxEventQ) with orchestration automation, source and target compatibility diagnostics and remediation, and a unified user experience.

  DBMS_AQMIGTOOL

- The `DBMS_XMLSCHEMA_UTIL` package provides an interface for XML schema validation.

  DBMS_XMLSCHEMA_UTIL

- The `DBMS_SPACE` package provides an interface for analyzing and shrinking a bigfile tablespace.

  DBMS_SPACE

- DBMS_CACHEUTIL has new `TRUE_CACHE_KEEP` and `TRUE_CACHE_UNKEEP` procedures to manage assignments to the `KEEP` buffer pool for Oracle True Cache.

- The following new procedures are introduced in the `DBMS_APP_CONT` package to enhance the Application Continuity functionality:

  - `APPLY_REPLAY_RULE`
  - `GET_REPLAY_RULES`
  - `RESET_REPLAY_RULES`

- The following new procedures are introduced in the `DBMS_APP_CONT_ADMIN` package to enhance the Application Continuity protection check:
  - `ENABLE_AC`
  - `ENABLE_TAC`
  - `ACCHK_SET_FILTER`
  - `ACCHK_SHOW_FILTERS`
  - `ACCHK_CLEAR_FILTER`
  - `DISABLE_FAILOVER`
  - `ENABLE_RESET_STATE`
  - `MODIFY_SERVICE`
  - `SET_DRAINING`
  - `SET_LOAD_BALANCING`

Deprecated Features

Review this list of deprecated features in Oracle Database release 23ai *Oracle Database PL/SQL Packages and Types Reference*.

Oracle recommends that you do not use deprecated features or values in new applications. Support for deprecated features is for backward compatibility only. For more information about deprecated features, see *Oracle Database Upgrade Guide*.

Deprecation of the mkstore wallet management command line tool

The `mkstore` wallet management command line tool is deprecated with Oracle Database 23ai, and can be removed in a future release.
To manage wallets, Oracle recommends that you use the `orapki` command line tool.

Deprecation of DBMS_RESULT_CACHE Function Names

Oracle is changing the names of several DBMS_RESULT_CACHE function names in Oracle Database 23ai.
The following functions and procedures are deprecated:

- `BLACK_LIST` function. Use `BLOCK_LIST` function.
- BLACK_LIST_ADD procedure. Use BLOCK_LIST_ADD procedure.
- BLACK_LIST_CLEAR procedure. Use BLOCK_LIST_CLEAR procedure
- BLACK_LIST_REMOVE procedure. Use BLOCK_LIST_REMOVE procedure
- OBJECT_BLACK_LIST function. Use OBJECT_BLOCK_LIST function
- OBJECT_BLACK_LIST_ADD procedure. Use OBJECT_BLOCK_LIST_ADD procedure.
- OBJECT_BLACK_LIST_CLEAR procedure. Use OBJECT BLOCK LIST_CLEAR procedure.
- OBJECT_BLACK_LIST_REMOVE procedure. Use OBJECT_BLOCK_LIST_REMOVE procedure.

Deprecation of DBMS_XMLSTORE

The PL/SQL package `DBMS_XMLSTORE` is deprecated in Oracle Database 23ai.

DBMS_XMLSTORE is a non-standard Oracle-proprietary package that enables you to store and manipulate XML data in Oracle Database. This package is deprecated, and can be desupported in a future release. Oracle recommends that you use regular SQL DML and with standard XQuery and SQL/XML to store and manage XML data. Using standard functionality provides future-proof way to store and manipulate XML data.

Deprecation of DBMS_XMLGEN PL/SQL Package

The PL/SQL package DBMS_XMLGEN is deprecated in Oracle Database 23ai.
DBMS_XMLGEN is a non-standard Oracle-proprietary package that is provided to generate and convert XML documents from SQL queries or with PL/SQL. This package is deprecated, and can be desupported in a future release. Oracle recommends that you use SQL/XML operators to generate XML from relational columns instead. Using ANSI SQL/XML operators for any generation and modification of XML documents provides a standardized and future-proof way to work with XML documents.

Deprecation of XML DB Repository

The Oracle XML DB Repository is deprecated with Oracle Database 23ai.
Oracle recommends that you replace any functionality used in XML DB Repository with alternative technologies.

Deprecation of DBMS_HANG_MANAGER Package

The DBMS_HANG_MANAGER package is deprecated in Oracle Database 23ai. Use DBMS_BLOCKER_RESOLVER instead.
The DBMS_HANG_MANAGER package provides a method of changing some configuration parameters and constraints to address session issues. This package is being replaced with DBMS_BLOCKER_RESOLVER. DBMS_HANG_MANAGER can be removed in a future release.

Deprecation of Traditional Auditing Packages and Functions

Traditional auditing packages and functions are deprecated in Oracle Database 23ai.

With the desupport of traditional auditing, the PL/SQL packages and functions associated with traditional auditing are deprecated, This deprecation includes the packages and functions INIT_CLEANUP, DEINIT_CLEANUP, and IS_CLEANUP_INITIALIZED. While these packages or functions continue to operate in Oracle Database 23ai, you can neither add to or modify traditional auditing configurations.

Deprecation of Oracle OLAP

Analytic workspaces, the OLAP DML programming language, financial reporting, and the OLAP Java API continue to be deprecated in Oracle Database 23ai.

Be aware that OLAP will not be supported beyond the term of the current release (Oracle Database 23ai) premier support. Oracle strongly recommends that you do not start new projects using OLAP and begin migrating applications using OLAP to alternatives now. If your application requires an in-database dimensional model, then consider using Oracle Analytic Views. Analytic views provide a dimensional semantic model, calculations, and query semantics using data in Oracle Database. When used with columnar tables, analytic views provide query performance similar to the OLAP Option. If your application requires support for advanced dimensional analytics, what-if analysis, or forecasting, then consider Oracle Essbase. Oracle Essbase is a multidimensional database management system with support for complex dimensional business analytics.

Related Topics

- Oracle Database Changes Desupports and Deprecations in *Oracle Database Upgrade Guide*

Desupported Features

This section lists the desupported features in Oracle Database Release 23ai *Oracle Database PL/SQL Packages and Types Reference*.

For more information about desupported features in this release, see *Oracle Database Upgrade Guide*

The following features are desupported in this release:

Desupport of Oracle Enterprise Manager Database Express

Oracle Enterprise Manager Database Express (EM Express) is desupported in Oracle Database Release 23ai.
EM Express is a web-based database management tool that is built inside Oracle Database. It supports key performance management and basic database administration functions. EM Express was deprecated in Oracle Database 21c. Many of EM Express's capabilities are now available in Oracle Cloud Infrastructure (OCI) Database Management service, Oracle Enterprise Manager Cloud Control, or Oracle SQL Developer.

Desupport of Service Attribute Value SESSION_STATE_CONSISTENCY = STATIC Parameter

The service attribute values `FAILOVER_TYPE = TRANSACTION` with `SESSION_STATE_CONSISTENCY = STATIC` are no longer a supported service attribute combination.
In previous releases, you could use the service parameter `SESSION_STATE_CONSISTENCY` to manage session state automatically using Application Continuity by setting `SESSION_STATE_CONSISTENCY` to `DYNAMIC` or `STATIC`. However, starting with Oracle Database 23ai, you can no longer use the `STATIC` option. Instead, use one of the following failover options:

- `FAILOVER_TYPE = AUTO` with `SESSION_STATE_CONSISTENCY = AUTO`
- `FAILOVER_TYPE = TRANSACTION` with `SESSION_STATE_CONSISTENCY = DYNAMIC`

These configurations enforce session state tracking in Oracle Database, ensuring that session state is preserved at session migration and session failover.

Desupport of Oracle Wallet Manager (OWM)

Starting with Oracle Database 23ai, the Oracle Wallet Manager (OWM) is desupported. Oracle recommends using the `orapki` command line tool to replace OWM.

Related Topics

- Oracle Database Changes Desupports and Deprecations in *Oracle Database Upgrade Guide*

1

Introduction to Oracle Supplied PL/SQL Packages & Types

Oracle supplies many PL/SQL packages with the Oracle server to extend database functionality and provide PL/SQL access to SQL features. You can use the supplied packages when creating your applications or for ideas in creating your own stored procedures.

This manual covers the packages provided with the Oracle database server. Packages supplied with other products, such as Oracle Developer or the Oracle Application Server, are not covered.

Note that not every package or type described in this manual or elsewhere in the Oracle Database Documentation Library is installed by default. In such cases, the documentation states this and explains how to install the object. Run this query as a suitably privileged user:

```
SELECT DISTINCT Owner, Object_Type, Object_Name FROM DBA_Objects_AE
     WHERE Owner IN (
       'SYS', 'OUTLN', 'SYSTEM', 'CTXSYS', 'DBSNMP',
       'LOGSTDBY_ADMINISTRATOR', 'ORDSYS',
       'ORDPLUGINS', 'OEM_MONITOR', 'MDSYS', 'LBACSYS',
       'DMSYS', 'WMSYS', 'OLAPDBA', 'OLAPSVR', 'OLAP_USER',
       'OLAPSYS', 'EXFSYS', 'SYSMAN', 'MDDATA',
       'SI_INFORMTN_SCHEMA', 'XDB', 'ODM')
     AND Object_Type IN ('PACKAGE', 'TYPE')
     ORDER BY Owner, Object_Type, Object_Name
```

This lists every Oracle-supplied package and type that is currently installed in the database. Note that it lists a number of objects not mentioned in the Oracle Database Documentation Library. This is deliberate. Some of the Oracle-supplied packages and types are intended to be used only by other Oracle-supplied components. Any package or type that is not described in the Oracle Database Documentation Library is not supported for direct customer use.

This chapter contains the following topics:

- Package Overview
- Summary of Oracle Supplied PL/SQL Packages and Types

> **See Also:**
>
> *Oracle Database Development Guide* for information on how to create your own packages

1.1 Package Overview

A *package* is an encapsulated collection of related program objects stored together in the database. Program objects are procedures, functions, variables, constants, cursors, and exceptions.

1.1.1 Package Components

PL/SQL packages have two parts: the specification and the body, although sometimes the body is unnecessary. The specification is the interface to your application; it declares the types, variables, constants, exceptions, cursors, and subprograms available for use. The body fully defines cursors and subprograms, and so implements the specification.

Unlike subprograms, packages cannot be called, parameterized, or nested. However, the formats of a package and a subprogram are similar:

```
CREATE PACKAGE name AS  -- specification (visible part)
   -- public type and item declarations
   -- subprogram specifications
END [name];

CREATE PACKAGE BODY name AS  -- body (hidden part)
   -- private type and item declarations
   -- subprogram bodies
[BEGIN
   -- initialization statements]
END [name];
```

The specification holds public declarations that are visible to your application. The body holds implementation details and private declarations that are hidden from your application. You can debug, enhance, or replace a package body without changing the specification. You can change a package body without recompiling calling programs because the implementation details in the body are hidden from your application.

1.1.2 Using Oracle Supplied Packages

Most Oracle supplied packages are automatically installed when the database is created. Certain packages are not installed automatically. Special installation instructions for these packages are documented in the individual chapters.

To call a PL/SQL function from SQL, you must either own the function or have `EXECUTE` privileges on the function. To select from a view defined with a PL/SQL function, you must have `SELECT` privileges on the view. No separate `EXECUTE` privileges are needed to select from the view. Instructions on special requirements for packages are documented in the individual chapters.

1.1.3 Creating New Packages

To create packages and store them permanently in an Oracle database, use the `CREATE PACKAGE` and `CREATE PACKAGE BODY` statements. You can execute these statements interactively from SQL*Plus or Enterprise Manager.

To create a new package, do the following:

1. Create the package specification with the `CREATE PACKAGE` statement.

   You can declare program objects in the package specification. Such objects are called *public* objects. Public objects can be referenced outside the package, as well as by other objects in the package.

> **Note:**
>
> It is often more convenient to add the `OR REPLACE` clause in the `CREATE PACKAGE` statement. But note that `CREATE PACKAGE` warns you if you are about to overwrite an existing package with the same name while `CREATE OR REPLACE` just overwrites it with no warning.

2. Create the package body with the `CREATE PACKAGE BODY` statement.

   You can declare and define program objects in the package body.

   - You must define public objects declared in the package specification.
   - You can declare and define additional package objects, called *private* objects. Private objects are declared in the package body rather than in the package specification, so they can be referenced only by other objects in the package. They cannot be referenced outside the package.

> **See Also:**
>
> – *Oracle Database PL/SQL Language Reference*
> – *Oracle Database Development Guide* for more information on creating new packages
> – *Oracle Database Concepts*
>
> for more information on storing and executing packages

1.1.3.1 Separating the Specification and Body

The specification of a package declares the public types, variables, constants, and subprograms that are visible outside the immediate scope of the package. The body of a package defines the objects declared in the specification, as well as private objects that are not visible to applications outside the package.

Oracle stores the specification and body of a package separately in the database. Other schema objects that call or reference public program objects depend only on the package specification, not on the package body. Using this distinction, you can change the definition of a program object in the package body without causing Oracle to invalidate other schema objects that call or reference the program object. Oracle invalidates dependent schema objects only if you change the declaration of the program object in the package specification.

1.1.3.1.1 Creating a New Package: Example

The following example shows a package specification for a package named `EMPLOYEE_MANAGEMENT`. The package contains one stored function and two stored procedures.

```
CREATE PACKAGE employee_management AS
   FUNCTION hire_emp (name VARCHAR2, job VARCHAR2,
      mgr NUMBER, hiredate DATE, sal NUMBER, comm NUMBER,
      deptno NUMBER) RETURN NUMBER;
   PROCEDURE fire_emp (emp_id NUMBER);
   PROCEDURE sal_raise (emp_id NUMBER, sal_incr NUMBER);
END employee_management;
```

The body for this package defines the function and the procedures:

```
CREATE PACKAGE BODY employee_management AS
   FUNCTION hire_emp (name VARCHAR2, job VARCHAR2,
      mgr NUMBER, hiredate DATE, sal NUMBER, comm NUMBER,
      deptno NUMBER) RETURN NUMBER IS
```

The function accepts all arguments for the fields in the employee table except for the employee number. A value for this field is supplied by a sequence. The function returns the sequence number generated by the call to this function.

```
       new_empno    NUMBER(10);

   BEGIN
      SELECT emp_sequence.NEXTVAL INTO new_empno FROM dual;
      INSERT INTO emp VALUES (new_empno, name, job, mgr,
         hiredate, sal, comm, deptno);
      RETURN (new_empno);
   END hire_emp;

   PROCEDURE fire_emp(emp_id IN NUMBER) AS
```

The procedure deletes the employee with an employee number that corresponds to the argument `emp_id`. If no employee is found, then an exception is raised.

```
   BEGIN
      DELETE FROM emp WHERE empno = emp_id;
      IF SQL%NOTFOUND THEN
      raise_application_error(-20011, 'Invalid Employee
         Number: ' || TO_CHAR(emp_id));
   END IF;
END fire_emp;

PROCEDURE sal_raise (emp_id IN NUMBER, sal_incr IN NUMBER) AS
```

The procedure accepts two arguments. `Emp_id` is a number that corresponds to an employee number. `Sal_incr` is the amount by which to increase the employee's salary.

```
   BEGIN

   -- If employee exists, then update salary with increase.

      UPDATE emp
         SET sal = sal + sal_incr
         WHERE empno = emp_id;
      IF SQL%NOTFOUND THEN
         raise_application_error(-20011, 'Invalid Employee
            Number: ' || TO_CHAR(emp_id));
      END IF;
   END sal_raise;
END employee_management;
```

> **Note:**
>
> If you want to try this example, then first create the sequence number `emp_sequence`. You can do this using the following SQL*Plus statement:
>
> ```
> SQL> CREATE SEQUENCE emp_sequence
>    > START WITH 8000 INCREMENT BY 10;
> ```

1.1.4 Referencing Package Contents

When you want to reference the types, items, and subprograms declared in a package specification, use the dot notation.

For example:

```
package_name.type_name
package_name.item_name
package_name.subprogram_name
```

1.2 Summary of Oracle Supplied PL/SQL Packages and Types

PL/SQL packages are supplied with the Oracle server that extend database functionality and provide PL/SQL access to SQL features.

These packages, listed in the following table, run as the invoking user, rather than the package owner. Unless otherwise noted, the packages are callable through public synonyms of the same name.

> **Note:**
>
> - The procedures and functions provided in these packages and their external interfaces are reserved by Oracle and are subject to change.
> - Modifying Oracle supplied packages can cause internal errors and database security violations. Do not modify supplied packages.

Table 1-1 Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
`APEX_APPLICATION` through `APEX_ZIP`	For a complete description of all the packages that ship with Oracle Application Express, see *Oracle Application Express API Reference*
`CTX_ADM`	Lets you administer servers and the data dictionary
`CTX_ANL`	The `CTX_ANL` package is used with `AUTO_LEXER` and provides procedures for adding and dropping a custom dictionary from the lexer.
`CTX_CLS`	Lets you generate `CTXRULE` rules for a set of documents
`CTX_DDL`	Lets you create and manage the preferences, section lists and stopgroups required for Text indexes
`CTX_DOC`	Lets you request document services
`CTX_ENTITY`	The `CTX_ENTITY` package enables you to search for terms that are unknown to you without specifying a particular search text.
`CTX_OUTPUT`	Lets you manage the index log
`CTX_QUERY`	Lets you generate query feedback, count hits, and create stored query expressions
`CTX_REPORT`	Lets you create various index reports

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
CTX_THES	Lets you to manage and browse thesauri
CTX_ULEXER	For use with the user-lexer
DBMS_ADDM	Facilitates the use of Advisor functionality regarding the Automatic Database Diagnostic Monitor
DBMS_ADVANCED_REWRITE	Contains interfaces for advanced query rewrite users to create, drop, and maintain functional equivalence declarations for query rewrite
DBMS_ADVISOR	Part of the SQLAccess Advisor, an expert system that identifies and helps resolve performance problems relating to the execution of SQL statements
DBMS_AUTOIM	Provides functions to manage the execution of the AIM feature
DBMS_ALERT	Provides support for the asynchronous notification of database events
DBMS_APP_CONT	The DBMS_APP_CONT package provides an interface to determine if the in-flight transaction on a now unavailable session committed or not, and if the last call on that session completed or not.
DBMS_APP_CONT_ADMIN	This package provides a collection dba level admin operations in relation to Application Continuity
DBMS_APP_CONT_REPORT	This procedure generates a report of ACCHK trace collected while ACCHK_SET(TRUE)
DBMS_APPLICATION_INFO	Lets you register an application name with the database for auditing or performance tracking purposes
DBMS_APPLY_ADM	Provides administrative procedures to start, stop, and configure an apply process
DBMS_AQ	Lets you add a message (of a predefined object type) onto a queue or to dequeue a message
DBMS_AQADM	Lets you perform administrative functions on a queue or queue table for messages of a predefined object type
DBMS_AQELM	Provides procedures to manage the configuration of Advanced Queuing asynchronous notification by e-mail and HTTP
DBMS_AQIN	Plays a part in providing secure access to the Oracle JMS interfaces
DBMS_AQMIGTOOL	The DBMS_AQMIGTOOL package provides procedures for migrating AQ queues toTransactional Event Queues (TxEventQ)
DBMS_ASSERT	Provides an interface to validate properties of the input value
DBMS_AUDIT_UTIL	The DBMS_AUDIT_UTIL package provides functions that enable you to format the output of queries to the DBA_FGA_AUDIT_TRAIL, DBA_AUDIT_TRAIL, UNIFIED_AUDIT_TRAIL, and V$XML_AUDIT_TRAIL views.
DBMS_AUDIT_MGMT	The DBMS_AUDIT_MGMT package provides subprograms to manage audit trail records. These subprograms enable audit administrators to manage the audit trail.

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_AUTO_REPORT	The DBMS_AUTO_REPORT package provides an interface to view SQL Monitoring and Real-time Automatic Database Diagnostic Monitor (ADDM) data that has been captured into Automatic Workload Repository (AWR). It also provides subprograms to control the behavior of how these data are captured to AWR.
DBMS_AUTO_SQLTUNE	The DBMS_AUTO_SQLTUNE package is the interface for managing the Automatic SQL Tuning task
DBMS_AUTO_INDEX	The DBMS_AUTO_INDEX package provides the interface for managing auto indexes in an Oracle database
DBMS_AUTO_MV	The DBMS_AUTO_MV contains subprograms for configuring automatic materialized views.
DBMS_AUTO_TASK_ADMIN	Used by the DBA as well as Enterprise Manager to access the AUTOTASK controls
DBMS_AUTO_ZONEMAP	The DBMS_AUTO_ZONEMAP package provides autonomous maintenance and creation of zonemap
DBMS_AVTUNE	The DBMS_AVTUNE package analyzes query access levels and aggregation workloads for each auto-cache enabled analytic view (AV) and creates or drops AV auto-caches to improve the overall performance of the SQL queries on that AV.
DBMS_AW_STATS	Contains a subprogram that generates and stores optimizer statistics for cubes and dimensions
DBMS_BLOCKCHAIN_TABLE	A blockchain table is an append-only table designed for centralized blockchain applications.
DBMS_BLOCKER_RESOLVER	The DBMS_BLOCKER_RESOLVER package provides a method of changing some Blocker Resolver configuration parameters.
DBMS_CACHEUTIL	Provides an interface to control object-level cache fusion locking policies in an Oracle Real Application Clusters (Oracle RAC) environment and to manage assignments to the KEEP buffer pool for Oracle True Cache.
DBMS_CAPTURE_ADM	Describes administrative procedures to start, stop, and configure a capture process; used in Replication
DBMS_CLOUD	The DBMS_CLOUD package provides comprehensive support for working with data in ObjectStorage
DBMS_COMPARISON	Provides interfaces to compare and converge database objects at different databases
DBMS_COMPRESSION	Provides an interface to facilitate choosing the correct compression level for an application
DBMS_CONNECTION_POOL	Provides an interface to manage the Database Resident Connection Pool
DBMS_CQ_NOTIFICATION	Is part of a set of features that clients use to receive notifications when result sets of a query have changed. The package contains interfaces that can be used by mid-tier clients to register objects and specify delivery mechanisms.
DBMS_CREDENTIAL	The DBMS_CREDENTIAL package provides an interface for authenticating and impersonating EXTPROC callout functions, as well as external jobs, remote jobs and file watchers from the SCHEDULER.

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_CRYPTO	Lets you encrypt and decrypt stored data, can be used in conjunction with PL/SQL programs running network communications, and supports encryption and hashing algorithms
DBMS_CSX_ADMIN	Provides an interface to customize the setup when transporting a tablespace containing binary XML data
DBMS_CUBE	Contains subprograms that create OLAP cubes and dimensions, and that load and process the data for querying
DBMS_CUBE_ADVISE	Contains subprograms for evaluating cube materialized views to support log-based fast refresh and query rewrite
DBMS_CUBE_LOG	DBMS_CUBE_LOG contains subprograms for creating and managing logs for cubes and cube dimensions.
DBMS_DATA_MINING	Implements the Oracle Data Mining interface for creating, evaluating, and managing mining models
DBMS_DATA_MINING_TRANSFORM	Provides subroutines that can be used to prepare data for Oracle Data Mining
DBMS_DATAPUMP	Lets you move all, or part of, a database between databases, including both data and metadata
DBMS_DB_VERSION	The DBMS_DB_VERSION package specifies the Oracle version numbers and other information useful for simple conditional compilation selections based on Oracle versions.
DBMS_DBCOMP	The DBMS_DBCOMP package performs block comparison to detect lost writes or database inconsistencies between a primary database and one or more physical standby databases.
DBMS_DBFS_CONTENT	Provides an interface comprising a file system-like abstraction backed by one or more Store Providers
DBMS_DBFS_CONTENT_SPI	Provides the Application Programming Interface (API) specification for DBMS_DBFS_CONTENT service providers
DBMS_DBFS_HS	Provides users the ability to use tape or Amazon S3 Web service as a storage tier when doing Information Lifecycle Management for their database tables
DBMS_DBFS_SFS	Provides an interface to operate a SecureFile-based store (SFS) for the content interface described in the DBMS_DBFS_CONTENT package
DBMS_DDL	Provides access to some SQL DDL statements from stored procedures, and provides special administration operations not available as DDLs
DBMS_DEBUG	Implements server-side debuggers and provides a way to debug server-side PL/SQL program units
DBMS_DEBUG_JDWP	The DBMS_DEBUG_JDWP provides the interface to initiate and control the debugging of PL/SQL stored procedures and Java stored procedures over Java Debug Wire Protocol (JDWP).
DBMS_DEBUG_JDWP_CUSTOM	The DBMS_DEBUG_JDWP_CUSTOM package is a PL/SQL package that provides database users a means to perform custom handling of a debug connection request of a database session to a debugger using the Java Debug Wire Protocol (JDWP).

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_DESCRIBE	Describes the arguments of a stored procedure with full name translation and security checking
DBMS_DG	Allows applications to notify the primary database in an Oracle Data Guard broker environment to initiate a fast-start failover when the application encounters a condition that warrants a failover
DBMS_DICTIONARY_CHECK	Enables you to identify Oracle Database dictionary inconsistencies.
DBMS_DIMENSION	Enables you to verify dimension relationships and provides an alternative to the Enterprise Manager Dimension Wizard for displaying a dimension definition
DBMS_DISTRIBUTED_TRUST_ADMIN	Maintains the Trusted Database List, which is used to determine if a privileged database link from a particular server can be accepted
DBMS_DNFS	The DBMS_DNFS package provides an interface to assists in creating a database using files in the backup set.
DBMS_DST	The DBMS_DST package provides an interface to apply the Daylight Saving Time (DST) patch to the Timestamp with Time Zone datatype.
DBMS_EDITIONS_UTILITIES	The DBMS_EDITIONS_UTILITIES package provides helper functions for edition related operations.
DBMS_EPG	Implements the embedded PL/SQL gateway that enables a Web browser to invoke a PL/SQL stored procedure through an HTTP listener
DBMS_ERRLOG	Provides a procedure that enables you to create an error logging table so that DML operations can continue after encountering errors rather than abort and roll back
DBMS_FGA	Provides fine-grained security functions
DBMS_FILE_TRANSFER	Lets you copy a binary file within a database or to transfer a binary file between databases
DBMS_FLASHBACK	Lets you flash back to a version of the database at a specified wall-clock time or a specified system change number (SCN)
DBMS_FLASHBACK_ARCHIVE	Contains procedures for disassociation and re-association of a Flashback Data Archive (FDA) enabled table from/with its underlying FDA respectively
DBMS_FLASHBACK_ARCHIVE_MIGRATE	Enables you to export and import the Flashback Archive base tables along with their history to another database using the Transportable Tablespaces
DBMS_FREQUENT_ITEMSET	Enables frequent itemset counting
DBMS_FS	The DBMS_FS package for performing operations on an Oracle file system (make, mount, unmount and destroy operations) in an Oracle database.
DBMS_GOLDENGATE_ADM	The DBMS_GOLDENGATE_ADM package provides subprograms to configure and manage Oracle GoldenGate conflict detection and resolution
DBMS_GOLDENGATE_AUTH	The DBMS_GOLDENGATE_AUTH package provides subprograms for granting privileges to and revoking privileges from GoldenGate administrators.

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_HADOOP	The DBMS_HADOOP package provides a PL/SQL procedure called CREATE_EXTDDL_FOR_HIVE(), that creates an Oracle external table for a given hive table.
DBMS_HEAT_MAP	The DBMS_HEAT_MAP package provides an interface to externalize heatmaps at various levels of storage including block, extent, segment, object and tablespace.
DBMS_HIERARCHY	Contains functions and procedures for validating the data in tables used by analytic views and hierarchies and for creating a log table.
DBMS_HM	Contains constants and procedure declarations for health check management
DBMS_HPROF	Provides an interface for profiling the execution of PL/SQL applications
DBMS_HS_PARALLEL	Enables parallel processing for heterogeneous targets access
DBMS_HS_PASSTHROUGH	Lets you use Heterogeneous Services to send pass-through SQL statements to non-Oracle systems
DBMS_HYBRID_VECTOR	Contains a SEARCH API that lets you to perform hybrid searches by both vectors and keywords against hybrid vector indexes.
DBMS_ILM	The DBMS_ILM package provides an interface for implementing Information Lifecycle Management (ILM) strategies using Automatic Data Optimization (ADO) policies.
DBMS_ILM_ADMIN	The DBMS_ILM_ADMIN package provides an interface to customize Automatic Data Optimization (ADO) policy execution.
DBMS_IMMUTABLE_TABLE	The DBMS_IMMUTABLE_TABLE package allows you to delete the expired rows in an immutable table and add interval partitioning
DBMS_INMEMORY	The DBMS_INMEMORY package provides an interface for In-Memory Column Store (IM column store) functionality.
DBMS_INMEMORY_ADMIN	The DBMS_INMEMORY_ADMIN package provides interfaces for managing an In-Memory FastStart (IM FastStart) area and In-Memory Expressions (IM expressions).
DBMS_IOT	Creates a table into which references to the chained rows for an Index Organized Table can be placed using the ANALYZE command
DBMS_JAVA	Provides a PL/SQL interface for accessing database functionality from Java
DBMS_JAVASCRIPT	The DBMS_JAVASCRIPT package provides a PL/SQL procedure for running DBMS-residingJavaScript code from the location identified by its argument
DBMS_JOB	Schedules and manages jobs in the job queue (Deprecated)
DBMS_JSON_DUALITY	The DBMS_JSON_DUALITY package provides PL/SQL interface for accessing JSON relational duality functionality
DBMS_JSON	Provides an interface for data-guide operations
DBMS_KAFKA	The DBMS_KAFKA package provides a PL/SQL interface for enabling Oracle SQL access to topics in Kafka clusters

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_KAFKA_ADM	The DBMS_KAFKA_ADM package provides a PL/SQL interface to create cluster definitions, which you can then use to grant access to Kafka cluster data for applications
DBMS_LDAP	Provides functions and procedures to access data from LDAP servers
DBMS_LDAP_UTL	Provides the Oracle Extension utility functions for LDAP
DBMS_LIBCACHE	Prepares the library cache on an Oracle instance by extracting SQL and PL/SQL from a remote instance and compiling this SQL locally without execution
DBMS_LOB	Provides general purpose routines for operations on Oracle Large Object (LOBs) datatypes - BLOB, CLOB (read/write), and BFILEs (read-only)
DBMS_LOCK	Lets you request, convert and release locks through Oracle Lock Management services
DBMS_LOGMNR	Provides functions to initialize and run the log reader
DBMS_LOGMNR_D	Queries the dictionary tables of the current database, and creates a text based file containing their contents
DBMS_LOGSTDBY	Describes procedures for configuring and managing the logical standby database environment
DBMS_LOGSTDBY_CONTEXT	You can use the procedures provided in the DBMS_LOGSTDBY_CONTEXT package to set and retrieve various parameters associated with LSBY_APPLY_CONTEXT.
DBMS_MEMOPTIMIZE	The DBMS_MEMOPTIMIZE package provides the interface for managing the Memoptimized Rowstore data buffered in the large pool and the memoptimize pool
DBMS_MEMOPTIMIZE_ADMIN	The DBMS_MEMOPTIMIZE_ADMIN package provides the interface for flushing the Memoptimized Rowstore data buffered in the large pool to disk
DBMS_METADATA	Lets callers easily retrieve complete database object definitions (metadata) from the dictionary
DBMS_METADATA_DIFF	Contains the interfaces for comparing two metadata documents in SXML format. The result of the comparison is an SXML difference document. This document can be converted to other formats using the DBMS_METADATA submit interface and the CONVERT API.
DBMS_MGD_ID_UTL	Provides a set of utility subprograms
DBMS_MGWADM	Describes the Messaging Gateway administrative interface; used in Advanced Queuing
DBMS_MGWMSG	Describes object types (used by the canonical message types to convert message bodies) and helper methods, constants, and subprograms for working with the Messaging Gateway message types; used in Advanced Queuing.
DBMS_MLE	The DBMS_MLE package allows users to execute JavaScript code inside the Oracle Database and exchange data seamlessly between PL/SQL and JavaScript.
DBMS_MONITOR	Let you use PL/SQL for controlling additional tracing and statistics gathering

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_MVIEW	Lets you refresh snapshots that are not part of the same refresh group and purge logs. DBMS_SNAPSHOT is a synonym.
DBMS_MVIEW_STATS	The DBMS_MVIEW_STATS package provides an interface to manage the collection and retention of statistics for materialized view refresh operations.
DBMS_NETWORK_ACL_ADMIN	Provides the interface to administer the network Access Control List (ACL)
DBMS_NETWORK_ACL_UTILITY	The DBMS_NETWORK_ACL_UTILITY package provides the utility functions to facilitate the evaluation of access control list (ACL) assignments governing TCP connections to network hosts.
DBMS_ODCI	Returns the CPU cost of a user function based on the elapsed time of the function
DBMS_OPTIM_BUNDLE	This package is used to manage the installed but disabled module bug fixes that cause an execution plan change.
DBMS_OUTLN	Provides the interface for procedures and functions associated with management of stored outlines Synonymous with OUTLN_PKG
DBMS_OUTPUT	Accumulates information in a buffer so that it can be retrieved later
DBMS_PARALLEL_EXECUTE	Enables the user to incrementally update table data in parallel
DBMS_PART	The DBMS_PART package provides an interface for maintenance and management operations on partitioned objects.
DBMS_PCLXUTIL	Provides intra-partition parallelism for creating partition-wise local indexes
DBMS_PDB	The DBMS_PDB package provides an interface to examine and manipulate data about pluggable databases (PDBs) in a multitenant container database (CDB). It also contains an interface specify which database objects are application common objects.
DBMS_PDB_ALTER_SHARING	In an application container with a pre-installed application, the DBMS_PDB_ALTER_SHARING package provides an interface to set database objects as application common objects or to specify that a database object is not an application common object.
DBMS_PERF	The DBMS_PERF package provides and interface to generate active reports for monitoring database performance.
DBMS_PIPE	Provides a DBMS pipe service which enables messages to be sent between sessions
DBMS_PLSQL_CODE_COVERAGE	The DBMS_PLSQL_CODE_COVERAGE package provides an interface for the collection of code coverage data of PL/SQL applications at the basic block level.
DBMS_PREDICTIVE_ANALYTICS	Provides subroutines that implement automatic data mining operations for predict, explain, and profile
DBMS_PREPROCESSOR	Provides an interface to print or retrieve the source text of a PL/SQL unit in its post-processed form

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_PRIVILEGE_CAPTURE	The DBMS_PRIVILEGE_CAPTURE package provides an interface to database privilege analysis.
DBMS_PROCESS	The DBMS_PROCESS package provides an interface to manage the prespawned servers.
DBMS_PROFILER	Provides a Probe Profiler API to profile existing PL/SQL applications and identify performance bottlenecks
DBMS_PROPAGATION_ADM	Provides administrative procedures for configuring propagation from a source queue to a destination queue
DBMS_QOPATCH	The DBMS_QOPATCH package provides an interface to view the installed database patches.
DBMS_RANDOM	Provides a built-in random number generator
DBMS_REDACT	The DBMS_REDACT package provides an interface to Oracle Data Redaction, which enables you to mask (redact) data that is returned from queries issued by lowprivileged users or an application.
DBMS_REDEFINITION	Lets you perform an online reorganization of tables
DBMS_REFRESH	Lets you create groups of snapshots that can be refreshed together to a transactionally consistent point in time Requires the Distributed Option
DBMS_REPAIR	Provides data corruption repair procedures
DBMS_RESCONFIG	Provides an interface to operate on the Resource Configuration List, and to retrieve listener information for a resource
DBMS_RESOURCE_MANAGER	Maintains plans, consumer groups, and plan directives; it also provides semantics so that you may group together changes to the plan schema
DBMS_RESOURCE_MANAGER_PRIVS	Maintains privileges associated with resource consumer groups
DBMS_RESULT_CACHE	Provides an interface to operate on the Result Cache
DBMS_RESUMABLE	Lets you suspend large operations that run out of space or reach space limits after executing for a long time, fix the problem, and make the statement resume execution
DBMS_RLS	Provides row level security administrative interface
DBMS_ROLLING	The DBMS_ROLLING PL/SQL package is used to implement the Rolling Upgrade Using Active Data Guard feature, which streamlines the process of upgrading Oracle Database software in a Data Guard configuration in a rolling fashion.
DBMS_ROWID	Provides procedures to create rowids and to interpret their contents
DBMS_RULE	Describes the subprograms that enable the evaluation of a rule set for a specified event
DBMS_RULE_ADM	Provides the subprograms for creating and managing rules, rule sets, and rule evaluation contexts
DBMS_SAGA_ADM	The DBMS_SAGA_ADM package provides a collection of saga administration functions and procedures to define and manage saga participants, coordinators, and brokers.

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_SAGA	The DBMS_SAGA package provides a collection of saga functions and procedures to initiate and finalize sagas.
DBMS_SCHEDULER	Provides a collection of scheduling functions that are callable from any PL/SQL program
DBMS_SEARCH	Enables the indexing of multiple schema objects in a single index
DBMS_SERVER_ALERT	Lets you issue alerts when some threshold has been violated
DBMS_SERVICE	Lets you create, delete, activate and deactivate services for a single instance
DBMS_SESSION	Provides access to SQL ALTER SESSION statements, and other session information, from stored procedures
DBMS_SFW_ACL_ADMIN	The DBMS_SFW_ACL_ADMIN package provides interfaces for administering and managing access control policies for the "database service firewall" feature
DBMS_SHARED_POOL	Lets you keep objects in shared memory, so that they will not be aged out with the normal LRU mechanism
DBMS_SHARDED_DIRECTORY	Provides procedures to manage an Oracle Globally Distributed Database (sharded database) created with directory-based data distribution.
DBMS_SODA	Implements Simple Oracle Document Access (SODA) which allows you to use the Oracle Database as a NoSQL document store.
DBMS_SPACE	Provides segment space information not available through standard SQL
DBMS_SPACE_ADMIN	Provides tablespace and segment space administration not available through the standard SQL
DBMS_SPD	The DBMS_SPD package provides subprograms for managing SQL plan directives (SPD).
DBMS_SPM	Supports the SQL plan management feature by providing an interface for the DBA or other user to perform controlled manipulation of plan history and SQL plan baselines maintained for various SQL statements
DBMS_SQL	Lets you use dynamic SQL to access the database
DBMS_SQL_FIREWALL	Lets you monitor users and prevent them from performing SQL injection attacks
DBMS_SQL_MONITOR	The DBMS_SQL_MONITOR package provides information about Real-time SQL Monitoring and Real-time Database Operation Monitoring.
DBMS_SQL_TRANSLATOR	The DBMS_SQL_TRANSLATOR package provides an interface for creating, configuring, and using SQL translation profiles.
DBMS_SQLDIAG	Provides an interface to the SQL Diagnosability functionality
DBMS_SQLPA	Provides an interface to implement the SQL Performance Analyzer
DBMS_SQLQ	The DBMS_SQLQ package provides the interface for configuring quarantine thresholds for execution plans of SQL statements
DBMS_SQLSET	The DBMS_SQLSET package provides an interface to manage SQL tuning sets

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_SQLTUNE	Provides the interface to tune SQL statements
DBMS_STAT_FUNCS	Provides statistical functions
DBMS_STATS	Provides a mechanism for users to view and modify optimizer statistics gathered for database objects
DBMS_STORAGE_MAP	Communicates with FMON to invoke mapping operations
DBMS_SYNC_REFRESH	The DBMS_SYNC_REFRESH package provides an interface to perform a synchronous refresh of materialized views.
DBMS_TABLE_DATA	The DBMS_TABLE_DATA package provides procedures that can be used to retrieve the byte value of a single column, or a series of columns, in any table.
DBMS_TDB	Reports whether a database can be transported between platforms using the RMAN CONVERT DATABASE command. It verifies that databases on the current host platform are of the same endian format as the destination platform, and that the state of the current database does not prevent transport of the database.
DBMS_TF	The DBMS_TF package contains utilities for POLYMORPHIC TABLE functions (PTFs) implementation. You can use DBMS_TF subprograms to consume and produce data, and get information about its execution environment..
DBMS_TNS	The DBMS_TNS package provides the RESOLVE_TNSNAME function to resolve a TNS name and return the corresponding Oracle Net8 connection string.
DBMS_TRACE	Provides routines to start and stop PL/SQL tracing
DBMS_TRANSACTION	Provides access to SQL transaction statements from stored procedures and monitors transaction activities
DBMS_TRANSFORM	Provides an interface to the message format transformation features of Oracle Advanced Queuing
DBMS_TSDP_MANAGE	The DBMS_TSDP_MANAGE package provides an interface to import and manage sensitive columns and sensitive column types in the database.
DBMS_TSDP_PROTECT	The DBMS_TSDP_PROTECT package provides an interface to configure transparent sensitive data protection (TSDP) policies in conjunction with the DBMS_TSDP_MANAGE package. The DBMS_TSDP_PROTECT package is available with the Enterprise Edition only.
DBMS_TTS	Checks if the transportable set is self-contained
DBMS_TYPES	Consists of constants, which represent the built-in and user-defined types
DBMS_UMF	The DBMS_UMF package provides an interface for deploying the Remote Management Framework (RMF) for an Oracle Database. The RMF is used for collecting performance statistics for an Oracle Database.
DBMS_USER_CERTS	The DBMS_USER_CERTS package allows you add and delete certificates.
DBMS_UTILITY	Provides various utility routines.

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_VECTOR	Simplifies common operations with Oracle AI Vector Search, such as chunking and embedding data, generating text for prompts and images, or managing vector indexes.
DBMS_VECTOR_CHAIN	Enables advanced operations with Oracle AI Vector Search, such as chunking and embedding data, generating text for prompts and images along with text processing and end-to-end similarity search.
DBMS_WARNING	Provides the interface to query, modify and delete current system or session settings.
DBMS_WM	Describes how to use the programming interface to Oracle Database Workspace Manager to work with long transactions
DBMS_WORKLOAD_CAPTURE	Configures the Workload Capture system and produce the workload capture data.
DBMS_WORKLOAD_REPLAY	Provides an interface to replay and report on a record of a workload on a production or test system
DBMS_WORKLOAD_REPOSITORY	Lets you manage the Workload Repository, performing operations such as managing snapshots and baselines
DBMS_XA	Contains the XA/Open interface for applications to call XA interface in PL/SQL
DBMS_XDB	Describes Resource Management and Access Control interface for PL/SQL
DBMS_XDB_ADMIN	Provides an interface to implement XMLIndex administration operation
DBMS_XDB_CONFIG	The DBMS_XDB_CONFIG package provides an interface for configuring Oracle XML DB and its repository.
DBMS_XDB_CONSTANTS	The DBMS_XDB_CONSTANTS package provides an interface to commonly used constants.
DBMS_XDB_REPOS	The DBMS_XDB_REPOS package provides an interface to operate on the Oracle XML database Repository.
DBMS_XDBRESOURCE	Provides an interface to operate on the XDB resource's metadata and contents
DBMS_XDB_VERSION	Describes the versioning interface
DBMS_XDBZ	Controls the Oracle XML DB repository security, which is based on Access Control Lists (ACLs)
DBMS_XEVENT	Provides event-related types and supporting subprograms
DBMS_XMLDOM	Explains access to XMLType objects
DBMS_XMLGEN	Converts the results of a SQL query to a canonical XML format
DBMS_XMLINDEX	Provides an interface to implement asychronous indexing and apply node referencing
DBMS_XMLPARSER	Explains access to the contents and structure of XML documents
DBMS_XMLSCHEMA	Explains procedures to register and delete XML schemas
DBMS_XMLSCHEMA_ANNOTATE	The DBMS_XMLSCHEMA_ANNOTATE package provides an interface to manage and configure the structured storage model, mainly through the use of pre-registration schema annotations

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
DBMS_XMLSCHEMA_UTIL	The DBMS_XMLSCHEMA_UTIL package provides an interface for XML schema validation
DBMS_XMLSTORAGE_MANAGE	The DBMS_XMLSTORAGE_MANAGE package provides an interface to manage and modify XML storage after schema registration has been completed
DBMS_XMLSTORE	Provides the ability to store XML data in relational tables
DBMS_XMLTRANSLATIONS	Provides an interface to perform translations so that strings can be searched or displayed in various languages
DBMS_XPLAN	Describes how to format the output of the EXPLAIN PLAN command
DBMS_XSTREAM_ADM	The DBMS_XSTREAM_ADM package provides interfaces for streaming database changes between an Oracle database and other systems
DBMS_XSTREAM_AUTH	The DBMS_XSTREAM_AUTH package provides subprograms for granting privileges to and revoking privileges from XStream administrators
DBMS_XSLPROCESSOR	Explains access to the contents and structure of XML documents
DEBUG_EXTPROC	Lets you debug external procedures on platforms with debuggers that attach to a running process
HTF	Hypertext functions generate HTML tags
HTP	Hypertext procedures generate HTML tags
OWA_CACHE	Provides an interface that enables the PL/SQL Gateway cache to improve the performance of PL/SQL Web applications
OWA_COOKIE	Provides an interface for sending and retrieving HTTP cookies from the client's browser
OWA_CUSTOM	Provides a Global PLSQL Agent Authorization callback function
OWA_IMAGE	Provides an interface to access the coordinates where a user clicked on an image
OWA_OPT_LOCK	Contains subprograms that impose optimistic locking strategies so as to prevent lost updates
OWA_PATTERN	Provides an interface to locate text patterns within strings and replace the matched string with another string
OWA_SEC	Provides an interface for custom authentication
OWA_TEXT	Contains subprograms used by OWA_PATTERN for manipulating strings. They are externalized so you can use them directly.
OWA_UTIL	Contains utility subprograms for performing operations such as getting the value of CGI environment variables, printing the data that is returned to the client, and printing the results of a query in an HTML table
SDO_CS	Provides functions for coordinate system transformation
SDO_CSW_PROCESS	Contains subprograms for various processing operations related to support for Catalog Services for the Web (CSW)

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
SDO_GCDR	Contains the Oracle Spatial geocoding subprograms, which let you geocode unformatted postal addresses
SDO_GEOM	Provides functions implementing geometric operations on spatial objects
SDO_GEOR	Contains functions and procedures for the Spatial GeoRaster feature, which lets you store, index, query, analyze, and deliver raster image data and its associated Spatial vector geometry data and metadata
SDO_GEOR_ADMIN	Contains subprograms for administrative operations related to GeoRaster.
SDO_GEOR_AGGR	The SDO_GEOR_AGGR package provides an interface to the SDO_GEOR_AGGR package for performing aggregate operations on GeoRaster objects
SDO_GEOR_RA	The SDO_GEOR_RA package provides an interface to the SDO_GEOR_RA package for performing raster algebra and analytic operations related to GeoRaster
SDO_GEOR_UTL	Contains utility functions and procedures for the Spatial GeoRaster feature, including those related to using triggers with GeoRaster data
SDO_LRS	Provides functions for linear referencing system support
SDO_MIGRATE	Provides functions for migrating spatial data from previous releases
SDO_NET	Provides functions and procedures for working with data modeled as nodes and links in a network
SDO_NET_MEM	Contains functions and procedures for performing editing and analysis operations on network data using a network memory object
SDO_NFE	The SDO_NFE package contains functions and procedures for performing network feature editing
SDO_OLS	Contains functions and procedures for performing editing and analysis operations on network data using a network memory object
SDO_PC_PKG	Contains subprograms to support the use of point clouds in Spatial
SDO_SAM	Contains functions and procedures for spatial analysis and data mining
SDO_TIN_PKG	Contains subprograms to support the use of triangulated irregular networks (TINs) in Spatial
SDO_TOPO	Provides procedures for creating and managing Spatial topologies
SDO_TOPO_MAP	Contains subprograms for editing Spatial topologies using a cache (TopoMap object)
SDO_TUNE	Provides functions for selecting parameters that determine the behavior of the spatial indexing scheme used in Oracle Spatial
SDO_UTIL	Provides utility functions and procedures for Oracle Spatial
SDO_WFS_LOCK	Contains subprograms for WFS support for registering and unregistering feature tables

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
SDO_WFS_PROC	Provides utility functions and procedures for Oracle Spatial
SEM_APIS	Contains subprograms for working with the Resource Description Framework (RDF) and Web Ontology Language (OWL) in an Oracle database.
SEM_OLS	The SEM_OLS package provides an interface to the SEM_OLS package for providing triple-level security to RDF data, using Oracle Label Security (OLS)
SEM_PERF	Contains subprograms for examining and enhancing the performance of the Resource Description Framework (RDF) and Web Ontology Language (OWL) support in an Oracle database
SEM_RDFCTX	Contains subprograms for managing extractor policies and semantic indexes created for documents
SEM_RDFSA	Contains subprograms for providing fine-grained access control to RDF data, using either a virtual private database (VPD) or Oracle Label Security (OLS)
UTL_CALL_STACK	The UTL_CALL_STACK package provides an interface to provide information about currently executing subprograms
UTL_COLL	Enables PL/SQL programs to use collection locators to query and update
UTL_COMPRESS	Provides a set of data compression utilities
UTL_ENCODE	Provides functions that encode RAW data into a standard encoded format so that the data can be transported between hosts
UTL_FILE	Enables your PL/SQL programs to read and write operating system text files and provides a restricted version of standard operating system stream file I/O
UTL_HTTP	Enables HTTP callouts from PL/SQL and SQL to access data on the Internet or to call Oracle Web Server Cartridges
UTL_I18N	Provides a set of services (Oracle Globalization Service) that help developers build multilingual applications
UTL_INADDR	Provides a procedure to support internet addressing
UTL_IDENT	Specifies which database or client PL/SQL is running
UTL_LMS	Retrieves and formats error messages in different languages
UTL_MAIL	A utility for managing email which includes commonly used email features, such as attachments, CC, BCC, and return receipt
UTL_MATCH	The UTL_MATCH package facilitates matching two records. This is typically used to match names, such as two First Names or two Last Names
UTL_NLA	Exposes a subset of the BLAS and LAPACK (Version 3.0) operations on vectors and matrices represented as VARRAYs
UTL_RAW	Provides SQL functions for manipulating RAW datatypes
UTL_RECOMP	Recompiles invalid PL/SQL modules, invalid views, Java classes, indextypes and operators in a database, either sequentially or in parallels

Table 1-1 (Cont.) Summary of Oracle Supplied PL/SQL Packages

Package Name	Description
UTL_REF	Enables a PL/SQL program to access an object by providing a reference to the object
UTL_SMTP	Provides PL/SQL functionality to send emails
UTL_RPADV	Provides subprograms to collect and analyze statistics for the Oracle Replication components in a distributed database environment
UTL_TCP	Provides PL/SQL functionality to support simple TCP/IP-based communications between servers and the outside world
UTL_URL	Provides escape and unescape mechanisms for URL characters
WPG_DOCLOAD	Provides an interface to download files, both BLOBs and BFILEs
ANYDATA TYPE	A self-describing data instance type containing an instance of the type plus a description
ANYDATASET TYPE	Contains a description of a given type plus a set of data instances of that type
ANYTYPE TYPE	Contains a type description of any persistent SQL type, named or unnamed, including object types and collection types; or, it can be used to construct new transient type descriptions
Oracle Database Advanced Queuing Types	Describes the types used in Advanced Queuing
DBFS Content Interface Types	Describes public types defined to support the DBMS_DBFS_CONTENT interface.
Database URI Type	Contains URI Support, UriType Super Type, HttpUriType Subtype, DBUriType Subtype, XDBUriType Subtype, UriFactory Package
JMS TYPES	Describes JMS types so that a PL/SQL application can use JMS queues of JMS types
LOGICAL CHANGE RECORD TYPES	Describes LCR types, which are message payloads that contain information about changes to a database.
MG_ID Package Types	Provides an extensible framework that supports current RFID tags with the standard family of EPC bit encodings for the supported encoding types
POLYMORPHIC TABLE FUNCTION (PTF) Package Types	Describes types defined in the DBMS_TF package to support PTF
RULES TYPEs	Describes the types used with rules, rule sets, and evaluation contexts
SODA Types	Describes the SODA Types.
UTL Streams Types	Describes abstract streams types used with Oracle XML functionality
XMLType	Describes the types and functions used for native XML support in the server

2

Oracle Application Express Packages APEX_APPLICATION Through APEX_ZIP

The Oracle Application Express supplies PL/SQL packages for application developers who are building database-centric web applications using Oracle Application Express.

For a complete description of all the packages that ship with Oracle Application Express, see the *Oracle Application Express API Reference*.

3
CTX_ADM

The `CTX_ADM` package lets you administer the Oracle Text data dictionary.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_ADM` in the Oracle Text Reference.

4
CTX_ANL

The `CTX_ANL` package is used with AUTO_LEXER and provides procedures for adding and dropping a custom dictionary from the lexer.

A custom dictionary might be one that you develop for a special field of study or for your industry. In most cases, the dictionaries supplied for the supported languages with Oracle Text are more than sufficient to handle your requirements.

For a complete description of this package within the context of Oracle Text, see `CTX_ANL` in the Oracle Text Reference.

4.1 GET_COMPRESSION_RATIO Procedure

Use this procedure to estimate the storage space that you can save by enabling the compression feature for an existing SecureFile LOB. It analyzes the compression ratio of a table or an index and gives information about compressibility of the object. You can provide various parameters to selectively analyze different compression types.

In Oracle Database 23ai, this procedure has been enhanced to estimate the compression ratio faster for LOBs while using less space. Now you can also estimate the compression ratio for BasicFile LOBs. This helps you decide upfront whether you want to compress BasicFile LOBs, before migrating BasicFile LOBs to SecureFile LOBs. You can also estimate the compression ratio at the LOB byte level and the time taken, in hours, to compress the LOB data in the table.

The compression ratio is estimated for the number of rows in the LOB column that you specify. For example, let's consider that the compression ratio is 2.33. It indicates that after you enable the compression feature, you can save around half of the space for the sampled rows in the LOB column.

Disclaimer: The compression ratio is an approximate value, which is calculated based on the sampled rows in the LOB column. The actual space that you save when you enable compression for the complete table may be different.

Syntax

The syntax to get the compression ratio differs for objects, LOBs, IOTs, and indexes on a table.

- Syntax to get the compression ratio for an object (table or index, default is table).

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
   scratchtbsname        IN     VARCHAR2,
   ownname               IN     VARCHAR2,
   objname               IN     VARCHAR2,
   subobjname            IN     VARCHAR2,
   comptype              IN     NUMBER,
   blkcnt_cmp            OUT    PLS_INTEGER,
   blkcnt_uncmp          OUT    PLS_INTEGER,
   row_cmp               OUT    PLS_INTEGER,
   row_uncmp             OUT    PLS_INTEGER,
   cmp_ratio             OUT    NUMBER,
   comptype_str          OUT    VARCHAR2,
   block_compr_ratio     OUT    PLS_INTEGER,
   byte_comp_ratio       OUT    NUMBER,
```

```
   subset_numrows        IN     NUMBER  DEFAULT COMP_RATIO_MINROWS,
   objtype               IN     PLS_INTEGER DEFAULT OBJTYPE_TABLE);
```

- Syntax to get compression ratio for BasicFile and SecureFile LOBs:

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
   scratchtbsname        IN     VARCHAR2,
   tabowner              IN     VARCHAR2,
   tabname               IN     VARCHAR2,
   lobname               IN     VARCHAR2,
   partname              IN     VARCHAR2,
   comptype              IN     NUMBER,
   blkcnt_cmp            OUT    PLS_INTEGER,
   blkcnt_uncmp          OUT    PLS_INTEGER,
   lobcnt                OUT    PLS_INTEGER,
   cmp_ratio             OUT    NUMBER,
   comptype_str          OUT    VARCHAR2,
   byte_comp_ratio       OUT    NUMBER,
   total_time            OUT    NUMBER
   subset_numrows        IN     NUMBER DEFAULT COMP_RATIO_LOB_MAXROWS);
```

- Syntax to get the compression ratio for all indexes on a table. The compression ratios are returned as a collection.

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
  scratchtbsname        IN     VARCHAR2,
  ownname               IN     VARCHAR2,
  tabname               IN     VARCHAR2,
  comptype              IN     NUMBER,
  index_cr              OUT    DBMS_COMPRESSION.COMPRECLIST,
  comptype_str          OUT    VARCHAR2,
  subset_numrows        IN     NUMBER DEFAULT COMP_RATIO_INDEX_MINROWS);
```

- Syntax to get the compression ratio for IOTs.

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
     scratchtbsname       IN     VARCHAR2,
     ownname              IN     VARCHAR2,
     objname              IN     VARCHAR2,
     subobjname           IN     VARCHAR2,
     comptype             IN     NUMBER,
     iotcomp_cr           OUT    DBMS_COMPRESSION.COMPRECLIST,
     comptype_str         OUT    VARCHAR2,
     subset_numrows       IN     NUMBER DEFAULT COMP_RATIO_INDEX_MINROWS);
```

Parameters

Table 4-1 GET_COMPRESSION_RATIO Procedure Parameters

Parameter	Description
`scratchtbsname`	Temporary scratch tablespace that can be used for analysis
`ownname` / `tabowner`	Schema of the table to analyze
`tabname`	Name of the table to analyze
`objname`	Name of the object
`subobjname`	Name of the partition or sub-partition of the object

Table 4-1 (Cont.) GET_COMPRESSION_RATIO Procedure Parameters

Parameter	Description
comptype	Compression types for which analysis should be performed When the object is an index, only the following compression types are valid: COMP_INDEX_ADVANCED_HIGH (value 1024) and COMP_INDEX_ADVANCED_LOW (value 2048). **Note:** The following compression types cannot be specified in this parameter for any type of object: COMP_BLOCK (value 64) and COMP_BASIC (value 4096).
blkcnt_cmp	Number of blocks used by compressed sample of the table
blkcnt_uncmp	Number of blocks used by uncompressed sample of the table
row_cmp	Number of rows in a block in compressed sample of the table
row_uncmp	Number of rows in a block in uncompressed sample of the table
cmp_ratio	Compression ratio, blkcnt_uncmp divided by blkcnt_cmp. It provides the ratio of **blocks** occupied by the uncompressed data to the blocks occupied by the compressed data.
comptype_str	String describing the compression type
subset_numrows	Number of rows sampled to estimate compression ratio.
objtype	Type of the object, either OBJTYPE_TABLE or OBJTYPE_INDEX
lobname	Name of the LOB column
partname	In case of partitioned tables, the related partition name
lobcnt	Number of lobs actually sampled to estimate compression ratio
byte_comp_ratio	Provides the ratio of bytes of uncompressed data to the bytes of compressed data for LOBs.
index_cr	List of indexes and their estimated compression ratios
iotcomp_cr	Compression ratio for the IOT The first object contains the compression ratio for the whole IOT. The second object contains the compression ratio only for the top index section of the IOT (excludes the overflow segment).
total_time	Provides an estimate of the time taken, in hours, to compress the LOB data in the table.

Example: Estimate the compression ratio for inline and out-of-line LOBs

The following example shows how to estimate the compression ratio for LOBs.

```
SET SERVEROUTPUT ON
DECLARE
    bcmp                    INTEGER;
    buncmp                  INTEGER;
    lobcmp                  INTEGER;
    cr                      NUMBER;
    byte_cr                 NUMBER;
    cstr                    VARCHAR2(2000);
    total_time              NUMBER;
    l_segment_name          VARCHAR2(30);
    l_segment_size_blocks   NUMBER;
    l_segment_size_bytes    NUMBER;
    l_used_blocks           NUMBER;
```

```
    l_used_bytes            NUMBER;
    l_expired_blocks        NUMBER;
    l_expired_bytes         NUMBER;
    l_unexpired_blocks      NUMBER;
    l_unexpired_bytes       NUMBER;
BEGIN
   DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
     scratchtbsname        => 'LOBTBSP',
     tabowner              => 'CMPADV',
     tabname               => p_tablename,
     lobname               => 'C',
     partname              => NULL,
     comptype              => 256,
     blkcnt_cmp            => bcmp,
     blkcnt_uncmp          => buncmp,
     lobcnt                => lobcmp,
     cmp_ratio             => cr,
     comptype_str          => cstr,
     subset_numrows        => 1000,
     byte_comp_ratio       => byte_cr,
     total_time            => total_time
);
DBMS_OUTPUT.put_line('Estimated ratio of blocks used by the uncompressed data
to the compressed data :  ' || cr);
DBMS_OUTPUT.put_line('Estimated ratio of bytes used by the uncompressed data
to the compressed data  :  ' || byte_cr);
END;
/
```

To understand the output of this procedure, let's consider `tab_inline`, an inline table, and `tab_outofline`, an out-of-line table as shown in the following example.

```
CREATE TABLE tab_inline
(
    a NUMBER,
    c CLOB
)
LOB(c) STORE AS SECUREFILE  (ENABLE STORAGE IN ROW CACHE LOGGING);


CREATE TABLE tab_outofline
(
    a NUMBER,
    c CLOB
)
LOB(c) STORE AS SECUREFILE  (DISABLE STORAGE IN ROW CACHE LOGGING);
```

Data is stored in different ways in `tab_inline` and `tab_outofline`. In the `tab_inline` table, if the LOB is less than 4K, then data is stored in the table segment; otherwise, it is stored in the LOB segment. For the `tab_outofline` table, data of all sizes is stored in the LOB segment.

Let's consider that you have inserted 1000 LOBs of 3K each in both the tables, and then calculate the compression ratios. You can use the `dbms_space.space_usage` procedure to calculate the space used by the data that is stored in the LOB segments.

Sample output of compression ratio for inline LOBs.

```
Estimated block compression ratio : 1
Estimated byte compression ratio           : 57.6
Space used(in bytes)                       : 0
space used(in blocks)                      : 0
```

Sample output of compression ratio for out-of-line LOBs.

```
Estimated block compression ratio : 1
Estimated byte compression ratio           : 56.1
Space used(in bytes)                       : 8 MB
space used(in blocks)                      : 1000
```

In this example, even though the estimated byte and block compression ratios are almost the same for inline and out-of-line LOBs, the space that is used is different. In the case of `tab_inline`, LOB segment is not used so the space used is 0. In both cases, the data is approximately 3KB, which is small. Therefore, the data before and after compression uses the same number of blocks (that is 1 block), so the block compression ratio is 1. However, the byte level compression ratio, `byte_comp_ratio`, which compares the actual number of bytes used by the LOBs before and after compression is 57.6 or 56.1.

Example: Estimate the compression ratio for indexes on a table with low compression type

The following example shows how to estimate the compression ratio for advanced index compression (low):

```
SET SERVEROUTPUT ON
DECLARE
  l_blkcnt_cmp       PLS_INTEGER;
  l_blkcnt_uncmp   PLS_INTEGER
  l_row_cmp       PLS_INTEGER;
  l_row_uncmp       PLS_INTEGER;
  l_cmp_ratio       NUMBER;
  l_comptype_str   VARCHAR2(32767);
BEGIN
   DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
    scratchtbsname =>       'USERS' ,
    ownname       =>       'TEST' ,
    objname       =>       'SALES_IDX' ,
    subobjname     =>        NULL ,
    comptype       =>        DBMS_COMPRESSION.COMP_INDEX_ADVANCED_LOW,
    blkcnt_cmp     =>        l_blkcnt_cmp,
    blkcnt_uncmp   =>        l_blkcnt_uncmp,
    row_cmp       =>        l_row_cmp,
    row_uncmp      =>        l_row_uncmp,
    cmp_ratio      =>        l_cmp_ratio,
    comptype_str   =>        l_comptype_str,
    subset_numrows =>        DBMS_COMPRESSION.comp_ratio_minrows,
    objtype       =>        DBMS_COMPRESSION.objtype_index
   );
DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
object   :  ' || l_blkcnt_cmp);
DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
```

```
the object :  ' || l_blkcnt_uncmp);
DBMS_OUTPUT.put_line( 'Number of rows in a block in compressed sample of the
object    :  ' || l_row_cmp);
DBMS_OUTPUT.put_line( 'Number of rows in a block in uncompressed sample of
the object :  ' || l_row_uncmp);
DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of
Sample                          :  ' || l_cmp_ratio);
DBMS_OUTPUT.put_line( 'Compression Type                              :  ' ||
l_comptype_str);
END;
/
```

Output of compression advisor estimate for advanced index compression (Low):

```
Number of blocks used by the compressed sample of the object     : 243
Number of blocks used by the uncompressed sample of the object   : 539
Number of rows in a block in compressed sample of the object     : 499
Number of rows in a block in uncompressed sample of the object   : 145
Estimated Compression Ratio of Sample                 : 2.2
Compression Type                         : “Compress Advanced Low”
```

Example: Estimate the compression ratio for LOBs with medium compression type

The following example shows how to estimate the compression ratio for advanced LOB compression (medium):

```
SET SERVEROUTPUT ON
DECLARE
  l_blkcnt_cmp     PLS_INTEGER;
  l_blkcnt_uncmp   PLS_INTEGER;
  l_row_cmp       PLS_INTEGER;
  l_lobcnt       PLS_INTEGER;
  l_cmp_ratio      NUMBER;
  l_comptype_str   VARCHAR2(32767);
BEGIN
   DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
    scratchtbsname => 'USERS' ,
    tabowner       => 'TEST' ,
    tabname       => 'PARTS' ,
    lobname       =>  'PART_DESCRIPTION' ,
    partname       =>  NULL ,
    comptype       =>  DBMS_COMPRESSION.COMP_LOB_MEDIUM,
    blkcnt_cmp     => l_blkcnt_cmp,
    blkcnt_uncmp   => l_blkcnt_uncmp,
    row_cmp       => l_row_cmp,
    lobcnt       => l_lobcnt,
    cmp_ratio      => l_cmp_ratio,
    comptype_str   => l_comptype_str,
    subset_numrows => DBMS_COMPRESSION.comp_ratio_lob_maxrows
   );
DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
object    :  ' || l_blkcnt_cmp);
DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the object  :  ' || l_blkcnt_uncmp);
DBMS_OUTPUT.put_line( 'Number of rows in a block in compressed sample of the
object    :  ' || l_row_cmp);
DBMS_OUTPUT.put_line( 'Number of LOBS actually
```

```
sampled                        :  ' || 1_lobcnt);
DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of
Sample                              :  ' || 1_cmp_ratio);
DBMS_OUTPUT.put_line( 'Compression
Type                                                :  ' || 1_comptype_str);
END;
/
```

Output of compression advisor estimate for advanced LOB compression (Medium):

```
Number of blocks used by the compressed sample of the object   : 199
Number of blocks used by the uncompressed sample of the object     : 389
Number of rows in a block in compressed sample of the object       : 293
Number of LOBS actually sampled                              : 55
Estimated Compression Ratio of Sample                        : 1.9
Compression Type                                        : “Compress Medium”
```

Example: Estimate the compression ratio for IoTs

The following example shows how to estimate the compression ratio for IOTs:

```
SET SERVEROUTPUT ON
DECLARE
  bcmp       INTEGER;
  buncmp     INTEGER;
  rowcmp     INTEGER;
  rowuncmp   INTEGER;
  cr         NUMBER;
  cstr       VARCHAR2(2000);
  iotcomp_cr DBMS_COMPRESSION.COMPRECLIST;
 BEGIN
    DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
    scratchtbsname        => 'USERS',
    ownname               => 'TEST',
    objname               => 'SALES',
    subobjname            =>  NULL,
    comptype              =>  DBMS_COMPRESSION.COMP_INDEX_ADVANCED_LOW,
    iotcomp_cr            =>  iotcomp_cr,
    comptype_str          =>  cstr,
    subset_numrows        =>  DBMS_COMPRESSION.COMP_RATIO_ALLROWS
    );
 --information about the index and the overflow segment
 DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
IOT table                                     : ' || iotcomp_cr(1).blkcnt_cmp);
 DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the IOT table                                   : ' ||
iotcomp_cr(1).blkcnt_uncmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the compressed
sample of the IOT table                         : ' || iotcomp_cr(1).row_cmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the uncompressed
sample of the IOT table                      : ' || iotcomp_cr(1).row_uncmp);
 DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of the
sample                                                          : ' ||
iotcomp_cr(1).cmp_ratio);
 --information about the index segment
 DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
index segment of the IOT table               : ' || iotcomp_cr(2).blkcnt_cmp);
```

```
 DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the index segment of the IOT table             : ' ||
iotcomp_cr(2).blkcnt_uncmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the compressed
sample of the index segment of the IOT table   : ' || iotcomp_cr(2).row_cmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the uncompressed
sample of the index segment of the IOT table : ' || iotcomp_cr(2).row_uncmp);
 DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of the
sample                                                          : ' ||
iotcomp_cr(2).cmp_ratio);
 END;
 /
```

Output of the compression ratio for IOTs:

```
Number of blocks used by the compressed sample of the IOT
table                                    : 5027
Number of blocks used by the uncompressed sample of the IOT
table                                  : 7950
Average number of rows in a block in the compressed sample of the IOT
table                         : 199
Average number of rows in a block in the uncompressed sample of the IOT
table                       : 126
Estimated Compression Ratio of the
sample                                                        : 1.58
Number of blocks used by the compressed sample of the index segment of the IOT
table               : 3238
Number of blocks used by the uncompressed sample of the index segment of the IOT
table             : 6161
Average number of rows in a block in the compressed sample of the index segment of the
IOT table   : 309
Average number of rows in a block in the uncompressed sample of the index segment of the
IOT table : 162
Estimated Compression Ratio of the
sample                                                        : 1.9
```

Usage Notes

- The procedure creates different tables in the scratch tablespace and runs analysis on these objects. It does not modify anything in the user-specified tables.
- From 23ai onwards, this feature has been enhanced to estimate the compression ratio faster for LOBs while using less space. To get a more accurate result, run the following command to switch to the old method. The older method to calculate the compression ratio takes more time to return the results and uses more space.

  ```
  alter session set "_kdlf_new_compression_adv"= FALSE;
  ```

- To understand the impact of compression, use the value of the byte compression ratio for inline LOBs and for out-of-line LOBs, use the value of the block compression ratio and space used.
- You can get more benefits when you compress large volume of data as compared to small volumes of data. If you want to compress small volumes of data, look at the byte ratio instead of the block ratio to understand the impact of compression.

5
CTX_DDL

The `CTX_DDL` package lets you create and manage the preferences, section groups, and stoplists required for Text indexes.

Note that you must install this package in order to use it.

For complete description of this package within the context of Oracle Text, see `CTX_DDL` in the Oracle Text Reference.

6
CTX_CLS

The `CTX_CLS` package enables generation of CTXRULE rules for a set of documents.

For a complete description of this package within the context of Oracle Text, see `CTX_CLS` in the Oracle Text Reference.

7
CTX_DOC

The `CTX_DOC` package lets you request document services.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_DOC` in the Oracle Text Reference.

8
CTX_ENTITY

The `CTX_ENTITY` package enables you to search for terms that are unknown to you without specifying a particular search text.

It does this by identifying names, places, dates, and other objects when they are mentioned in a document and by tagging each occurrence (called a mention) with its type and subtype. This process enables you to produce a structured view of a document that can later be used for text and data mining and more comprehensive intelligence analysis.

For the complete description of this package within the context of Oracle Text, see `CTX_ENTITY` in the *Oracle Text Reference.*

9
CTX_OUTPUT

This Oracle Text package lets you manage the index log.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_OUTPUT` in the Oracle Text Reference.

10
CTX_QUERY

This Oracle Text package lets you generate query feedback, count hits, and create stored query expressions.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_QUERY` in the Oracle Text Reference.

11
CTX_REPORT

This Oracle Text package lets you create various index reports.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_REPORT` in the Oracle Text Reference.

12
CTX_THES

This Oracle Text package lets you to manage and browse thesauri.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_THES` in the Oracle Text Reference.

13
CTX_ULEXER

This Oracle Text package is for use with the user-lexer.

Note that you must install this package in order to use it.

For a complete description of this package within the context of Oracle Text, see `CTX_ULEXER` in the Oracle Text Reference.

14
DBMS_ACTIVITY

`DBMS_ACTIVITY` contains functions and procedures allowing authorized users to control the activity information captured by OATS.

This package is owned by `SYS`, so `EXECUTE` package privilege is required of the non-`SYS` users. Users with DBA role are granted the `EXECUTE` privilege on this package.

This chapter contains the following topics:

- Using DBMS_ACTIVITY
- Summary of DBMS_ACTIVITY Subprograms

14.1 Using DBMS_ACTIVITY

Object Activity Tracking System (OATS) is a generic RDBMS based tracking service that provides information about various types of activities associated with different database objects. An activity represents a user or system-initiated action such as scanning or loading a table. Most of the activities are tracked in the form of frequencies (i.e. counts over fixed time intervals) such as the number of scans of a table in 15 minutes.

The database objects whose activities are tracked include tables and materialized views. Different types of activities include row insert, delete and update, table scan, load and truncate, partition maintenance operations (create, drop, move, split, merge, and exchange), materialized view rewrite and refresh.

Another important class of activities is related to the usage or non-usage of auxiliary structures such as MVs, indexes and zone maps.

Depending on the type of activity the tracking technique can be precise, approximate, or probabilistic. Most of the activities are tracked approximately mostly for efficiency reasons.

Depending on the type of activity and its usage by different clients a certain tracking technique would be more suitable than others. For example, counters are useful for tracking the index and materialized view usage. Counters are maintained for fixed time intervals in order to capture the object usage pattern over time. The same technique can be used to track the update activity of table columns by maintaining update counters for each. Setting bits in a bitvector is another form of tracking technique which is suitable for registering the occurrence of a certain activity within a time interval without saying how many times that activity occurred.

Precise Tracking: an activity is accounted for with 100% accuracy. Some form of atomic update or some latching is required to get precise tracking in a multi-processor environment. OATS generally avoidS precise tracking, unless a client needs that kind of precision.

Approximate Tracking: an activity is accounted for in almost all cases but with rare exceptions. An example of approximate tracking is the dirty update of activity counters. Dirty updating is very efficient because no locking or latching is performed and atomic operations are not used when counters are changed. However, dirty updating can result in lost updates leading to under counting. If lost updates are rare then dirty updating is a very efficient technique for approximate tracking.

Probabilistic Tracking: an activity is sampled with a certain probability, and it is accounted for when it becomes part of the sample. Probabilistic tracking produces activity data that is less accurate but it prevents frequent tracking actions.

Occurrence Tracking: records certain activity as having occurred within a time interval. Specifically, occurrence tracking provides information that says that certain object activity occurred at least once in a given time interval without saying how many times it occurred. The use of bitvector is a popular technique for occurrence tracking. Depending on how the bits are set in a bitvector the occurrence tracking can be either precise (no collisions) or approximate (possible collisions).

Counter Effects of Transactions and Rollback: Except for the precise tracking counters, none of the other types of counters is transactional or affected by rollback. This means that if a transaction fails, the approximate counters may still record any actions from that transaction. And when the database is rolled back to a save point, the approximate counters are not set back to their previous state.

14.2 Summary of DBMS_ACTIVITY Subprograms

`DBMS_ACTIVITY` uses the `CONFIGURE`, `CREATE_SNAPSHOP`, and `DELETE_SNAPSHOTS` procedures.

Table 14-1 DBMS_ACTIVITY Package Subprograms

Subprogram	Description
CONFIGURE Procedure	Allows an authorized user to set configuration parameters for OATS in parameter/value format.
CREATE_SNAPSHOT Procedure	Allows an authorized user to manually create an activity snapshot on local instance or on all database instances by flushing the activity information maintained in memory to disk
DELETE_SNAPSHOT Procedure	Allows an authorized user to manually delete all older snapshots based on an input snapshot id called `BEFORE_SNAP_ID`.
DELETE_SNAPSHOTS	Allows an authorized user to manually delete all older snapshots based on an input timestamp value called `BEFORE_TIME`.

14.2.1 CONFIGURE Procedure

This procedure allows an authorized user to set configuration parameters for OATS in parameter/value format.

This function can be called numerous times, each time setting a different configuration parameter, or same parameter but with different value. The parameter setting applies to the indicated database or the local database.

Syntax

```
DBMS_ACTIVITY.CONFIGURE (
   PARAMETER_NAME       IN VARCHAR2
   PARAMETER_VALUE      IN NUMBER,
   CON_DBNAME           IN VARCHAR2 := NULL
);
```

Parameters

Table 14-2 CONFIGURE Procedure Parameters

Parameter	Description
PARAMETER_NAME	Name of the configuration parameter to set. Parameters available: • ACTIVITY_INTERVAL_MINUTES: The interval in minutes for maintaining activity information before it is flushed and reset. The default is 15. Other values: 30, 60, 120, 180, 240, 360, 480, 720, 1440 • ACTIVITY_RETENTION_DAYS: the number of days to maintain the activity information before it is purged. The default is 400. From 8 to 2000. • ACTIVITY_SPACE_PERCENT: soft limit in percent of the SYSAUX space for storing the activity information. The default is 5. From 1 to 25.
PARAMETER_VALUE	Value of the configuration parameter to use.
CON_DBNAME	Name of a container in the consolidated database (CDB). It is either root or a pluggable database. When omitted, the default is the local database.

> **Note:**
>
> The current CDB or PDB name is the only non-null value supported.

Usage Notes

The user must be SYS, or must have the DBA role, or granted the EXECUTE package privilege.

14.2.2 CREATE_SNAPSHOT Procedure

This procedure allows an authorized user to manually create an activity snapshot on local instance or on all database instances by flushing the activity information maintained in memory to disk.

Syntax

```
DBMS_ACTIVITY.CREATE_SNAPSHOT(
   ALL_INSTANCES       IN BOOLEAN := TRUE,
   CON_DBNAME          IN VARCHAR2 := NULL
);
```

Parameters

Table 14-3 CREATE_SNAPSHOT Procedure Parameters

Parameter	Description
ALL_INSTANCES	Specify FALSE if activity snapshot should be created only for the local instance. The default is to create snapshot on all database instances.

Table 14-3 (Cont.) CREATE_SNAPSHOT Procedure Parameters

Parameter	Description
CON_DBNAME	Name of a container in the consolidated database (CDB). It is either root or a pluggable database. When omitted, the default is the local database.

> **Note:**
>
> The current CDB or PDB name is the only non-null value supported.

Usage Notes

The SNAP_ID of snapshot created.

The user must be SYS, or must have the DBA role, or granted the EXECUTE package privilege.

14.2.3 DELETE_SNAPSHOT Procedure

This procedure allows an authorized user to manually delete all older snapshots based on an input snapshot id called BEFORE_SNAP_ID.

Syntax

```
DBMS_ACTIVITY.DELETE_SNAPSHOT (
   BEFORE_SNAP_ID      IN NUMBER,
   CON_DBNAME          IN VARCHAR2 := NULL
);
```

Parameters

Table 14-4 DELETE_SNAPSHOT Procedure Parameters

Parameter	Description
BEFORE_SNAP_ID	All snapshots with SNAP_ID value less than this argument value is removed from the disk storage (SYSAUX).
CON_DBNAME	Name of a container in the consolidated database (CDB). It is either root or a pluggable database. When omitted, the default is the local database.

> **Note:**
>
> The current CDB or PDB name is the only non-null value supported.

Usage Notes

Returns TRUE if one or more snapshots were deleted; FALSE otherwise.

The user must be SYS, or must have the DBA role, or granted the EXECUTE package privilege.

14.2.4 DELETE_SNAPSHOT Procedure

This procedure allows an authorized user to manually delete all older snapshots based on an input timestamp value called BEFORE_TIME.

Syntax

```
DBMS_ACTIVITY.DELETE_SNAPSHOT (
   BEFORE_TIME          IN TIMESTAMP,
   CON_DBNAME           IN VARCHAR2 := NULL
);
```

Parameters

Table 14-5 DELETE_SNAPSHOT Procedure Parameters

Parameter	Description
BEFORE_TIME	All snapshots associated with a time less than this argument value are removed from the disk storage (SYSAUX).
CON_DBNAME	Name of a container in the consolidated database (CDB). It is either root or a pluggable database. When omitted, the default is the local database.

> **Note:**
>
> The current CDB or PDB name is the only non-null value supported.

Usage Notes

Returns TRUE if one or more snapshots were deleted; FALSE otherwise.

The user must be SYS, or must have the DBA role, or granted the EXECUTE package privilege.

15
DBMS_ADDM

The `DBMS_ADDM` package facilitates the use of Advisor functionality regarding the Automatic Database Diagnostic Monitor.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_ADDM Subprograms

> **See Also:**
>
> - *Oracle Real Application Clusters Administration and Deployment Guide* for more information about "Automatic Workload Repository in Oracle Real Application Clusters Environments"
> - *Oracle Database Performance Tuning Guide* for more information about "Automatic Performance Diagnostics"

15.1 DBMS_ADDM Security Model

The `DBMS_ADDM` package runs with the caller's permission, not the definer's, and then applies the security constraints required by the `DBMS_ADVISOR` package.

> **See Also:**
>
> The DBMS_ADVISOR package for more information about "Security Model".

15.2 Summary of DBMS_ADDM Subprograms

The table in this topic lists and describes the DBMS_ADDM subprograms.

Table 15-1 DBMS_ADDM Package Subprograms

Subprogram	Description
ANALYZE_DB Procedure	Creates an ADDM task for analyzing in database analysis mode and executes it
ANALYZE_INST Procedure	Creates an ADDM task for analyzing in instance analysis mode and executes it.
ANALYZE_PARTIAL Procedure	Creates an ADDM task for analyzing a subset of instances in partial analysis mode and executes it
COMPARE_CAPTURE_REPLAY_REPORT Function	Produces a Compare Period ADDM report comparing the performance of a capture to a replay
COMPARE_DATABASES Function	Produces a Compare Period ADDM report for a database-wide performance comparison
COMPARE_INSTANCES Function	Produces a Compare Period ADDM report for an instance-level performance comparison
COMPARE_REPLAY_REPLAY_REPORT Function	Produces a Compare Period ADDM report comparing the performance of a replay to another replay
DELETE Procedure	Deletes an already created ADDM task (of any kind)
DELETE_FINDING_DIRECTIVE Procedure	Deletes a finding directive
DELETE_PARAMETER_DIRECTIVE Procedure	Deletes a parameter directive
DELETE_SEGMENT_DIRECTIVE Procedure	Deletes a segment directive
DELETE_SQL_DIRECTIVE Procedure	Deletes a SQL directive
FAILED_AUTO_TASKS_REPORT Function	Creates a plain text, user-readable report for the failed tasks registered in the `DBA_ADDM_PENDING_AUTOTASKS` views.
GET_ASH_QUERY Function	Returns a string containing the SQL text of an ASH query identifying the rows in ASH with impact for the finding
GET_REPORT Function	Retrieves the default text report of an executed ADDM task
INSERT_FINDING_DIRECTIVE Procedure	Creates a directive to limit reporting of a specific finding type.
INSERT_PARAMETER_DIRECTIVE Procedure	Creates a directive to prevent ADDM from creating actions to alter the value of a specific system parameter
INSERT_SEGMENT_DIRECTIVE Procedure	Creates a directive to prevent ADDM from creating actions to "run Segment Advisor" for specific segments
INSERT_SQL_DIRECTIVE Procedure	Creates a directive to limit reporting of actions on specific SQL
REAL_TIME_ADDM_REPORT Function	Produces a real-time report of ADDM activity
REEXECUTE_FAILED_AUTO_TASKS Procedure	Re-executes the tasks registered in the `DBA_ADDM_PENDING_AUTOTASKS` views.

15.2.1 ANALYZE_DB Procedure

This procedure creates an ADDM task for analyzing in database analysis mode and executes it.

Syntax

```
DBMS_ADDM.ANALYZE_DB (
   task_name                 IN OUT VARCHAR2,
   begin_snapshot            IN     NUMBER,
   end_snapshot              IN     NUMBER,
   read_only_type_override   IN     VARCHAR2,
   db_id                     IN     NUMBER := NULL);
```

Parameters

Table 15-2 ANALYZE_DB Procedure Parameters

Parameter	Description
task_name	Name of the task to be created.
begin_snapshot	Number of the snapshot that starts the analysis period.
end_snapshot	Number of the snapshot that ends the analysis period.
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data
db_id	Database ID for the database you to analyze. By default, this is the database currently connected.

Return Values

The name of the created task is returned in the task_name parameter. It may be different from the value that is given as input (only in cases that name is already used by another task).

Examples

To create an ADDM task in database analysis mode and execute it, with its name in variable tname:

```
var tname VARCHAR2(60);
BEGIN
  :tname := 'my_database_analysis_mode_task';
  DBMS_ADDM.ANALYZE_DB(:tname, 1, 2);
END
```

To see a report:

```
SET LONG 100000
SET PAGESIZE 50000
SELECT DBMS_ADDM.GET_REPORT(:tname) FROM DUAL;
```

Note that the return type of a report is a CLOB, formatted to fit line size of 80.

15.2.2 ANALYZE_INST Procedure

This procedure creates an ADDM task for analyzing in instance analysis mode and executes it.

Syntax

```
DBMS_ADDM.ANALYZE_INST (
   task_name                 IN OUT VARCHAR2,
   begin_snapshot            IN     NUMBER,
   end_snapshot              IN     NUMBER,
   cdb_type_override         IN     VARCHAR2,
   read_only_type_override   IN     VARCHAR2,
   instance_number           IN     NUMBER := NULL,
   db_id                     IN     NUMBER := NULL);
```

Parameters

Table 15-3 ANALYZE_INST Procedure Parameters

Parameter	Description
task_name	Name of the task to be created
begin_snapshot	Number of the snapshot that starts the analysis period
end_snapshot	Number of the snapshot that ends the analysis period
cdb_type_override	Overrides the type of CDB that ADDM determines for doing analysis. The possible values are: • AUTONOMOUS OLTP—autonomous OLTP inside a PDB • AUTONOMOUS DATA WAREHOUSE—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows ADDM to decide the type of CDB to override based on the data
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data
instance_number	Number of the instance to analyze. By default it is the instance currently connected
db_id	Database ID for the database you to analyze. By default, this is the database currently connected

Return Values

The name of the created task is returned in the `task_name` parameter. It may be different from the value that is given as input (only in cases that name is already used by another task).

Usage Notes

On single instance systems (when not using Oracle RAC) the resulting task is identical to using the `ANALYZE_DB` procedure.

Examples

To create an ADDM task in instance analysis mode and execute it, with its name in variable `tname`:

```
var tname VARCHAR2(60);
BEGIN
  :tname := 'my_instance_analysis_mode_task';
  DBMS_ADDM.ANALYZE_INST(:tname, 1, 2);
END
```

To see a report:

```
SET LONG 100000
SET PAGESIZE 50000
SELECT DBMS_ADDM.GET_REPORT(:tname) FROM DUAL;
```

Note that the return type of a report is a `CLOB`, formatted to fit line size of 80.

15.2.3 ANALYZE_PARTIAL Procedure

This procedure creates an ADDM task for analyzing a subset of instances in partial analysis mode and executes it.

Syntax

```
DBMS_ADDM.ANALYZE_PARTIAL (
   task_name                 IN OUT VARCHAR2,
   instance_numbers          IN     VARCHAR2,
   begin_snapshot            IN     NUMBER,
   end_snapshot              IN     NUMBER,
   cdb_type_override         IN     VARCHAR2,
   read_only_type_override   IN     VARCHAR2,
   db_id                     IN     NUMBER := NULL);
```

Parameters

Table 15-4 ANALYZE_PARTIAL Procedure Parameters

Parameter	Description
`task_name`	Name of the task to be created
`instance_numbers`	Comma separated list of instance numbers to analyze
`begin_snapshot`	Number of the snapshot that starts the analysis period
`end_snapshot`	Number of the snapshot that ends the analysis period

Table 15-4 (Cont.) ANALYZE_PARTIAL Procedure Parameters

Parameter	Description
cdb_type_override	Overrides the type of CDB that ADDM determines for doing analysis. The possible values are: • AUTONOMOUS OLTP—autonomous OLTP inside a PDB • AUTONOMOUS DATA WAREHOUSE—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows ADDM to decide the type of CDB to override based on the data
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data
db_id	Database ID for the database you to analyze. By default, this is the database currently connected

Return Values

The name of the created task is returned in the task_name parameter. It may be different from the value that is given as input (only in cases that name is already used by another task).

Examples

To create an ADDM task in partial analysis mode and execute it, with its name in variable tname:

```
var tname VARCHAR2(60);
BEGIN
  :tname := 'my_partial_analysis_modetask';
  DBMS_ADDM.ANALYZE_PARTIAL(:tname, '1,2,3', 1, 2);
END
```

To see a report:

```
SET LONG 100000
SET PAGESIZE 50000
SELECT DBMS_ADDM.GET_REPORT(:tname) FROM DUAL;
```

Note that the return type of a report is a CLOB, formatted to fit line size of 80.

15.2.4 COMPARE_CAPTURE_REPLAY_REPORT Function

This function produces a Compare Period ADDM report comparing the performance of a capture to a replay.

The AWR data must reside in the same database, but it can originate from different databases. The function generates a report in either XML or HTML(Active Report) format.

Syntax

```
DBMS_ADDM.COMPARE_CAPTURE_REPLAY_REPORT (
   replay_id                 IN NUMBER,
   cdb_type_override         IN VARCHAR2,
   read_only_type_override   IN VARCHAR2,
   report_type               IN VARCHAR2 := 'HTML')
  RETURN CLOB;
```

Parameters

Table 15-5 COMPARE_CAPTURE_REPLAY_REPORT Function Parameters

Parameter	Description
`replay_id`	Replay ID to use as the base period. The base period is the baseline period to compare in order to determine improvement or regression.
`cdb_type_override`	Overrides the type of `CDB` that `ADDM` determines for doing analysis. The possible values are: • `AUTONOMOUS OLTP`—autonomous OLTP inside a PDB • `AUTONOMOUS DATA WAREHOUSE`—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows `ADDM` to decide the type of `CDB` to override based on the data
`read_only_type_override`	Overrides the type of `CDB` `ADDM` determines for analysis. The possible values are: • `READ-WRITE`—a regular database or the primary database in a data guard configuration • `READ-ONLY`—a database open in read-only mode, such as an active data guard standby • `AUTO`—allows `ADDM` to decide the type of `CDB` to override based on the data
`report_type`	`HTML` (the default) for an HTML active report, `'XML'` for an XML report

Return Values

A `CLOB` containing a compare period ADDM report

15.2.5 COMPARE_DATABASES Function

This function produces a Compare Period ADDM report comparing the performance of a database over two different time periods or the performance of two different databases over two different time periods.

The AWR data must reside in the same database, but it can originate from different databases. The function generates a report in either XML or HTML(Active Report) format.

Syntax

```
DBMS_ADDM.COMPARE_DATABASES (
   base_dbid             IN NUMBER := NULL,
   base_begin_snap_id    IN NUMBER,
   base_end_snap_id      IN NUMBER,
   comp_dbid             IN NUMBER := NULL,
   comp_begin_snap_id    IN NUMBER,
   comp_end_snap_id      IN NUMBER,
   cdb_type_override          IN     VARCHAR2,
   read_only_type_override    IN     VARCHAR2,
   report_type           IN VARCHAR2 := 'HTML')
  RETURN CLOB;
```

Parameters

Table 15-6 COMPARE_DATABASES Function Parameters

Parameter	Description
base_dbid	Database id (DBID) of the base period. The base period is the baseline period that we compare to in order to determine improvement or regression.
base_begin_snap_ids	Begin AWR snapshot ID of the base period.
base_end_snap_id	End AWR snapshot ID of the base period.
comp_dbid	Database id (DBID) of the comparison period. The comparison period is the period we compare to the base period.
comp_begin_snap_id	Begin AWR snapshot ID of the comparison period
comp_end_snap_id	End AWR snapshot ID of the comparison period
cdb_type_override	Overrides the type of CDB that ADDM determines for doing analysis. The possible values are: • AUTONOMOUS OLTP—autonomous OLTP inside a PDB • AUTONOMOUS DATA WAREHOUSE—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows ADDM to decide the type of CDB to override based on the data

Table 15-6 (Cont.) COMPARE_DATABASES Function Parameters

Parameter	Description
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data
report_type	'HTML' (the default) for an HTML active report, 'XML' for an XML report

Return Values

A CLOB containing a compare period ADDM report

15.2.6 COMPARE_INSTANCES Function

This function produces a Compare Period ADDM report comparing the performance of a single instance over two different time periods or the performance of two different instances over two different time periods.

The AWR data must reside in the same database, but it can originate from different databases. The function generates a report in either XML or HTML(Active Report) format.

Syntax

```
DBMS_ADDM.COMPARE_INSTANCES (
   base_dbid                 IN NUMBER := NULL,
   base_instance_id          IN NUMBER,
   base_begin_snap_id        IN NUMBER,
   base_end_snap_id          IN NUMBER,
   comp_dbid                 IN NUMBER := NULL,
   comp_instance_id          IN NUMBER,
   comp_begin_snap_id        IN NUMBER,
   comp_end_snap_id          IN NUMBER,
   cdb_type_override         IN VARCHAR2,
   read_only_type_override   IN VARCHAR2,
   report_type               IN VARCHAR2 := 'HTML')
  RETURN CLOB;
```

Parameters

Table 15-7 COMPARE_INSTANCES Function Parameters

Parameter	Description
base_dbid	Database id (DBID) of the base period. The base period is the baseline period that we compare to in order to determine improvement or regression.
base_instance_id	Instance number of the database instance to include from the base period
base_begin_snap_id	Begin AWR snapshot ID of the base period.

Table 15-7 (Cont.) COMPARE_INSTANCES Function Parameters

Parameter	Description
base_end_snap_id	End AWR snapshot ID of the base period.
comp_dbid	Database id (DBID) of the comparison period. The comparison period is the period we compare to the base period.
comp_instance_id	Instance number of the database instance to include from the comparison period
comp_begin_snap_id	Begin AWR snapshot ID of the comparison period
comp_end_snap_id	End AWR snapshot ID of the comparison period
cdb_type_override	Overrides the type of CDB that ADDM determines for doing analysis. The possible values are: • AUTONOMOUS OLTP—autonomous OLTP inside a PDB • AUTONOMOUS DATA WAREHOUSE—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows ADDM to decide the type of CDB to override based on the data
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data
report_type	'HTML' (the default) for an HTML active report, 'XML' for an XML report

Return Values

A CLOB containing a compare period ADDM report

15.2.7 COMPARE_REPLAY_REPLAY_REPORT Function

This function produces a Compare Period ADDM report comparing the performance of a replay to another replay.

The AWR data must reside in the same database, but it can originate from different databases. The function generates a report in either XML or HTML(Active Report) format.

Syntax

```
DBMS_ADDM.COMPARE_CAPTURE_REPLAY_REPORT (
   replay_id1                IN     NUMBER,
   replay_id2                IN     NUMBER,
```

```
  cdb_type_override          IN     VARCHAR2,
  read_only_type_override    IN     VARCHAR2,
  report_type                IN     VARCHAR2 := 'HTML')
 RETURN CLOB;
```

Parameters

Table 15-8 COMPARE_REPLAY_REPLAY_REPORT Function Parameters

Parameter	Description
replay_id1	Replay ID to use as the base period. The base period is the baseline period to compare in order to determine improvement or regression.
replay_id2	Replay ID to use as the comparison period. The comparison period is the period to compare to the base period in order to determine improvement or regression.
cdb_type_override	Overrides the type of CDB that ADDM determines for doing analysis. The possible values are: • AUTONOMOUS OLTP—autonomous OLTP inside a PDB • AUTONOMOUS DATA WAREHOUSE—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows ADDM to decide the type of CDB to override based on the data
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data
report_type	'HTML' (the default) for an HTML active report, 'XML' for an XML report

Return Values

A CLOB containing a compare period ADDM report

15.2.8 DELETE Procedure

This procedure deletes an already created ADDM task (of any kind). For database analysis mode and partial analysis mode this deletes the local tasks associated with the main task.

Syntax

```
DBMS_ADDM.DELETE (
   task_name           IN VARCHAR2);
```

Parameters

Table 15-9 DELETE Procedure Parameters

Parameter	Description
task_name	Name of the task to be deleted

Examples

```
BEGIN
  DBMS_ADDM.DELETE ('my_partial_analysis_mode_task');
END
```

15.2.9 DELETE_FINDING_DIRECTIVE Procedure

This procedure deletes a finding directive.

Syntax

```
DBMS_ADDM.DELETE_FINDING_DIRECTIVE (
   task_name           IN VARCHAR2,
   dir_name            IN VARCHAR2);
```

Parameters

Table 15-10 DELETE_FINDING_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it is a system directive.
dir_name	Name of the directive. All directives must be given unique names.

15.2.10 DELETE_PARAMETER_DIRECTIVE Procedure

This procedure deletes a parameter directive. This removes a specific system directive for parameters. Subsequent ADDM tasks are not affected by this directive.

Syntax

```
DBMS_ADDM.DELETE_PARAMETER_DIRECTIVE (
   task_name           IN VARCHAR2,
   dir_name            IN VARCHAR2);
```

Parameters

Table 15-11 DELETE_PARAMETER_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it is a system directive.
dir_name	Name of the directive. All directives must be given unique names.

Examples

```
BEGIN
   DBMS_ADDM.DELETE_PARAMETER_DIRECTIVE (NULL,'my Parameter directive');
END;
```

15.2.11 DELETE_SEGMENT_DIRECTIVE Procedure

This procedure deletes a segment directive.

Syntax

```
DBMS_ADDM.DELETE_SEGMENT_DIRECTIVE (
   task_name           IN VARCHAR2,
   dir_name            IN VARCHAR2);
```

Parameters

Table 15-12 DELETE_SEGMENT_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it is a system directive.
dir_name	Name of the directive. All directives must be given unique names.

15.2.12 DELETE_SQL_DIRECTIVE Procedure

This procedure deletes a SQL directive.

Syntax

```
DBMS_ADDM.DELETE_SQL_DIRECTIVE (
   task_name           IN VARCHAR2,
   dir_name            IN VARCHAR2);
```

Parameters

Table 15-13 DELETE_SQL_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it is a system directive.
dir_name	Name of the directive. All directives must be given unique names.

15.2.13 FAILED_AUTO_TASKS_REPORT Function

This function creates a plain text, user-readable report for the failed tasks registered in the `DBA_ADDM_PENDING_AUTOTASKS` views.

Syntax

```
DBMS_ADDM.FAILED_AUTO_TASKS_REPORT(
   instance_number     IN   NUMBER :=NULL,
   begin_snapshot      IN   NUMBER :=NULL,
   end_snapshot        IN   NUMBER :=NULL,
   dbid                IN   NUMBER :=NULL)
```

Parameters

Table 15-14 FAILED_AUTO_TASKS_REPORT Function Parameters

Parameter	Description
`instance_number`	Instance number for the tasks to be reported.
`begin_snapshot`	Earliest begin snapshot ID for the tasks to be reported.
`end_snapshot`	Latest end snapshot ID for the tasks to be reported.
`dbid`	Database ID for the tasks to be reported.

15.2.14 GET_ASH_QUERY Function

The function returns a string containing the SQL text of an ASH query identifying the rows in ASH with impact for the finding.

For most types of findings this identifies the exact rows in ASH corresponding to the finding. For some types of findings the query is an approximation and should not be used for exact identification of the finding's impact or the finding's specific activity.

Syntax

```
DBMS_ADDM.GET_ASH_QUERY (
   task_name           IN   VARCHAR2,
   finding_id          IN   NUMBER)
  RETURN VARCHAR2;
```

Parameters

Table 15-15 GET_ASH_QUERY Function Parameters

Parameter	Description
`task_name`	Name of the task
`finding`	ID of the finding within the task

Return Values

A `VARCHAR` containing an ASH query identifying the rows in ASH with impact for the finding

15.2.15 GET_REPORT Function

This function retrieves the default text report of an executed ADDM task.

Syntax

```
DBMS_ADDM.GET_REPORT (
   task_name           IN VARCHAR2)
  RETURN CLOB;
```

Parameters

Table 15-16 GET_REPORT Function Parameters

Parameter	Description
task_name	Name of the task

Examples

```
Set long 1000000
Set pagesize 50000
SELECT DBMS_ADDM.GET_REPORT('my_partial_analysis_mode_task') FROM DUAL;
```

15.2.16 INSERT_FINDING_DIRECTIVE Procedure

This procedure creates a directive to limit reporting of a specific finding type. The directive can be created for a specific task (only when the task is in INITIAL status), or for all subsequently created ADDM tasks (such as a system directive).

Syntax

```
DBMS_ADDM.INSERT_FINDING_DIRECTIVE (
   task_name             IN VARCHAR2,
   dir_name              IN VARCHAR2,
   finding_name          IN VARCHAR2,
   min_active_sessions   IN NUMBER := 0,
   min_perc_impact       IN NUMBER := 0);
```

Parameters

Table 15-17 INSERT_FINDING_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it applies to all subsequently created ADDM Tasks.
dir_name	Name of the directive. All directives must be given unique names.
finding_name	Name of an ADDM finding to which this directive applies. All valid findings names appear in the NAME column of view DBA_ADVISOR_FINDING_NAMES.
min_active_sessions	Minimal number of active sessions for the finding. If a finding has less than this number, it is filtered from the ADDM result.

Table 15-17 (Cont.) INSERT_FINDING_DIRECTIVE Procedure Parameters

Parameter	Description
min_perc_impact	Minimal number for the "percent impact" of the finding relative to total database time in the analysis period. If the finding's impact is less than this number, it is filtered from the ADDM result.

Examples

A new ADDM task is created to analyze a local instance. However, it has special treatment for 'Undersized SGA' findings. The result of GET_REPORT shows only an 'Undersized SGA' finding if the finding is responsible for at least 2 average active sessions during the analysis period, and this constitutes at least 10% of the total database time during that period.

```
var tname VARCHAR2(60);
BEGIN
  DBMS_ADDM.INSERT_FINDING_DIRECTIVE(
   NULL,
   'Undersized SGA directive',
   'Undersized SGA',
   2,
   10);
  :tname := 'my_instance_analysis_mode_task';
  DBMS_ADDM.ANALYZE_INST(:tname, 1, 2);
END;
```

To see a report containing 'Undersized SGA' findings regardless of the directive:

```
SELECT DBMS_ADVISOR.GET_TASK_REPORT(:tname, 'TEXT', 'ALL') FROM DUAL;
```

15.2.17 INSERT_PARAMETER_DIRECTIVE Procedure

This procedure creates a directive to prevent ADDM from creating actions to alter the value of a specific system parameter. The directive can be created for a specific task (only when the task is in INITIAL status), or for all subsequently created ADDM tasks (such as a system directive).

Syntax

```
DBMS_ADDM.INSERT_PARAMETER_DIRECTIVE (
   task_name             IN VARCHAR2,
   dir_name              IN VARCHAR2,
   parameter_name        IN VARCHAR2);
```

Parameters

Table 15-18 INSERT_PARAMETER_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it applies to all subsequently created ADDM Tasks.
dir_name	Name of the directive. All directives must be given unique names.
parameter_name	Specifies the parameter to use. Valid parameter names appear in V$PARAMETER.

Examples

A new ADDM task is created to analyze a local instance. However, it has special treatment for all actions that recommend modifying the parameter 'sga_target'. The result of GET_REPORT does not show these actions.

```
var tname varchar2(60);
BEGIN
  DBMS_ADDM.INSERT_PARAMETER_DIRECTIVE(
   NULL,
   'my Parameter directive',
   'sga_target');
  :tname := 'my_instance_analysis_mode_task';
  DBMS_ADDM.ANALYZE_INST(:tname, 1, 2);
END;
```

To see a report containing all actions regardless of the directive:

```
SELECT DBMS_ADVISOR.GET_TASK_REPORT(:tname, 'TEXT', 'ALL') FROM DUAL;
```

15.2.18 INSERT_SEGMENT_DIRECTIVE Procedure

This procedure creates a directive to prevent ADDM from creating actions to "run Segment Advisor" for specific segments. The directive can be created for a specific task (only when the task is in INITIAL status), or for all subsequently created ADDM tasks (such as a system directive).

Syntax

```
DBMS_ADDM.INSERT_SEGMENT_DIRECTIVE (
   task_name             IN VARCHAR2,
   dir_name              IN VARCHAR2,
   owner_name            IN VARCHAR2,
   object_name           IN VARCHAR2 := NULL,
   sub_object_name       IN VARCHAR2 := NULL);


DBMS_ADDM.INSERT_SEGMENT_DIRECTIVE (
   task_name             IN VARCHAR2,
   dir_name              IN VARCHAR2,
   object_number         IN NUMBER);
```

Parameters

Table 15-19 INSERT_SEGMENT_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it applies to all subsequently created ADDM Tasks.
dir_name	Name of the directive. All directives must be given unique names.
owner_name	Specifies the owner of the segment/s to be filtered. A wildcard is allowed in the same syntax used for "like" constraints.
object_name	Name of the main object to be filtered. Again, wildcards are allowed. The default value of NULL is equivalent to a value of '%'.

Table 15-19 (Cont.) INSERT_SEGMENT_DIRECTIVE Procedure Parameters

Parameter	Description
sub_object_name	Name of the part of the main object to be filtered. This could be a partition name, or even sub partitions (separated by a '.'). Again, wildcards are allowed. The default value of NULL is equivalent to a value of '%'.
object_number	Object number of the SEGMENT that this directive is to filter, found in views DBA_OBJECTS or DBA_SEGMENTS

Examples

A new ADDM task is created to analyze a local instance. However, it has special treatment for all segments that belong to user SCOTT. The result of GET_REPORT does not show actions for running Segment advisor for segments that belong to SCOTT.

```
var tname VARCHAR2(60);
BEGIN
  DBMS_ADDM.INSERT_SEGMENT_DIRECTIVE(NULL,
                                     'my Segment directive',
                                     'SCOTT');
  :tname := 'my_instance_analysis_mode_task';
  DBMS_ADDM.ANALYZE_INST(:tname, 1, 2);
END;
```

To see a report containing all actions regardless of the directive:

```
SELECT DBMS_ADVISOR.GET_TASK_REPORT(:tname, 'TEXT', 'ALL') FROM DUAL;
```

15.2.19 INSERT_SQL_DIRECTIVE Procedure

This procedure creates a directive to limit reporting of actions on specific SQL. The directive can be created for a specific task (only when the task is in INITIAL status), or for all subsequently created ADDM tasks (such as a system directive).

Syntax

```
DBMS_ADDM.INSERT_SQL_DIRECTIVE (
   task_name             IN VARCHAR2,
   dir_name              IN VARCHAR2,
   sql_id                IN VARCHAR2,
   min_active_sessions   IN NUMBER := 0,
   min_response_time     IN NUMBER := 0);
```

Parameters

Table 15-20 INSERT_SQL_DIRECTIVE Procedure Parameters

Parameter	Description
task_name	Name of the task this directive applies to. If the value is NULL, it applies to all subsequently created ADDM Tasks.
dir_name	Name of the directive. All directives must be given unique names.
sql_id	Identifies which SQL statement to filter. A valid value contains exactly 13 characters from '0' to '9' and 'a' to 'z'.

Table 15-20 (Cont.) INSERT_SQL_DIRECTIVE Procedure Parameters

Parameter	Description
min_active_sessions	Minimal number of active sessions for the SQL. If a SQL action has less than this number, it is filtered from the ADDM result.
min_response_time	Minimal value for response time of the SQL (in microseconds). If the SQL had lower response time, it is filtered from the ADDM result.

Examples

A new ADDM task is created to analyze a local instance. However, it has special treatment for SQL with id 'abcd123456789'. The result of GET_REPORT shows only actions for that SQL (actions to tune the SQL, or to investigate application using it) if the SQL is responsible for at least 2 average active sessions during the analysis period, and the average response time was at least 1 second.

```
var tname VARCHAR2(60);
BEGIN
  DBMS_ADDM.INSERT_SQL_DIRECTIVE(
      NULL,
      'my SQL directive',
      'abcd123456789',
      2,
      1000000);
  :tname := 'my_instance_analysis_mode_task';
  DBMS_ADDM.ANALYZE_INST(:tname, 1, 2);
END;
```

To see a report containing all actions regardless of the directive:

```
SELECT DBMS_ADVISOR.GET_TASK_REPORT(:tname, 'TEXT', 'ALL') FROM DUAL;
```

15.2.20 REAL_TIME_ADDM_REPORT Function

This function produces a real-time ADDM report for ADDM-related activity for the last five minutes. In an Oracle Real Application Clusters (Oracle RAC) environment, the function assumes that executing SQL over GV$ is possible.

Syntax

```
DBMS_ADDM.REAL_TIME_ADDM_REPORT (
   cdb_type_override         IN     VARCHAR2,
   read_only_type_override   IN     VARCHAR2)
 RETURN CLOB;
```

Parameters

Table 15-21 REAL_TIME_ADDM_REPORT Function Parameters

Parameter	Description
cdb_type_override	Overrides the type of CDB that ADDM determines for doing analysis. The possible values are: • AUTONOMOUS OLTP—autonomous OLTP inside a PDB • AUTONOMOUS DATA WAREHOUSE—autonomous data warehouse (ADWH) inside a PDB • PDB —a regular PDB • CDB ROOT —the root of a CDB • NON-CDB —a system that is not CDB or PDB • AUTO —allows ADDM to decide the type of CDB to override based on the data
read_only_type_override	Overrides the type of CDB ADDM determines for analysis. The possible values are: • READ-WRITE—a regular database or the primary database in a data guard configuration • READ-ONLY—a database open in read-only mode, such as an active data guard standby • AUTO—allows ADDM to decide the type of CDB to override based on the data

Return Values

CLOB containing a real-time ADDM report

15.2.21 REEXECUTE_FAILED_AUTO_TASKS Procedure

This procedure re-executes the tasks registered in the DBA_ADDM_PENDING_AUTOTASKS views.

This API can be called from both the CDB and PDB level. The user specifies the range for the snapshots or a time interval. The duration of the re-execution is time constrained by a user-specified parameter or a default value. Once a task has executed successfully, its corresponding row is removed from the DBA_ADDM_PENDING_AUTOTASKS view.

Syntax

```
DBMS_ADDM.REXECUCUTE_FAILED_AUTO_TASKS(
   instance_number      IN   NUMBER :=NULL,
   begin_snapshot       IN   NUMBER :=NULL,
   end_snapshot         IN   NUMBER :=NULL,
   dbid                 IN   NUMBER :=NULL,
   time_budget_in_sec   IN   NUMBER :=NULL,
   max_attempts         IN   NUMBER :=NULL);
```

Parameters

Table 15-22 REEXECUTE_FAILED_AUTO_TASKS Function Parameters

Parameter	Description
instance_number	Instance number for the tasks to be re executed.
begin_snapshot	Earliest begin snapshot ID for the tasks to be re executed.
end_snapshot	Latest end snapshot ID for the tasks to be re executed.
dbid	Database ID for the tasks to be re executed.
time_budget_in_sec	Forced time out to run the procedure after exiting.
max_attempts	Maximum number of attempts to re execute a task.

16
DBMS_ADVANCED_REWRITE

DBMS_ADVANCED_REWRITE contains interfaces for advanced query rewrite users. Using this package, you can create, drop, and maintain functional equivalence declarations for query rewrite.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information about query rewrite

This chapter contains the following topics:

- DBMS_ADVANCED_REWRITE Security Model
- Summary of DBMS_ADVANCED_REWRITE Subprograms

16.1 DBMS_ADVANCED_REWRITE Security Model

Default privileges are not granted to anyone for access to DBMS_ADVANCED_REWRITE Security Model procedures. To gain access to these procedures, you must connect as SYSDBA and explicitly grant execute access to the desired database administrators.

You can control security on this package by granting the EXECUTE privilege to selected database administrators or roles. For example, the user er can be given access to use this package by the following statement, executed as SYSDBA:

```
GRANT EXECUTE ON DBMS_ADVANCED_REWRITE TO er;
```

You may want to write a separate cover package on top of this package for restricting the alert names used. Instead of granting the EXECUTE privilege on the DBMS_ADVANCED_REWRITE package directly, you can then grant it to the cover package.

In addition, similar to the privilege required for regular materialized views, the user should be granted the privilege to create an equivalence. For example, the user er can be granted this privilege by executing the following statement as SYSDBA:

```
GRANT CREATE MATERIALIZED VIEW TO er;
```

16.2 Summary of DBMS_ADVANCED_REWRITE Subprograms

This table lists the DBMS_ADVANCED_REWRITE subprograms and briefly describes them.

Table 16-1 DBMS_ADVANCED_REWRITE Package Subprograms

Subprogram	Description
ALTER_REWRITE_EQUIVALENCE Procedure	Changes the mode of the rewrite equivalence declaration to the mode you specify

Table 16-1 (Cont.) DBMS_ADVANCED_REWRITE Package Subprograms

Subprogram	Description
BUILD_SAFE_REWRITE_EQUIVALENCE Procedure	Enables the rewrite of top-level materialized views using submaterialized views. Oracle Corporation does not recommend you directly use this procedure
DECLARE_REWRITE_EQUIVALENCE Procedures	Creates a declaration indicating that `source_stmt` is functionally equivalent to `destination_stmt` for as long as the equivalence declaration remains enabled, and that `destination_stmt` is more favorable in terms of performance
DROP_REWRITE_EQUIVALENCE Procedure	Drops the specified rewrite equivalence declaration
VALIDATE_REWRITE_EQUIVALENCE Procedure	Validates the specified rewrite equivalence declaration using the same validation method as described with the `validate` parameter

16.2.1 ALTER_REWRITE_EQUIVALENCE Procedure

This table list the all the package subprograms in alphabetical order.

Syntax

```
DBMS_ADVANCED_REWRITE.ALTER_REWRITE_EQUIVALENCE (
   name            VARCHAR2,
   rewrite_mode    VARCHAR2);
```

Parameters

Table 16-2 ALTER_REWRITE_EQUIVALENCE Procedure Parameters

Parameter	Description
`name`	A name for the equivalence declaration to alter. The name can be of the form `owner.name`, where `owner` complies with the rules for a schema name, and `name` compiles with the rules for a table name. Alternatively, a simple name that complies with the rules for a table name can be specified. In this case, the rewrite equivalence is altered in the current schema. The invoker must have the appropriate alter materialized view privileges to alter an equivalence declaration outside their own schema.

Table 16-2 (Cont.) ALTER_REWRITE_EQUIVALENCE Procedure Parameters

Parameter	Description
rewrite_mode	The following modes are supported, in increasing order of power: disabled: Query rewrite does not use the equivalence declaration. Use this mode to temporarily disable use of the rewrite equivalence declaration. text_match: Query rewrite uses the equivalence declaration only in its text match modes. This mode is useful for simple transformations. general: Query rewrite uses the equivalence declaration in all of its transformation modes against the incoming request queries. However, query rewrite makes no attempt to rewrite the specified destination_query. recursive: Query rewrite uses the equivalence declaration in all of its transformation modes against the incoming request queries. Moreover, query rewrite further attempts to rewrite the specified destination_query for further performance enhancements whenever it uses the equivalence declaration. Oracle recommends you use the least powerful mode that is sufficient to solve your performance problem.

16.2.2 BUILD_SAFE_REWRITE_EQUIVALENCE Procedure

This procedure enables the rewrite and refresh of top-level materialized views using submaterialized views. It is provided for the exclusive use by scripts generated by the DBMS_ADVISOR.TUNE_MVIEW procedure.

It is required to enable query rewrite and fast refresh when DBMS_ADVISOR.TUNE_MVIEW decomposes a materialized view into a top-level materialized view and one or more submaterialized views.

Oracle does not recommend you directly use the BUILD_SAFE_REWRITE_EQUIVALENCE procedure. You should use either the DBMS_ADVISOR.TUNE_MVIEW or the DBMS_ADVANCED_REWRITE.CREATE_REWRITE_EQUIVALENCE procedure as appropriate.

16.2.3 DECLARE_REWRITE_EQUIVALENCE Procedures

This procedure creates a declaration indicating that source_stmt is functionally equivalent to destination_stmt for as long as the equivalence declaration remains enabled, and that destination_stmt is more favorable in terms of performance.

The scope of the declaration is system wide. The query rewrite engine uses such declarations to perform rewrite transformations in QUERY_REWRITE_INTEGRITY = trusted and stale_tolerated modes.

Because the underlying equivalences between the source and destination statements cannot be enforced by the query rewrite engine, queries can be only rewritten in trusted and stale_tolerated integrity modes.

Syntax

```
DBMS_ADVANCED_REWRITE.DECLARE_REWRITE_EQUIVALENCE (
   name                 VARCHAR2,
   source_stmt          VARCHAR2,
   destination_stmt     VARCHAR2,
   validate             BOOLEAN    := TRUE,
```

```
   rewrite_mode         VARCHAR2   := 'TEXT_MATCH');

DBMS_ADVANCED_REWRITE.DECLARE_REWRITE_EQUIVALENCE (
   name                 VARCHAR2,
   source_stmt          CLOB,
   destination_stmt     CLOB,
   validate             BOOLEAN    := TRUE,
   rewrite_mode         VARCHAR2   := 'TEXT_MATCH');
```

Parameters

Table 16-3 DECLARE_REWRITE_EQUIVALENCE Procedure Parameters

Parameter	Description
`name`	A name for the equivalence declaration. The name can be of the form `owner.name`, where owner complies with the rules for a schema name, and name compiles with the rules for a table name. Alternatively, a simple name that complies with the rules for a table name can be specified. In this case, the rewrite equivalence is created in the current schema. The invoker must have the appropriate `CREATE MATERIALIZED VIEW` privileges to alter an equivalence declaration.
`source_stmt`	A sub-`SELECT` expression in either `VARCHAR2` or `CLOB` format. This is the query statement that is the target of optimization.
`destination_stmt`	A sub-`SELECT` expression in either `VARCHAR2` or `CLOB` format.
`validate`	A Boolean indicating whether to validate that the specified `source_stmt` is functionally equivalent to the specified `destination_stmt`. If `validate` is specified as `TRUE`, `DECLARE_REWRITE_EQUIVALENCE` evaluates the two sub-`SELECT`s and compares their results. If the results are not the same, `DECLARE_REWRITE_EQUIVALENCE` does not create the rewrite equivalence and returns an error condition. If `FALSE`, `DECLARE_REWRITE_EQUIVALENCE` does not validate the equivalence.
`rewrite_mode`	The following modes are supported, in increasing order of power: • `disabled`: Query rewrite does not use the equivalence declaration. Use this mode to temporarily disable use of the rewrite equivalence declaration. • `text_match`: Query rewrite uses the equivalence declaration only in its text match modes. This mode is useful for simple transformations. • `general`: Query rewrite uses the equivalence declaration in all of its transformation modes against the incoming request queries. However, query rewrite makes no attempt to rewrite the specified `destination_query`. • `recursive`: Query rewrite uses the equivalence declaration in all of its transformation modes against the incoming request queries. Moreover, query rewrite further attempts to rewrite the specified `destination_query` for further performance enhancements whenever it uses the equivalence declaration. Oracle recommends you use the least powerful mode that is sufficient to solve your performance problem.

Exceptions

Table 16-4 DECLARE_REWRITE_EQUIVALENCE Procedure Exceptions

Exception	Description
ORA-30388	Name of the rewrite equivalence is not specified
ORA-30391	The specified rewrite equivalence does not exist
ORA-30392	The checksum analysis for the rewrite equivalence failed
ORA-30393	A query block in the statement did not write
ORA-30396	Rewrite equivalence procedures require the COMPATIBLE parameter to be set to 10.1 or greater

Usage Notes

Query rewrite using equivalence declarations occurs simultaneously and in concert with query rewrite using materialized views. The same query rewrite engine is used for both. The query rewrite engine uses the same rewrite rules to rewrite queries using both equivalence declarations and materialized views. Because the rewrite equivalence represents a specific rewrite crafted by a sophisticated user, the query rewrite engine gives priority to rewrite equivalences over materialized views when it is possible to perform a rewrite with either a materialized view or a rewrite equivalence. For this same reason, the cost-based optimizer (specifically, cost-based rewrite) will not choose an unrewritten query plan over a query plan that is rewritten to use a rewrite equivalence even if the cost of the un-rewritten plan appears more favorable. Query rewrite matches properties of the incoming request query against the equivalence declaration's source_stmt or the materialized view's defining statement, respectively, and derives an equivalent relational expression in terms of the equivalence declaration's destination_stmt or the materialized view's container table, respectively.

16.2.4 DROP_REWRITE_EQUIVALENCE Procedure

This procedure drops the specified rewrite equivalence declaration.

Syntax

```
DBMS_ADVANCED_REWRITE.DROP_REWRITE_EQUIVALENCE (
   name         VARCHAR2);
```

Parameters

Table 16-5 DROP_REWRITE_EQUIVALENCE Procedure Parameters

Parameter	Description
name	A name for the equivalence declaration to drop. The name can be of the form owner.name, where owner complies with the rules for a schema name, and name compiles with the rules for a table name. Alternatively, a simple name that complies with the rules for a table name can be specified. In this case, the rewrite equivalence is dropped in the current schema. The invoker must have the appropriate drop materialized view privilege to drop an equivalence declaration outside their own schema.

16.2.5 VALIDATE_REWRITE_EQUIVALENCE Procedure

This procedure validates the specified rewrite equivalence declaration.

It uses the same validation method as described with the `VALIDATE` parameter in "VALIDATE_REWRITE_EQUIVALENCE Procedure".

Syntax

```
DBMS_ADVANCED_REWRITE.VALIDATE_REWRITE_EQUIVALENCE (
   name         VARCHAR2);
```

Parameters

Table 16-6 VALIDATE_REWRITE_EQUIVALENCE Procedure Parameters

Parameter	Description
`name`	A name for the equivalence declaration to validate. The name can be of the form owner.name, where owner complies with the rules for a schema name, and name compiles with the rules for a table name. Alternatively, a simple name that complies with the rules for a table name can be specified. In this case, the rewrite equivalence is validated in the current schema. The invoker must have sufficient privileges to execute both the `source_stmt` and `destination_stmt` of the specified equivalence declaration.

17
DBMS_ADVISOR

`DBMS_ADVISOR` is part of the server manageability suite of advisors, a set of expert systems that identifies and helps resolve performance problems relating to database server components.

Some advisors have their own packages. For these advisors, Oracle recommends that you use the advisor-specific package rather than `DBMS_ADVISOR`. Each of the following advisors has its own package, tailored to its specific functionality:

- Automatic Database Diagnostic Monitor (`DBMS_ADDM`)
- SQL Performance Analyzer (`DBMS_SQLPA`)
- SQL Repair Advisor (`DBMS_SQLDIAG`)
- SQL Tuning Advisor (`DBMS_SQLTUNE`)
- Compression Advisor (`DBMS_COMPRESSION.GET_COMPRESSION_RATIO`)

SQL Access Advisor and Segment Advisor are the only advisors with common use cases for `DBMS_ADVISOR`. Undo Advisor and Compression Advisor do not support `DBMS_ADVISOR` subprograms.

This chapter contains the following topics:

- DBMS_ADVISOR Deprecated Subprograms
- DBMS_ADVISOR Security Model
- Summary of DBMS_ADVISOR Subprograms

> **See Also:**
>
> - *Oracle Database Administrator's Guide* to learn about Segment Advisor
> - *Oracle Database Get Started with Performance Tuning* to learn how to use SQL Access Advisor in Enterprise Manager
> - *Oracle Database SQL Tuning Guide* to learn more about SQL Access Advisor

17.1 DBMS_ADVISOR Deprecated Subprograms

The section lists programs that are deprecated with Oracle Database 11*g*.

> **Note:**
>
> Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

The following subprograms are deprecated:

- ADD_SQLWKLD_REF Procedure
- CREATE_SQLWKLD Procedure
- DELETE_SQLWKLD Procedure
- DELETE_SQLWKLD_REF Procedure
- DELETE_SQLWKLD_STATEMENT Procedure
- IMPORT_SQLWKLD_SCHEMA Procedure
- IMPORT_SQLWKLD_SQLCACHE Procedure
- IMPORT_SQLWKLD_STS Procedure
- IMPORT_SQLWKLD_SUMADV Procedure
- IMPORT_SQLWKLD_USER Procedure
- RESET_SQLWKLD Procedure
- SET_SQLWKLD_PARAMETER Procedure
- UPDATE_SQLWKLD_ATTRIBUTES Procedure
- UPDATE_SQLWKLD_STATEMENT Procedure

17.2 DBMS_ADVISOR Security Model

The `ADVISOR` privilege is required to use the DBMS_ADVISOR package.

17.3 Summary of DBMS_ADVISOR Subprograms

This topic lists and describes the subprograms in the DBMS_ADVISOR package.

In the following table, the `Used in` column lists advisors relevant for each subprogram, but excludes ADDM, SQL Performance Analyzer, SQL Repair Advisor, and SQL Tuning Advisor because these advisors have their own packages.

Table 17-1 DBMS_ADVISOR Package Subprograms

Subprogram	Description	Used in
ADD_SQLWKLD_REF Procedure	Adds a workload reference to an Advisor task (Caution: Deprecated Subprogram)	SQL Access Advisor
ADD_SQLWKLD_STATEMENT Procedure	Adds a single statement to a workload	SQL Access Advisor
ADD_STS_REF Procedure	Establishes a link between the current SQL Access Advisor task and a SQL tuning set	SQL Access Advisor
CANCEL_TASK Procedure	Cancels a currently executing task operation	Segment Advisor, SQL Access Advisor
COPY_SQLWKLD_TO_STS Procedure	Copies the contents of a SQL workload object to a SQL tuning set	SQL Access Advisor
CREATE_FILE Procedure	Creates an external file from a PL/SQL CLOB variable, which is useful for creating scripts and reports	SQL Access Advisor

Table 17-1 (Cont.) DBMS_ADVISOR Package Subprograms

Subprogram	Description	Used in
CREATE_OBJECT Procedure	Creates a new task object	Segment Advisor
CREATE_SQLWKLD Procedure	Creates a new workload object (Caution: Deprecated Subprogram)	SQL Access Advisor
CREATE_TASK Procedures	Creates a new Advisor task in the repository	Segment Advisor, SQL Access Advisor
DELETE_SQLWKLD Procedure	Deletes an entire workload object (Caution: Deprecated Subprogram)	SQL Access Advisor
DELETE_SQLWKLD_REF Procedure	Deletes an entire workload object (Caution: Deprecated Subprogram)	SQL Access Advisor
DELETE_SQLWKLD_STATEMENT Procedure	Deletes one or more statements from a workload (Caution: Deprecated Subprogram)	SQL Access Advisor
DELETE_STS_REF Procedure	Removes a link between the current SQL Access Advisor task and a SQL tuning set object	SQL Access Advisor
DELETE_TASK Procedure	Deletes the specified task from the repository	SQL Access Advisor
EXECUTE_TASK Procedure	Executes the specified task	Segment Advisor, SQL Access Advisor
GET_REC_ATTRIBUTES Procedure	Retrieves specific recommendation attributes from a task	SQL Access Advisor
GET_TASK_REPORT Function	Creates and returns a report for the specified task	
GET_TASK_SCRIPT Function	Creates and returns an executable SQL script of the Advisor task's recommendations in a buffer	SQL Access Advisor
IMPLEMENT_TASK Procedure	Implements the recommendations for a task	SQL Access Advisor
IMPORT_SQLWKLD_SCHEMA Procedure	Imports data into a workload from the current SQL cache (Caution: Deprecated Subprogram)	SQL Access Advisor
IMPORT_SQLWKLD_SQLCACHE Procedure	Imports data into a workload from the current SQL cache (Caution: Deprecated Subprogram)	SQL Access Advisor
IMPORT_SQLWKLD_STS Procedure	Imports data from a SQL tuning set into a SQL workload data object (Caution: Deprecated Subprogram)	SQL Access Advisor
IMPORT_SQLWKLD_SUMADV Procedure	Imports data into a workload from the current SQL cache (Caution: Deprecated Subprogram)	SQL Access Advisor
IMPORT_SQLWKLD_USER Procedure	Imports data into a workload from the current SQL cache (Caution: Deprecated Subprogram)	SQL Access Advisor
INTERRUPT_TASK Procedure	Stops a currently executing task, ending its operations as it would at a normal exit, so that the recommendations are visible	Segment Advisor, SQL Access Advisor

Table 17-1 (Cont.) DBMS_ADVISOR Package Subprograms

Subprogram	Description	Used in
MARK_RECOMMENDATION Procedure	Sets the `annotation_status` for a particular recommendation	Segment Advisor, SQL Access Advisor
QUICK_TUNE Procedure	Performs an analysis on a single SQL statement	SQL Access Advisor
RESET_SQLWKLD Procedure	Resets a workload to its initial starting point (Caution: Deprecated Subprogram)	SQL Access Advisor
RESET_TASK Procedure	Resets a task to its initial state	Segment Advisor, SQL Access Advisor
SET_DEFAULT_SQLWKLD_PARAMETER Procedure	Imports data into a workload from schema evidence	SQL Access Advisor
SET_DEFAULT_TASK_PARAMETER Procedure	Modifies a default task parameter	Segment Advisor, SQL Access Advisor
SET_SQLWKLD_PARAMETER Procedure	Sets the value of a workload parameter	SQL Access Advisor
SET_TASK_PARAMETER Procedure	Sets the specified task parameter value	Segment Advisor, SQL Access Advisor
TUNE_MVIEW Procedure	Shows how to decompose a materialized view into two or more materialized views or to restate the materialized view in a way that is more advantageous for fast refresh and query rewrite	SQL Access Advisor
UPDATE_OBJECT Procedure	Updates a task object	Segment Advisor
UPDATE_REC_ATTRIBUTES Procedure	Updates an existing recommendation for the specified task	SQL Access Advisor
UPDATE_SQLWKLD_ATTRIBUTES Procedure	Updates a workload object	SQL Access Advisor
UPDATE_SQLWKLD_STATEMENT Procedure	Updates one or more SQL statements in a workload	SQL Access Advisor
UPDATE_TASK_ATTRIBUTES Procedure	Updates a task's attributes	Segment Advisor, SQL Access Advisor

17.3.1 ADD_SQLWKLD_REF Procedure

This procedure establishes a link between the current SQL Access Advisor task and a SQL Workload object.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

The link allows an advisor task to access interesting data for doing an analysis. The link also provides a stable view of the data. Once a connection between a SQL Access Advisor task and a SQL Workload object is made, the workload is protected from removal or modification.

Users should use `ADD_STS_REF` instead of `ADD_SQLWKLD_REF` for all SQL tuning set-based advisor runs. This function is only provided for backward compatibility.

Syntax

```
DBMS_ADVISOR.ADD_SQLWKLD_REF (
   task_name              IN VARCHAR2,
   workload_name          IN VARCHAR2,
   is_sts                 IN NUMBER :=0);
```

Parameters

Table 17-2 ADD_SQLWKLD_REF Procedure Parameters

Parameter	Description
`task_name`	The SQL Access Advisor task name that uniquely identifies an existing task.
`workload_name`	The name of the workload object to be linked. Once a object has been linked to a task, it becomes read-only and cannot be deleted. There is no limit to the number of links to workload objects. To remove the link to the workload object, use the procedure `DELETE_REFERENCE`.
`is_sts`	Indicates the type of workload source. Possible values are: • 0 - SQL workload object • 1 - SQL tuning set

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name, 1);
```

```
END;
/
```

17.3.2 ADD_SQLWKLD_STATEMENT Procedure

This procedure adds a single statement to the specified workload.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT (
   workload_name          IN VARCHAR2,
   module                 IN VARCHAR2,
   action                 IN VARCHAR2,
   cpu_time               IN NUMBER := 0,
   elapsed_time           IN NUMBER := 0,
   disk_reads             IN NUMBER := 0,
   buffer_gets            IN NUMBER := 0,
   rows_processed         IN NUMBER := 0,
   optimizer_cost         IN NUMBER := 0,
   executions             IN NUMBER := 1,
   priority               IN NUMBER := 2,
   last_execution_date    IN DATE := 'SYSDATE',
   stat_period            IN NUMBER := 0,
   username               IN VARCHAR2,
   sql_text               IN CLOB);
```

Parameters

Table 17-3 ADD_SQLWKLD_STATEMENT Procedure Parameters

Parameter	Description
workload_name	The workload name that uniquely identifies an existing workload.
module	An optional business application module that will be associated with the SQL statement.
action	An optional application action that will be associated with the SQL statement.
cpu_time	The total CPU time in seconds that is consumed by the SQL statement.
elapsed_time	The total elapsed time in seconds that is consumed by the SQL statement.
disk_reads	The total disk-read operations that are consumed by the SQL statement.
buffer_gets	The total buffer-get operations that are consumed by the SQL statement.
rows_processed	The average number of rows processed by the SQL statement.
optimizer_cost	The cost value calculated by the optimizer.
executions	The total execution count of the SQL statement. This value should be greater than zero.

Table 17-3 (Cont.) ADD_SQLWKLD_STATEMENT Procedure Parameters

Parameter	Description
priority	The relative priority of the SQL statement. The value must be one of the following: 1-HIGH, 2-MEDIUM, or 3-LOW.
last_execution_date	The date and time at which the SQL statement last executed. If the value is NULL, then the database uses the current date and time.
stat_period	Time interval in seconds from which statement statistics were calculated.
username	The database user that executed the SQL statement. Because a user name is an Oracle identifier, the username value must be entered exactly as it is stored in the server. For example, if the user SCOTT is the executing user, then you must provide the user identifier SCOTT in all uppercase letters. It will not recognize the user scott or Scott as a match for SCOTT.
sql_text	The complete SQL statement. To increase the quality of a recommendation, the SQL statement should not contain bind variables.

Usage Notes

You cannot modify or delete a workload when it is currently referenced by an active task. A task is considered active if it is not in its initial state. See RESET_TASK Procedure for directions on setting a task to its initial state.

The ADD_SQLWKLD_STATEMENT procedure accepts several parameters that may be ignored by the caller. The database only uses the disk_reads, buffer_gets, and optimizer_cost parameters to sort workload data when actual analysis occurs. Therefore, actual values are only necessary when the order_list task parameter references a particular statistic.

To determine what statistics to provide when adding a new SQL statement to a workload, examine or set the task parameter order_list. The order_list parameter accepts any combination of the keys:

- cpu_time
- elapsed_time
- buffer_gets
- optimizer_cost
- disk_reads
- executions
- priority

The optimizer_cost key, which is a typical setting of priority, indicates that SQL Access Advisor sorts the workload data by priority and optimizer_cost, and processes the highest cost statements first. Any statements that you add to the workload must include appropriate priority and optimizer_cost values. All other statistics can be defaulted or set to zero.

For the statistical keys referenced by the order_list task parameter, the actual parameter values should be reasonably accurate since they will be compared to other statements in the workload. If the caller is unable to estimate values, then choose values that would determine its importance relative to other statements in the workload. For example, if the current statement is considered the most critical query in your business, then an appropriate value would be anything greater than all other values for the same statistic found in the workload.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT
AVG(amount_sold) FROM sh.sales');
END;
/
```

17.3.3 ADD_STS_REF Procedure

This procedure establishes a link between the current SQL Access Advisor task and a SQL tuning set.

The link enables an advisor task to access data for the purpose of doing an analysis. The link also provides a stable view of the data. Once a connection between a SQL Access Advisor task and a SQL tuning set is made, the STS is protected from removal or modification.

Use `ADD_STS_REF` for any STS-based advisor runs. The older method of using `ADD_SQLWKLD_REF` with parameter `IS_STS=1` is only supported for backward compatibility. Furthermore, the `ADD_STS_REF` function accepts a SQL tuning set owner name, whereas `ADD_SQLWKLD_REF` does not.

Syntax

```
DBMS_ADVISOR.ADD_STS_REF(
  task_name       IN VARCHAR2 NOT NULL,
  sts_owner       IN VARCHAR2,
  workload_name   IN VARCHAR2 NOT NULL);
```

Parameters

Table 17-4 ADD_STS_REF Procedure Parameters

Parameter	Description
`task_name`	The SQL Access Advisor task name that uniquely identifies an existing task.
`sts_owner`	The owner of the SQL tuning set. The value of this parameter may be `NULL`, in which case the advisor assumes the SQL tuning set to be owned by the currently logged-in user.

Table 17-4 (Cont.) ADD_STS_REF Procedure Parameters

Parameter	Description
workload_name	The name of the workload to be linked. A workload consists of one or more SQL statements, plus statistics and attributes that fully describe each statement. The database stores a workload as a SQL tuning set. After a workload has been linked to a task, it becomes read-only and cannot be deleted. There is no limit to the number of links to workloads. To remove the link to the workload, use the procedure DBMS_ADVISOR.DELETE_STS_REF.

Examples

```
DBMS_ADVISOR.ADD_STS_REF ('My Task', 'SCOTT', 'My Workload');
```

17.3.4 CANCEL_TASK Procedure

This procedure causes a currently executing operation to terminate.

This call performs a soft interrupt. It will not break into a low-level database access call like a hard interrupt such as Ctrl-C. The SQL Access Advisor periodically checks for soft interrupts and acts appropriately. As a result, this operation may take a few seconds to respond to a call.

Syntax

```
DBMS_ADVISOR.CANCEL_TASK (
   task_name      IN  VARCHAR2);
```

Parameters

Table 17-5 CANCEL_TASK Procedure Parameter

Parameter	Description
task_name	A valid Advisor task name that uniquely identifies an existing task.

Usage Notes

A cancel command restores the task to its condition prior to the start of the canceled operation. Therefore, a canceled task or data object cannot be resumed.

Because all Advisor task procedures are synchronous, to cancel an operation, you must use a separate database session.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
BEGIN
```

```
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.CANCEL_TASK('My Task');
END;
/
```

17.3.5 COPY_SQLWKLD_TO_STS Procedure

This procedure copies the contents of a SQL workload object to a SQL tuning set.

Syntax

To use this procedure, the caller must have privileges to create and modify a SQL tuning set.

```
DBMS_ADVISOR.COPY_SQLWKLD_TO_STS (
   workload_name          IN VARCHAR2,
   sts_name               IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW');
```

Parameters

Table 17-6 COPY_SQLWKLD_TO_STS Procedure Parameter

Parameter	Description
workload_name	The SQL Workload object name to copy.
sts_name	The SQL tuning set name into which the SQL Workload object will be copied.
import_mode	Specifies the handling of the target SQL tuning set. Possible values are: • APPEND Causes SQL Workload data to be appended to the target SQL tuning set. • NEW Indicates the SQL tuning set can only contain the copied contents. If the SQL tuning set exists and has data, an error will be reported. • REPLACE Causes any existing data in the target SQL tuning set to be purged prior to the workload copy. In all cases, if the specified SQL tuning set does not exist, it will be created.

Usage Notes

To use this procedure, the caller must have privileges to create and modify a SQL tuning set.

Examples

```
BEGIN
   DBMS_ADVISOR.COPY_SQLWKLD_TO_STS('MY_OLD_WORKLOAD', 'MY_NEW_STS', 'NEW');
END;
/
```

17.3.6 CREATE_FILE Procedure

This procedure creates an external file from a PL/SQL `CLOB` variable, which is used for creating scripts and reports.

Syntax

```
DBMS_ADVISOR.CREATE_FILE (
   buffer       IN  CLOB,
   location     IN  VARCHAR2,
   filename     IN  VARCHAR2);
```

Parameters

Table 17-7 CREATE_FILE Procedure Parameters

Parameter	Description
`buffer`	A `CLOB` buffer containing report or script information.
`location`	The name of the directory that will contain the output file. You must use the alias as defined by the `CREATE DIRECTORY` statement. The Advisor translates the alias into the actual directory location.
`filename`	The name of the output file. The file name can only contain the name and an optional file type of the form `filename.filetype`.

Usage Notes

You must embed all formatting within the `CLOB`.

The database restricts file access within stored procedures. This means that file locations and names must adhere to the known file permissions in the server.

Examples

```
CREATE DIRECTORY MY_DIR as '/homedir/user4/gssmith';
GRANT READ,WRITE ON DIRECTORY MY_DIR TO PUBLIC;

DECLARE
  v_task_id NUMBER;
  v_task_name VARCHAR2(30);
  v_workload_name VARCHAR2(30);
BEGIN
  v_task_name := 'My Task';
  v_workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(
      advisor_name => DBMS_ADVISOR.SQLACCESS_ADVISOR
  ,   task_id      => v_task_id
  ,   task_name    => v_task_name );
  DBMS_ADVISOR.CREATE_SQLWKLD(
      workload_name => v_workload_name
  ,   description   => 'My Workload' );
```

```
  DBMS_ADVISOR.ADD_SQLWKLD_REF(
      task_name     => v_task_name
  ,   workload_name => v_workload_name);
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(
      workload_name       => v_workload_name
  ,   module              => 'MONTHLY'
  ,   action              => 'ROLLUP'
  ,   cpu_time            => 100
  ,   elapsed_time        => 400
  ,   disk_reads          => 5041
  ,   buffer_gets         => 103
  ,   rows_processed      => 640445
  ,   optimizer_cost      => 680000
  ,   executions          => 2
  ,   priority            => 1
  ,   last_execution_date => SYSDATE
  ,   stat_period         => 1
  ,   username            => 'SH'
  ,   sql_text            => 'SELECT AVG(amount_sold) FROM sh.sales' );
  DBMS_ADVISOR.EXECUTE_TASK(v_task_name);
  DBMS_ADVISOR.CREATE_FILE(
      buffer   => DBMS_ADVISOR.GET_TASK_SCRIPT(v_task_name)
  ,   location => 'MY_DIR'
  ,   filename => 'script.sql' );
END;
/
```

17.3.7 CREATE_OBJECT Procedure

This procedure creates a new task object.

Syntax

```
DBMS_ADVISOR.CREATE_OBJECT (
   task_name         IN VARCHAR2,
   object_type       IN VARCHAR2,
   attr1             IN VARCHAR2 :=  NULL,
   attr2             IN VARCHAR2 :=  NULL,
   attr3             IN VARCHAR2 :=  NULL,
   attr4             IN CLOB     :=  NULL,
   attr5             IN VARCHAR2 :=  NULL,
   object_id         OUT NUMBER,
   attr6             IN VARCHAR2 :=  NULL,
   attr7             IN VARCHAR2 :=  NULL,
   attr8             IN VARCHAR2 :=  NULL,
   attr9             IN VARCHAR2 :=  NULL,
   attr10            IN VARCHAR2 :=  NULL);
```

Parameters

Table 17-8 CREATE_OBJECT Procedure Parameters

Parameter	Description
task_name	A valid Advisor task name that uniquely identifies an existing task.
object_type	Specifies the external object type.
attr1	Advisor-specific data.
attr2	Advisor-specific data.
attr3	Advisor-specific data.
attr4	Advisor-specific data.
attr5	Advisor-specific data.
object_id	The advisor-assigned object identifier.
attr6	Advisor-specific data.
attr7	Advisor-specific data.
attr8	Advisor-specific data.
attr9	Advisor-specific data.
attr10	Advisor-specific data.

The attribute parameters have different values depending upon the object type. See *Oracle Database Administrator's Guide* for details regarding these parameters and object types.

Return Values

Returns the new object identifier.

Usage Notes

Task objects are typically used as input data for a particular advisor. Segment advice can be generated at the object, segment, or tablespace level. If for the object level, advice is generated on all partitions of the object (if the object is partitioned). The advice is not cascaded to any dependent objects. If for the segment level, advice can be obtained on a single segment, such as the partition or subpartition of a table, index, or `LOB` column. If for a tablespace level, target advice for every segment in the tablespace will be generated.

See *Oracle Database Administrator's Guide* for further information regarding the Segment Advisor.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  obj_id NUMBER;
BEGIN
  task_name := 'My Task';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.CREATE_OBJECT (task_name,'SQL',NULL,NULL,NULL, 'SELECT * FROM
```

```
TEST_TAB',NULL,obj_id,NULL,NULL,NULL,NULL,NULL);
END;
/
```

17.3.8 CREATE_SQLWKLD Procedure

This procedure creates a new private SQL Workload object for the user.

A SQL Workload object manages a SQL workload on behalf of the SQL Access Advisor. A SQL Workload object must exist prior to performing any other SQL Workload operations, such as importing or updating SQL statements.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.CREATE_SQLWKLD (
   workload_name            IN OUT VARCHAR2,
   description              IN VARCHAR2 := NULL,
   template                 IN VARCHAR2 := NULL,
   is_template              IN VARCHAR2 := 'FALSE');
```

Parameters

Table 17-9 CREATE_SQLWKLD Procedure Parameters

Parameter	Description
`workload_name`	A name that uniquely identifies the created workload. If not specified, the system will generate a unique name. Names can be up to 30 characters long.
`description`	Specifies an optional workload description. Descriptions can be up to 256 characters.
`template`	An optional SQL Workload name of an existing workload data object or data object template.
`is_template`	An optional value that enables you to set the newly created workload as a template. Valid values are `TRUE` and `FALSE`.

Return Values

The SQL Access Advisor returns a unique workload object identifier number that must be used for subsequent activities within the new SQL Workload object.

Usage Notes

By default, workload objects are created using built-in default settings. To create a workload using the parameter settings of an existing workload or workload template, the user may specify an existing workload name.

After a SQL Workload object is present, it can then be referenced by one or more SQL Access Advisor tasks using the `ADD_SQLWKLD_REF` procedure.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
END;
/
```

17.3.9 CREATE_TASK Procedures

This procedure creates a new Advisor task in the repository.

Syntax

```
DBMS_ADVISOR.CREATE_TASK (
   advisor_name          IN VARCHAR2,
   task_id               OUT NUMBER,
   task_name             IN OUT VARCHAR2,
   task_desc             IN VARCHAR2 := NULL,
   template              IN VARCHAR2 := NULL,
   is_template           IN VARCHAR2 := 'FALSE',
   how_created           IN VARCHAR2 := NULL);

DBMS_ADVISOR.CREATE_TASK (
   advisor_name          IN VARCHAR2,
   task_name             IN VARCHAR2,
   task_desc             IN VARCHAR2 := NULL,
   template              IN VARCHAR2 := NULL,
   is_template           IN VARCHAR2 := 'FALSE',
   how_created           IN VARCHAR2 := NULL);

DBMS_ADVISOR.CREATE_TASK (
   parent_task_name      IN VARCHAR2,
   rec_id                IN NUMBER,
   task_id               OUT NUMBER,
   task_name             IN OUT VARCHAR2,
   task_desc             IN VARCHAR2,
   template              IN VARCHAR2);
```

Parameters

Table 17-10 CREATE_TASK Procedure Parameters

Parameter	Description
advisor_name	Specifies the unique advisor name as defined in the view DBA_ADVISOR_DEFINITIONS.
task_id	A number that uniquely identifies the created task. The number is generated by the procedure and returned to the user.

Table 17-10 (Cont.) CREATE_TASK Procedure Parameters

Parameter	Description
task_name	Specifies a new task name. Names must be unique among all tasks for the user. When using the second form of the CREATE_TASK syntax listed above (with OUT), a unique name can be generated. Names can be up to 30 characters long.
task_desc	Specifies an optional task description. Descriptions can be up to 256 characters in length.
template	An optional task name of an existing task or task template. To specify built-in SQL Access Advisor templates, use the template name as described earlier.
is_template	An optional value that allows the user to set the newly created task as template. Valid values are: TRUE and FALSE.
how_created	An optional value that identifies how the source was created.

Return Values

Returns a unique task ID number and a unique task name if one is not specified.

Usage Notes

A task must be associated with an advisor, and once the task has been created, it is permanently associated with the original advisor. By default, tasks are created using built-in default settings. To create a task using the parameter settings of an existing task or task template, the user may specify an existing task name.

For the SQL Access Advisor, use the identifier DBMS_ADVISOR.SQLACCESS_ADVISOR as the advisor_name.

The SQL Access Advisor provides three built-in task templates, using the following constants:

- DBMS_ADVISOR.SQLACCESS_OLTP

  Parameters are preset to favor an OLTP application environment.

- DBMS_ADVISOR.SQLACCESS_WAREHOUSE

  Parameters are preset to favor a data warehouse application environment.

- DBMS_ADVISOR.SQLACCESS_GENERAL

  Parameters are preset to favor a hybrid application environment where both OLTP and data warehouse operations may occur. For the SQL Access Advisor, this is the default template.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';
  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
```

```
END;
/
```

17.3.10 DELETE_SQLWKLD Procedure

This procedure deletes an existing SQL Workload object from the repository.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.DELETE_SQLWKLD (
   workload_name         IN VARCHAR2);
```

Parameters

Table 17-11 DELETE_SQLWKLD Procedure Parameters

Parameter	Description
`workload_name`	The workload object name that uniquely identifies an existing workload. The wildcard `%` is supported as a `WORKLOAD_NAME`. The rules of use are identical to the `LIKE` operator. For example, to delete all tasks for the current user, use the wildcard `%` as the `WORKLOAD_NAME`. If a wildcard is provided, the `DELETE_SQLWKLD` operation will not delete any workloads marked as `READ_ONLY` or `TEMPLATE`.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See the RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.DELETE_SQLWKLD(workload_name);
END;
/
```

17.3.11 DELETE_SQLWKLD_REF Procedure

This procedure removes a link between the current SQL Access task and a SQL Workload data object.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Use `DELETE_STS_REF` instead of `DELETE_SQLWKLD_REF` for all SQL tuning set-based advisor runs. This function is only provided for backward compatibility.

Syntax

```
DBMS_ADVISOR.DELETE_SQLWKLD_REF (
   task_name              IN VARCHAR2,
   workload_name          IN VARCHAR2,
   is_sts                 IN NUMBER :=0);
```

Parameters

Table 17-12 DELETE_SQLWKLD_REF Procedure Parameters

Parameter	Description
`task_name`	The SQL Access task name that uniquely identifies an existing task.
`workload_name`	The name of the workload object to be unlinked. The wildcard `%` is supported as a `workload_name`. The rules of use are identical to the `LIKE` operator. For example, to remove all links to workload objects, use the wildcard `%` as the `workload_name`.
`is_sts`	Indicates the type of workload source. Possible values are: • 0 - SQL workload object • 1 - SQL tuning set

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.DELETE_SQLWKLD_REF(task_name, workload_name);
END;
/
```

17.3.12 DELETE_SQLWKLD_STATEMENT Procedure

This procedure deletes one or more statements from a workload.

> **Note:**
>
> This procedure has been deprecated.

Syntax

```
DBMS_ADVISOR.DELETE_SQLWKLD_STATEMENT (
   workload_name        IN VARCHAR2,
   sql_id               IN NUMBER);

DBMS_ADVISOR.DELETE_SQLWKLD_STATEMENT (
   workload_name        IN VARCHAR2,
   search               IN VARCHAR2,
   deleted              OUT NUMBER);
```

Parameters

Table 17-13 DELETE_SQLWKLD_STATEMENT Procedure Parameters

Parameter	Description
workload_name	The workload object name that uniquely identifies an existing workload.
sql_id	The Advisor-generated identifier number that is assigned to the statement. To specify all workload statements, use the constant ADVISOR_ALL.
search	Disabled.
deleted	Returns the number of statements deleted by the searched deleted operation.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See the RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
  deleted NUMBER;
  id NUMBER;
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'YEARLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
```

```
                                        FROM sh.sales');

   SELECT sql_id INTO id FROM USER_ADVISOR_SQLW_STMTS
   WHERE workload_name = 'My Workload';

  DBMS_ADVISOR.DELETE_SQLWKLD_STATEMENT(workload_name, id);
END;
/
```

17.3.13 DELETE_STS_REF Procedure

This procedure removes a link between the current SQL Access Advisor task and a SQL tuning set.

Use `DELETE_STS_REF` for any STS-based advisor runs. The older method of using `DELETE_SQLWKLD_REF` with parameter `IS_STS=1` is only supported for backward compatibility. Furthermore, the `DELETE_STS_REF` function accepts an STS owner name, whereas `DELETE_SQLWKLD_REF` does not.

Syntax

```
DBMS_ADVISOR.DELETE_STS_REF (
  task_name      IN VARCHAR2 NOT NULL,
  sts_owner      IN VARCHAR2,
  workload_name  IN VARCHAR2 NOT NULL);
```

Parameters

Table 17-14 DELETE_STS_REF Procedure Parameters

Parameter	Description
`task_name`	The SQL Access Advisor task name that uniquely identifies an existing task.
`sts_owner`	The owner of the SQL tuning set. The value of this parameter may be `NULL`, in which case the advisor assumes the SQL tuning set to be owned by the currently logged-in user.
`workload_name`	The name of the workload to be unlinked. A workload consists of one or more SQL statements, plus statistics and attributes that fully describe each statement. The database stores a workload as a SQL tuning set. The wildcard `%` is supported as a workload name. The rules of use are identical to the SQL `LIKE` operator. For example, to remove all links to SQL tuning set objects, use the wildcard `%` as the `STS_NAME`.

Examples

```
DBMS_ADVISOR.DELETE_STS_REF ('My task', 'SCOTT', 'My workload');
```

17.3.14 DELETE_TASK Procedure

This procedure deletes an existing task from the repository.

Syntax

```
DBMS_ADVISOR.DELETE_TASK (
   task_name          IN VARCHAR2);
```

Parameters

Table 17-15 DELETE_TASK Procedure Parameters

Parameter	Description
task_name	A single Advisor task name that will be deleted from the repository. The wildcard % is supported as a TASK_NAME. The rules of use are identical to the LIKE operator. For example, to delete all tasks for the current user, use the wildcard % as the TASK_NAME. If a wildcard is provided, the DELETE_TASK operation will not delete any tasks marked as READ_ONLY or TEMPLATE.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.DELETE_TASK(task_name);
END;
/
```

17.3.15 EXECUTE_TASK Procedure

This procedure performs the Advisor analysis or evaluation for the specified task. The procedure is overloaded.

The execution-related arguments are optional and you do not need to set them for advisors that do not allow their tasks to be executed multiple times.

Advisors can execute a task multiple times and use the results for further processing and analysis.

Syntax

```
DBMS_ADVISOR.EXECUTE_TASK (
   task_name          IN VARCHAR2);

DBMS_ADVISOR.EXECUTE_TASK (
```

```
   task_name         IN VARCHAR2,
   execution_type    IN VARCHAR2             := NULL,
   execution_name    IN VARCHAR2             := NULL,
   execution_params  IN dbms_advisor.argList := NULL,
   execution_desc    IN VARCHAR2             := NULL,
RETURN VARCHAR2;
```

Parameters

Table 17-16 EXECUTE_TASK Procedure Parameters

Parameter	Description
task_name	The task name that uniquely identifies an existing task.
execution_type	The type of action to be performed by the function. If NULL, it will default to the value of the DEFAULT_EXECUTION_TYPE parameter. As an example, the SQL Performance Analyzer accepts the following possible values: • EXPLAIN PLAN: Generate an explain plan for a SQL statement. This is similar to an EXPLAIN PLAN command. The resulting plans will be stored in the advisor framework in association with the task. • TEST EXECUTE: Test execute the SQL statement and collect its execute plan and statistics. The resulting plans and statistics are stored in the advisor framework. • ANALYZE PERFORMANCE: Analyze and compare two versions of SQL performance data. The performance data is generated by test executing a SQL statement or generating its explain plan.
execution_name	A name to qualify and identify an execution. If not specified, it will be generated by the Advisor and returned by function.
execution_params	A list of parameters (name, value) for the specified execution. Note that execution parameters are real task parameters, but they affect only the execution they are specified for. As an example, consider the following: `DBMS_ADVISOR.ARGLIST('time_limit', 12, 'username', 'hr')`
execution_desc	A 256-length string describing the execution.

Usage Notes

Task execution is a synchronous operation. Control will not be returned to the caller until the operation has completed, or a user-interrupt was detected.

Upon return, you can check the DBA_ADVISOR_LOG table for the execution status.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';
```

```
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.EXECUTE_TASK(task_name);
END;
/
```

17.3.16 GET_REC_ATTRIBUTES Procedure

This procedure retrieves a specified attribute of a new object as recommended by Advisor analysis.

Syntax

```
DBMS_ADVISOR.GET_REC_ATTRIBUTES (
   workload_name         IN VARCHAR2,
   rec_id                IN NUMBER,
   action_id             IN NUMBER,
   attribute_name        IN VARCHAR2,
   value                 OUT VARCHAR2,
   owner_name            IN VARCHAR2 := NULL);
```

Parameters

Table 17-17 GET_REC_ATTRIBUTES Procedure Parameters

Parameter	Description
task_name	The task name that uniquely identifies an existing task.
rec_id	The Advisor-generated identifier number that is assigned to the recommendation.
action_id	The Advisor-generated action identifier that is assigned to the particular command.
attribute_name	Specifies the attribute to change.
value	The buffer to receive the requested attribute value.
owner_name	Optional owner name of the target task. This permits access to task data not owned by the current user.

Return Values

The requested attribute value is returned in the VALUE argument.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
  attribute VARCHAR2(100);
BEGIN
  task_name := 'My Task';
```

```
  workload_name := 'My Workload';
  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales WHERE promo_id = 10');
  DBMS_ADVISOR.EXECUTE_TASK(task_name);
  DBMS_ADVISOR.GET_REC_ATTRIBUTES(task_name, 1, 1, 'NAME', attribute);
END;
/
```

17.3.17 GET_TASK_REPORT Function

This function creates and returns a report for the specified task.

Syntax

```
DBMS_ADVISOR.GET_TASK_REPORT (
   task_name      IN VARCHAR2,
   type           IN VARCHAR2 := 'TEXT',
   level          IN VARCHAR2 := 'TYPICAL',
   section        IN VARCHAR2 := 'ALL',
   owner_name     IN VARCHAR2 := NULL,
   execution_name IN VARCHAR2 := NULL,
   object_id      IN NUMBER   := NULL)
RETURN CLOB;
```

Parameters

Table 17-18 GET_TASK_REPORT Function Parameters

Parameter	Description
task_name	The name of the task from which the script will be created.
type	The only valid value is TEXT.
level	The possible values are BASIC, TYPICAL, and ALL.
section	Advisor-specific report sections.
owner_name	Owner of the task. If specified, the system will check to see if the current user has read privileges to the task data.
execution_name	An identifier of a specific execution of the task. It is needed only for advisors that allow their tasks to be executed multiple times.
object_id	An identifier of an advisor object that can be targeted by the script.

Return Values

Returns the buffer receiving the script.

17.3.18 GET_TASK_SCRIPT Function

This function creates a SQL*Plus-compatible SQL script and sends the output to a file.

The output script contains all of the accepted recommendations from the specified task.

Syntax

```
DBMS_ADVISOR.GET_TASK_SCRIPT (
   task_name          IN VARCHAR2
   type               IN VARCHAR2 := 'IMPLEMENTATION',
   rec_id             IN NUMBER   := NULL,
   act_id             IN NUMBER   := NULL,
   owner_name         IN VARCHAR2 := NULL,
   execution_name     IN VARCHAR2 := NULL,
   object_id          IN NUMBER   := NULL)
RETURN CLOB;
```

Parameters

Table 17-19 GET_TASK_SCRIPT Function Parameters

Parameter	Description
`task_name`	The task name that uniquely identifies an existing task.
`type`	Specifies the type of script to generate. The possible values are `IMPLEMENTATION` and `UNDO`.
`rec_id`	An optional recommendation identifier number that can be used to extract a subset of the implementation script. A zero or the value `DBMS_ADVISOR.ADVISOR_ALL` indicates all accepted recommendations would be included. The default is to include all accepted recommendations for the task.
`act_id`	Optional action identifier number that can be used to extract a single action as a DDL command. A zero or the value `DBMS_ADVISOR.ADVISOR_ALL` indicates all actions for the recommendation would be included. The default is to include all actions for a recommendation.
`owner_name`	An optional task owner name.
`execution_name`	An identifier of a specific execution of the task. It is needed only for advisors that allow their tasks to be executed multiple times.
`object_id`	An identifier of an advisor object that can be targeted by the script.

Return Values

Returns the script as a `CLOB` buffer.

Usage Notes

Though the script is ready to execute, Oracle recommends that the user review the script for acceptable locations for new materialized views and indexes.

For a recommendation to appear in a generated script, it must be marked as accepted.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
  buf CLOB;
BEGIN
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales');
  DBMS_ADVISOR.EXECUTE_TASK(task_name);
    buf := DBMS_ADVISOR.GET_TASK_SCRIPT(task_name);
END;
/
```

17.3.19 IMPLEMENT_TASK Procedure

This procedure implements the recommendations of the specified Advisor task.

Syntax

```
DBMS_ADVISOR.IMPLEMENT_TASK (
   task_name          IN VARCHAR2,
   rec_id             IN NUMBER := NULL,
   exit_on_error      IN BOOLEAN := NULL);
```

Parameters

Table 17-20 IMPLEMENT_TASK Procedure Parameters

Parameter	Description
task_name	The name of the task.
rec_id	An optional recommendation ID.
exit_on_error	An optional Boolean to exit on the first error.

17.3.20 IMPORT_SQLWKLD_SCHEMA Procedure

This procedure constructs and loads a SQL workload based on schema evidence. The workload is also referred to as a hypothetical workload.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.IMPORT_SQLWKLD_SCHEMA (
   workload_name          IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW',
   priority               IN NUMBER := 2,
   saved_rows             OUT NUMBER,
   failed_rows            OUT NUMBER);
```

Parameters

Table 17-21 IMPORT_SQLWKLD_SCHEMA Procedure Parameters

Parameter	Description
workload_name	The workload object name that uniquely identifies an existing workload.
import_mode	Specifies the action to be taken when storing the workload. Possible values are: • APPEND Indicates that the collected workload will be added to any existing workload in the task. • NEW Indicates that the collected workload will be the exclusive workload for the task. If an existing workload is found, an exception will be thrown. • REPLACE Indicates the collected workload will be the exclusive workload for the task. If an existing workload is found, it will be deleted prior to saving the new workload. The default value is NEW.
priority	Specifies the application priority for each statement that is saved in the workload object. The value must be one of the following: 1-HIGH, 2-MEDIUM, or 3-LOW.
failed_rows	Returns the number or rows that were not saved due to syntax or validation errors
saved_rows	Returns the number of rows actually saved in the repository.

Return Values

This call returns the number of rows saved and failed as output parameters.

Usage Notes

To successfully import a hypothetical workload, the target schemas must contain dimensions.

If the VALID_TABLE_LIST parameter is not set, the search space may become very large and require a significant amount of time to complete. Oracle recommends that you limit your search space to specific set of tables.

If a task contains valid recommendations from a prior run, adding or modifying task will mark the task as invalid, preventing the viewing and reporting of potentially valuable recommendation data.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
  saved NUMBER;
  failed NUMBER;
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.SET_SQLWKLD_PARAMETER(workload_name,'VALID_TABLE_LIST','SH.%');
  DBMS_ADVISOR.IMPORT_SQLWKLD_SCHEMA(workload_name, 'REPLACE', 1, saved,
     failed);
END;
/
```

17.3.21 IMPORT_SQLWKLD_SQLCACHE Procedure

This procedure creates a SQL workload from the current contents of the server's SQL cache.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.IMPORT_SQLWKLD_SQLCACHE (
   workload_name          IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW',
   priority               IN NUMBER := 2,
   saved_rows             OUT NUMBER,
   failed_rows            OUT NUMBER);
```

Parameters

Table 17-22 IMPORT_SQLWKLD_SQLCACHE Procedure Parameters

Parameter	Description
workload_name	The workload object name that uniquely identifies an existing workload.

Table 17-22 (Cont.) IMPORT_SQLWKLD_SQLCACHE Procedure Parameters

Parameter	Description
import_mode	Specifies the action to be taken when storing the workload. Possible values are: • APPEND Indicates that the collected workload will be added to any existing workload in the task. • NEW Indicates that the collected workload will be the exclusive workload for the task. If an existing workload is found, an exception will be thrown. • REPLACE Indicates the collected workload will be the exclusive workload for the task. If an existing workload is found, it will be deleted prior to saving the new workload. The default value is NEW.
priority	Specifies the application priority for each statement that is saved in the workload object. The value must be one of the following 1-HIGH, 2-MEDIUM, or 3-LOW.
saved_rows	Returns the number of rows saved as output parameters.
failed_rows	Returns the number of rows that were not saved due to syntax or validation errors.

Return Values

This call returns the number of rows saved and failed as output parameters.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
  saved NUMBER;
  failed NUMBER;
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.SET_SQLWKLD_PARAMETER(workload_name,'VALID_TABLE_LIST','SH.%');
  DBMS_ADVISOR.IMPORT_SQLWKLD_SQLCACHE(workload_name, 'REPLACE', 1, saved,
     failed);
END;
/
```

17.3.22 IMPORT_SQLWKLD_STS Procedure

This procedure loads a SQL workload from an existing SQL tuning set. A SQL tuning set is typically created from the server workload repository using various time and data filters.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.IMPORT_SQLWKLD_STS (
   workload_name          IN VARCHAR2,
   sts_name               IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW',
   priority               IN NUMBER := 2,
   saved_rows             OUT NUMBER,
   failed_rows            OUT NUMBER);


DBMS_ADVISOR.IMPORT_SQLWKLD_STS (
   workload_name          IN VARCHAR2,
   sts_owner              IN VARCHAR2,
   sts_name               IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW',
   priority               IN NUMBER := 2,
   saved_rows             OUT NUMBER,
   failed_rows            OUT NUMBER);
```

Parameters

Table 17-23 IMPORT_SQLWKLD_STS Procedure Parameters

Parameter	Description
`workload_name`	The workload object name that uniquely identifies an existing workload.
`sts_owner`	The optional owner of the SQL tuning set.
`sts_name`	The name of an existing SQL tuning set workload from which the data will be imported. If the `sts_owner` value is not provided, the owner will default to the current user.
`import_mode`	Specifies the action to be taken when storing the workload. Possible values are: • `APPEND` Indicates that the collected workload will be added to any existing workload in the task. • `NEW` Indicates that the collected workload will be the exclusive workload for the task. If an existing workload is found, an exception will be thrown. • `REPLACE` Indicates the collected workload will be the exclusive workload for the task. If an existing workload is found, it will be deleted prior to saving the new workload. The default value is `NEW`.

Table 17-23 (Cont.) IMPORT_SQLWKLD_STS Procedure Parameters

Parameter	Description
priority	Specifies the application priority for each statement that is saved in the workload object. The value must be one of the following: 1-HIGH, 2-MEDIUM, or 3-LOW. The default value is 2.
saved_rows	Returns the number of rows actually saved in the repository.
failed_rows	Returns the number of rows that were not saved due to syntax or validation errors.

Return Values

This call returns the number of rows saved and failed as output parameters.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
  saved NUMBER;
  failed NUMBER;
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.SET_SQLWKLD_PARAMETER(workload_name,'VALID_TABLE_LIST','SH.%');
  DBMS_ADVISOR.IMPORT_SQLWKLD_STS(workload_name, 'MY_SQLSET', 'REPLACE', 1,
     saved, failed);
END;
/
```

17.3.23 IMPORT_SQLWKLD_SUMADV Procedure

This procedure collects a SQL workload from a Summary Advisor workload.

This procedure is intended to assist Oracle9*i* Database Summary Advisor users in the migration to SQL Access Advisor.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.IMPORT_SQLWKLD_SUMADV (
   workload_name          IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW',
   priority               IN NUMBER := 2,
   sumadv_id              IN NUMBER,
   saved_rows             OUT NUMBER,
   failed_rows            OUT NUMBER);
```

Parameters

Table 17-24 IMPORT_SQLWKLD_SUMADV Procedure Parameters

Parameter	Description
workload_name	The workload object name that uniquely identifies an existing workload.
import_mode	Specifies the action to be taken when storing the workload. Possible values are: • APPEND Indicates that the collected workload will be added to any existing workload in the task. • NEW Indicates that the collected workload will be the exclusive workload for the task. If an existing workload is found, an exception will be thrown. • REPLACE Indicates the collected workload will be the exclusive workload for the task. If an existing workload is found, it will be deleted prior to saving the new workload. The default value is NEW.
priority	Specifies the default application priority for each statement that is saved in the workload object. If a Summary Advisor workload statement contains a priority of zero, the default priority will be applied. If the workload statement contains a valid priority, then the Summary Advisor priority will be converted to a comparable SQL Access Advisor priority. The value must be one of the following: 1-HIGH, 2-MEDIUM, or 3-LOW.
sumadv_id	Specifies the Summary Advisor workload identifier number.
saved_rows	Returns the number of rows actually saved in the repository.
failed_rows	Returns the number of rows that were not saved due to syntax or validation errors.

Return Values

This call returns the number of rows saved and failed as output parameters.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
```

```
  saved NUMBER;
  failed NUMBER;
  sumadv_id NUMBER;
BEGIN
  workload_name := 'My Workload';
  sumadv_id := 394;

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.SET_SQLWKLD_PARAMETER(workload_name,'VALID_TABLE_LIST','SH.%');
  DBMS_ADVISOR.IMPORT_SQLWKLD_SUMADV(workload_name, 'REPLACE', 1, sumadv_id,
    saved, failed);
END;
/
```

17.3.24 IMPORT_SQLWKLD_USER Procedure

This procedure collects a SQL workload from a specified user table.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.IMPORT_SQLWKLD_USER (
   workload_name          IN VARCHAR2,
   import_mode            IN VARCHAR2 := 'NEW',
   owner_name             IN VARCHAR2,
   table_name             IN VARCHAR2,
   saved_rows             OUT NUMBER,
   failed_rows            OUT NUMBER);
```

Parameters

Table 17-25 IMPORT_SQLWKLD_USER Procedure Parameters

Parameter	Description
`workload_name`	The workload object name that uniquely identifies an existing workload.
`import_mode`	Specifies the action to be taken when storing the workload. Possible values are: • `APPEND` Indicates that the collected workload will be added to any existing workload in the task. • `NEW` Indicates that the collected workload will be the exclusive workload for the task. If an existing workload is found, an exception will be thrown. • `REPLACE` Indicates the collected workload will be the exclusive workload for the task. If an existing workload is found, it will be deleted prior to saving the new workload. The default value is `NEW`.
`owner_name`	Specifies the owner name of the table or view from which workload data will be collected.

Table 17-25 (Cont.) IMPORT_SQLWKLD_USER Procedure Parameters

Parameter	Description
table_name	Specifies the name of the table or view from which workload data will be collected.
saved_rows	Returns the number of rows actually saved in the workload object.
failed_rows	Returns the number of rows that were not saved due to syntax or validation errors.

Return Values

This call returns the number of rows saved and failed as output parameters.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
  saved NUMBER;
  failed NUMBER;
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.SET_SQLWKLD_PARAMETER(workload_name,'VALID_TABLE_LIST','SH.%');
  DBMS_ADVISOR.IMPORT_SQLWKLD_USER(workload_name, 'REPLACE', 'SH',
    'USER_WORKLOAD', saved, failed);
END;
/
```

17.3.25 INTERRUPT_TASK Procedure

This procedure stops a currently executing task.

The task will end its operations as it would at a normal exit. The user will be able to access any recommendations that exist to this point.

Syntax

```
DBMS_ADVISOR.INTERRUPT_TASK (
   task_name          IN VARCHAR2);
```

Parameters

Table 17-26 INTERRUPT_TASK Procedure Parameters

Parameter	Description
task_name	A single Advisor task name that will be interrupted.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.EXECUTE_TASK(task_name);
END;
/
```

While this session is executing its task, you can interrupt the task from a second session using the following statement:

```
BEGIN
  DBMS_ADVISOR.INTERRUPT_TASK('My Task');
END;
/
```

17.3.26 MARK_RECOMMENDATION Procedure

This procedure marks a recommendation for import or implementation.

Syntax

```
DBMS_ADVISOR.MARK_RECOMMENDATION (
   task_name          IN VARCHAR2
   id                 IN NUMBER,
   action             IN VARCHAR2);
```

Parameters

Table 17-27 MARK_RECOMMENDATION Procedure Parameters

Parameter	Description
task_name	Name of the task.
id	The recommendation identifier number assigned by the Advisor.

Table 17-27 (Cont.) MARK_RECOMMENDATION Procedure Parameters

Parameter	Description
action	The recommendation action setting. The possible actions are: • ACCEPT Marks the recommendation as accepted. With this setting, the recommendation will appear in implementation and undo scripts. • IGNORE Marks the recommendation as ignore. With this setting, the recommendation will not appear in an implementation or undo script. • REJECT Marks the recommendation as rejected. With this setting, the recommendation will not appear in any implementation or undo scripts.

Usage Notes

For a recommendation to be implemented, it must be marked as accepted. By default, all recommendations are considered accepted and will appear in any generated scripts.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
  attribute VARCHAR2(100);
  rec_id NUMBER;
BEGIN
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales WHERE promo_id = 10');
  DBMS_ADVISOR.EXECUTE_TASK(task_name);

  rec_id := 1;
  DBMS_ADVISOR.MARK_RECOMMENDATION(task_name, rec_id, 'REJECT');
END;
/
```

17.3.27 QUICK_TUNE Procedure

This procedure performs an analysis and generates recommendations for a single SQL statement.

This provides a shortcut method of all necessary operations to analyze the specified SQL statement. The operation creates a task using the specified task name. The task will be created using a specified Advisor task template. Finally, the task will be executed and the results will be saved in the repository.

Syntax

```
DBMS_ADVISOR.QUICK_TUNE (
   advisor_name           IN VARCHAR2,
   task_name              IN VARCHAR2,
   attr1                  IN CLOB,
   attr2                  IN VARCHAR2 := NULL,
   attr3                  IN NUMBER := NULL,
   template               IN VARCHAR2 := NULL,
   implement              IN BOOLEAN := FALSE,
   description            IN VARCHAR2 := NULL);
```

Parameters

Table 17-28 QUICK_TUNE Procedure Parameters

Parameter	Description
advisor_name	Name of the Advisor that will perform the analysis.
task_name	Name of the task.
attr1	Advisor-specific attribute in the form of a CLOB variable.
attr2	Advisor-specific attribute in the form of a VARCHAR2 variable.
attr3	Advisor-specific attribute in the form of a NUMBER.
template	Name of an existing task or template from which the initial settings need to be copied.
implement	Flag specifying whether to implement the task.
description	Description of the task.

Usage Notes

If indicated by the user, the final recommendations can be implemented by the procedure.

The task will be created using either a specified SQL Access task template or the built-in default template of SQLACCESS_GENERAL. The workload will only contain the specified statement, and all task parameters will be defaulted.

attr1 must be the single SQL statement to tune. For the SQL Access Advisor, attr2 is the user who would execute the single statement. If omitted, the current user will be used.

Examples

```
DECLARE
  task_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';

  DBMS_ADVISOR.QUICK_TUNE(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_name,
                  'SELECT AVG(amount_sold) FROM sh.sales WHERE promo_id=10');
END;
/
```

17.3.28 RESET_SQLWKLD Procedure

This procedure resets a workload to its initial starting point.

Resetting the workload has the effect of removing all journal and log messages, and recalculating necessary volatility and usage statistics.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.RESET_SQLWKLD (
   workload_name        IN VARCHAR2);
```

Parameters

Table 17-29 RESET_SQLWKLD Procedure Parameters

Parameter	Description
workload_name	The SQL Workload object name that uniquely identifies an existing workload.

Usage Notes

`RESET_SQLWKLD` should be executed after any workload adjustments such as adding or removing SQL statements.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales WHERE promo_id = 10');

  DBMS_ADVISOR.RESET_SQLWKLD(workload_name);
END;
/
```

17.3.29 RESET_TASK Procedure

This procedure re-initializes the metadata for the specified task. The task status will be set to `INITIAL`.

Syntax

```
DBMS_ADVISOR.RESET_TASK (
   task_name          IN VARCHAR2);
```

Parameters

Table 17-30 RESET_TASK Procedure Parameters

Parameter	Description
`task_name`	The task name that uniquely identifies an existing task.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.EXECUTE_TASK(task_name);
  DBMS_ADVISOR.RESET_TASK(task_name);
END;
/
```

17.3.30 SET_DEFAULT_SQLWKLD_PARAMETER Procedure

This procedure modifies the default value for a user parameter within a SQL Workload object or SQL Workload object template.

A user parameter is a simple variable that stores various attributes that affect workload collection, tuning decisions and reporting. When a default value is changed for a parameter, workload objects will inherit the new value when they are created.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.SET_DEFAULT_SQLWKLD_PARAMETER (
   parameter            IN VARCHAR2,
   value                IN VARCHAR2);

DBMS_ADVISOR.SET_DEFAULT_SQLWKLD_PARAMETER (
   parameter            IN VARCHAR2,
   value                IN NUMBER);
```

Parameters

Table 17-31 SET_DEFAULT_SQLWKLD_PARAMETER Procedure Parameters

Parameter	Description
parameter	The name of the data parameter to be modified. Parameter names are not case sensitive. Parameter names are unique to the workload object type, but not necessarily unique to all workload object types. Various object types may use the same parameter name for different purposes.
value	The value of the specified parameter. The value can be specified as a string or a number. If the value is DBMS_ADVISOR.DEFAULT, the value will be reset to the default value.

Usage Notes

A parameter will only affect operations that modify the workload collection. Therefore, parameters should be set prior to importing or adding new SQL statements to a workload. If a parameter is set after data has been placed in a workload object, it will have no effect on the existing data.

Examples

```
BEGIN
  DBMS_ADVISOR.SET_DEFAULT_SQLWKLD_PARAMETER('VALID_TABLE_LIST','SH.%');
END;
/
```

17.3.31 SET_DEFAULT_TASK_PARAMETER Procedure

This procedure modifies the default value for a user parameter within a task or a template.

A user parameter is a simple variable that stores various attributes that affect various Advisor operations. When a default value is changed for a parameter, tasks will inherit the new value when they are created.

A default task is different from a regular task. The default value is the initial value that will be inserted into a newly created task, while setting a task parameter with SET_TASK_PARAMETER sets the local value only. Thus, SET_DEFAULT_TASK_PARAMETER has no effect on an existing task.

Syntax

```
DBMS_ADVISOR.SET_DEFAULT_TASK_PARAMETER (
   advisor_name         IN VARCHAR2
   parameter            IN VARCHAR2,
   value                IN VARCHAR2);

DBMS_ADVISOR.SET_DEFAULT_TASK_PARAMETER (
   advisor_name         IN VARCHAR2
   parameter            IN VARCHAR2,
   value                IN NUMBER);
```

Parameters

Table 17-32 SET_DEFAULT_TASK_PARAMETER Procedure Parameters

Parameter	Description
advisor_name	Specifies the unique advisor name as defined in the view DBA_ADVISOR_DEFINITIONS.
parameter	The name of the task parameter to be modified. Parameter names are not case sensitive. Parameter names are unique to the task type, but not necessarily unique to all task types. Various task types may use the same parameter name for different purposes.
value	The value of the specified task parameter. The value can be specified as a string or a number.

Examples

```
BEGIN
  DBMS_ADVISOR.SET_DEFAULT_TASK_PARAMETER(DBMS_ADVISOR.SQLACCESS_ADVISOR,
   'VALID_TABLE_LIST', 'SH.%');
END;
/
```

17.3.32 SET_SQLWKLD_PARAMETER Procedure

This procedure modifies a user parameter within a SQL Workload object or SQL Workload object template.

A user parameter is a simple variable that stores various attributes that affect workload collection, tuning decisions and reporting.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.SET_SQLWKLD_PARAMETER (
   workload_name        IN VARCHAR2,
   parameter            IN VARCHAR2,
   value                IN VARCHAR2);

DBMS_ADVISOR.SET_SQLWKLD_PARAMETER (
   workload_name        IN VARCHAR2,
   parameter            IN VARCHAR2,
   value                IN NUMBER);
```

Parameters

Table 17-33 SET_SQLWKLD_PARAMETER Procedure Parameters

Parameter	Description
workload_name	The SQL Workload object name that uniquely identifies an existing workload.
parameter	The name of the data parameter to be modified. Parameter names are not case sensitive.
value	The value of the specified parameter. The value can be specified as a string or a number. If the value is DBMS_ADVISOR.DEFAULT, the value will be reset to the default value.

Usage Notes

A parameter will only affect operations that modify the workload collection. Therefore, parameters should be set prior to importing or adding new SQL statements to a workload. If a parameter is set after data has been placed in a workload object, it will have no effect on the existing data.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.SET_SQLWKLD_PARAMETER(workload_name,
'VALID_TABLE_LIST','SH.%');
END;
/
```

17.3.33 SET_TASK_PARAMETER Procedure

This procedure modifies a user parameter within an Advisor task or a template. A user parameter is a simple variable that stores various attributes that affect workload collection, tuning decisions and reporting.

Syntax

```
DBMS_ADVISOR.SET_TASK_PARAMETER (
   task_name           IN VARCHAR2
   parameter           IN VARCHAR2,
   value               IN VARCHAR2);

DBMS_ADVISOR.SET_TASK_PARAMETER (
   task_name           IN VARCHAR2
   parameter           IN VARCHAR2,
   value               IN NUMBER);
```

Parameters

Table 17-34 SET_TASK_PARAMETER Procedure Parameters

Parameter	Description
task_name	The Advisor task name that uniquely identifies an existing task.
parameter	The name of the task parameter to be modified. Parameter names are not case sensitive. Parameter names are unique to the task type, but not necessarily unique to all task types. Various task types may use the same parameter name for different purposes.
value	The value of the specified task parameter. The value can be specified as a string or a number. If the value is DEFAULT, the value will be reset to the default value.

Usage Notes

A task cannot be modified unless it is in its initial state. See RESET_TASK Procedure to set a task to its initial state. See your Advisor-specific documentation for further information on using this procedure.

SQL Access Advisor Task Parameters

Table 17-35 lists SQL Access Advisor task parameters.

Table 17-35 SQL Access Advisor Task Parameters

Parameter	Description
ANALYSIS_SCOPE	A comma-separated list that specifies the tuning artifacts to consider during analysis. The possible values are: • ALL Short name for specifying INDEX, MVIEW, TABLE, and PARTITION. • EVALUATION Causes a read-only evaluation of the specified workload. No new recommendations will be made. Can only be specified alone. • INDEX Allows the SQL Access Advisor to recommend index structure changes. • MVIEW Allows the SQL Access Advisor to recommend materialized view and log changes. • PARTITION Allows the SQL Access Advisor to recommend partition options. Use this in conjunction with the INDEX, MVIEW, and TABLE options. • TABLE Allows the SQL Access Advisor to make base-table recommendations. In this release, the only base-table recommendation is partitioning. Using the new keywords, the following combinations are valid: • INDEX • MVIEW • INDEX, PARTITION • INDEX, MVIEW, PARTITION • INDEX, TABLE, PARTITION • MVIEW, PARTITION • MIVEW, TABLE, PARTITION • INDEX, MVIEW, TABLE, PARTITION • TABLE, PARTITION • EVALUATION The default value is INDEX. The data type is STRINGLIST.
CREATION_COST	When set to true (default), the SQL Access Advisor will weigh the cost of creation of the access structure (index or materialized view) against the frequency of the query and potential improvement in the query execution time. When set to false, the cost of creation is ignored. The data type is STRING.
DAYS_TO_EXPIRE	Specifies the expiration time in days for the current SQL Access Advisor task. The value is relative to the last modification date. Once the task expires, it will become a candidate for removal by an automatic purge operation. Specifies the expiration time in days for the current Access Advisor task. The value is relative to the last modification date. The data type is NUMBER. Once the task expires, it becomes a candidate for removal by an automatic purge operation. The possible values are: • an integer in the range of 0 to 2147483647 • ADVISOR_UNLIMITED • ADVISOR_UNUSED The default value is 30.
DEF_EM_TEMPLATE	Contains the default task or template name from which the Enterprise Manager SQL Access Advisor Wizard reads its initial values. The default value is SQLACCESS_EMTASK. The data type is STRING.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
DEF_INDEX_OWNER	Specifies the default owner for new index recommendations. When a script is created, this value will be used to qualify the index name. Possible values are: • Existing schema name. Quoted identifiers are supported. • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRING.
DEF_INDEX_TABLESPACE	Specifies the default tablespace for new index recommendations. When a script is created, this value will be used to specify a tablespace clause. Possible values are: • Existing tablespace name. Quoted identifiers are supported. • ADVISOR_UNUSED No tablespace clause will be present in the script for indexes. The default value is ADVISOR_UNUSED. The data type is STRING.
DEF_MVIEW_OWNER	Specifies the default owner for new materialized view recommendations. When a script is created, this value will be used to qualify the materialized view name. Possible values are: • Existing schema name. Quoted identifiers are supported. • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRING.
DEF_MVIEW_TABLESPACE	Specifies the default tablespace for new materialized view recommendations. When a script is created, this value will be used to specify a tablespace clause. Possible values are • Existing tablespace name. Quoted identifiers are supported. • ADVISOR_UNUSED. No tablespace clause will be present in the script for materialized view logs. The default value is ADVISOR_UNUSED. The data type is STRING.
DEF_MVLOG_TABLSPACE	Specifies the default tablespace for new materialized view log recommendations. When a script is created, this value will be used to specify a tablespace clause. Possible values are: • Existing tablespace name. Quoted identifiers are supported. • ADVISOR_UNUSED. No tablespace clause will be present in the script for materialized view logs. The default value is ADVISOR_UNUSED. The data type is STRING.
DEF_PARTITION_TABLESPACE	Specifies the default tablespace for new partitioning recommendations. When a script is created, this value will be used to specify a tablespace clause. Possible values are: • Existing tablespace name. Quoted identifiers are supported. • ADVISOR_UNUSED. No tablespace clause will be present in the script for materialized views. The default value is ADVISOR_UNUSED. The data type is STRING.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
DML_VOLATILITY	When set to TRUE, the SQL Access Advisor will consider the impact of index maintenance and materialized view refresh in determining the recommendations. It will limit the access structure recommendations involving columns or tables that are frequently updated. For example, if there are too many DML statements on a column, then it may favor a B-tree index over a bitmap index on that column. For this process to be effective, the workload must include DML (insert/update/delete/merge/direct path inserts) statements that represent the update behavior of the application. The data type is STRING. See the related parameter refresh_mode.
END_TIME	Specifies an end time for selecting SQL statements. If the statement did not execute on or before the specified time, it will not be processed. Each date must be in the standard Oracle form of MM-DD-YYYY HH24:MI:SS, where: • DD is the numeric date • MM is the numeric month • YYYY is the numeric year • HH is the hour in 24 hour format • MI is the minute • SS is the second The data type is STRING.
EVALUATION_ONLY	This parameter is maintained for backward compatibility. All values will be translated and placed into the ANALYSIS_SCOPE task parameter. If set to TRUE, causes SQL Access Advisor to analyze the workload, but only comment on how well the current configuration is supporting it. No tuning recommendations will be generated. Possible values are: • FALSE • TRUE The default value is FALSE. The data type is STRING.
EXECUTION_TYPE	This parameter is maintained for backward compatibility. All values will be translated and placed into the ANALYSIS_SCOPE task parameter. The translated values are: • FULL => FULL • INDEX_ONLY => INDEX • MVIEW_ONLY => MVIEW • MVIEW_LOG_ONLY => MVIEW_LOG_ONLY The type of recommendations that is desired. Possible values: • FULL All supported recommendation types will be considered. • INDEX_ONLY The SQL Access Advisor will only consider index solutions as recommendations. • MVIEW_ONLY The SQL Access Advisor will consider materialized view and materialized view log solutions as recommendations. • MVIEW_LOG_ONLY The SQL Access Advisor will only consider materialized view log solutions as recommendations. The default value is FULL. The data type is STRINGLIST.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
IMPLEMENT_EXIT_ON_ERROR	When performing an IMPLEMENT_TASK operation, this parameter will control behavior when an action fails to implement. If set to TRUE, IMPLEMENT_TASK will stop on the first unexpected error. The possible values are: • TRUE • FALSE The default value is TRUE. The data type is STRING.
INDEX_NAME_TEMPLATE	Specifies the method by which new index names are formed. If the *TASK_ID* is omitted from the template, names generated by two concurrently executing SQL Access Advisor tasks may conflict and cause undesirable effects. So it is recommended that you include the *TASK_ID* in the template. Once formatted, the maximum size of a name is 30 characters. Valid keywords are: • Any literal value up to 22 characters. • TABLE Causes the parent table name to be substituted into the index name. If the name is too long, it will be trimmed to fit. • TASK_ID Causes the current task identifier number to be inserted in hexadecimal form. • SEQ Causes a sequence number to be inserted in hexadecimal form. Because this number is used to guarantee uniqueness, it is a required token. The default template is *table*_IDX$$_*task_idsequence*. The data type is STRING.
INVALID_ACTION_LIST	Contains a fully qualified list of actions that are not eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. An action can be any string. If an action is not quoted, it will be changed to uppercase lettering and stripped of leading and trailing spaces. An action string is not scanned for correctness. During a task execution, if a SQL statement's action matches a name in the action list, it will not be processed by the task. An action name is case sensitive. The possible values are: • single action • comma-delimited action list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.
INVALID_MODULE_LIST	Contains a fully qualified list of modules that are not eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. A module can be any string. If a module is not quoted, it will be changed to uppercase lettering and stripped of leading and trailing spaces. A module string is not scanned for correctness. During a task execution, if a SQL statement's module matches a name in the list, it will not be processed by the task. A module name is case sensitive. The possible values are: • single application • comma-delimited module list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
INVALID_SQLSTRING_LIST	Contains a fully qualified list of text strings that are not eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted values are supported. A SQL string can be any string. If a string is not quoted, it will be changed to uppercase lettering and stripped of leading and trailing spaces. A SQL string is not scanned for correctness. During a task execution, if a SQL statement contains a string in the SQL string list, it will not be processed by the task. The possible values are: • single string • comma-delimited string list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.
INVALID_USERNAME_LIST	Contains a fully qualified list of user names that are not eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. During a task execution, if a SQL statement's user name matches a name in the user name list, it will not be processed by the task. A user name is not case sensitive unless it is quoted. The possible values are: • single user name • comma-delimited user name list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.
JOURNALING	Controls the logging of messages to the journal (DBA_ADVISOR_JOURNAL and USER_ADVISOR_JOURNAL views). The higher the setting, the more information is logged to the journal. Possible values are: • UNUSED: no journal messages • FATAL: explanation of fatal conditions • ERROR: explanation of errors • WARNING: explanation of warnings • INFORMATION: information message • INFORMATION2: common information • INFORMATION3: common information • INFORMATION4: common information • INFORMATION5: common information • INFORMATION6: common information Each journal value represents all recorded messages at that level or lower. For example, when choosing WARNING, all messages marked WARNING as well as ERROR and FATAL will be recorded in the repository. INFORMATION6 represents the most thorough message recording and UNUSED is the least. The default value is INFORMATION. The data type is NUMBER.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
LIMITED_PARTITION_SCHEMES	User can suggest that the Partition Expert cut off the number of partitioning schemes to investigate. This can help with cutting down the run time of the advisor. Possible values are: • An integer in the range of 1 to 10 • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is NUMBER.
MAX_NUMBER_PARTITIONS	Limits the number of partitions the advisor will recommend for any base table, index, or materialized view. Possible values are: • An integer in the range of 1 to 4294967295 • ADVISOR_UNLIMITED • ADVISOR_UNUSED The default value is ADVISOR_UNLIMITED. The data type is NUMBER.
MODE	Specifies the mode by which Access Advisor will operate during an analysis. Valid values are: • LIMITED Indicates the Advisor will attempt to a quick job by limiting the search-space of candidate recommendations, and correspondingly, the results may be of a low quality. • COMPREHENSIVE Indicates the Advisor will search a large pool of candidates that may take long to run, but the resulting recommendations will be of the highest quality. The default value is COMPREHENSIVE. The data type is STRING.
MVIEW_NAME_TEMPLATE	Specifies the method by which new materialized view names are formed. If the *TASK_ID* is omitted from the template, names generated by two concurrently executing SQL Access Advisor tasks may conflict and cause undesirable effects. So it is recommended that you include the *TASK_ID* in the template. The format is any combination of keyword tokens and literals. However, once formatted, the maximum size of a name is 30 characters. Valid tokens are: • Any literal value up to 22 characters. • *TASK_ID* Causes the current task identifier number to be inserted in hexadecimal form. • *SEQ* Causes a sequence number to be inserted in hexadecimal form. Because this number is used to guarantee uniqueness, it is a required token. The default template is: MV$$_*task_idsequence*. The data type is STRING.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
ORDER_LIST	This parameter has been deprecated. Contains the primary natural order in which the Access Advisor processes workload elements during the analysis operation. To determine absolute natural order, Access Advisor sorts the workload using ORDER_LIST values. A comma must separate multiple order keys. Possible values are: • BUFFER_GETS Sets the order using the SQL statement's buffer-get count value. • CPU_TIME Sets the order using the SQL statement's CPU time value. • DISK_READS Sets the order using the SQL statement's disk-read count value. • ELAPSED_TIME Sets the order using the SQL statement's elapsed time value. • EXECUTIONS Sets the order using the SQL statement's execution frequency value. • OPTIMIZER_COST Sets the order using the SQL statement's optimizer cost value. • I/O Sets the order using the SQL statement's I/O count value. • PRIORITY Sets the order using the user-supplied business priority value. All values are accessed in descending order, where a high value is considered more interesting than a low value. The default value is PRIORITY, OPTIMIZER_COST. The data type is STRINGLIST.
PARTITION_NAME_TEMPLATE	Specifies the method by which new partition names are formed. The format is any combination of keyword tokens and literals. However, once formatted, the maximum size of a name is 30 characters. Valid tokens are: • Any literal value up to 22 characters. • *table* - Causes the parent table name to be substituted into the partition name. If the name is too long, it will be trimmed to fit. • *task_id* - Causes the current task identifier number to be inserted in hexadecimal form. • *sequence* - Causes a sequence number to be inserted in hexadecimal form. Because this number is used to guarantee uniqueness, it is a required token. The default template is PTN$$_*table_task_idsequence*. The data type is STRING.
PARTITIONING_GOAL	Specifies the approach used to make partitioning recommendations. One possible value is PERFORMANCE, which is the default. The data type is STRING.
PARTITIONING_TYPES	Specifies the type of partitioning used. Possible values are RANGE and HASH. The data type is STRING.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
RANKING_MEASURE	Contains the primary natural order in which the SQL Access Advisor processes workload elements during the analysis operation. To determine absolute natural order, SQL Access Advisor sorts the workload using RANKING_MEASURE values. A comma must separate multiple order keys. Possible values are: • BUFFER_GETS Sets the order using the SQL statement's buffer-get count value. • CPU_TIME Sets the order using the SQL statement's CPU time value. • DISK_READS Sets the order using the SQL statement's disk-read count value. • ELAPSED_TIME Sets the order using the SQL statement's elapsed time value. • EXECUTIONS Sets the order using the SQL statement's elapsed time value. • OPTIMIZER_COST Sets the order using the SQL statement's optimizer cost value. • PRIORITY Sets the order using the user-supplied business priority value. All values are accessed in descending order, where a high value is considered more interesting than a low value. The default value is PRIORITY, OPTIMIZER_COST. The data type is STRINGLIST.
RECOMMEND_MV_EXACT_TEXT_MATCH	When considering candidate materialized views, exact text match solutions will only be included if this parameter contains TRUE. The possible values are: • TRUE • FALSE The default value is TRUE. The data type is STRING.
RECOMMENDED_TABLESPACES	Allows the SQL Access Advisor to recommend optimal tablespaces for any partitioning scheme. If this is not set, the SQL Access Advisor will simply recommend a partitioning method but give no advice on physical storage. Possible values are: • TRUE • FALSE (the default) The data type is STRING.
REFRESH_MODE	Specifies whether materialized views are refreshed ON_DEMAND or ON_COMMIT. This will be used to weigh the impact of materialized view refresh when the parameter dml_volatility is set to TRUE. Possible values are: • ON_DEMAND • ON_COMMIT The default value is ON_DEMAND. The data type is STRING.
REPORT_DATE_FORMAT	This is the default date and time formatting template. The default format is DD/MM/YYYYHH24:MI. The data type is STRING.
SHOW_RETAINS	Controls the display of RETAIN actions within an implementation script and the SQL Access Advisor wizard. The possible values are: • TRUE • FALSE The default value is TRUE. The data type is STRING.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
SQL_LIMIT	Specifies the number of SQL statements to be analyzed. The SQL_LIMIT filter is applied after all other filters have been applied. For example, if only statements referencing the table hr.employees are to be accepted, the SQL_LIMIT value will be only apply to those statements. When used in conjunction with the parameter ORDER_LIST, SQL Access Advisor will process the most interesting SQL statements by ordering the statements according to the specified sort keys. The possible values are: • An integer in the range of 1 to 2147483647 • ADVISOR_UNLIMITED • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is NUMBER.
START_TIME	Specifies a start time for selecting SQL statements. If the statement did not execute on or before the specified time, it will not be processed. Each date must be in the standard Oracle form of MM-DD-YYYY HH24:MI:SS, where: • DD is the numeric date • MM is the numeric month • YYYY is the numeric year • HH is the hour in 24 hour format • MI is the minute • SS is the second The data type is STRING.
STORAGE_CHANGE	Contains the amount of space adjustment that can be consumed by SQL Access Advisor recommendations. Zero or negative values are only permitted if the workload scope is marked as FULL. When the SQL Access Advisor produces a set of recommendations, the resultant physical structures must be able to fit into the budgeted space. A space budget is computed by adding the STORAGE_CHANGE value to the space quantity currently used by existing access structures. A negative STORAGE_CHANGE value may force SQL Access Advisor to remove existing structures in order to shrink space demand. Possible values: • Any valid integer including negative values, zero and positive values. The default value is ADVISOR_UNLIMITED. The data type is NUMBER.
TIME_LIMIT	Specifies the time in minutes that the SQL Access Advisor can use to perform an analysis operation. If the SQL Access Advisor reaches a specified recommendation quality or all input data has been analyzed, processing will terminate regardless of any remaining time. Possible values: • An integer in the range of 1 to 10,000 • ADVISOR_UNLIMITED The default value is 720 (12 hours). The data type is NUMBER. Note that specifying ADVISOR_UNLIMITED has the same effect as setting the parameter to the maximum of 10,000 (about one week). The SQL Access Advisor will never run for more than 10,000 minutes.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
VALID_ACTION_LIST	Contains a fully qualified list of actions that are eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. An action can be any string. If an action is not quoted, it will be changed to uppercase lettering and stripped of leading and trailing spaces. An action string is not scanned for correctness. During a task execution, if a SQL statement's action does not match a name in the action list, it will not be processed by the task. An action name is case sensitive. The possible values are: • single action • comma-delimited action list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.
VALID_MODULE_LIST	Contains a fully qualified list of application modules that are eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. A module can be any string. If a module is not quoted, it will be changed to uppercase lettering and stripped of leading and trailing spaces. A module string is not scanned for correctness. During a task execution, if a SQL statement's module does not match a name in the module list, it will not be processed by the task. A module name is case sensitive. The possible values are: • single application • comma-delimited module list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.
VALID_SQLSTRING_LIST	Contains a fully qualified list of text strings that are eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. A SQL string can be any string. If a string is not quoted, it will be changed to uppercase lettering and stripped of leading and trailing spaces. A SQL string is not scanned for correctness. During a task execution, if a SQL statement does not contain string in the SQL string list, it will not be processed by the task. The possible values are: • single string • comma-delimited string list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.

Table 17-35 (Cont.) SQL Access Advisor Task Parameters

Parameter	Description
VALID_TABLE_LIST	Contains a fully qualified list of tables that are eligible for tuning. The list elements are comma-delimited, and quoted identifiers are supported. Wildcard specifications are supported for tables. The default value is all tables within the user's scope are eligible for tuning. Supported wildcard character is %. A % wildcard matches any set of consecutive characters. When a SQL statement is processed, it will not be accepted unless at least one referenced table is specified in the valid table list. If the list is unused, then all table references within a SQL statement are considered valid. The valid syntax for a table reference is: • schema.table • schema • schema.% (equivalent to schema) • comma-delimited action list • ADVISOR_UNUSED The possible values are: • single table reference • comma-delimited reference list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is TABLELIST.
VALID_USERNAME_LIST	Contains a fully qualified list of user names that are eligible for processing in a SQL workload object. The list elements are comma-delimited, and quoted names are supported. During a task execution, if a SQL statement's user name does not match a name in the user name list, it will not be processed by the task. A user name is not case sensitive unless it is quoted. The possible values are: • single user name • comma-delimited user name list • ADVISOR_UNUSED The default value is ADVISOR_UNUSED. The data type is STRINGLIST.
WORKLOAD_SCOPE	Describes the level of application coverage the workload represents. Possible values are FULL and PARTIAL. FULL Should be used if the workload contains all interesting application SQL statements for the targeted tables. PARTIAL (default) Should be used if the workload contains anything less than a full representation of the interesting application SQL statements for the targeted tables. The data type is STRING.

Segment Advisor Parameters

Table 17-36 lists the input task parameters that can be set in the Segment Advisor using the SET_TASK_PARAMETER procedure.

Table 17-36 Segment Advisor Task Parameters

Parameter	Description
MODE	The data to use for analysis. The default value is COMPREHENSIVE, and the possible values are: • LIMITED: Analysis restricted to statistics available in the Automatic Workload Repository • COMPREHENSIVE: Analysis based on sampling and Automatic Workload Repository statistics
TIME_LIST	The time limit for which the Advisor should run. It is specified in seconds, and the default and possible values are UNLIMITED.
RECOMMEND_ALL	Whether to generate recommendations for all segments. The default value is TRUE. If set to TRUE, it generates recommendations all segments specified by the user. If set to FALSE, it generates recommendations for only those objects that are eligible for shrink.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';

   DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
   DBMS_ADVISOR.SET_TASK_PARAMETER(task_name, 'VALID_TABLELIST',
      'SH.%,SCOTT.EMP');
END;
/
```

Undo Advisor Task Parameters

Table 17-37 lists the input task parameters that can be set in the Undo Advisor using the SET_TASK_PARAMETER procedure.

Table 17-37 Undo Advisor Task Parameters

Parameter	Description
TARGET_OBJECTS	The undo tablespace of the system. There is no default value, and the possible value is UNDO_TBS.
START_SNAPSHOT	The starting time for the system to perform analysis using the snapshot numbers in the AWR repository. There is no default value and the possible values are the valid snapshot numbers in the AWR repository.
END_SNAPSHOT	The ending time for the system to perform analysis using the snapshot numbers in the AWR repository. There is no default value and the possible values are the valid snapshot numbers in the AWR repository.
BEGIN_TIME_SEC	The number of seconds between the beginning time of the period and now. Describes a period of time for the system to perform analysis. BEGIN_TIME_SEC should be greater than END_TIME_SEC. There is no default value and the possible values are any positive integer.

Table 17-37 (Cont.) Undo Advisor Task Parameters

Parameter	Description
END_TIME_SEC	The number of seconds between the ending time of the period and now. END_TIME_SEC should be less than BEGIN_TIME_SEC. There is no default value and the possible values are any positive integer.

Examples

```
DECLARE
   tname  VARCHAR2(30);
   oid    NUMBER;
   BEGIN
     DBMS_ADVISOR.CREATE_TASK('Undo Advisor', tid, tname, 'Undo Advisor
Task');
     DBMS_ADVISOR.CREATE_OBJECT(tname, 'UNDO_TBS', null, null, null, 'null',
oid);
     DBMS_ADVISOR.SET_TASK_PARAMETER(tname, 'TARGET_OBJECTS', oid);
     DBMS_ADVISOR.SET_TASK_PARAMETER(tname, 'START_SNAPSHOT', 1);
     DBMS_ADVISOR.SET_TASK_PARAMETER(tname, 'END_SNAPSHOT', 2);
     DBMS_ADVISOR.SET_TASK_PARAMETER(tname, 'INSTANCE', 1);
     DBMS_ADVISOR.EXECUTE_TASK(tname);
   END;
/
```

Automatic Database Diagnostic Monitor (ADDM) Task Parameters

Table 17-38 lists the input task parameters that can be set in ADDM using the SET_TASK_PARAMETER procedure.

Table 17-38 ADDM Task Parameters

Parameter	Description
START_SNAPSHOT	The starting time for the system to perform analysis using the snapshot numbers in the AWR repository. There is no default value, and the possible values are the valid snapshot numbers in the AWR repository.
END_SNAPSHOT	The ending time for the system to perform analysis using the snapshot numbers in the AWR repository. There is no default value, and the possible values are the valid snapshot numbers in the AWR repository.
DB_ID	The database for START_SNAPSHOT and END_SNAPSHOT. The default value is the current database ID.
INSTANCE	The instance for START_SNAPSHOT and END_SNAPSHOT. The default value is 0 or UNUSED, and the possible values are all positive integers. By default, all instances are analyzed.
INSTANCES	If the INSTANCE parameter has been set, INSTANCES is ignored. The default value is UNUSED, and the possible values are comma-separated list of instance numbers (for example, "1, 3, 5"). By default, all instances are analyzed.
DBIO_EXPECTED	The average time to read the database block in microseconds. The default value is 10 milliseconds, and the possible values are system-dependent.

Examples

The following creates and executes an ADDM task for the current database and an AWR snapshot range between 19 and 26. Note that this example will analyze all instances, whether you have only one or an Oracle RAC database.

```
DECLARE
   tid     NUMBER;
   tname VARCHAR2(30) := 'ADDM_TEST';
BEGIN
   DBMS_ADVISOR.CREATE_TASK('ADDM', tid, tname, 'my test');
   DBMS_ADVISOR.SET_TASK_PARAMETER(tname, 'START_SNAPSHOT', '19');
   DBMS_ADVISOR.SET_TASK_PARAMETER(tname, 'END_SNAPSHOT', '26');
   DBMS_ADVISOR.EXECUTE_TASK(tname);
END;
/
```

See Also:

- *Oracle Database Performance Tuning Guide* to learn more about using ADDM
- The DBMS_ADDM package for details on how to create and execute ADDM tasks

SQL Tuning Advisor Task Parameters

See the DBMS_SQLTUNE package and *Oracle Database SQL Tuning Guide* for more information.

17.3.34 TUNE_MVIEW Procedure

This procedure shows how to decompose a materialized view into multiple views and to restate the materialized view to be optimized for fast refresh and query rewrite. It also shows how to fix materialized view logs and to enable query rewrite.

Syntax

```
DBMS_ADVISOR.TUNE_MVIEW (
   task_name         IN OUT VARCHAR2,
   mv_create_stmt    IN     [CLOB | VARCHAR2]);
```

Parameters

Table 17-39 TUNE_MVIEW Procedure Parameters

Parameter	Description
task_name	The task name for querying the results in a catalog view. If not specified, the database generates a task name, and then returns.
mv_create_stmt	The original materialized view creation statement.

Usage Notes

Executing TUNE_MVIEW generates two sets of output results: one for the implementation, and the other for undoing the implementation. The output is accessible through USER_TUNE_MVIEW and DBA_TUNE_MVIEW views. You can also use DBMS_ADVISOR.GET_TASK_SCRIPT and DBMS_ADVISOR.CREATE_FILE to print the TUNE_MVIEW results into a script file for later execution.

Table 17-40 USER_TUNE_MVIEW and DBA_TUNE_MVIEW Views

Column Name	Column Description
OWNER	The name of the materialized view owner.
TASK_NAME	The name of the task. This name serves as a key to access the set of recommendations.
SCRIPT_TYPE	Recommendation ID that indicates whether the row is for the IMPLEMENTATION or UNDO script.
ACTION_ID	Action ID used as the command order number.
STATEMENT	For TUNE_MVIEW output, this column represents the following statements, and includes statement properties such as REFRESH and REWRITE options: • CREATE MATERIALIZED VIEW LOG • ALTER MATERIALIZED VIEW LOG FORCE • [CREATE | DROP] MATERIALIZED VIEW

Examples

The following example shows how to use TUNE_MVIEW to optimize a CREATE MATERIALIZED VIEW statement:

```
DECLARE
  v_tname VARCHAR2(30);
BEGIN
  v_tname := 'mview_task';
  DBMS_ADVISOR.TUNE_MVIEW(
      task_name      => v_tname
  ,   mv_create_stmt =>
       'CREATE MATERIALIZED VIEW omv REFRESH WITH ROWID AS SELECT * FROM
orders');
END;
```

You can view the results by querying USER_TUNE_MVIEW or DBA_TUNE_MVIEW as the following example (sample output included):

```
SET LINESIZE 120
COL TASK_NAME FORMAT a20
COL STATEMENT FORMAT a40

SELECT *
FROM   USER_TUNE_MVIEW
WHERE  TASK_NAME='mview_task'
AND    SCRIPT_TYPE='IMPLEMENTATION';

TASK_NAME             ACTION_ID SCRIPT_TYPE    STATEMENT
-------------------- ---------- --------------
```

```
----------------------------------------
mview_task                     1 IMPLEMENTATION CREATE MATERIALIZED VIEW LOG
ON "OE"."OR
                                                DERS" WITH ROWID

mview_task                     2 IMPLEMENTATION ALTER MATERIALIZED VIEW LOG
FORCE ON "OE
                                                "."ORDERS" ADD ROWID

mview_task                     3 IMPLEMENTATION CREATE MATERIALIZED VIEW
OE.OMV REFRESH
                                                FAST WITH ROWID DISABLE QUERY
REWRITE
```

Alternatively, you can save the output results in an external script file as in the following example:

```
CREATE DIRECTORY TUNE_RESULTS_DIR AS  '/tmp';
GRANT READ, WRITE ON DIRECTORY TUNE_RESULTS_DIR TO PUBLIC;
BEGIN
  DBMS_ADVISOR.CREATE_FILE(
      buffer     => DBMS_ADVISOR.GET_TASK_SCRIPT( task_name => 'mview_task')
  ,   location   => 'TUNE_RESULTS_DIR'
  ,   filename   => 'mview_create.sql' );
END;
```

The preceding statement will save the results in `/tmp/mview_create.sql`.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide*for more information about using the `TUNE_MVIEW` procedure

17.3.35 UPDATE_OBJECT Procedure

This procedure updates an existing task object.

Task objects are typically used as input data for a particular advisor. Segment advice can be generated at the object, segment, or tablespace level.

Syntax

```
DBMS_ADVISOR.UPDATE_OBJECT (
   task_name       IN VARCHAR2
   object_id       IN NUMBER,
   attr1           IN VARCHAR2 := NULL,
   attr2           IN VARCHAR2 := NULL,
   attr3           IN VARCHAR2 := NULL,
   attr4           IN CLOB := NULL,
   attr5           IN VARCHAR2 := NULL);
```

Parameters

Table 17-41 UPDATE_OBJECT Procedure Parameters

Parameter	Description
task_name	A valid advisor task name that uniquely identifies an existing task.
object_id	The advisor-assigned object identifier.
attr1	Advisor-specific data. If set to NULL, there will be no effect on the target object.
attr2	Advisor-specific data. If set to NULL, there will be no effect on the target object.
attr3	Advisor-specific data. If set to NULL, there will be no effect on the target object.
attr4	Advisor-specific data. If set to NULL, there will be no effect on the target object.
attr5	Advisor-specific data. If set to null, there will be no effect on the target object.

The attribute parameters have different values depending upon the object type. See *Oracle Database Administrator's Guide* for details regarding these parameters and object types.

Usage Notes

If for the object level, advice is generated on all partitions of the object (if the object is partitioned). The advice is not cascaded to any dependent objects. If for the segment level, advice can be obtained on a single segment, such as the partition or subpartition of a table, index, or lob column. If for a tablespace level, target advice for every segment in the tablespace will be generated.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  obj_id NUMBER;
BEGIN
  task_name := 'My Task';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
task_name);
  DBMS_ADVISOR.CREATE_OBJECT (task_name,'SQL',NULL,NULL,NULL,
                              'SELECT * FROM SH.SALES',obj_id);
  DBMS_ADVISOR.UPDATE_OBJECT (task_name, obj_id,NULL,NULL,NULL,
                              'SELECT count(*) FROM SH.SALES');
END;
/
```

> **See Also:**
>
> *Oracle Database Administrator's Guide* for further information regarding the Segment Advisor

17.3.36 UPDATE_REC_ATTRIBUTES Procedure

This procedure updates the owner, name, and tablespace for a recommendation.

Syntax

```
DBMS_ADVISOR.UPDATE_REC_ATTRIBUTES (
   task_name            IN VARCHAR2
   rec_id               IN NUMBER,
   action_id            IN NUMBER,
   attribute_name       IN VARCHAR2,
   value                IN VARCHAR2);
```

Parameters

Table 17-42 UPDATE_REC_ATTRIBUTES Procedure Parameters

Parameter	Description
`task_name`	The task name that uniquely identifies an existing task.
`rec_id`	The Advisor-generated identifier number that is assigned to the recommendation.
`action_id`	The Advisor-generated action identifier that is assigned to the particular command.
`attribute_name`	Name of the attribute to be changed. The valid values are: • `owner` The new owner of the object. • `name` The new name of the object. • `tablespace` The new tablespace for the object.
`value`	Specifies the new value for the recommendation attribute.

Usage Notes

Recommendation attributes cannot be modified unless the task has successfully executed.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
  workload_name VARCHAR2(30);
  attribute VARCHAR2(100);
BEGIN
  task_name := 'My Task';
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
```

```
task_name);
  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_REF(task_name, workload_name);
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales WHERE promo_id = 10');
  DBMS_ADVISOR.EXECUTE_TASK(task_name);

attribute := 'SH';

  DBMS_ADVISOR.UPDATE_REC_ATTRIBUTES(task_name, 1, 3, 'OWNER', attribute);
END;
/
```

17.3.37 UPDATE_SQLWKLD_ATTRIBUTES Procedure

This procedure changes various attributes of a SQL Workload object or template.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.UPDATE_SQLWKLD_ATTRIBUTES (
   workload_name        IN VARCHAR2,
   new_name             IN VARCHAR2 := NULL,
   description          IN VARCHAR2 := NULL,
   read_only            IN VARCHAR2 := NULL,
   is_template          IN VARCHAR2 := NULL,
   how_created          IN VARCHAR2 := NULL);
```

Parameters

Table 17-43 UPDATE_SQLWKLD_ATTRIBUTES Procedure Parameters

Parameter	Description
workload_name	The workload object name that uniquely identifies an existing workload.
new_name	The new workload object name. If the value is NULL or contains the value ADVISOR_UNUSED, the workload will not be renamed. A task name can be up to 30 characters long.
description	A new workload description. If the value is NULL or contains the value ADVISOR_UNUSED, the description will not be changed. Names can be up to 256 characters long.
read_only	Set to TRUE so it cannot be changed.
is_template	TRUE if workload is to be used as a template.

Table 17-43 (Cont.) UPDATE_SQLWKLD_ATTRIBUTES Procedure Parameters

Parameter	Description
how_created	Indicates a source application name that initiated the workload creation. If the value is NULL or contains the value ADVISOR_UNUSED, the source will not be changed.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales WHERE promo_id = 10');
  DBMS_ADVISOR.UPDATE_SQLWKLD_ATTRIBUTES(workload_name,'New workload name');
END;
/
```

17.3.38 UPDATE_SQLWKLD_STATEMENT Procedure

This procedure updates an existing SQL statement in a specified SQL workload.

> **Note:**
>
> This procedure is deprecated starting in Oracle Database 11*g*.

Syntax

```
DBMS_ADVISOR.UPDATE_SQLWKLD_STATEMENT (
   workload_name     IN VARCHAR2,
   sql_id            IN NUMBER,
   application       IN VARCHAR2 := NULL,
   action            IN VARCHAR2 := NULL,
   priority          IN NUMBER := NULL,
   username          IN VARCHAR2 := NULL);

DBMS_ADVISOR.UPDATE_SQLWKLD_STATEMENT (
   workload_name     IN VARCHAR2,
   search            IN VARCHAR2,
   updated           OUT NUMBER,
   application       IN VARCHAR2 := NULL,
   action            IN VARCHAR2 := NULL,
   priority          IN NUMBER := NULL,
   username          IN VARCHAR2 := NULL);
```

Parameters

Table 17-44 UPDATE_SQLWKLD_STATEMENT Procedure Parameters

Parameter	Description
workload_name	The SQL Workload object name that uniquely identifies an existing workload.
sql_id	The Advisor-generated identifier number that is assigned to the statement. To specify all workload statements, use the constant DBMS_ADVISOR.ADVISOR_ALL.
updated	Returns the number of statements changed by a searched update.
application	Specifies a business application name that will be associated with the SQL statement. If the value is NULL or contains the value ADVISOR_UNUSED, then the column will not be updated in the repository.
action	Specifies the application action for the statement. If the value is NULL or contains the value ADVISOR_UNUSED, then the column will not be updated in the repository.
priority	The relative priority of the SQL statement. The value must be one of the following: 1 - HIGH, 2 - MEDIUM, or 3 - LOW. If the value is NULL or contains the value ADVISOR_UNUSED, then the column will not be updated in the repository.
username	The Oracle user name that executed the SQL statement. If the value is NULL or contains the value ADVISOR_UNUSED, then the column will not be updated in the repository. Because a user name is an Oracle identifier, the username value must be entered exactly like it is stored in the database. For example, if the user SCOTT is the executing user, then you must provide the user identifier SCOTT in all uppercase letters. The database does not recognize the user scott as a match for SCOTT.
search	Disabled.

Usage Notes

A workload cannot be modified or deleted if it is currently referenced by an active task. A task is considered active if it is not in its initial state. See RESET_TASK Procedure to set a task to its initial state.

Examples

```
DECLARE
  workload_name VARCHAR2(30);
  updated NUMBER;
  id NUMBER;
BEGIN
  workload_name := 'My Workload';

  DBMS_ADVISOR.CREATE_SQLWKLD(workload_name, 'My Workload');
  DBMS_ADVISOR.ADD_SQLWKLD_STATEMENT(workload_name, 'MONTHLY', 'ROLLUP',
                                     100,400,5041,103,640445,680000,2,
                                     1,SYSDATE,1,'SH','SELECT AVG(amount_sold)
                                     FROM sh.sales WHERE promo_id = 10');

   SELECT sql_id INTO id FROM USER_ADVISOR_SQLW_STMTS
```

```
   WHERE workload_name = 'My Workload';

  DBMS_ADVISOR.UPDATE_SQLWKLD_STATEMENT(workload_name, id);
END;
/
```

17.3.39 UPDATE_TASK_ATTRIBUTES Procedure

This procedure changes various attributes of a task or a task template.

Syntax

```
DBMS_ADVISOR.UPDATE_TASK_ATTRIBUTES (
   task_name          IN VARCHAR2
   new_name           IN VARCHAR2 := NULL,
   description        IN VARCHAR2 := NULL,
   read_only          IN VARCHAR2 := NULL,
   is_template        IN VARCHAR2 := NULL,
   how_created        IN VARCHAR2 := NULL);
```

Parameters

Table 17-45 UPDATE_TASK_ATTRIBUTES Procedure Parameters

Parameter	Description
task_name	The Advisor task name that uniquely identifies an existing task.
new_name	The new Advisor task name. If the value is NULL or contains the value ADVISOR_UNUSED, the task will not be renamed. A task name can be up to 30 characters long.
description	A new task description. If the value is NULL or contains the value ADVISOR_UNUSED, the description will not be changed. Names can be up to 256 characters long.
read_only	Sets the task to read-only. Possible values are: TRUE and FALSE. If the value is NULL or contains the value ADVISOR_UNUSED, the setting will not be changed.
is_template	Marks the task as a template. Physically, there is no difference between a task and a template; however, a template cannot be executed. Possible values are: TRUE and FALSE. If the value is NULL or contains the value ADVISOR_UNUSED, the setting will not be changed.
how_created	Indicates a source application name that initiated the task creation. If the value is NULL or contains the value ADVISOR_UNUSED, the source will not be changed.

Examples

```
DECLARE
  task_id NUMBER;
  task_name VARCHAR2(30);
BEGIN
  task_name := 'My Task';

  DBMS_ADVISOR.CREATE_TASK(DBMS_ADVISOR.SQLACCESS_ADVISOR, task_id,
```

```
task_name);
  DBMS_ADVISOR.UPDATE_TASK_ATTRIBUTES(task_name,'New Task Name');
  DBMS_ADVISOR.UPDATE_TASK_ATTRIBUTES('New Task Name',NULL,'New description');
END;
/
```

18
DBMS_AUTOIM

This API provides functions to manage the execution of the AIM feature.

18.1 Using DBMS_AUTOIM

Automatic In-Memory Management (AIM) is an In-Memory column store feature that optimizes the column store through actions such evicting cold data, populating frequently accessed data, and compressing cold columns.

The DBMS_AUTOIM package provides APIs to manage the execution of the AIM feature. These are procedures to get and set AIM parameters as well as a reporting function.

18.2 Summary of DBMS_AUTOIM Subprograms

The table in this topic lists and describes the DBMS_AUTOIM subprograms.

Table 18-1 DBMS_AUTOIM Subprograms

Subprogram	Description
SET_PARAMETER	Sets the value of one of the available parameters: `AUTO_IM_PERF_TASK`, `INTERVAL`, or `AIM_STATWINDOW_DAYS`.
ACTIVITY_REPORT	Reports AIM task activity within the specified time interval for the given feature or for all features.

18.2.1 SET_PARAMETER Procedure

Use this DBMS_AUTOIM procedure to set configuration parameter values for AIM features.

Syntax

```
DBMS_AUTOIM.SET_PARAMETER( parameter_name IN VARCHAR2, parameter_value IN
VARCHAR2 );
```

Parameters

The parameters for DBMS_AUTOIM.SET_PARAMETER are as follows.

Table 18-2 SET_PARAMETER Procedure Parameters

Parameter	Description
`aim_task`	Turns the automatic IM automatic performance feature creation task on or off. Values: `enable`, `disable`.

Table 18-2 (Cont.) SET_PARAMETER Procedure Parameters

Parameter	Description
interval	Controls the statistics time window used by AIM to compute decisions. Value: Any positive integer representing the time window in seconds. The default is 900 seconds.

18.2.2 ACTIVITY_REPORT Function

Use this function to generate a report on the automatic creation of IM performance features. These features include Autonomous Join Groups, Autonomous Bloom Filter Optimization, Autonomous Optimized Arithmetic, and Autonomous Vector Optimization.

Syntax

```
DBMS_AUTOIM.ACTIVITY_REPORT( feature_id IN NUMBER, start_time IN TIMESTAMP,
end_time IN TIMESTAMP, type IN VARCHAR2, level IN VARCHAR2 )
```

Table 18-3 ACTIVITY_REPORT Function Parameters

Parameter	Type	Description
feature_id	integer	A NULL feature_id generates report on all AIM features. The default is NULL. You can get the feature_id for each feature from the view V$AIM_FEATURES .
start_time	timestamp	Start of the time window in which AIM activities are considered for the report. A NULL start_time returns the report for last execution. Default is NULL.
end_time	timestamp	End of time window in which AIM activities are considered for the report. A NULL end_time returns the report for last execution. Default is NULL.
type	varchar2	Type of the report. Possible values are: 'TEXT', 'HTML', 'XML'. Default is 'TEXT'.
level	varchar2	Level of verbosity. Possible values are 'BASIC' and 'DETAILED'. Default is 'BASIC'.

Example 18-1 ACTIVY_REPORT function

Create an AIM activity report for the previous day in TEXT format. The return type is CLOB.

```
declare
   report clob := null;
   begin
       report := DBMS_AUTOIM.activity_report();
   end;
```

> **Note:**
>
> See the view V$AIM_FEATURES to obtain the `feature_id` for each feature.

19
DBMS_ALERT

DBMS_ALERT supports asynchronous notification of database events (alerts). By appropriate use of this package and database triggers, an application can notify itself whenever values of interest in the database are changed.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Restrictions
- Exceptions
- Operational Notes
- Examples
- Summary of DBMS_ALERT Subprograms

19.1 DBMS_ALERT Overview

This scenario describes a possible use of the DBMS_ALERT package.

Suppose a graphics tool is displaying a graph of some data from a database table. The graphics tool can, after reading and graphing the data, wait on a database alert (WAITONE) covering the data just read. The tool automatically wakes up when the data is changed by any other user. All that is required is that a trigger be placed on the database table, which performs a signal (SIGNAL) whenever the trigger is fired.

19.2 DBMS_ALERT Security Model

Security on this package can be controlled by granting EXECUTE on this package to selected users or roles. You might want to write a cover package on top of this one that restricts the alert names used. EXECUTE privilege on this cover package can then be granted rather than on this package.

19.3 DBMS_ALERT Constants

The DBMS_ALERT package uses the constants listed and described in tis topic.

Table 19-1 DBMS_ALERT Constants

Name	Type	Value	Description
MAXWAIT	INTEGER	86400000	The maximum time to wait for an alert (1000 days which is essentially forever).

19.4 DBMS_ALERT Restrictions

Because database alerters issue commits, they cannot be used with Oracle Forms. For more information on restrictions on calling stored procedures while Oracle Forms is active, refer to your Oracle Forms documentation.

19.5 DBMS_ALERT Exceptions

`DBMS_ALERT` raises the application error -20000 on error conditions.

Table 19-2 shows the messages and the procedures that can raise them.

19.6 DBMS_ALERT Operational Notes

This topic lists notes related to general and specific applications. Also, a list of DBMS_ALERT error messages is provided.

- Alerts are transaction-based. This means that the waiting session is not alerted until the transaction signalling the alert commits. There can be any number of concurrent signalers of a given alert, and there can be any number of concurrent waiters on a given alert.
- A waiting application is blocked in the database and cannot do any other work.
- An application can register for multiple events and can then wait for any of them to occur using the `WAITANY` procedure.
- An application can also supply an optional `timeout` parameter to the `WAITONE` or `WAITANY` procedures. A `timeout` of 0 returns immediately if there is no pending alert.
- The signalling session can optionally pass a message that is received by the waiting session.
- Alerts can be signalled more often than the corresponding application wait calls. In such cases, the older alerts are discarded. The application always gets the latest alert (based on transaction commit times).
- If the application does not require transaction-based alerts, the `DBMS_PIPE` package may provide a useful alternative.

> **See Also:**
>
> DBMS_PIPE

- If the transaction is rolled back after the call to `SIGNAL`, no alert occurs.
- It is possible to receive an alert, read the data, and find that no data has changed. This is because the data changed after the *prior* alert, but before the data was read for that *prior* alert.
- Usually, Oracle is event-driven; this means that there are no polling loops. There are two cases where polling loops can occur:
  - Shared mode. If your database is running in shared mode, a polling loop is required to check for alerts from another instance. The polling loop defaults to five seconds and can be set by the `SET_DEFAULTS` procedure.

- `WAITANY` procedure. If you use the `WAITANY` procedure, and if a signalling session does a signal but does not commit within one second of the signal, a polling loop is required so that this uncommitted alert does not camouflage other alerts. The polling loop begins at a one second interval and exponentially backs off to 30-second intervals.

Table 19-2 DBMS_ALERT Error Messages

Error Message	Procedure
ORU-10001 lock request error, status: N	SIGNAL
ORU-10015 error: N waiting for pipe status	WAITANY
ORU-10016 error: N sending on pipe 'X'	SIGNAL
ORU-10017 error: N receiving on pipe 'X'	SIGNAL
ORU-10019 error: N on lock request	WAIT
ORU-10020 error: N on lock request	WAITANY
ORU-10021 lock request error; status: N	REGISTER
ORU-10022 lock request error, status: N	SIGNAL
ORU-10023 lock request error; status N	WAITONE
ORU-10024 there are no alerts registered	WAITANY
ORU-10025 lock request error; status N	REGISTER
ORU-10037 attempting to wait on uncommitted signal from same session	WAITONE

19.7 DBMS_ALERT Examples

In this example, suppose that you want to graph average salaries by department, for all employees. Your application needs to know whenever `EMP` is changed.

Your application would look similar to this code:

```
DBMS_ALERT.REGISTER('emp_table_alert');
    <<readagain>>:
   /* ... read the emp table and graph it */
      DBMS_ALERT.WAITONE('emp_table_alert', :message, :status);
      if status = 0 then goto <<readagain>>; else
      /* ... error condition */
```

The `EMP` table would have a trigger similar to this:

```
CREATE TRIGGER emptrig AFTER INSERT OR UPDATE OR DELETE ON emp
    BEGIN
      DBMS_ALERT.SIGNAL('emp_table_alert', 'message_text');
   END;
```

When the application is no longer interested in the alert, it makes this request:

```
DBMS_ALERT.REMOVE('emp_table_alert');
```

This reduces the amount of work required by the alert signaller. If a session exits (or dies) while registered alerts exist, the alerts are eventually cleaned up by future users of this package.

The example guarantees that the application always sees the latest data, although it may not see every intermediate value.

19.8 Summary of DBMS_ALERT Subprograms

This table lists the `DBMS_ALERT` subprograms and briefly describes them.

Table 19-3 DBMS_ALERT Package Subprograms

Subprogram	Description
REGISTER Procedure	Receives messages from an alert
REMOVE Procedure	Disables notification from an alert
REMOVEALL Procedure	Removes all alerts for this session from the registration list
SET_DEFAULTS Procedure	Sets the polling interval
SIGNAL Procedure	Signals an alert (send message to registered sessions)
WAITANY Procedure	Waits `timeout` seconds to receive alert message from an alert registered for session
WAITONE Procedure	Waits `timeout` seconds to receive message from named alert

19.8.1 REGISTER Procedure

This procedure lets a session register interest in an alert.

Syntax

```
DBMS_ALERT.REGISTER (
   name      IN  VARCHAR2,
   cleanup   IN  BOOLEAN DEFAULT TRUE);
```

Parameters

Table 19-4 REGISTER Procedure Parameters

Parameter	Description
`name`	Name of the alert in which this session is interested
`cleanup`	Specifies whether to perform cleanup of any extant orphaned pipes used by the `DBMS_ALERT` package. This cleanup is only performed on the first call to `REGISTER` for each package instantiation. The default for the parameter is `TRUE`.

> **WARNING:**
>
> Alert names beginning with 'ORA$' are reserved for use for products provided by Oracle. Names must be 30 bytes or less. The name is case insensitive.

Usage Notes

A session can register interest in an unlimited number of alerts. Alerts should be deregistered when the session no longer has any interest, by calling `REMOVE`.

19.8.2 REMOVE Procedure

This procedure enables a session that is no longer interested in an alert to remove that alert from its registration list. Removing an alert reduces the amount of work done by signalers of the alert.

Syntax

```
DBMS_ALERT.REMOVE (
   name  IN  VARCHAR2);
```

Parameters

Table 19-5 REMOVE Procedure Parameters

Parameter	Description
name	Name of the alert (case-insensitive) to be removed from registration list.

Usage Notes

Removing alerts is important because it reduces the amount of work done by signalers of the alert. If a session dies without removing the alert, that alert is eventually (but not immediately) cleaned up.

19.8.3 REMOVEALL Procedure

This procedure removes all alerts for this session from the registration list. You should do this when the session is no longer interested in any alerts.

This procedure is called automatically upon first reference to this package during a session. Therefore, no alerts from prior sessions which may have terminated unusually can affect this session.

This procedure always performs a commit.

Syntax

```
DBMS_ALERT.REMOVEALL;
```

19.8.4 SET_DEFAULTS Procedure

In case a polling loop is required, use the `SET_DEFAULTS` procedure to set the polling interval.

Syntax

```
DBMS_ALERT.SET_DEFAULTS (
   sensitivity  IN  NUMBER);
```

Parameters

Table 19-6 SET_DEFAULTS Procedure Parameters

Parameter	Description
sensitivity	Polling interval, in seconds, to sleep between polls. The default interval is five seconds.

19.8.5 SIGNAL Procedure

This procedure signals an alert. The effect of the SIGNAL call only occurs when the transaction in which it is made commits. If the transaction rolls back, SIGNAL has no effect.

All sessions that have registered interest in this alert are notified. If the interested sessions are currently waiting, they are awakened. If the interested sessions are not currently waiting, they are notified the next time they do a wait call.

Multiple sessions can concurrently perform signals on the same alert. Each session, as it signals the alert, blocks all other concurrent sessions until it commits. This has the effect of serializing the transactions.

Syntax

```
DBMS_ALERT.SIGNAL (
   name     IN  VARCHAR2,
   message  IN  VARCHAR2);
```

Parameters

Table 19-7 SIGNAL Procedure Parameters

Parameter	Description
name	Name of the alert to signal.
message	Message, of 1800 bytes or less, to associate with this alert. This message is passed to the waiting session. The waiting session might be able to avoid reading the database after the alert occurs by using the information in the message.

19.8.6 WAITANY Procedure

Call this procedure to wait for an alert to occur for any of the alerts for which the current session is registered.

Syntax

```
DBMS_ALERT.WAITANY (
   name      OUT  VARCHAR2,
   message   OUT  VARCHAR2,
   status    OUT  INTEGER,
   timeout   IN   NUMBER DEFAULT MAXWAIT);
```

Parameters

Table 19-8 WAITANY Procedure Parameters

Parameter	Description
name	Returns the name of the alert that occurred.
message	Returns the message associated with the alert. This is the message provided by the SIGNAL call. If multiple signals on this alert occurred before WAITANY, the message corresponds to the most recent SIGNAL call. Messages from prior SIGNAL calls are discarded.
status	Values returned: 0 - alert occurred 1 - timeout occurred
timeout	Maximum time to wait for an alert. If no alert occurs before timeout seconds, this returns a status of 1.

Usage Notes

An implicit COMMIT is issued before this procedure is executed. The same session that waits for the alert may also first signal the alert. In this case remember to commit after the signal and before the wait; otherwise, DBMS_LOCK.REQUEST (which is called by DBMS_ALERT) returns status 4.

Exceptions

```
-20000, ORU-10024: there are no alerts registered.
```

19.8.7 WAITONE Procedure

This procedure waits for a specific alert to occur.

An implicit COMMIT is issued before this procedure is executed. A session that is the first to signal an alert can also wait for the alert in a subsequent transaction. In this case, remember to commit after the signal and before the wait; otherwise, DBMS_LOCK.REQUEST (which is called by DBMS_ALERT) returns status 4.

Syntax

```
DBMS_ALERT.WAITONE (
   name      IN   VARCHAR2,
   message   OUT  VARCHAR2,
   status    OUT  INTEGER,
   timeout   IN   NUMBER DEFAULT MAXWAIT);
```

Parameters

Table 19-9 WAITONE Procedure Parameters

Parameter	Description
name	Name of the alert to wait for.

Table 19-9 (Cont.) WAITONE Procedure Parameters

Parameter	Description
`message`	Returns the message associated with the alert. This is the message provided by the `SIGNAL` call. If multiple signals on this alert occurred before `WAITONE`, the message corresponds to the most recent `SIGNAL` call. Messages from prior `SIGNAL` calls are discarded.
`status`	Values returned: 0 - alert occurred 1 - timeout occurred
`timeout`	Maximum time to wait for an alert. If the named alert does not occurs before `timeout` seconds, this returns a status of 1.

20
DBMS_APP_CONT

The `DBMS_APP_CONT` package provides an interface to determine if the in-flight transaction on a now unavailable session committed or not, and if the last call on that session completed or not.

> **See Also:**
>
> *Oracle Database Development Guide* for explanations of application continuity and Transaction Guard, and the relationship between these two features:
>
> - "Using Transaction Guard"

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_APP_CONT Subprograms

20.1 DBMS_APP_CONT Overview

The DBMS_APP_CONT package can be used to solve this example issue.

Problem Description

One of the fundamental problems for recovering applications after an outage is that the commit message that is sent back to the client is not durable. If there is a break between the client and the server, the client sees an error message indicating that the communication failed. This error does not inform the application whether the submission executed any commit operations or if a procedural call, ran to completion executing all expected commits and session state changes or failed part way through or yet worse, is still running disconnected from the client.

GET_LTXID_OUTCOME

The purpose of the GET_LTXID_OUTCOME Procedure is to determine if the in-flight transaction on a now unavailable session completed or not. It is used when the original session returned an error due to unavailability. Situations that can cause such session unavailability may occur at the session, instance, server, or network, and result from planned or unplanned outages. When such an outage occurs, the application receives a disconnection error. Such an error provides no insight as to whether the transaction committed. It also does not reveal what the application might have been expecting from that commit if it had returned.

> **See Also:**
>
> *Oracle Database Concepts* for explanation of Logical Transaction ID

20.2 DBMS_APP_CONT Security Model

Applications must have the `EXECUTE` privilege on the `DBMS_APP_CONT` package.

To grant this privilege, ask your database administrator to run the following SQL statement:

```
GRANT execute on DBMS_APP_CONT to application user ;
```

20.3 Summary of DBMS_APP_CONT Subprograms

The `DBMS_APP_CONT` package contains the `GET_LTXID_OUTCOME` Procedure.

Table 20-1 DBMS_APP_CONT Package Subprograms

Subprogram	Description
ADD_SQL_CONNECTION_TEST Procedure	This procedure adds a new connection test that is used during draining sessions before planned maintenance begins. Use this procedure when the SQL connection test is not covered by standard tests. The test is enabled when added.
APPLY_REPLAY_RULE Procedure	This procedure activates a new replay rule that is inherited by any sub-routine unless it is overwritten by a new rule applied to that sub-routine.
DELETE_SQL_CONNECTION_TEST Procedure	This procedure deletes a connection test that is no longer needed for planned draining. Removing a test applies immediately to all RAC instances where the PDB is open.
DISABLE_CONNECTION_TEST Procedure	This procedure disables usage of a connection test during draining of sessions. Disabling a test applies immediately to all RAC instances where the PDB is open.
ENABLE_CONNECTION_TEST Procedure	This procedure enables usage of a connection test for draining database sessions before planned maintenance. Enabling a test applies immediately to all RAC instances where the PDB is open.
GET_LTXID_OUTCOME Procedure	Lets customer applications and third party application servers determine the transactional status of the last session when that session becomes unavailable.
GET_REPLAY_RULES Function	This function returns a list of the replay rules. From the results, you can test if any rule target is replayable with bit and function.
RESET_REPLAY_RULES Procedure	This procedure resets all existing replay rules. When you call this procedure, it clears all locally defined rules in the specified scope and restores them to the original state when they were inherited from their parent scope, or the service if parent scope is not specified.

20.3.1 ADD_SQL_CONNECTION_TEST Procedure

This procedure adds a new connection test that is used during draining sessions before planned maintenance begins. Use this procedure when the SQL connection test is not covered by standard tests. The test is enabled when added. If the optional service name qualifier is provided, the test only applies only to that service name.

Syntax

```
DBMS_SESSION.ADD_SQL_CONNECTION_TEST (
   connection_test          IN VARCHAR2
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 20-2 ADD_SQL_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST	The SQL text used to test and drain connections.
SERVICE_NAME	Optional service name qualifier.

Usage Notes

The ADD_SQL_CONNECTION_TEST Procedure adds a connection test for the purpose of draining sessions before planned maintenance begins. The connection test is used by the application to test connections that are marked for draining. Sessions are set for draining at stop and relocate operations for services or PDBs. When set the RDBMS closes the connection while draining so the application sees no errors during planned maintenance. You can enter as many CONNECTION TESTs as needed. They are used only during planned maintenance. The tests apply to all RAC instances.

Check online documentation for latest updates on service qualifier availability.

Added connection can be viewed by querying the view DBA_CONNECTION_TESTS.

This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level.

20.3.2 APPLY_REPLAY_RULE Procedure

This procedure activates a new replay rule that is inherited by any sub-routine unless it is overwritten by a new rule applied to that sub-routine.

A replay rule activated in a PL/SQL block loses its effect when reset_replay_rules is called or the block returns.

> **Note:**
>
> Calling this procedure in a sub-block of a procedure, function, or anonymous block applies the rule to its parent scope. When the sub-block returns, the rule will not lose its effect. Thus, you are not recommended to call this procedure in a sub-block.

Syntax

```
DBMS_APP_CONT.APPLY_REPLAY_RULE (
 replayable IN  BOOLEAN,
 targets    IN  BINARY_INTEGER,
 scope      IN  BINARY_INTEGER DEFAULT DBMS_APP_CONT.SCOPE_CURRENT);
```

Parameters

Table 20-3 *APPLY_REPLAY_RULE Procedure Parameters*

Parameter	Description
replayable	TRUE or FALSE, depending on weather the rule is replayable or not.

Table 20-3 (Cont.) *APPLY_REPLAY_RULE Procedure Parameters*

Parameter	Description
targets	Target effects on which this rule is applied. If multiple effects are desired, pass them as bit or as individual effects, such as: dbms_app_cont.side_effects + dbms_app_cont.autonomous_transactions
scope	Scope of the rule, either dbms_app_cont.scope_current (default) or dbms_app_cont.scope_parent.

20.3.3 DELETE_SQL_CONNECTION_TEST Procedure

This procedure deletes a connection test that is no longer needed for planned draining. Removing a test applies immediately to all RAC instances where the PDB is open.

Syntax

```
DBMS_SESSION.DELETE_SQL_CONNECTION_TEST (
   connection_test          IN VARCHAR2
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 20-4 DELETE_SQL_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST	The SQL text used to test and drain connections.
SERVICE_NAME	Optional service name qualifier. If the optional SERVICE_NAME qualifier is provided, only the test for that service name is deleted.

Usage Notes

If you are not certain if a test should be deleted, you can disable the test using DISABLE_CONNECTION_TEST Procedure. Only custom SQL tests can be deleted. Predefined tests cannot be deleted. Check for latest updates on service qualifier availability.

This procedure is owned by SYS at CDB$ROOT or PDB level, or SYS for when not multitenant.

Connection tests and their status can be checked by querying the view DBA_CONNECTION_TESTS.

20.3.4 DISABLE_CONNECTION_TEST Procedure

This procedure disables usage of a connection test during draining of sessions. Disabling a test applies immediately to all RAC instances where the PDB is open.

Syntax

```
DBMS_SESSION.DISABLE_CONNECTION_TEST (
   connection_test_type     IN VARCHAR2,
   connection_test          IN VARCHAR2,
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 20-5 DISABLE_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST_TYPE	The permitted values are: • DBMS_SESSION.SQL_TEST • DBMS_SESSION.PING_TEST • DBMS_SESSION.ENDREQUEST_TEST
CONNECTION_TEST	The SQL text used to test and drain connections. This parameter is allowed only if the value of CONNECTION_TEST_TYPE is SQL_TEST.
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, only the test for that service name is enabled. A disable at service name level takes precedence over an enable at PDB level. That is the PDB can be enabled, and the service disabled.

Usage Notes

This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level.

Connection tests and their status can be checked by querying the view DBA_CONNECTION_TESTS.

20.3.5 ENABLE_CONNECTION_TEST Procedure

This procedure enables usage of a connection test for draining database sessions before planned maintenance. Enabling a test applies immediately to all RAC instances where the PDB is open.

Syntax

```
DBMS_SESSION.ENABLE_CONNECTION_TEST (
   connection_test_type     IN VARCHAR2,
   connection_test          IN VARCHAR2,
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 20-6 ENABLE_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST_TYPE	The connection type used when managing connection tests for draining before planned maintenance. See ADD, DELETE, ENABLE, DISABLE procedures for connection tests. The permitted values are: • DBMS_SESSION.SQL_TEST • DBMS_SESSION.PING_TEST • DBMS_SESSION.ENDREQUEST_TEST

Table 20-6 (Cont.) ENABLE_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST	The SQL text used to test and drain connections at the RDBMS before planned maintenance starts. This parameter is allowed only if the value of CONNECTION_TEST_TYPE is SQL_TEST.
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, only the test for that service name is enabled. An enable at service name level overrides any higher-level disables. That is, the PDB can be disabled, and the service enabled.

Usage Notes

- This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level
- ENABLE_CONNECTION_TEST enables a connection test for draining sessions during planned maintenance. The enable operation applies to all RAC instances where the PDB is open. It persists across database restarts.
- This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level.

20.3.6 GET_LTXID_OUTCOME Procedure

This procedure lets customer applications and third party application servers determine the transactional status of the last session when that session becomes unavailable.

Syntax

```
DBMS_APP_CONT.GET_LTXID_OUTCOME (
 client_ltxid          IN    RAW,
 committed             OUT   BOOLEAN,
 user_call_completed   OUT   BOOLEAN)
```

Parameters

Table 20-7 *GET_LTXID_OUTCOME Procedure Parameters*

Parameter	Description
client_ltxid	Client-side logical transaction ID. Obtain the LTXID from the previous failed session using the client driver provided APIs - getLTXID for JDBC, and LogicalTransactionId for ODP.net., and OCI_ATTR_GET with LTXID for OCI.
committed	Returns TRUE if the transaction with the named logical LTXID has COMMITTED. Returns FALSE if the logical LTXID has not COMMITTED. When returning FALSE, the procedure blocks the LTXID from further use so that there is no possibility of previous in-flight work committing this LTXID.

Table 20-7 (Cont.) *GET_LTXID_OUTCOME Procedure Parameters*

Parameter	Description
user_call_completed	Whether all information has been returned to the client. Examples of such messages are the number of rows processed when using autocommit or commit on success, parameter and function results when calling PL/SQL, or PL/SQL with more work to do after the COMMIT. Applications that expect to use data returned from the commit in order to function correctly must look at this second parameter.

Exceptions

Table 20-8 GET_LTXID_OUTCOME Procedure Exceptions

Exception	Description
ORA-14950 - SERVER_AHEAD	The server is ahead so the transaction is both an old transaction and one which has already committed. This is an error as the application is passing an older LTXID that is the not the last used for that session. The purpose of GET_LTXID_OUTCOME is to return the current transaction outcome for that session after a recoverable outage.
ORA-14951 - CLIENT_AHEAD	The client is ahead of the server. This can happen if the server has been flashed backed, recovered using media recovery, or is a standby that has opened earlier with data loss.
ORA-14906 - SAME_SESSION	Executing GET_LTXID_OUTCOME is not supported on the session owning the LTXID as it blocks further processing on that session after a recoverable outage.
ORA-14909 - COMMIT_BLOCKED	Your session has been blocked from committing by another user with the same username using GET_LTXID_OUTCOME. GET_LTXID_OUTCOME should only be called on dead sessions. Please check with your application administrator.
ORA-14952 - ERROR	The outcome cannot be determined. During processing an error happened. The error stack shows the error detail.

20.3.7 GET_REPLAY_RULES Function

This function returns a list of the replay rules. From the results, you can test if any rule target is replayable with bit and function.

Syntax

```
DBMS_APP_CONT.GET_REPLAY_RULES (
 replayable BOOLEAN,
 scope      IN  BINARY_INTEGER DEFAULT DBMS_APP_CONT.SCOPE_CURRENT)
 RETURN     BINARY_INTEGER;
```

Parameters

Table 20-9 *GET_REPLAY_RULES Procedure Parameters*

Parameter	Description
`replayable`	`TRUE` or `FALSE`, depending on weather the rule is replayable or not.
`scope`	Scope of the rules, either `dbms_app_cont.scope_current` (default) or `dbms_app_cont.scope_parent`.

20.3.8 RESET_REPLAY_RULES Procedure

This procedure resets all existing replay rules. When you call this procedure, it clears all locally defined rules in the specified scope and restores them to the original state when they were inherited from their parent scope, or the service if parent scope is not specified.

Syntax

```
DBMS_APP_CONT.RESET_REPLAY_RULES (
 targets  IN  BINARY_INTEGER DEFAULT NULL,
 scope    IN  BINARY_INTEGER DEFAULT DBMS_APP_CONT.SCOPE_CURRENT);
```

Parameters

Table 20-10 *RESET_REPLAY_RULES Procedure Parameters*

Parameter	Description
`targets`	Target effects on which you want to reset the rules.
`scope`	Scope of the rule, either `dbms_app_cont.scope_current` (default) or `dbms_app_cont.scope_parent`.

21
DBMS_APP_CONT_ADMIN

This package provides a collection dba level admin operations in relation to Application Continuity.

This chapter contains the following topics:

- DBMS_APP_CONT_ADMIN Security Model
- Summary of DBMS_APP_CONT_ADMIN Subprograms

21.1 DBMS_APP_CONT_ADMIN Security Model

Applications must have the `EXECUTE` privilege on the `DBMS_APP_CONT_ADMIN` package.

21.2 Summary of DBMS_APP_CONT_ADMIN Subprograms

This topic lists the `DBMS_APP_CONT_ADMIN` subprograms in alphabetical order and briefly describes them.

Table 21-1 DBMS_APP_CONT_ADMIN Package Subprograms

Subprogram	Description
ACCHK_CLEAR_FILTER Procedure	This procedure clears an ACCHK filter by providing a filter type and filter name.
ACCHK_SET Procedure	This procedure enables or disables data collection for acchk protection for your application when using Application Continuity or Transparent Application Continuity. The `DISABLE_TIME` parameter is used to extend the runtime. Default is `5` minutes.
ACCHK_SET_FILTER Procedure	This procedure sets ACCHK filtering options for `acchk` protection analysis by service name, module name, program name, and machine name.
ACCHK_SHOW_FILTERS Procedure	This procedure uses a cursor to show ACCHK filters ordered by service, program, module, and machine name.
ADD_SQL_CONNECTION_TEST Procedure	This procedure adds a new connection test that is used during draining sessions before planned maintenance begins.
DELETE_SQL_CONNECTION_TEST Procedure	This procedure deletes a connection test that is no longer needed for planned draining. Removing a test applies immediately to all RAC instances where the PDB is open.
DISABLE_CONNECTION_TEST Procedure	This procedure disables usage of a connection test during draining of sessions.
DISABLE_FAILOVER Procedure	This procedure disables failover on a given service.
ENABLE_AC Procedure	This procedure enables Application Continuity (AC) on a given service.
ENABLE_CONNECTION_TEST Procedure	This procedure enables usage of a connection test for draining database sessions before planned maintenance. Enabling a test applies immediately to all RAC instances where the PDB is open.

Table 21-1 (Cont.) DBMS_APP_CONT_ADMIN Package Subprograms

Subprogram	Description
ENABLE_RESET_STATE Procedure	This procedure enables clearing the session state usage between requests, so that each new request starts clean (usage web and stateless applications).
ENABLE_TAC Procedure	This procedure enables Transparent Application Continuity (TAC) on a given service.
MODIFY_SERVICE Procedure	This procedure modifies a database service using the given parameters.
SET_DRAINING Procedure	This procedure configures draining options for your service, such as timeout value and stop option.
SET_LOAD_BALANCING Procedure	This procedure configures load balancing options for your service.

21.2.1 ACCHK_CLEAR_FILTER Procedure

This procedure clears an ACCHK filter by providing a filter type, filter name, or all the filters.

- You should call `ACCHK_CLEAR_FILTER` before `ACCHK_SET`, which loads the filters.
- `ACCHK_CLEAR_FILTER` is set at the PDB level.

Syntax

```
DBMS_APP_CONT_ADMIN.ACCHK_CLEAR_FILTER (
   filter_type IN NUMBER   DEFAULT NULL,
   filter_name IN VARCHAR2 DEFAULT NULL,
   purge_all   IN BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 21-2 ACCHK_CLEAR_FILTER Procedure Parameters

Parameter	Description
`filter_type`	Type of the filter, which can be one of the constants `SERVICE_FILTER`, `PROGRAM_FILTER`, `MODULE_FILTER`, or `MACHINE_FILTER` from the `DBMS_APP_CONT_ADMIN` package.
`filter_name`	Name of the filter that you want to delete when `purge_all` is FALSE.
`purge_all`	This parameter is used to delete all the filters. • `TRUE`- enables deletion of all the filters. • `FALSE`- disables deletion of all the filters. When this parameter is set to `FALSE`, then only the specified filter is deleted.

Error Messages

Table 21-3 ACCHK_CLEAR_FILTER Procedure Error Messages

Error Code	Description
ORA-20000	Provided filter name exceeds maximum number of characters.
ORA-20000	Invalid filter type value.
ORA-20000	Filter name is not specified.

Examples

The following examples illustrate how to clear an ACCHK service filter:

```
SQL> execute
dbms_app_cont_admin.acchk_clear_filter(DBMS_APP_CONT_ADMIN.SERVICE_FILTER,
'ORACLE.Service1');
```

21.2.2 ACCHK_SET Procedure

This procedure enables or disables data collection for `acchk` protection reports for your application when using Application Continuity or Transparent Application Continuity. The `DISABLE_TIME` parameter is used to extend the runtime. Default value is `10` minutes.

Enabling or disabling data collection applies to the level connected, that is, a CDB or a PDB.

- `acchk` is enabled/disabled at your CDB if connected to the container.
- `acchk` is enabled/disabled at your PDB only when connected to the PDB.

Data collection applies to new sessions only.

Once enabled, data is collected for the workload run under this service. You can then view this data in the `ACCHK_REPORT`s and can also be mined in the `ACCHK` views.

Syntax

```
DBMS_APP_CONT_ADMIN.ACCHK_SET (
   enabled                 IN BOOLEAN,
   disable_time_in_seconds  DEFAULT 600);
```

Parameters

Table 21-4 ACCHK_SET Procedure Parameters

Parameter	Description
enabled	This parameter is used to enable or disable data collection at a CDB or PDB level. • `TRUE`-enables data collection at this level. • `FALSE`-explicitly disables data collection.

Table 21-4 (Cont.) ACCHK_SET Procedure Parameters

Parameter	Description
disable_time_in_seconds	Optional parameter used to disable ACCHK tracing automatically in a given number of seconds. The maximum value that you can specify is 3600 seconds. The default value is 600 seconds.

Usage Notes

- This procedure is owned by SYS at CDB$ROOT or PDB level, or by SYS when not multitenant
- The acchk activation is enabled across all instances of RAC supporting that service.
- Enabling is persistent to allow for failover and restart tests, that is, implementation uses ALTER SESSION SET EVENTS ….. SCOPE=BOTH. The enable is per database. For Data Guard, enable and disable must be at each database.

Examples

Application Continuity Protection Check is not enabled by default. Follow this procedure to enable or disable ACCHK and generate reports to check protection level for the applications.

1. Grant read access to the users, who will run the Application Continuity Protection Check report and views, using the ACCHK_READ role:

   ```
   GRANT ACCHK_READ TO USER;
   ```

2. Enable Application Continuity tracing for your applications using the dbms_app_cont_admin.acchk_set(true) procedure:

   ```
   SQL> execute dbms_app_cont_admin.acchk_set(true);
   ```

   By default, ACCHK is disabled automatically after 600 seconds. You can specify a lower number to reduce the auto disable time. For example, to disable ACCHK after 300 seconds:

   ```
   SQL> dbms_app_cont_admin.acchk_set(true,300);
   ```

   The dbms_app_cont_admin.acchk_set(true) procedure enables Application Continuity tracing at the database level to which you are connected. If you are connected at the CDB level, then tracing is enabled for the CDB, and if you are connected at the PDB level, then tracing is enabled for the PDB.

   > **Note:**
   >
   > Set the COMPATIBLE parameter to 12.2.0 or greater.

3. To disable Application Continuity tracing for new sessions in your applications:

   ```
   SQL> execute dbms_app_cont_admin.acchk_set(false);
   ```

> **Note:**
>
> The tracing will not be disabled for the current sessions until the sessions are terminated.

21.2.3 ACCHK_SET_FILTER Procedure

This procedure sets ACCHK filtering options for `acchk` protection analysis by service name, module name, program name, and machine name.

Filter options are cumulative by multiple calls of this procedure.

- You should call `ACCHK_SET_FILTER` before `ACCHK_SET`, which loads the filters.
- `ACCHK_SET_FILTER` is set at the PDB level.

Syntax

```
DBMS_APP_CONT_ADMIN.ACCHK_SET_FILTER (
   filter_type IN NUMBER,
   filter_name IN VARCHAR2);
```

Parameters

Table 21-5 ACCHK_SET_FILTER Procedure Parameters

Parameter	Description
`filter_type`	Type of the filter, which can be one of the constants`SERVICE_FILTER`, `PROGRAM_FILTER`, `MODULE_FILTER`, or `MACHINE_FILTER` from the `DBMS_APP_CONT_ADMIN` package.
`filter_name`	Name of the filter, which must match a service name, program name, module name, or machine name.

Error Messages

Table 21-6 ACCHK_SET_FILTER Procedure Error Messages

Error Code	Description
`ORA-20000`	Provided filter name exceeds maximum number of characters.
`ORA-20000`	Service `service_name` does not exist.
`ORA-20000`	Invalid filter type value.
`ORA-20000`	Filter name is not specified.
`ORA-20000`	The limit has been reached, only 1024 filters are allowed.
`ORA-20000`	Filter name already exists.

Examples

The following examples illustrate how to set filters for ACCHK constants:

- To set a service filter:

```
SQL> execute
dbms_app_cont_admin.acchk_set_filter(DBMS_APP_CONT_ADMIN.SERVICE_FILTER,
'ORACLE.Service1');
```

- To set a program filter:

```
SQL> execute
dbms_app_cont_admin.acchk_set_filter(DBMS_APP_CONT_ADMIN.PROGRAM_FILTER,
'Oracle.Program');
```

- To set a module filter:

```
SQL> execute
dbms_app_cont_admin.acchk_set_filter(DBMS_APP_CONT_ADMIN.MODULE_FILTER,
'Oracle.Module');
```

21.2.4 ACCHK_SHOW_FILTERS Procedure

This procedure uses a cursor to show ACCHK filters ordered by service, program, module, and machine name.

Syntax

```
DBMS_APP_CONT_ADMIN.ACCHK_SHOW_FILTERS;
```

Examples

The following examples illustrate how to show ACCHK filters:

```
SQL> execute dbms_app_cont_admin.acchk_show_filters;
ResultSet TYPE  FILTER_NAME  SERVICE          PROGRAM            MODULE
#1                           oracle.service1  oracle.program     oracle.module
```

21.2.5 ADD_SQL_CONNECTION_TEST Procedure

This procedure adds a new connection test that is used during draining sessions before planned maintenance begins. Use this procedure when the SQL connection test is not covered by standard tests. The test is enabled when added. If the optional service name qualifier is provided, the test only applies only to that service name.

Syntax

```
DBMS_APP_CONT_ADMIN.ADD_SQL_CONNECTION_TEST (
   connection_test          IN VARCHAR2
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 21-7 ADD_SQL_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST	The SQL text used to test and drain connections.

Table 21-7 (Cont.) ADD_SQL_CONNECTION_TEST Procedure Parameters

Parameter	Description
SERVICE_NAME	Optional service name qualifier.

Usage Notes

The ADD_SQL_CONNECTION_TEST Procedure adds a connection test for the purpose of draining sessions before planned maintenance begins. The connection test is used by the application to test connections that are marked for draining. Sessions are set for draining at stop and relocate operations for services or PDBs. When set the RDBMS closes the connection while draining so the application sees no errors during planned maintenance. You can enter as many CONNECTION TESTs as needed. They are used only during planned maintenance. The tests apply to all RAC instances.

Check online documentation for latest updates on service qualifier availability.

Added connection can be viewed by querying the view DBA_CONNECTION_TESTS.

This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level.

21.2.6 DELETE_SQL_CONNECTION_TEST Procedure

This procedure deletes a connection test that is no longer needed for planned draining. Removing a test applies immediately to all RAC instances where the PDB is open.

Syntax

```
DBMS_APP_CONT_ADMIN.DELETE_SQL_CONNECTION_TEST (
   connection_test          IN VARCHAR2
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 21-8 DELETE_SQL_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST	The SQL text used to test and drain connections.
SERVICE_NAME	Optional service name qualifier. If the optional SERVICE_NAME qualifier is provided, only the test for that service name is deleted.

Usage Notes

If you are not certain if a test should be deleted, you can disable the test using DISABLE_CONNECTION_TEST Procedure. Only custom SQL tests can be deleted. Predefined tests cannot be deleted. Check for latest updates on service qualifier availability.

This procedure is owned by SYS at CDB$ROOT or PDB level, or SYS for when not multitenant.

Connection tests and their status can be checked by querying the view DBA_CONNECTION_TESTS.

21.2.7 DISABLE_CONNECTION_TEST Procedure

This procedure disables usage of a connection test during draining of sessions. Disabling a test applies immediately to all RAC instances where the PDB is open.

Syntax

```
DBMS_APP_CONT_ADMIN.DISABLE_CONNECTION_TEST (
   connection_test_type     IN VARCHAR2,
   connection_test          IN VARCHAR2,
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 21-9 DISABLE_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST_TYPE	The permitted values are: • DBMS_APP_CONT_ADMIN.SQL_TEST • DBMS_APP_CONT_ADMIN.PING_TEST • DBMS_APP_CONT_ADMIN.ENDREQUEST_TEST
CONNECTION_TEST	The SQL text used to test and drain connections. This parameter is allowed only if the value of CONNECTION_TEST_TYPE is SQL_TEST.
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, only the test for that service name is enabled. A disable at service name level takes precedence over an enable at PDB level. That is the PDB can be enabled, and the service disabled.

Usage Notes

This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level.

Connection tests and their status can be checked by querying the view DBA_CONNECTION_TESTS.

21.2.8 DISABLE_FAILOVER Procedure

This procedure disables failover on a given service.

Syntax

```
DBMS_APP_CONT_ADMIN.DISABLE_FAILOVER (
    service_name IN VARCHAR2);
```

Parameters

Table 21-10 DISABLE_FAILOVER Procedure Parameters

Parameter	Description
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, then failover is disabled only for that service name.

Usage Notes

- You must have the PDBADMIN user permissions to use this procedure.
- Use the full service name on which you want to disable the failover.

Examples

The following example illustrates how to disable failover for a service:

```
SQL> execute dbms_app_cont_admin.disable_failover('TPURGENT');
```

21.2.9 ENABLE_AC Procedure

This procedure enables Application Continuity (AC) on a given service.

Syntax

```
DBMS_APP_CONT_ADMIN.ENABLE_AC (
   service_name              IN VARCHAR2,
   failover_restore          IN VARCHAR2       DEFAULT 'LEVEL1'
   replay_initiation_timeout IN BINARY_INTEGER DEFAULT 300);
```

Parameters

Table 21-11 ENABLE_AC Procedure Parameters

Parameter	Description
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, AC is enabled only for the specified service. An enable at service name level overrides any higher-level disables. That is, the PDB can be disabled, and the service enabled.
FAILOVER_RESTORE	Failover restore. Possible values are NONE or LEVEL1.
REPLAY_INITIATION_TIMEOUT	Replay timeout that specifies how many seconds after a request is submitted to allow that request to replay.

Usage Notes

- You must have the PDBADMIN user permissions to use this procedure.
- Use the full service name on which you want to enable AC.

Examples

The following example illustrates how to enable Application Continuity for your service:

```
SQL> execute dbms_app_cont_admin.enable_ac('TPURGENT', 'LEVEL1', 600);
```

21.2.10 ENABLE_CONNECTION_TEST Procedure

This procedure enables usage of a connection test for draining database sessions before planned maintenance. Enabling a test applies immediately to all RAC instances where the PDB is open.

Syntax

```
DBMS_APP_CONT_ADMIN.ENABLE_CONNECTION_TEST (
   connection_test_type     IN VARCHAR2,
   connection_test          IN VARCHAR2,
   service_name             IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 21-12 ENABLE_CONNECTION_TEST Procedure Parameters

Parameter	Description
CONNECTION_TEST_TYPE	The connection type used when managing connection tests for draining before planned maintenance. See ADD, DELETE, ENABLE, DISABLE procedures for connection tests. The permitted values are: • DBMS_APP_CONT_ADMIN.SQL_TEST • DBMS_APP_CONT_ADMIN.PING_TEST • DBMS_APP_CONT_ADMIN.ENDREQUEST_TEST • DBMS_APP_CONT_ADMIN.BEGINREQUEST_TEST
CONNECTION_TEST	The SQL text used to test and drain connections at the RDBMS before planned maintenance starts. This parameter is allowed only if the value of CONNECTION_TEST_TYPE is SQL_TEST.
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, only the test for that service name is enabled. An enable at service name level overrides any higher-level disables. That is, the PDB can be disabled, and the service enabled.

Usage Notes

- This procedure is owned by SYS and is granted to users for execution at CDB$ROOT or PDB levels, or when not multitenant, at dictionary level
- ENABLE_CONNECTION_TEST enables a connection test for draining sessions during planned maintenance. The enable operation applies to all RAC instances where the PDB is open. It persists across database restarts.

- This procedure is owned by `SYS` and is granted to users for execution at `CDB$ROOT` or `PDB` levels, or when not multitenant, at dictionary level.

21.2.11 ENABLE_RESET_STATE Procedure

This procedure enables clearing the session state usage between requests, so that each new request starts clean (usage web and stateless applications).

Syntax

```
DBMS_APP_CONT_ADMIN.ENABLE_RESET_STATE (
   service_name IN VARCHAR2,
   level        IN VARCHAR2 DEFAULT AUTO);
```

Parameters

Table 21-13 ENABLE_RESET_STATE Procedure Parameters

Parameter	Description
`SERVICE_NAME`	Optional service name qualifier. If the optional service name qualifier is provided, `RESET_STATE` is enabled only for the specified service.
`LEVEL`	Level of the Oracle Database reset session state configuration. Available options are `LEVEL1`, `LEVEL2`, or `AUTO`.

Usage Notes

- You must have the `PDBADMIN` user permissions to use this procedure.
- Use the full service name for which you want to clear the session state.

Examples

The following example illustrates how to enable reset state for your service:

```
SQL> execute dbms_app_cont_admin.enable_reset_state('TPURGENT', 'AUTO');
```

21.2.12 ENABLE_TAC Procedure

This procedure enables Transparent Application Continuity (TAC) on a given service.

Syntax

```
DBMS_APP_CONT_ADMIN.ENABLE_TAC (
   service_name              IN VARCHAR2,
   failover_restore          IN VARCHAR2       DEFAULT 'AUTO'
   replay_initiation_timeout IN BINARY_INTEGER DEFAULT 300
   session_state_consistency IN VARCHAR2       DEFAULT 'AUTO');
```

Parameters

Table 21-14 ENABLE_TAC Procedure Parameters

Parameter	Description
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, TAC is enabled only for the specified service. An enable at service name level overrides any higher-level disables. That is, the PDB can be disabled, and the service enabled.
FAILOVER_RESTORE	Failover restore. Possible values are AUTO or LEVEL1.
REPLAY_INITIATION_TIMEOUT	Replay timeout that specifies how many seconds after a request is submitted to allow that request to replay.
SESSION_STATE_CONSISTENCY	Session State Consistency. Possible values are AUTO or HYBRID.

Exceptions

- When FAILOVER_RESTORE is set to NULL or FAILOVER_RESTORE is not set to AUTO or LEVEL1: ORA-20000 - Invalid failover_restore parameter.
- When SESSION_STATE_CONSISTENCY is set to NULL or SESSION_STATE_CONSISTENCY is not set as AUTO: ORA-20000 - Invalid session_state_consistency parameter.

Usage Notes

- You must have the PDBADMIN user permissions to use this procedure.
- Use the full service name on which you want to enable TAC.

Examples

The following example illustrates how to enable Transparent Application Continuity for your service:

```
SQL> execute dbms_app_cont_admin.enable_tac('TPURGENT', 'AUTO', 600, 'AUTO');
```

21.2.13 MODIFY_SERVICE Procedure

This procedure modifies a database service using the given parameters.

Syntax

```
DBMS_APP_CONT_ADMIN.MODIFY_SERVICE(
   service_name   IN VARCHAR2,
   service_params IN svc_parameter_array);
```

Parameters

Table 21-15 MODIFY_SERVICE Procedure Parameters

Parameter	Description
service_name	Name of the service, limited to 64 characters in the Data Dictionary
parameter_array	Associative array with name/value pairs of the service attributes. Supported names: • aq_ha_notifications • auto_connection_rebalance • clb_goal • commit_fast_path • commit_outcome • drain_timeout • edition • failover_delay • failover_method • failover_restore • failover_retries • failover_type • goal • replay_initiation_timeout • reset_state • retention_timeout • session_state_consistency • sql_translation_profile • stop_option • template_timeout • true_cache_service
aq_ha_notifications	Determines whether Fast Application Notification (FAN) is enabled for OCI/OCCI/ODP. In Oracle Database12*c*, FAN uses Oracle Notification Services (ONS). This parameter is still used to enable FAN. FAN is recommended for all High Availability systems, and is on by default for Application Continuity
auto_connection_rebalance	Automatically balances the workload management.
clb_goal	Connection Load Balancing goal, either LONG or SHORT.
commit_fast_path	Enable or disable Oracle Database transaction guard.
commit_outcome	Determines whether transaction COMMIT outcome is accessible after the COMMIT has executed. While the database guarantees that COMMIT is durable, this ensures that the outcome of the COMMIT is durable.Applications use the feature to probe the status of the commit last executed after an outage, and is available to applications to determine an outcome. Note: • Invoking the GET_LTXID_OUTCOME Procedure of the DBMS_APP_CONT package requires that the commit_outcome attribute be set. • commit_outcome has no effect on active Data Guard and read-only databases. • commit_outcome is allowed only on user-defined database services

Table 21-15 (Cont.) MODIFY_SERVICE Procedure Parameters

Parameter	Description
drain_timeout	When this parameter is set, all sessions connected to that service are drained by the client drivers and pools using Fast Connection Failover (FCF). The drain_timeout can be set on the service, to stop and relocate drains for this time by default.
edition	If this argument has a non-NULL value, this provides the initial session edition for subsequent database connections using this service that do not specify an edition. If no value is specified, this argument has no effect. During service creation or modification, no validation is performed on this parameter. At connection time, if the connecting user does not have USE privilege on the edition, or the edition does not exist, this raises the error ORA-38802 (edition does not exist).
failover_delay	Delay in seconds between connection retries for Application Continuity and TAF. The default is 10 seconds for Application Continuity. Do not use a 0-second delay if the service needs time to failover and register. Long delays are good for planned outages and to failover to Data Guard. Short delays work well with Oracle RAC when the service is already available.
failover_method	Failover TYPE for the service for Application Continuity and TAF. If the failover_type is set to TRANSACTION on the service, this automatically sets COMMIT_OUTCOME to TRUE. JDBC Replay Driver uses the FAILOVER_TYPE service attribute setting of TRANSACTION for TRANSACTION failover. OCI uses the older settings of SELECT and SESSION. The server only accepts FAILOVER_METHOD = BASIC with the TRANSACTION setting.
failover_restore	For Application Continuity, when the failover_restore parameter is set, the session states are restored before replaying for ODP.NET and Java. Use LEVEL1 for ODP.NET and Java with Application Continuity to restore the initial state. For AC OCI, use NONE for applications that are not STATIC.
failover_retries	Number of connection retries for Application Continuity and TAF. Using the failover_retries and failover_delay parameters, the failover can be delayed until the service is next available. This parameter is for connecting. It does not control the number of failovers, which is 3 for each incident for Application Continuity.
failover_type	Failover TYPE for the service for Application Continuity and TAF.
goal	Workload management goal directive for the service. Valid values: • DBMS_SERVICE.GOAL_SERVICE_TIME • DBMS_SERVICE.GOAL_THROUGHPUT • DBMS_SERVICE.GOAL_NONE
replay_initiation_timeout	For Application Continuity, replay_initiation_timeout is the difference between the time of original execution of first operation of a request, and the time that the replay is ready to start after a successful reconnect. Replay initiation time is measured from the time that the request was originally submitted until the time that replay has connected and is ready to replay. When replay is expected, keep this value high. Default is 900 seconds.

Table 21-15 (Cont.) MODIFY_SERVICE Procedure Parameters

Parameter	Description
reset_state	Use to clean session state automatically by Oracle Database between requests.
retention_timeout	Used in conjunction with commit_outcome, it determines the amount of time (in seconds) that the COMMIT_OUTCOME is retained. Default is 24 hours (86400). Maximum value is 30 days (2592000).
session_state_consistency	Describes how nontransactional is changed during a request. This parameter is considered only if failover_type is set to TRANSACTION for Application Continuity. Examples of session state are NLS settings, optimizer preferences, event settings, PL/SQL global variables, temporary tables, advanced queues, LOBs, and result cache. If these values change after the request starts, set to DYNAMIC (default). Almost all applications should use DYNAMIC mode. If you are unsure, use DYNAMIC mode.
sql_translation_profile	Name of SQL translation profile.
stop_option	Stop options for the service.
template_timeout	Template timeout time in seconds.
true_cache_service	Name of the True Cache service being registered with the primary service.

Examples

```
DECLARE
params dbms_app_cont_admin.svc_parameter_array;
BEGIN
   params('FAILOVER_TYPE')             :='TRANSACTION';
   params('FAILOVER_RESTORE')          :='LEVEL2';
   params('RESET_STATE')               :='NONE';
   params('SESSION_STATE_CONSISTENCY') :='DYNAMIC';
   DBMS_APP_CONT_ADMIN.MODIFY_SERVICE('ernie.example.com', params);
END;
```

21.2.14 SET_DRAINING Procedure

This procedure configures draining options for your service, such as timeout value and stop option.

Syntax

```
DBMS_APP_CONT_ADMIN.SET_DRAINING (
   service_name  IN VARCHAR2,
   drain_timeout IN BINARY_INTEGER DEFAULT 300
   stop_option   IN VARCHAR2       DEFAULT 'NONE');
```

Parameters

Table 21-16 SET_DRAINING Procedure Parameters

Parameter	Description
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, draining is set only for that service.
DRAIN_TIMEOUT	Specify the time, in seconds, allowed for resource draining to be completed. Accepted values are an empty string (""), 0, or any positive integer.
STOP_OPTION	Specify the method of stopping the service. Available options are NONE, IMMEDIATE, or TRANSACTIONAL.

Examples

The following example illustrates how to set draining options for your service:

```
SQL> execute dbms_app_cont_admin.set_draining('TPURGENT', 300, 'IMMEDIATE');
```

21.2.15 SET_LOAD_BALANCING Procedure

This procedure configures load balancing options for your service.

Syntax

```
DBMS_APP_CONT_ADMIN.SET_LOAD_BALANCING (
   service_name IN VARCHAR2,
   goal         IN VARCHAR2);
```

Parameters

Table 21-17 SET_LOAD_BALANCING Procedure Parameters

Parameter	Description
SERVICE_NAME	Optional service name qualifier. If the optional service name qualifier is provided, load balancing is set only for the specified service.
GOAL	Load balancing goal. Possible values are: • CLBGOAL- Connection Load Balancing Goal. • RLBGOAL- Runtime Load Balancing Goal.

Examples

The following example illustrates how to set load balancing goal for your service:

```
SQL> execute dbms_app_cont_admin.set_load_balancing('TPURGENT', 'CLBGOAL');
```

22
DBMS_APP_CONT_REPORT

This procedure generates a report of ACCHK trace collected while `ACCHK_SET(TRUE)`.

This chapter contains the following topics:

- DBMS_APP_CONT_REPORT Security Model
- Summary of DBMS_APP_CONT_REPORT Subprograms

22.1 DBMS_APP_CONT_REPORT Security Model

Applications must have the `EXECUTE` privilege on the `DBMS_APP_CONT_REPORT` package.

22.2 Summary of DBMS_APP_CONT_REPORT Subprograms

This topic lists the `DBMS_APP_CONT_REPORT` subprograms in alphabetical order and briefly describes them.

Table 22-1 *DBMS_APP_CONT_REPORT Package Subprograms*

Subprogram	Description
ACCHK_REPORT Procedure	This procedure provides a report of your run. The data is derived from the `acchk` views.

22.2.1 ACCHK_REPORT Procedure

This procedure generates a report of your run. The data is derived from the `acchk` views.

Syntax

```
DBMS_APP_CONT_REPORT.ACCHK_REPORT (
   level        IN NUMBER    DEFAULT SUMMARY,
   service_name IN VARCHAR2  DEFAULT NULL,
   start_time   IN DATE      DEFAULT NULL,
   duration     IN NUMBER    DEFAULT NULL,
   source       IN NUMBER    DEFAULT FROM_TABLES);
```

Parameters

Table 22-2 ACCHK_REPORT Procedure Parameters

Parameter	Description
`level`	These are constants. • `SUMMARY`-summary report only (default) • `WARNING`-summary and warnings • `FULL`-summary and all workload capture, including where there are no warnings

Table 22-2 (Cont.) ACCHK_REPORT Procedure Parameters

Parameter	Description
service_name	The service name qualifier (Optional). The value NULL applies to all service_names.
start_time	Optional parameter for the start time to filter events and statistics. By default, starts at first available row in tables/trace with ACCHK.
duration	Optional parameter for the amount of time in seconds to filter events and statistics. By default, uses available rows in tables/traces starting from the start_time value. **Note:** If both start_time and duration are not specified, then all available rows in tables/traces are used for reporting.
source	Optional parameter for obtaining the statistics and events from trace files when capture was in read-only mode by providing the constant value DBMS_APP_CONT_REPORT.FROM_TRACES. DBMS_APP_CONT_REPORT.FROM_TABLES is the default value and uses the views CDB_ACCHK_STATISTICS and CDB_ACCHK_EVENTS for the report. There are two constants from DBMS_APP_CONT_REPORT: • FROM_TABLES- The procedure uses CDB_ACCHK_STATISTICS and CDB_ACCHK_EVENTS to show the report. • FROM_TRACES- The procedure uses traces to show the report.

Procedure Exceptions

The exceptions of DBMS_APP_CONT_REPORT.ACCHK_REPORT procedure are as follows:

- 20000- Invalid report level, use DBMS_APP_CONT_REPORT with SUMMARY or WARNING or FULL.
- 20000- Invalid duration, set a valid value.
- 20000- Invalid source, use DBMS_APP_CONT_REPORT with either FROM_TABLES or FROM_TRACES.

Example

To get the recorded statistics and events starting from 10 minutes ago:

```
execute dbms_app_cont_report.acchk_report(dbms_app_cont_report.FULL,
start_time => current_date -INTERVAL '10' MINUTE);
```

To get the recorded statistics and events starting from 10 minutes ago up to 420 seconds (7 minutes):

```
execute dbms_app_cont_report.acchk_report(dbms_app_cont_report.FULL,
start_time  => current_date -INTERVAL '10' MINUTE, duration => 420);
```

To get the recorded statistics and events of the first 60 seconds starting from the first recorded statistic or event:

```
execute dbms_app_cont_report.acchk_report(dbms_app_cont_report.FULL, duration
=> 60);
```

To get the report by using tables:

```
execute dbms_app_cont_report.acchk_report(source =>
dbms_app_cont_report.FROM_TABLES);
```

To get the report by using traces. Use this option when in read-only mode:

```
execute dbms_app_cont_report.acchk_report(source =>
dbms_app_cont_report.FROM_TRACES);
```

To get only the summary report:

```
execute dbms_app_cont_report.acchk_report(dbms_app_cont_report.SUMMARY);
```

23
DBMS_APPLICATION_INFO

Application developers can use the `DBMS_APPLICATION_INFO` package with Oracle Trace and the SQL trace facility to record names of executing modules or transactions in the database for later use when tracking the performance of various modules and debugging.

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Summary of DBMS_APPLICATION_INFO Subprograms

23.1 DBMS_APPLICATION_INFO Overview

Registering the application allows system administrators and performance tuning specialists to track performance by module. System administrators can also use this information to track resource use by module. When an application registers with the database, its name and actions are recorded in the `V$SESSION` and `V$SQLAREA` views.

23.2 DBMS_APPLICATION_INFO Security Model

No further privileges are required. The `DBMSAPIN.SQL` script is already run as a part of standard database creation.

> **Note:**
>
> The public synonym for `DBMS_APPLICATION_INFO` is not dropped before creation so that you can redirect the public synonym to point to your own package.

23.3 DBMS_APPLICATION_INFO Operational Notes

Your applications should set the name of the module and name of the action automatically each time a user enters that module. The module name could be the name of a form in an Oracle Forms application, or the name of the code segment in an Oracle Precompilers application. The action name should usually be the name or description of the current transaction within a module.

If you want to gather your own statistics based on module, you can implement a wrapper around this package by writing a version of this package in another schema that first gathers statistics and then calls the `SYS` version of the package. The public synonym for `DBMS_APPLICATION_INFO` can then be changed to point to the DBA's version of the package.

23.4 Summary of DBMS_APPLICATION_INFO Subprograms

This table lists and describes the DBMS_APPLICATION_INFO package subprograms.

Table 23-1 DBMS_APPLICATION_INFO Package Subprograms

Subprogram	Description
READ_CLIENT_INFO Procedure	Reads the value of the client_info field of the current session
READ_MODULE Procedure	Reads the values of the module and action fields of the current session
SET_ACTION Procedure	Sets the name of the current action within the current module
SET_CLIENT_INFO Procedure	Sets the client_info field of the session
SET_MODULE Procedure	Sets the name of the module that is currently running to a new module
SET_SESSION_LONGOPS Procedure	Sets a row in the V$SESSION_LONGOPS table

23.4.1 READ_CLIENT_INFO Procedure

This procedure reads the value of the client_info field of the current session.

Syntax

```
DBMS_APPLICATION_INFO.READ_CLIENT_INFO (
   client_info OUT VARCHAR2);
```

Parameters

Table 23-2 *READ_CLIENT_INFO Procedure Parameters*

Parameter	Description
client_info	Last client information value supplied to the SET_CLIENT_INFO procedure.

23.4.2 READ_MODULE Procedure

This procedure reads the values of the module and action fields of the current session.

Syntax

```
DBMS_APPLICATION_INFO.READ_MODULE (
   module_name OUT VARCHAR2,
   action_name OUT VARCHAR2);
```

Parameters

Table 23-3 READ_MODULE Procedure Parameters

Parameter	Description
module_name	Last value that the module name was set to by calling SET_MODULE.

Table 23-3 (Cont.) READ_MODULE Procedure Parameters

Parameter	Description
action_name	Last value that the action name was set to by calling SET_ACTION or SET_MODULE.

Usage Notes

Module and action names for a registered application can be retrieved by querying V$SQLAREA or by calling the READ_MODULE procedure. Client information can be retrieved by querying the V$SESSION view, or by calling the READ_CLIENT_INFO Procedure.

Examples

The following sample query illustrates the use of the MODULE and ACTION column of the V$SQLAREA.

```
SELECT sql_text, disk_reads, module, action
FROM v$sqlarea
WHERE module = 'add_employee';

SQL_TEXT DISK_READS MODULE ACTION
------------------- ---------- ------------------ ----------------
INSERT INTO emp 1 add_employee insert into emp
(ename, empno, sal, mgr, job, hiredate, comm, deptno)
VALUES
(name, next.emp_seq, manager, title, SYSDATE, commission, department)

1 row selected.
```

23.4.3 SET_ACTION Procedure

This procedure sets the name of the current action within the current module.

Syntax

```
DBMS_APPLICATION_INFO.SET_ACTION (
   action_name IN VARCHAR2);
```

Parameters

Table 23-4 SET_ACTION Procedure Parameters

Parameter	Description
action_name	The name of the current action within the current module. When the current action terminates, call this procedure with the name of the next action if there is one, or NULL if there is not. Names longer than 32 bytes are truncated.

Usage Notes

The action name should be descriptive text about the current action being performed. You should probably set the action name before the start of every transaction.

Set the transaction name to NULL after the transaction completes, so that subsequent transactions are logged correctly. If you do not set the transaction name to NULL, subsequent transactions may be logged with the previous transaction's name.

Example

The following is an example of a transaction that uses the registration procedure:

```
CREATE OR REPLACE PROCEDURE bal_tran (amt IN NUMBER(7,2)) AS
BEGIN

-- balance transfer transaction

   DBMS_APPLICATION_INFO.SET_ACTION(
      action_name => 'transfer from chk to sav');
   UPDATE chk SET bal = bal + :amt
      WHERE acct# = :acct;
   UPDATE sav SET bal = bal - :amt
      WHERE acct# = :acct;
   COMMIT;
   DBMS_APPLICATION_INFO.SET_ACTION(null);

END;
```

23.4.4 SET_CLIENT_INFO Procedure

This procedure supplies additional information about the client application.

Syntax

```
DBMS_APPLICATION_INFO.SET_CLIENT_INFO (
   client_info IN VARCHAR2);
```

Parameters

Table 23-5 SET_CLIENT_INFO Procedure Parameters

Parameter	Description
client_info	Supplies any additional information about the client application. This information is stored in the V$SESSION view. Information exceeding 64 bytes is truncated.

> **Note:**
>
> CLIENT_INFO is readable and writable by any user. For storing secured application attributes, you can use the application context feature.

23.4.5 SET_MODULE Procedure

This procedure sets the name of the current application or module.

Syntax

```
DBMS_APPLICATION_INFO.SET_MODULE (
   module_name IN VARCHAR2,
   action_name IN VARCHAR2);
```

Parameters

Table 23-6 SET_MODULE Procedure Parameters

Parameter	Description
module_name	Name of module that is currently running. When the current module terminates, call this procedure with the name of the new module if there is one, or NULL if there is not. Names longer than 48 bytes are truncated.
action_name	Name of current action within the current module. If you do not want to specify an action, this value should be NULL. Names longer than 32 bytes are truncated.

Usage Notes

Example

```
CREATE or replace PROCEDURE add_employee(
  name VARCHAR2,
  salary NUMBER,
  manager NUMBER,
  title VARCHAR2,
  commission NUMBER,
  department NUMBER) AS
BEGIN
  DBMS_APPLICATION_INFO.SET_MODULE(
    module_name => 'add_employee',
    action_name => 'insert into emp');
  INSERT INTO emp
    (ename, empno, sal, mgr, job, hiredate, comm, deptno)
    VALUES (name, emp_seq.nextval, salary, manager, title, SYSDATE,
            commission, department);
  DBMS_APPLICATION_INFO.SET_MODULE(null,null);
END;
```

23.4.6 SET_SESSION_LONGOPS Procedure

This procedure sets a row in the V$SESSION_LONGOPS view. This is a view that is used to indicate the on-going progress of a long running operation. Some Oracle functions, such as parallel execution and Server Managed Recovery, use rows in this view to indicate the status of, for example, a database backup.

Applications may use the SET_SESSION_LONGOPS procedure to advertise information on the progress of application specific long running tasks so that the progress can be monitored by way of the V$SESSION_LONGOPS view.

Syntax

```
DBMS_APPLICATION_INFO.SET_SESSION_LONGOPS (
   rindex      IN OUT BINARY_INTEGER,
   slno        IN OUT BINARY_INTEGER,
   op_name     IN     VARCHAR2       DEFAULT NULL,
   target      IN     BINARY_INTEGER DEFAULT 0,
   context     IN     BINARY_INTEGER DEFAULT 0,
   sofar       IN     NUMBER         DEFAULT 0,
   totalwork   IN     NUMBER         DEFAULT 0,
   target_desc IN     VARCHAR2       DEFAULT 'unknown target',
   units       IN     VARCHAR2       DEFAULT NULL)

set_session_longops_nohint constant BINARY_INTEGER := -1;
```

Parameters

Table 23-7 SET_SESSION_LONGOPS Procedure Parameters

Parameter	Description
rindex	A token which represents the v$session_longops row to update. Set this to set_session_longops_nohint to start a new row. Use the returned value from the prior call to reuse a row.
slno	Saves information across calls to set_session_longops: It is for internal use and should not be modified by the caller.
op_name	Specifies the name of the long running task. It appears as the OPNAME column of v$session_longops. The maximum length is 64 bytes.
target	Specifies the object that is being worked on during the long running operation. For example, it could be a table ID that is being sorted. It appears as the TARGET column of v$session_longops.
context	Any number the client wants to store. It appears in the CONTEXT column of v$session_longops.
sofar	Any number the client wants to store. It appears in the SOFAR column of v$session_longops. This is typically the amount of work which has been done so far.
totalwork	Any number the client wants to store. It appears in the TOTALWORK column of v$session_longops. This is typically an estimate of the total amount of work needed to be done in this long running operation.
target_desc	Specifies the description of the object being manipulated in this long operation. This provides a caption for the target parameter. This value appears in the TARGET_DESC field of v$session_longops. The maximum length is 32 bytes.
units	Specifies the units in which sofar and totalwork are being represented. It appears as the UNITS field of v$session_longops. The maximum length is 32 bytes.

Example

This example performs a task on 10 objects in a loop. As the example completes each object, Oracle updates V$SESSION_LONGOPS on the procedure's progress.

```
DECLARE
        rindex    BINARY_INTEGER;
        slno      BINARY_INTEGER;
```

```
  totalwork number;
  sofar     number;
  obj       BINARY_INTEGER;

BEGIN
  rindex := dbms_application_info.set_session_longops_nohint;
  sofar := 0;
  totalwork := 10;

  WHILE sofar < 10 LOOP
    -- update obj based on sofar
    -- perform task on object target

    sofar := sofar + 1;
    dbms_application_info.set_session_longops(rindex, slno,
      "Operation X", obj, 0, sofar, totalwork, "table", "tables");
  END LOOP;
END;
```

24
DBMS_APPLY_ADM

The `DBMS_APPLY_ADM` package provides subprograms to configure and manage Oracle Apply processes, XStream outbound servers, and XStream inbound servers.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Summary of DBMS_APPLY_ADM Subprograms

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about XStream outbound servers and inbound servers

24.1 DBMS_APPLY_ADM Overview

The `DBMS_APPLY_ADM` package provides interfaces to start, stop, and configure Oracle Apply processes, XStream outbound servers, and XStream inbound servers.

This package includes subprograms for configuring apply handlers, setting enqueue destinations for messages, and specifying execution directives for messages. This package also provides administrative subprograms that set the instantiation SCN for objects at a destination database. This package also includes subprograms for managing apply errors.

XStream inbound servers and outbound servers can be used in an XStream configuration in a multitenant container database (CDB). A CDB is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs).

Note:

- For simplicity, this chapter refers to apply processes, XStream outbound servers, and XStream inbound servers as **apply components**. This chapter identifies a specific type of apply component when necessary.
- Using XStream requires purchasing a license for the Oracle GoldenGate product.

See Also:

- *Oracle Database XStream Guide*
- *Oracle Database Concepts* for more information about CDBs and PDBs

24.2 DBMS_APPLY_ADM Security Model

Security on this package can be controlled by either granting `EXECUTE` on this package to selected users or roles, or by granting `EXECUTE_CATALOG_ROLE` to selected users or roles..

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

When the `DBMS_APPLY_ADM` package is used to manage an Oracle Replication configuration, it requires that the user is granted the privileges of an Oracle Replication administrator.

When the `DBMS_APPLY_ADM` package is used to manage an XStream configuration, it requires that the user is granted the privileges of an XStream administrator.

Note:

The user must be granted additional privileges to perform some administrative tasks using the subprograms in this package, such as setting an apply user. If additional privileges are required for a subprogram, then the privileges are documented in the section that describes the subprogram.

See Also:

Oracle Database XStream Guide for information about configuring an XStream administrator

24.3 DBMS_APPLY_ADM Deprecated Subprograms

The `NONE` value for the `commit_serialization` apply component parameter is deprecated. It is replaced by the `DEPENDENT_TRANSACTIONS` value.

> **Note:**
>
> Oracle recommends that you do not use deprecated apply component parameter values. Support for deprecated features is for backward compatibility only.

> **See Also:**
>
> SET_PARAMETER Procedure

24.4 Summary of DBMS_APPLY_ADM Subprograms

This table topic lists and describes the `DBMS_APPLY_ADM` subprograms.

Table 24-1 DBMS_APPLY_ADM Package Subprograms

Subprogram	Description
ALTER_APPLY Procedure	Alters an apply component
CLEAR_KEY_COLUMNS Procedure	Removes the key columns that were used as the substitute primary key by the `SET_KEY_COLUMNS` procedure
COMPARE_OLD_VALUES Procedure	Specifies whether to compare the old value of one or more columns in a row logical change record (row LCR) with the current value of the corresponding columns at the destination site during apply
CREATE_APPLY Procedure	Creates an apply component
CREATE_OBJECT_DEPENDENCY Procedure	Creates an object dependency
DELETE_ALL_ERRORS Procedure	Deletes all the error transactions for the specified apply component
DELETE_ERROR Procedure	Deletes the specified error transaction
DROP_APPLY Procedure	Drops an apply component
DROP_OBJECT_DEPENDENCY Procedure	Drops an object dependency
EXECUTE_ALL_ERRORS Procedure	Reexecutes the error transactions for the specified apply component
EXECUTE_ERROR Procedure	Reexecutes the specified error transaction
GET_ERROR_MESSAGE Function	Returns the message payload from the error queue for the specified message number and transaction identifier
HANDLE_COLLISIONS	Enables or disables basic conflict resolution for an apply process and a table

Table 24-1 (Cont.) DBMS_APPLY_ADM Package Subprograms

Subprogram	Description
SET_DML_CONFLICT_HANDLER Procedure	Adds, modifies, or removes a prebuilt DML conflict handler for `INSERT`, `UPDATE`, or `DELETE` conflicts on the specified object
SET_DML_HANDLER Procedure	Sets a user procedure as a procedure DML handler for a specified operation on a specified database object for a single apply component or for all apply components in the database
SET_ENQUEUE_DESTINATION Procedure	Sets the queue where the apply component automatically enqueues a message that satisfies the specified rule
SET_EXECUTE Procedure	Specifies whether a message that satisfies the specified rule is executed by an apply component
SET_GLOBAL_INSTANTIATION_SCN Procedure	Records the specified instantiation SCN for the specified source database and, optionally, for the schemas at the source database and the tables owned by these schemas
SET_KEY_COLUMNS Procedures	Records the set of columns to be used as the substitute primary key for local apply purposes and removes existing substitute primary key columns for the specified object if they exist
SET_PARAMETER Procedure	Sets an apply parameter to the specified value
SET_REPERROR_HANDLER Procedure	Specifies how a particular error is handled based on its error number
SET_SCHEMA_INSTANTIATION_SCN Procedure	Records the specified instantiation SCN for the specified schema in the specified source database and, optionally, for the tables owned by the schema at the source database
SET_TABLE_INSTANTIATION_SCN Procedure	Records the specified instantiation SCN for the specified table in the specified source database
SET_UPDATE_CONFLICT_HANDLER Procedure	Adds, updates, or drops an update conflict handler for the specified object
SET_VALUE_DEPENDENCY Procedure	Sets or removes a value dependency
START_APPLY Procedure	Directs the apply component to start applying messages
STOP_APPLY Procedure	Stops the apply component from applying any messages and rolls back any unfinished transactions being applied

Note:

All procedures commit unless specified otherwise. However, the `GET_ERROR_MESSAGE` function does not commit.

24.4.1 ALTER_APPLY Procedure

This procedure alters an apply component.

Syntax

```
DBMS_APPLY_ADM.ALTER_APPLY(
   apply_name                IN  VARCHAR2,
   rule_set_name             IN  VARCHAR2  DEFAULT NULL,
   remove_rule_set           IN  BOOLEAN   DEFAULT FALSE,
   message_handler           IN  VARCHAR2  DEFAULT NULL
   remove_message_handler    IN  BOOLEAN   DEFAULT FALSE,
   ddl_handler               IN  VARCHAR2  DEFAULT NULL,
   remove_ddl_handler        IN  BOOLEAN   DEFAULT FALSE,
   apply_user                IN  VARCHAR2  DEFAULT NULL,
   apply_tag                 IN  RAW       DEFAULT NULL,
   remove_apply_tag          IN  BOOLEAN   DEFAULT FALSE,
   precommit_handler         IN  VARCHAR2  DEFAULT NULL,
   remove_precommit_handler  IN  BOOLEAN   DEFAULT FALSE,
   negative_rule_set_name    IN  VARCHAR2  DEFAULT NULL,
   remove_negative_rule_set  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 24-2 ALTER_APPLY Procedure Parameters

Parameter	Description
apply_name	The name of the apply component being altered. You must specify the name of an existing apply component. Do not specify an owner.
rule_set_name	The name of the positive rule set for the apply component. The positive rule set contains the rules that instruct the apply component to apply messages. If you want to use a positive rule set for the apply component, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a positive rule set in the hr schema named job_apply_rules, enter hr.job_apply_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. If you specify NULL and the remove_rule_set parameter is set to FALSE, then this procedure retains any existing positive rule set for the specified apply component. If you specify NULL and the remove_rule_set parameter is set to TRUE, then this procedure removes any existing positive rule set from the specified apply component.

Table 24-2 (Cont.) ALTER_APPLY Procedure Parameters

Parameter	Description
remove_rule_set	If TRUE, then the procedure removes the positive rule set for the specified apply component. If you remove the positive rule set for an apply component, and the apply component does not have a negative rule set, then the apply component dequeues all messages in its queue. If you remove the positive rule set for an apply component, and a negative rule set exists for the apply component, then the apply component dequeues all messages in its queue that are not discarded by the negative rule set. If FALSE, then the procedure retains the positive rule set for the specified apply component. If the rule_set_name parameter is non-NULL, then this parameter should be set to FALSE.
message_handler	A user-defined procedure that processes non-LCR messages in the queue for the apply component. See "Usage Notes" in the CREATE_APPLY Procedure for more information about a message handler procedure.
remove_message_handler	If TRUE, then the procedure removes the message handler for the specified apply component. If FALSE, then the procedure retains any message handler for the specified apply component. If the message_handler parameter is non-NULL, then this parameter should be set to FALSE.
ddl_handler	A user-defined procedure that processes DDL logical change records (DDL LCRs) in the queue for the apply component. All applied DDL LCRs commit automatically. Therefore, if a DDL handler calls the EXECUTE member procedure of a DDL LCR, then a commit is performed automatically. See "Usage Notes" in the CREATE_APPLY Procedure for more information about a DDL handler procedure.
remove_ddl_handler	If TRUE, then the procedure removes the DDL handler for the specified apply component. If FALSE, then the procedure retains any DDL handler for the specified apply component. If the ddl_handler parameter is non-NULL, then this parameter should be set to FALSE.

Table 24-2 (Cont.) ALTER_APPLY Procedure Parameters

Parameter	Description
apply_user	The user in whose security domain an apply component dequeues messages that satisfy its rule sets, applies messages directly to database objects, runs custom rule-based transformations, and runs apply handlers. If NULL, then the apply user is not changed. If a non-NULL value is specified to change the apply user, then the user who invokes the ALTER_APPLY procedure must be granted the DBA role. Only the SYS user can set the apply_user to SYS. If you change the apply user, then this procedure grants the new apply user dequeue privilege on the queue used by the apply component. It also configures the user as a secure queue user of the queue. In addition to the privileges granted by this procedure, you also should grant the following privileges to the apply user: • The necessary privileges to perform DML and DDL changes on the apply objects • EXECUTE privilege on the rule sets used by the apply component • EXECUTE privilege on all rule-based transformation functions used in the rule set • EXECUTE privilege on all apply handler procedures These privileges can be granted directly to the apply user, or they can be granted through roles. In addition, the apply user must be granted the EXECUTE privilege on all packages, including Oracle-supplied packages, that are invoked in subprograms run by the apply component. These privileges must be granted directly to the apply user. They cannot be granted through roles. By default, this parameter is set to the user who created the apply component by running either the CREATE_APPLY procedure in this package. **Note:** If the apply user for an apply component is dropped using DROP USER. . . CASCADE, then the apply component is also dropped automatically.
apply_tag	A binary tag that is added to redo entries generated by the specified apply component. The tag is a binary value that can be used to track LCRs. The tag is relevant only if a capture process at the database where the apply component is running captures changes made by the apply component. If so, then the captured changes include the tag specified by this parameter. If NULL, the default, then the apply tag for the apply component is not changed. The following is an example of a tag with a hexadecimal value of 17: HEXTORAW('17')

Table 24-2 (Cont.) ALTER_APPLY Procedure Parameters

Parameter	Description
remove_apply_tag	If TRUE, then the procedure sets the apply tag for the specified apply component to NULL, and the apply component generates redo entries with NULL tags. If FALSE, then the procedure retains any apply tag for the specified apply component. If the apply_tag parameter is non-NULL, then this parameter should be set to FALSE.
precommit_handler	A user-defined procedure that can receive internal commit directives in the queue for the apply component before they are processed by the apply component. Typically, precommit handlers are used for auditing commit information for transactions processed by an apply component. An internal commit directive is enqueued in the following ways: • When a capture process captures row LCRs, the capture process enqueues the commit directive for the transaction that contains the row LCRs. • When a user or application enqueues messages and then issues a COMMIT statement, the commit directive is enqueued automatically. For a captured row LCR, a commit directive contains the commit SCN of the transaction from the source database. For a user message, the commit SCN is generated by the apply component. The precommit handler procedure must conform to the following restrictions: • Any work that commits must be an autonomous transaction. • Any rollback must be to a named savepoint created in the procedure. If a precommit handler raises an exception, then the entire apply transaction is rolled back, and all of the messages in the transaction are moved to the error queue. See "Usage Notes" in the CREATE_APPLY Procedure for more information about a precommit handler procedure.
remove_precommit_handler	If TRUE, then the procedure removes the precommit handler for the specified apply component. If FALSE, then the procedure retains any precommit handler for the specified apply component. If the precommit_handler parameter is non-NULL, then this parameter should be set to FALSE.

Table 24-2 (Cont.) ALTER_APPLY Procedure Parameters

Parameter	Description
negative_rule_set_name	The name of the negative rule set for the apply component. The negative rule set contains the rules that instruct the apply component to discard messages. If you want to use a negative rule set for the apply component, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a negative rule set in the hr schema named neg_apply_rules, enter hr.neg_apply_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. If you specify NULL and the remove_negative_rule_set parameter is set to FALSE, then the procedure retains any existing negative rule set. If you specify NULL and the remove_negative_rule_set parameter is set to TRUE, then the procedure removes any existing negative rule set. If you specify both a positive and a negative rule set for an apply component, then the negative rule set is always evaluated first.
remove_negative_rule_set	If TRUE, then the procedure removes the negative rule set for the specified apply component. If you remove the negative rule set for an apply component, and the apply component does not have a positive rule set, then the apply component dequeues all messages in its queue. If you remove the negative rule set for an apply component, and a positive rule set exists for the apply component, then the apply component dequeues all messages in its queue that are not discarded by the positive rule set. If FALSE, then the procedure retains the negative rule set for the specified apply component. If the negative_rule_set_name parameter is non-NULL, then this parameter should be set to FALSE.

Usage Notes

The following usage notes apply to this procedure:

- Automatic Restart of Apply Components
- The ALTER_APPLY Procedure and XStream Outbound Servers
- The ALTER_APPLY Procedure and XStream Inbound Servers

Automatic Restart of Apply Components

An apply component is stopped and restarted automatically when you change the value of one or more of the following ALTER_APPLY procedure parameters:

- message_handler
- ddl_handler
- apply_user
- apply_tag
- precommit_handler

The ALTER_APPLY Procedure and XStream Outbound Servers

The following usage notes apply to this procedure and XStream outbound servers:

- The `apply_user` parameter can change the connect user for an outbound server.
- You cannot specify an apply handler for an outbound server. An outbound server ignores the settings for the following parameters: `message_handler`, `ddl_handler`, and `precommit_handler`.

  The client application can perform custom processing of the LCRs instead if necessary.
- An outbound server cannot set an apply tag for the changes it processes. An outbound server ignores the setting for the `apply_tag` parameter.

The ALTER_APPLY Procedure and XStream Inbound Servers

Inbound servers can use apply handlers and process only DML and DDL LCRs. Therefore, inbound servers ignore message handlers specified in the `message_handler` parameter.

24.4.2 CLEAR_KEY_COLUMNS Procedure

This procedure removes the key columns that were used as the substitute primary key by the `SET_KEY_COLUMNS` procedure.

Syntax

```
DBMS_APPLY_ADM.CLEAR_KEY_COLUMNS(
   apply_name  IN  VARCHAR2  DEFAULT NULL);
```

Parameter

Table 24-3 CLEAR_KEY_COLUMNS Procedure Parameter

Parameter	Description
`apply_name`	The apply component name.

Usage Notes

The following usage notes apply to this procedure:

The CLEAR_KEY_COLUMNS Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The CLEAR_KEY_COLUMNS Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

The CLEAR_KEY_COLUMNS Procedure and CDBs

This procedure removes the columns that are used as a substitute primary key. You must perform the `CLEAR_KEY_COLUMNS` procedure in the appropriate PDB.

24.4.3 COMPARE_OLD_VALUES Procedure

This procedure specifies whether to compare the old values of one or more columns in a row logical change record (row LCR) with the current values of the corresponding columns at the destination site during apply.

This procedure is relevant only for `UPDATE` and `DELETE` operations because only these operations result in old column values in row LCRs. The default is to compare old values for all columns.

This procedure is overloaded. The `column_list` and `column_table` parameters are mutually exclusive.

Syntax

```
DBMS_APPLY_ADM.COMPARE_OLD_VALUES(
   object_name         IN VARCHAR2,
   column_list         IN VARCHAR2,
   operation           IN VARCHAR2 DEFAULT 'UPDATE',
   compare             IN BOOLEAN  DEFAULT TRUE,
   apply_database_link IN VARCHAR2 DEFAULT NULL);

DBMS_APPLY_ADM.COMPARE_OLD_VALUES(
   object_name         IN VARCHAR2,
   column_table        IN DBMS_UTILITY.LNAME_ARRAY,
   operation           IN VARCHAR2 DEFAULT 'UPDATE',
   compare             IN BOOLEAN  DEFAULT TRUE,
   apply_database_link IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 24-4 COMPARE_OLD_VALUES Procedure Parameters

Parameter	Description
`object_name`	The name of the source table specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default.
`column_list`	A comma-delimited list of column names in the table. There must be no spaces between entries. Specify `*` to include all nonkey columns.
`column_table`	A PL/SQL associative array of type `DBMS_UTILITY.LNAME_ARRAY` that contains names of columns in the table. The first column name should be at position 1, the second at position 2, and so on. The table does not need to be `NULL` terminated.
`operation`	The name of the operation, which can be specified as: • `UPDATE` for `UPDATE` operations • `DELETE` for `DELETE` operations • `*` for both `UPDATE` and `DELETE` operations
`compare`	If `compare` is `TRUE`, the old values of the specified columns are compared during apply. If `compare` is `FALSE`, the old values of the specified columns are not compared during apply.

Table 24-4 (Cont.) COMPARE_OLD_VALUES Procedure Parameters

Parameter	Description
apply_database_link	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database is a non-Oracle database.

Usage Notes

The following usage notes apply to this procedure:

- Conflict Detection
- The COMPARE_OLD_VALUES Procedure and XStream Outbound Servers
- The COMPARE_OLD_VALUES Procedure and XStream Inbound Servers

Conflict Detection

By default, an apply component uses the old column values in a row LCR to detect conflicts. You can choose not to compare old column values to avoid conflict detection for specific tables. For example, if you do not want to compare the old values for a set of columns during apply, then, using the `COMPARE_OLD_VALUES` procedure, specify the set of columns in the `column_list` or `column_table` parameter, and set the `compare` parameter to `FALSE`.

In addition, when the `compare_key_only` apply component parameter is set to `Y`, automatic conflict detection is disabled, and the apply component only uses primary key and unique key columns to identify the table row for a row LCR. When the `compare_key_only` apply component parameter is set to `N`, automatic conflict detection is enabled, and the apply component uses all of the old values in a row LCR to identify the table row for a row LCR.

Note:

- An apply component compares old values for non-key columns when they are present in a row LCR and when the apply component parameter `compare_key_only` is set to `N`.
- This procedure raises an error if a key column is specified in `column_list` or `column_table` and the `compare` parameter is set to `FALSE`.

See Also:

SET_PARAMETER Procedure for more information about the `compare_key_only` apply component parameter

The COMPARE_OLD_VALUES Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The COMPARE_OLD_VALUES Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.4 CREATE_APPLY Procedure

This procedure creates an apply component.

Syntax

```
DBMS_APPLY_ADM.CREATE_APPLY(
   queue_name              IN  VARCHAR2,
   apply_name              IN  VARCHAR2,
   rule_set_name           IN  VARCHAR2  DEFAULT NULL,
   message_handler         IN  VARCHAR2  DEFAULT NULL,
   ddl_handler             IN  VARCHAR2  DEFAULT NULL,
   apply_user              IN  VARCHAR2  DEFAULT NULL,
   apply_database_link     IN  VARCHAR2  DEFAULT NULL,
   apply_tag               IN  RAW       DEFAULT '00',
   apply_captured          IN  BOOLEAN   DEFAULT FALSE,
   precommit_handler       IN  VARCHAR2  DEFAULT NULL,
   negative_rule_set_name  IN  VARCHAR2  DEFAULT NULL,
   source_database         IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-5 CREATE_APPLY Procedure Parameters

Parameter	Description
queue_name	The name of the queue from which the apply component dequeues messages. You must specify an existing queue in the form [*schema_name.*]*queue_name*. For example, to specify a queue in the hr schema named streams_queue, enter hr.streams_queue. If the schema is not specified, then the current user is the default. **Note:** The queue_name setting cannot be altered after the apply component is created.
apply_name	The name of the apply component being created. A NULL specification is not allowed. Do not specify an owner. The specified name must not match the name of an existing apply component or messaging client. **Note:** The apply_name setting cannot be altered after the apply component is created.
rule_set_name	The name of the positive rule set for the apply component. The positive rule set contains the rules that instruct the apply component to apply messages. If you want to use a positive rule set for the apply component, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a positive rule set in the hr schema named job_apply_rules, enter hr.job_apply_rules. If the schema is not specified, then the current user is the default. If you specify NULL, and no negative rule set is specified, then the apply component applies either all captured messages or all messages in the persistent queue, depending on the setting of the apply_captured parameter. An error is returned if the specified rule set does not exist.

Table 24-5 (Cont.) CREATE_APPLY Procedure Parameters

Parameter	Description
message_handler	A user-defined procedure that processes non-LCR messages in the queue for the apply component. See "Usage Notes" for more information about a message handler procedure.
ddl_handler	A user-defined procedure that processes DDL logical change record (DDL LCRs) in the queue for the apply component. All applied DDL LCRs commit automatically. Therefore, if a DDL handler calls the EXECUTE member procedure of a DDL LCR, then a commit is performed automatically. See "Usage Notes" for more information about a DDL handler procedure.
apply_user	The user who applies all DML and DDL changes that satisfy the apply component rule sets and who runs user-defined apply handlers. If NULL, then the user who runs the CREATE_APPLY procedure is used. The apply user is the user in whose security domain an apply component dequeues messages that satisfy its rule sets, applies messages directly to database objects, runs custom rule-based transformations configured for apply component rules, and runs apply handlers configured for the apply component. This user must have the necessary privileges to apply changes. This procedure grants the apply user dequeue privilege on the queue used by the apply component and configures the user as a secure queue user of the queue. In addition to the privileges granted by this procedure, you also should grant the following privileges to the apply user: • The necessary privileges to perform DML and DDL changes on the apply objects • EXECUTE privilege on the rule sets used by the apply component • EXECUTE privilege on all rule-based transformation functions used in the rule set • EXECUTE privilege on all apply handler procedures These privileges can be granted directly to the apply user, or they can be granted through roles. In addition, the apply user must be granted EXECUTE privilege on all packages, including Oracle-supplied packages, that are invoked in subprograms run by the apply component. These privileges must be granted directly to the apply user. They cannot be granted through roles. You can use the DBMS_XSTREAM_AUTH package to grant and revoke administrative privileges in XStream configuration. These packages do not configure the necessary privileges to perform DML or DDL changes on the apply objects. **Note:** If the apply user for an apply component is dropped using DROP USER . . . CASCADE, then the apply component is also dropped automatically. See "Usage Notes" for more information about this parameter.

Table 24-5 (Cont.) CREATE_APPLY Procedure Parameters

Parameter	Description
apply_database_link	The database at which the apply component applies messages. This parameter is used by an apply component when applying changes from Oracle to non-Oracle systems, such as Sybase. Set this parameter to NULL to specify that the apply component applies messages at the local database. **Note:** The apply_database_link setting cannot be altered after the apply component is created.
apply_tag	A binary tag that is added to redo entries generated by the specified apply component. The tag is a binary value that can be used to track LCRs. The tag is relevant only if a capture process at the database where the apply component is running captures changes made by the apply component. If so, then the captured changes include the tag specified by this parameter. By default, the tag for an apply component is the hexadecimal equivalent of '00' (double zero). The following is an example of a tag with a hexadecimal value of 17: HEXTORAW('17') If NULL, then the apply component generates redo entries with NULL tags.
apply_captured	Either TRUE or FALSE. If TRUE, then the apply component applies only the captured LCRs in the queue. Captured LCRs are LCRs that were captured by an Oracle Replication capture process. If FALSE, then the apply component applies only the messages in a persistent queue. These are messages that were not captured by an Oracle Replication capture process, such as persistent LCRs or user messages. To apply both captured LCRs and messages in a persistent queue, you must create at least two apply components. **Note:** The apply_captured setting cannot be altered after the apply component is created.

Table 24-5 (Cont.) CREATE_APPLY Procedure Parameters

Parameter	Description
`precommit_handler`	A user-defined procedure that can receive internal commit directives in the queue for the apply component before they are processed by the apply component. Typically, precommit handlers are used for auditing commit information for transactions processed by an apply component. An internal commit directive is enqueued in the following ways: • When a capture process captures row LCRs, the capture process enqueues the commit directive for the transaction that contains the row LCRs. • When a synchronous capture captures row LCRs, the persistent LCRs that were enqueued by the synchronous capture are organized into a message group. The synchronous capture records the transaction identifier in each persistent LCR in a transaction. • When a user or application enqueues messages and then issues a `COMMIT` statement, the commit directive is enqueued automatically. For a row LCR captured by a capture process or synchronous capture, a commit directive contains the commit SCN of the transaction from the source database. For a message enqueued by a user or application, the commit SCN is generated by the apply component. The precommit handler procedure must conform to the following restrictions: • Any work that commits must be an autonomous transaction. • Any rollback must be to a named savepoint created in the procedure. If a precommit handler raises an exception, then the entire apply transaction is rolled back, and all of the messages in the transaction are moved to the error queue. See "Usage Notes" for more information about a precommit handler procedure.
`negative_rule_set_name`	The name of the negative rule set for the apply component. The negative rule set contains the rules that instruct the apply component to discard messages. If you want to use a negative rule set for the apply component, then you must specify an existing rule set in the form `[`*`schema_name`*`.]`*`rule_set_name`*. For example, to specify a negative rule set in the `hr` schema named `neg_apply_rules`, enter `hr.neg_apply_rules`. If the schema is not specified, then the current user is the default. If you specify `NULL`, and no positive rule set is specified, then the apply component applies either all captured LCRs or all of the messages in the persistent queue, depending on the setting of the `apply_captured` parameter. An error is returned if the specified rule set does not exist. If you specify both a positive and a negative rule set for an apply component, then the negative rule set is always evaluated first.

Table 24-5 (Cont.) CREATE_APPLY Procedure Parameters

Parameter	Description
source_database	The global name of the source database for the changes that will be applied by the apply component. The source database is the database where the changes originated. If an apply component applies captured messages, then the apply component can apply messages from only one capture process at one source database. If NULL, then the source database name of the first LCR received by the apply component is used for the source database. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is NET, then the procedure specifies DBS1.NET automatically. The rules in the apply component rule sets determine which messages are dequeued by the apply component. If the apply component dequeues an LCR with a source database that is different than the source database for the apply component, then an error is raised. You can determine the source database for an apply component by querying the DBA_APPLY_PROGRESS data dictionary view.

Usage Notes

The following sections describe usage notes for this procedure:

- DBA Role Requirement
- Handler Procedure Names
- Message Handler and DDL Handler Procedure
- Precommit Handler Procedure
- The CREATE_APPLY Procedure and XStream Outbound Servers
- The CREATE_APPLY Procedure and XStream Inbound Servers

DBA Role Requirement

If the user who invokes this procedure is different from the user specified in the apply_user parameter, then the invoking user must be granted the DBA role. If the user who invokes this procedure is the same as the user specified in the apply_user parameter, then the DBA role is not required for the invoking user. Only the SYS user can set the apply_user to SYS.

Handler Procedure Names

For the message_handler, ddl_handler, and precommit_handler parameters, specify an existing procedure in one of the following forms:

- *[schema_name.]procedure_name*
- *[schema_name.]package_name.procedure_name*

If the procedure is in a package, then the package_name must be specified. For example, to specify a procedure in the apply_pkg package in the hr schema named process_ddls, enter hr.apply_pkg.process_ddls. An error is returned if the specified procedure does not exist.

The user who invokes the CREATE_APPLY procedure must have EXECUTE privilege on a specified handler procedure. Also, if the *schema_name* is not specified, then the user who invokes the CREATE_APPLY procedure is the default.

Message Handler and DDL Handler Procedure

The procedure specified in both the message_handler parameter and the ddl_handler parameter must have the following signature:

```
PROCEDURE handler_procedure (
   parameter_name  IN  ANYDATA);
```

Here, *handler_procedure* stands for the name of the procedure and *parameter_name* stands for the name of the parameter passed to the procedure. For the message handler, the parameter passed to the procedure is a ANYDATA encapsulation of a user message. For the DDL handler procedure, the parameter passed to the procedure is a ANYDATA encapsulation of a DDL LCR.

> **See Also:**
>
> Logical Change Record TYPEs for information about DDL LCRs

Precommit Handler Procedure

The procedure specified in the precommit_handler parameter must have the following signature:

```
PROCEDURE handler_procedure (
   parameter_name   IN  NUMBER);
```

Here, *handler_procedure* stands for the name of the procedure and *parameter_name* stands for the name of the parameter passed to the procedure. The parameter passed to the procedure is the commit SCN of a commit directive.

The CREATE_APPLY Procedure and XStream Outbound Servers

This procedure cannot create an XStream outbound server. To create an XStream outbound server, use the DBMS_XSTREAM_ADM package.

The CREATE_APPLY Procedure and XStream Inbound Servers

The following usage notes apply to this procedure and XStream inbound servers:

- The CREATE_APPLY procedure always creates an apply process. The apply process remains an apply process if it receives messages from a source other than an XStream client application, such as a capture process. The apply process can become an inbound server if an XStream client application attaches to it before it receives messages from any other source. After the initial contact, an apply process cannot be changed into an inbound server, and an inbound server cannot be changed into an apply process.
- When creating an inbound server using the CREATE_APPLY procedure, set the apply_captured parameter to TRUE. Inbound servers only process LCRs captured by a capture process.
- Inbound servers can use apply handlers. Inbound servers process only DML and DDL LCRs. Therefore, inbound servers ignore message handlers specified in the message_handler parameter.

24.4.5 CREATE_OBJECT_DEPENDENCY Procedure

This procedure creates an object dependency. An object dependency is a virtual dependency definition that defines a parent-child relationship between two objects at a destination database.

An apply component schedules execution of transactions that involve the child object after all transactions with a lower commit system change number (commit SCN) that involve the parent object have been committed. An apply component uses the object identifier of the objects in the logical change records (LCRs) to detect dependencies. The apply component does not use column values in the LCRs to detect dependencies.

> **Note:**
>
> An error is raised if `NULL` is specified for either of the procedure parameters.

> **See Also:**
>
> DROP_OBJECT_DEPENDENCY Procedure

Syntax

```
DBMS_APPLY_ADM.CREATE_OBJECT_DEPENDENCY(
   object_name         IN  VARCHAR2,
   parent_object_name  IN  VARCHAR2);
```

Parameters

Table 24-6 CREATE_OBJECT_DEPENDENCY Procedure Parameters

Parameter	Description
`object_name`	The name of the child database object, specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default.
`parent_object_name`	The name of the parent database object, specified as `[schema_name.]object_name`. For example, `hr.departments`. If the schema is not specified, then the current user is the default.

Usage Notes

The following usage notes apply to this procedure:

- The CREATE_OBJECT_DEPENDENCY Procedure and XStream Outbound Servers
- The CREATE_OBJECT_DEPENDENCY Procedure and XStream Inbound Servers

The CREATE_OBJECT_DEPENDENCY Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The CREATE_OBJECT_DEPENDENCY Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.6 DELETE_ALL_ERRORS Procedure

This procedure deletes all the error transactions for the specified apply component.

Syntax

```
DBMS_APPLY_ADM.DELETE_ALL_ERRORS(
   apply_name  IN  VARCHAR2 DEFAULT NULL);
```

Parameter

Table 24-7 DELETE_ALL_ERRORS Procedure Parameter

Parameter	Description
`apply_name`	The name of the apply component that raised the errors while processing the transactions. Do not specify an owner. If `NULL`, then all error transactions for all apply components are deleted.

Usage Notes

The following usage notes apply to this procedure:

- The DELETE_ALL_ERRORS Procedure and XStream Outbound Servers
- The DELETE_ALL_ERRORS Procedure and XStream Inbound Servers

The DELETE_ALL_ERRORS Procedure and XStream Outbound Servers

Outbound servers do not enqueue error transactions into an error queue. This procedure has no effect on XStream outbound servers.

The DELETE_ALL_ERRORS Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.7 DELETE_ERROR Procedure

This procedure deletes the specified error transaction.

Syntax

```
DBMS_APPLY_ADM.DELETE_ERROR(
   local_transaction_id  IN  VARCHAR2);
```

Parameter

Table 24-8 DELETE_ERROR Procedure Parameter

Parameter	Description
`local_transaction_id`	The identification number of the error transaction to delete. If the specified transaction does not exist in the error queue, then an error is raised.

Usage Notes

The following usage notes apply to this procedure:

The DELETE_ERROR Procedure and XStream Outbound Servers

Outbound servers do not enqueue error transactions into an error queue. This procedure has no effect on XStream outbound servers.

The DELETE_ERROR Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.8 DROP_APPLY Procedure

This procedure drops an apply component.

Syntax

```
DBMS_APPLY_ADM.DROP_APPLY(
   apply_name             IN  VARCHAR2,
   drop_unused_rule_sets  IN  BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 24-9 DROP_APPLY Procedure Parameters

Parameter	Description
`apply_name`	The name of the apply component being dropped. You must specify an existing apply component name. Do not specify an owner.
`drop_unused_rule_sets`	If `TRUE`, then the procedure drops any rule sets, positive and negative, used by the specified apply component if these rule sets are not used by any other Oracle Replication or XStream component. These components include capture processes, propagations, apply processes, inbound servers, and messaging clients. If this procedure drops a rule set, then this procedure also drops any rules in the rule set that are not in another rule set. If `FALSE`, then the procedure does not drop the rule sets used by the specified apply component, and the rule sets retain their rules.

Usage Notes

The following usage notes apply to this procedure:

- The DROP_APPLY Procedure and Rules
- The DROP_APPLY Procedure and XStream Outbound Servers
- The DROP_APPLY Procedure and XStream Inbound Servers

The DROP_APPLY Procedure and Rules

When you use this procedure to drop an apply component, information about rules created for the apply component is removed from the data dictionary views for rules. Information about such a rule is removed even if the rule is not in either the positive or negative rule set for the apply component. The following are the data dictionary views for rules:

- `ALL_STREAMS_GLOBAL_RULES`

- DBA_STREAMS_GLOBAL_RULES
- ALL_STREAMS_SCHEMA_RULES
- DBA_STREAMS_SCHEMA_RULES
- ALL_STREAMS_TABLE_RULES
- DBA_STREAMS_TABLE_RULES

The DROP_APPLY Procedure and XStream Outbound Servers

When the DROP_APPLY procedure is executed on an outbound server, it runs the DROP_OUTBOUND procedure in the DBMS_XSTREAM_ADM package. Therefore, it might also drop the outbound server's capture process and queue.

The DROP_APPLY Procedure and XStream Inbound Servers

When the DROP_APPLY procedure is executed on an inbound server, it runs the DROP_INBOUND procedure in the DBMS_XSTREAM_ADM package. Therefore, it might also drop the inbound server's queue.

24.4.9 DROP_OBJECT_DEPENDENCY Procedure

This procedure drops an object dependency. An object dependency is a virtual dependency definition that defines a parent-child relationship between two objects at a destination database.

> **Note:**
>
> - An error is raised if an object dependency does not exist for the specified database objects.
> - An error is raised if NULL is specified for either of the procedure parameters.

> **See Also:**
>
> CREATE_OBJECT_DEPENDENCY Procedure

Syntax

```
DBMS_APPLY_ADM.DROP_OBJECT_DEPENDENCY(
   object_name         IN  VARCHAR2,
   parent_object_name  IN  VARCHAR2);
```

Parameters

Table 24-10 DROP_OBJECT_DEPENDENCY Procedure Parameters

Parameter	Description
object_name	The name of the child database object, specified as [*schema_name.*]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default.

Table 24-10 (Cont.) DROP_OBJECT_DEPENDENCY Procedure Parameters

Parameter	Description
parent_object_name	The name of the parent database object, specified as [*schema_name.*]*object_name*. For example, hr.departments. If the schema is not specified, then the current user is the default.

Usage Notes

The following usage notes apply to this procedure:

The DROP_OBJECT_DEPENDENCY Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The DROP_OBJECT_DEPENDENCY Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.10 EXECUTE_ALL_ERRORS Procedure

This procedure re-executes the error transactions in the error queue for the specified apply component.

The transactions are re-executed in commit SCN order. Error re-execution stops if an error is raised.

Syntax

```
DBMS_APPLY_ADM.EXECUTE_ALL_ERRORS(
   apply_name       IN  VARCHAR2  DEFAULT NULL,
   execute_as_user  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 24-11 EXECUTE_ALL_ERRORS Procedure Parameters

Parameter	Description
apply_name	The name of the apply component that raised the errors while processing the transactions. Do not specify an owner. If NULL, then all error transactions for all apply components are re-executed.
execute_as_user	If TRUE, then the procedure re-executes the transactions in the security context of the current user. If FALSE, then the procedure re-executes each transaction in the security context of the original receiver of the transaction. The original receiver is the user who was processing the transaction when the error was raised. The DBA_APPLY_ERROR data dictionary view lists the original receiver for each error transaction. The user who executes the transactions must have privileges to perform DML and DDL changes on the apply objects and to run any apply handlers. This user must also have dequeue privileges on the queue used by the apply component.

Usage Notes

The following usage notes apply to this procedure:

- The EXECUTE_ALL_ERRORS Procedure and XStream Outbound Servers
- The EXECUTE_ALL_ERRORS Procedure and XStream Inbound Servers

The EXECUTE_ALL_ERRORS Procedure and XStream Outbound Servers

Outbound servers do not enqueue error transactions into an error queue. This procedure cannot be used with XStream outbound servers.

The EXECUTE_ALL_ERRORS Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.11 EXECUTE_ERROR Procedure

This procedure re-executes the specified error transaction in the error queue.

Syntax

```
DBMS_APPLY_ADM.EXECUTE_ERROR(
   local_transaction_id  IN  VARCHAR2,
   execute_as_user       IN  BOOLEAN   DEFAULT FALSE,
   user_procedure        IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-12 EXECUTE_ERROR Procedure Parameters

Parameter	Description
`local_transaction_id`	The identification number of the error transaction to execute. If the specified transaction does not exist in the error queue, then an error is raised.
`execute_as_user`	If `TRUE`, then the procedure re-executes the transaction in the security context of the current user. If `FALSE`, then the procedure re-executes the transaction in the security context of the original receiver of the transaction. The original receiver is the user who was processing the transaction when the error was raised. The `DBA_APPLY_ERROR` data dictionary view lists the original receiver for each error transaction. The user who executes the transaction must have privileges to perform DML and DDL changes on the apply objects and to run any apply handlers. This user must also have dequeue privileges on the queue used by the apply component.
`user_procedure`	A user-defined procedure that modifies the error transaction so that it can be successfully executed. Specify `NULL` to execute the error transaction without running a user procedure. **See Also:** "Usage Notes" for more information about the user procedure

Usage Notes

The following usage notes apply to this procedure:

- The User Procedure
- The EXECUTE_ERROR Procedure and XStream Outbound Servers
- The EXECUTE_ERROR Procedure and XStream Inbound Servers

The User Procedure

You must specify the full procedure name for the `user_procedure` parameter in one of the following forms:

- `[`*`schema_name`*`.]`*`package_name`*`.`*`procedure_name`*
- `[`*`schema_name`*`.]`*`procedure_name`*

If the procedure is in a package, then the *`package_name`* must be specified. The user who invokes the `EXECUTE_ERROR` procedure must have `EXECUTE` privilege on the specified procedure. Also, if the *`schema_name`* is not specified, then the user who invokes the `EXECUTE_ERROR` procedure is the default.

For example, suppose the `procedure_name` has the following properties:

- `strmadmin` is the *`schema_name`*.
- `fix_errors` is the *`package_name`*.
- `fix_hr_errors` is the *`procedure_name`*.

In this case, specify the following:

```
strmadmin.fix_errors.fix_hr_errors
```

The procedure you create for error handling must have the following signature:

```
PROCEDURE user_procedure (
     in_anydata                    IN      ANYDATA,
     error_record                  IN      DBA_APPLY_ERROR%ROWTYPE,
     error_message_number          IN      NUMBER,
     messaging_default_processing  IN OUT  BOOLEAN,
     out_anydata                   OUT     ANYDATA);
```

The user procedure has the following parameters:

- `in_anydata`: The `ANYDATA` encapsulation of a message that the apply component passes to the procedure. A single transaction can include multiple messages. A message can be a row logical change record (row LCR), a DDL logical change record (DDL LCR), or a user message.
- `error_record`: The row in the `DBA_APPLY_ERROR` data dictionary view that identifies the transaction
- `error_message_number`: The message number of the `ANYDATA` object in the `in_anydata` parameter, starting at 1
- `messaging_default_processing`: If `TRUE`, then the apply component continues processing the message in the `in_anydata` parameter, which can include executing DML or DDL statements and invoking apply handlers.

  If `FALSE`, then the apply component skips processing the message in the `in_anydata` parameter and moves on to the next message in the `in_anydata` parameter.
- `out_anydata`: The `ANYDATA` object processed by the user procedure and used by the apply component if `messaging_default_processing` is `TRUE`.

If an LCR is executed using the EXECUTE LCR member procedure in the user procedure, then the LCR is executed directly, and the messaging_default_processing parameter should be set to FALSE. In this case, the LCR is not passed to any apply handlers.

Processing an error transaction with a user procedure results in one of the following outcomes:

- The user procedure modifies the transaction so that it can be executed successfully.
- The user procedure fails to make the necessary modifications, and an error is raised when transaction execution is attempted. In this case, the transaction is rolled back and remains in the error queue.

The following restrictions apply to the user procedure:

- Do not execute COMMIT or ROLLBACK statements. Doing so can endanger the consistency of the transaction.
- Do not modify LONG, LONG RAW or LOB column data in an LCR.
- If the ANYDATA object in the in_anydata parameter is a row LCR, then the out_anydata parameter must be row LCR if the messaging_default_processing parameter is set to TRUE.
- If the ANYDATA object in the in_anydata parameter is a DDL LCR, then the out_anydata parameter must be DDL LCR if the messaging_default_processing parameter is set to TRUE.
- The user who runs the user procedure must have the SELECT or READ privilege on the DBA_APPLY_ERROR data dictionary view.

> **Note:**
>
> LCRs containing transactional directives, such as COMMIT and ROLLBACK, are not passed to the user procedure.

The EXECUTE_ERROR Procedure and XStream Outbound Servers

Outbound servers do not enqueue error transactions into an error queue. This procedure cannot be used with XStream outbound servers.

The EXECUTE_ERROR Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.12 GET_ERROR_MESSAGE Function

This function returns the message payload from the error queue for the specified message number and transaction identifier. The message can be a logical change record (LCR) or a non-LCR message.

This function is overloaded. One version of this function contains two OUT parameters. These OUT parameters contain the destination queue into which the message should be enqueued, if one exists, and whether the message should be executed. The destination queue is specified using the SET_ENQUEUE_DESTINATION procedure, and the execution directive is specified using the SET_EXECUTE procedure.

> **See Also:**
>
> - SET_ENQUEUE_DESTINATION Procedure
> - SET_EXECUTE Procedure

Syntax

```
DBMS_APPLY_ADM.GET_ERROR_MESSAGE(
   message_number          IN   NUMBER,
   local_transaction_id    IN   VARCHAR2,
   destination_queue_name  OUT  VARCHAR2,
   execute                 OUT  BOOLEAN)
RETURN ANYDATA;

DBMS_APPLY_ADM.GET_ERROR_MESSAGE(
   message_number          IN   NUMBER,
   local_transaction_id    IN   VARCHAR2)
RETURN ANYDATA;
```

Parameters

Table 24-13 GET_ERROR_MESSAGE Function Parameters

Parameter	Description
message_number	The identification number of the message. This number identifies the position of the message in the transaction. Query the DBA_APPLY_ERROR data dictionary view to view the message number of each apply error.
local_transaction_id	Identifier of the error transaction for which to return a message
destination_queue_name	Contains the name of the queue into which the message should be enqueued. If the message should not be enqueued into a queue, then this parameter contains NULL.
execute	Contains TRUE if the message should be executed Contains FALSE if the message should not be executed

Usage Notes

The following usage notes apply to this procedure:

- The GET_ERROR_MESSAGE Procedure and XStream Outbound Servers
- The GET_ERROR_MESSAGE Procedure and XStream Inbound Servers

The GET_ERROR_MESSAGE Procedure and XStream Outbound Servers

Outbound servers do not enqueue error transactions into an error queue. This procedure cannot be used with XStream outbound servers.

The GET_ERROR_MESSAGE Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.13 HANDLE_COLLISIONS

This procedure enables or disables basic conflict resolution for an apply process and a table.

Syntax

```
DBMS_APPLY_ADM.HANDLE_COLLISIONS(
   apply_name    IN  VARCHAR2,
   enable        IN  BOOLEAN,
   object        IN  VARCHAR2,
   source_object IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 24-14 HANDLE_COLLISIONS Procedure Parameters

Parameter	Description
apply_name	The name of the apply process.
enable	If TRUE, then the following conflict resolution methods are used: • When a conflict is detected for a row that exists in the table, the data in the row LCR overwrites the data in the table. • When a conflict is detected for a row that does not exist in the table, the data in the row LCR is ignored. If FALSE then it disables conflict resolution set by this procedure for the specified apply process and object. If NULL, then removes any explicit table-level setting for collision handling for the specified apply process and object.
object	The schema and name of the target table, specified as [*schema_name.*]*table_name* for the change of the setting. For example, if you are changing the setting for table employees owned by user hr, then specify hr.employees. If the schema is not specified, then the current user is the default.
source_object	The schema and object name of the source table, specified as [*schema_name.*]*table_name* for the table where the change originated. For example, if the change originated at the employees table owned by user hr, then specify hr.employees. If the schema is not specified, then the current user is the default.

24.4.14 SET_DML_CONFLICT_HANDLER Procedure

This procedure adds, modifies, or removes a prebuilt DML conflict handler for INSERT, UPDATE, or DELETE conflicts on the specified object.

This procedure is overloaded. The column_list and column_table parameters are mutually exclusive.

Syntax

```
DBMS_APPLY_ADM.SET_DML_CONFLICT_HANDLER(
  apply_name             IN  VARCHAR2,
  conflict_handler_name  IN  VARCHAR2,
  object                 IN  VARCHAR2  DEFAULT NULL,
  operation_name         IN  VARCHAR2  DEFAULT NULL,
```

```
  conflict_type          IN  VARCHAR2  DEFAULT NULL,
  method_name            IN  VARCHAR2  DEFAULT NULL,
  column_list            IN  VARCHAR2  DEFAULT NULL,
  resolution_column      IN  VARCHAR2  DEFAULT NULL,
  source_object          IN  VARCHAR2  DEFAULT NULL);

DBMS_APPLY_ADM.SET_DML_CONFLICT_HANDLER(
  apply_name             IN  VARCHAR2,
  conflict_handler_name  IN  VARCHAR2,
  object                 IN  VARCHAR2  DEFAULT NULL,
  operation_name         IN  VARCHAR2  DEFAULT NULL,
  conflict_type          IN  VARCHAR2  DEFAULT NULL,
  method_name            IN  VARCHAR2  DEFAULT NULL,
  column_table           IN  DBMS_UTILITY.LNAME_ARRAY,
  resolution_column      IN  VARCHAR2  DEFAULT NULL,
  source_object          IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-15 SET_DML_CONFLICT_HANDLER Procedure Parameters

Parameter	Description
apply_name	The name of the apply process.
conflict_handler_name	The name of the conflict handler.
object	The schema and name of the target table, specified as [*schema name.*]*table name* for which a conflict handler is being added, modified, or removed. For example, if an update conflict handler is being added for table employees owned by user hr, then specify hr.employees. If the schema is not specified, then the current user is the default.
operation_name	The name of the operation, which can be specified as: • INSERT • UPDATE • DELETE In order to set up conflict handlers for different operations on the same table, you must make one call per operation.
conflict_type	Type of update conflict handler to create. You can specify one of the prebuilt handlers, which determine whether the column list from the source database is applied for the row or whether the values in the row at the destination database are retained: • ROW_EXISTS: A row with the same primary key already exists in the database for an insert or update. • ROW_MISSING:A row with the same primary key cannot be found for an update or delete.

Table 24-15 (Cont.) SET_DML_CONFLICT_HANDLER Procedure Parameters

Parameter	Description
method_name	Type of update conflict handler to create. You can specify one of the prebuilt handlers, which determine whether the column list from the source database is applied for the row or whether the values in the row at the destination database are retained: • DELTA: If the conflict type is ROW_EXISTS and the LCR is an update, then take the difference between the old and new values of the column in the LCR and add it to the current value of the column in the target database. The columns in the column group must be of type NUMBER. • IGNORE: Silently ignores the LCR and can be used for all conflict types. • MAXIMUM: Applies the column list from the source database if it has the greater value for the resolution column. Otherwise, retains the values at the destination database. This resolution method is only supported for ROW_EXISTS and only applies to inserts and updates. • MINIMUM: Applies the column list from the source database if it has the lesser value for the resolution column. Otherwise, retains the values at the destination database. This resolution method is only supported for ROW_EXISTS and only applies to inserts and updates. • OVERWRITE: Applies the column list from the source database, overwriting the column values at the destination database. An INSERT with ROW_EXISTS is converted to an UPDATE. An UPDATE with ROW_MISSING is converted to an INSERT. A DELETE with ROW_MISSING is ignored. • RECORD: Enqueue the LCR into the error queue. Can be used for all conflict types and can only be specified for a column group that contains all the columns in the table. • MAX_AND_EQUALS: Applies the column list from the source database if the value of resolution column is greater than or equal to the value of the column in the database. • MIN_AND_EQUALS: Applies the column list from the source database if the value of resolution column is less than or equal to the value of the column in the database. If NULL, then the procedure removes any existing conflict handler with the same object_name, resolution_group, and conflict_type. If a conflict handler already exists with the same object_name and resolution_column and conflict_type, then the existing handler is replaced.

Table 24-15 (Cont.) SET_DML_CONFLICT_HANDLER Procedure Parameters

Parameter	Description
column_list	A comma-separated list of the column names for which the conflict handler is called. The same column cannot be in more than one column list (for a given apply_name, object_name, operation_name and conflict_type). Specify * for the default column group, which includes all the columns in the table that are not already specified in another column list (for a given apply_name, object_name, operation_name and conflict_type). If a conflict occurs for one or more of the columns in the list when an apply component tries to apply a row logical change record (row LCR), then the conflict handler is called to resolve the conflict. The conflict handler is not called if a conflict occurs only for columns that are not in the list. You cannot use a column_list if you use a * in the object_name. The only time you can use multiple column groups is when you are specifying a conflict handler for insert or update for ROW_EXISTS. **Note:** Prebuilt conflict handlers do not support LOB, LONG, LONG RAW, user-defined type, and Oracle-suppled type columns. Therefore, you should not include these types of columns in the column_list parameter. This parameter must be set to '*' in the following cases: • The operation_name is DELETE. • The method_name is RECORD. • The operation_name is UPDATE and the conflict_type is ROW_MISSING.
column_table	An array of column names for which the conflict handler is called. This parameter is the same as the column_list parameter, but it uses an array instead of a list for the column names. **Note:** The column_list and column_table parameters are mutually exclusive.
resolution_column	For the MAXIMUM and MINIMUM prebuilt methods, the resolution column is the one tested to determine whether the current row or the LCR has the smaller value. The resolution column must be one of the columns listed in the column_list or column_table parameter. You can specify NULL for other resolution methods.
source_object	The schema and object name of the source table, specified as [*schema_name.*]*table_name* for the table where the change originated. For example, if the change originated at the employees table owned by user hr, then specify hr.employees. If the schema is not specified, then the current user is the default.

Usage Notes

The following usage notes apply to this procedure:

- Modifying an Existing Conflict Handler
- Removing an Existing Conflict Handler
- Series of Actions for Conflicts

- Procedure DML Handlers for Conflicts
- A Column Can Be in Only One Column List
- The SET_DML_CONFLICT_HANDLER Procedure and XStream Outbound Servers
- The SET_DML_CONFLICT_HANDLER Procedure and XStream Inbound Servers
- Table 24-16
- Example

Modifying an Existing Conflict Handler

If you want to modify an existing conflict handler, then you specify the `object`, `conflict_type`, and `resolution_column` of an the existing conflict handler. You can modify the `method_name` or the `column_list`.

Removing an Existing Conflict Handler

If you want to remove an existing conflict handler, then specify `NULL` for the `method_name` and specify the `object`, `conflict_type`, and `resolution_column` of the existing conflict handler.

Series of Actions for Conflicts

If an conflict occurs, then Oracle completes the following series of actions:

1. Calls the appropriate conflict handler to resolve the conflict
2. If no conflict handler is specified or if the conflict handler cannot resolve the conflict, then calls the appropriate error handler for the apply component, object name, and operation name to handle the error
3. If no error handler is specified or if the error handler cannot resolve the error, then raises an error and moves the transaction containing the row LCR that caused the error to the error queue

> **See Also:**
>
> "Signature of a DML Handler Procedure or Error Handler Procedure" for information about setting an error handler

Procedure DML Handlers for Conflicts

If you cannot use a prebuilt conflict handler to meet your requirements, then you can create a PL/SQL procedure to use as a custom conflict handler. You use the `SET_DML_HANDLER` procedure to designate one or more custom conflict handlers for a particular table. In addition, a custom conflict handler can process LOB columns and use LOB assembly.

> **See Also:**
>
> SET_DML_HANDLER Procedure

A Column Can Be in Only One Column List

When a column is in a column list, and you try to add the same column to another column list, this procedure returns the following error:

```
ORA-00001: UNIQUE CONSTRAINT (SYS.APPLY$_CONF_HDLR_COLUMNS_UNQ1) VIOLATED
```

The SET_DML_CONFLICT_HANDLER Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The SET_DML_CONFLICT_HANDLER Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

Table 24-16 Valid Combinations of Parameters

Operation	Conflict Type	Method
INSERT	ROW_EXISTS	OVERWRITE
		RECORD
		IGNORE
		MAXIMUM
		MINIMUM
UPDATE	ROW_EXISTS	OVERWRITE
		RECORD
		IGNORE
		MAXIMUM
		MINIMUM
		DELTA
UPDATE	ROW_MISSING	OVERWRITE
		RECORD
		IGNORE
DELETE	ROW_EXISTS	OVERWRITE
		RECORD
		IGNORE
DELETE	ROW_MISSING	RECORD
		IGNORE

Example

The following is an example for setting a conflict handler for the `employees` table in the `hr` schema:

```
DECLARE
  cols  DBMS_UTILITY.NAME_ARRAY;
BEGIN
  cols(1) := 'salary';
  cols(2) := 'commission_pct';
  DBMS_APPLY_ADM.SET_DML_CONFLICT_HANDLER(
    apply_name            =>  'appl1',
    conflict_handler_name =>  'emp_handler_update',
    object                =>  'hr.employees',
    operation_name        =>  'UPDATE',
    conflict_type         =>  'ROW_EXISTS',
    method_name           =>  'MAXIMUM',
    resolution_column     =>  'salary',
    column_table          =>  cols);
```

```
END;
/
```

This example sets a conflict handler named `emp_handler_update` that is called if a conflict occurs for the `salary` or `commission_pct` column in the `hr.employees` table. If such a conflict occurs, then the `salary` column is evaluated to resolve the conflict. If a conflict occurs only for a column that is not in the column list, such as the `job_id` column, then this conflict handler is not called.

24.4.15 SET_DML_HANDLER Procedure

This procedure sets or unsets a user procedure as a procedure DML handler for a specified operation on a specified database object for a single apply component or for all apply components in the database. The user procedure alters the apply behavior for the specified operation on the specified object.

Syntax

```
DBMS_APPLY_ADM.SET_DML_HANDLER(
   object_name          IN  VARCHAR2,
   object_type          IN  VARCHAR2,
   operation_name       IN  VARCHAR2,
   error_handler        IN  BOOLEAN   DEFAULT FALSE,
   user_procedure       IN  VARCHAR2,
   apply_database_link  IN  VARCHAR2  DEFAULT NULL,
   apply_name           IN  VARCHAR2  DEFAULT NULL,
   assemble_lobs        IN  BOOLEAN   DEFAULT TRUE);
```

Parameters

Table 24-17 SET_DML_HANDLER Procedure Parameters

Parameter	Description
`object_name`	The name of the source object specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default. The specified object does not need to exist when you run this procedure.
`object_type`	The type of the source object. Currently, `TABLE` is the only possible source object type.
`operation_name`	The name of the operation, which can be specified as: • `INSERT` • `UPDATE` • `DELETE` • `LOB_UPDATE` • `DEFAULT` The procedure must be run for each operation individually. Specify `DEFAULT` to set the procedure as the default procedure DML handler for the database object. In this case, the procedure DML handler is used for any `INSERT`, `UPDATE`, `DELETE`, and `LOB_WRITE` on the database object, if another procedure DML handler is not specifically set for the operation on the database object.

Table 24-17 (Cont.) SET_DML_HANDLER Procedure Parameters

Parameter	Description
error_handler	If TRUE, then the specified user procedure is run when a row logical change record (row LCR) involving the specified operation on the specified object raises an apply error. You can code the user procedure to resolve possible error conditions, notify administrators of the error, log the error, or any combination of these actions. If FALSE, then the handler being set is run for all row LCRs involving the specified operation on the specified object.
user_procedure	A user-defined procedure that is invoked during apply for the specified operation on the specified object. If the procedure is a procedure DML handler, then it is invoked instead of the default apply performed by Oracle. If the procedure is an error handler, then it is invoked when an apply error is encountered. Specify NULL to unset a procedure DML handler that is set for the specified operation on the specified object.
apply_database_link	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database is a non-Oracle database.
apply_name	The name of the apply component that uses the procedure DML handler or error handler. If NULL, then the procedure sets the procedure DML handler or error handler as a general handler for all apply components in the database. If the user_procedure parameter is set to NULL to unset a handler, and the handler being unset is set for a specific apply component, then use the apply_name parameter to specify the apply component to unset the handler.
assemble_lobs	If TRUE, then LOB assembly is used for LOB columns in LCRs processed by the handler. LOB assembly combines multiple LCRs for a LOB column resulting from a single row change into one row LCR before passing the LCR to the handler. Database compatibility must be 10.2.0 or higher to use LOB assembly. If FALSE, then LOB assembly is not used for LOB columns in LCRs processed by the handler.

Usage Notes

The following usage notes apply to this procedure:

- Run the SET_DML_HANDLER Procedure at the Destination Database
- Procedure DML Handlers and Error Handlers
- The apply_name Parameter
- Signature of a DML Handler Procedure or Error Handler Procedure
- LOB Assembly
- The SET_DML_HANDLER Procedure and XStream Outbound Servers
- The SET_DML_HANDLER Procedure and XStream Inbound Servers

Run the SET_DML_HANDLER Procedure at the Destination Database

Run this procedure at the destination database. The SET_DML_HANDLER procedure provides a way for users to apply logical change records containing DML changes (row LCRs) using a customized apply.

Procedure DML Handlers and Error Handlers

If the error_handler parameter is set to TRUE, then it specifies that the user procedure is an error handler. An error handler is invoked only when a row LCR raises an apply error. Such an error can result from a data conflict if no conflict handler is specified or if the update conflict handler cannot resolve the conflict. If the error_handler parameter is set to FALSE, then the user procedure is a procedure DML handler, not an error handler, and a procedure DML handler is always run instead of performing the specified operation on the specified object.

This procedure either sets a procedure DML handler or an error handler for a particular operation on an object. It cannot set both a procedure DML handler and an error handler for the same object and operation.

> **Note:**
>
> Currently, setting an error handler for an apply component that is applying changes to a non-Oracle database is not supported.

The apply_name Parameter

If the apply_name parameter is non-NULL, then the procedure DML handler or error handler is set for the specified apply component. In this case, this handler is not invoked for other apply components at the local destination database. If the apply_name parameter is NULL, the default, then the handler is set as a general handler for all apply components at the destination database. When a handler is set for a specific apply component, then this handler takes precedence over any general handlers. For example, consider the following scenario:

- A procedure DML handler named handler_hr is specified for an apply component named apply_hr for UPDATE operations on the hr.employees table.
- A general procedure DML handler named handler_gen also exists for UPDATE operations on the hr.employees table.

In this case, the apply_hr apply component uses the handler_hr procedure DML handler for UPDATE operations on the hr.employees table.

At the source database, you must specify an unconditional supplemental log group for the columns needed by a DML or error handler.

Signature of a DML Handler Procedure or Error Handler Procedure

You can use the SET_DML_HANDLER procedure to set either a procedure DML handler or an error handler for row LCRs that perform a specified operation on a specified object. The signatures of a DML handler procedure and of an error handler procedure are described following this section.

In either case, you must specify the full procedure name for the user_procedure parameter in one of the following forms:

- [*schema_name.*]*package_name.procedure_name*
- [*schema_name.*]*procedure_name*

If the procedure is in a package, then the *package_name* must be specified. The user who invokes the SET_DML_HANDLER procedure must have EXECUTE privilege on the specified procedure. Also, if the *schema_name* is not specified, then the user who invokes the SET_DML_HANDLER procedure is the default.

For example, suppose the procedure_name has the following properties:

- hr is the *schema_name*.
- apply_pkg is the *package_name*.
- employees_default is the *procedure_name*.

In this case, specify the following:

```
hr.apply_pkg.employees_default
```

The following restrictions apply to the user procedure:

- Do not execute COMMIT or ROLLBACK statements. Doing so can endanger the consistency of the transaction that contains the LCR.
- If you are manipulating a row using the EXECUTE member procedure for the row LCR, then do not attempt to manipulate more than one row in a row operation. You must construct and execute manually any DML statements that manipulate more than one row.
- If the command type is UPDATE or DELETE, then row operations resubmitted using the EXECUTE member procedure for the LCR must include the entire key in the list of old values. The key is the primary key or the smallest unique index that has at least one NOT NULL column, unless a substitute key has been specified by the SET_KEY_COLUMNS procedure. If there is no specified key, then the key consists of all non LOB, non LONG, and non LONG RAW columns.
- If the command type is INSERT, then row operations resubmitted using the EXECUTE member procedure for the LCR should include the entire key in the list of new values. Otherwise, duplicate rows are possible. The key is the primary key or the smallest unique index that has at least one NOT NULL column, unless a substitute key has been specified by the SET_KEY_COLUMNS procedure. If there is no specified key, then the key consists of all of the table columns, except for columns of the following data types: LOB, LONG, LONG RAW, user-defined types (including object types, REFs, varrays, nested tables), and Oracle-supplied types (including Any types, XML types, spatial types, and media types).

The procedure specified in the user_procedure parameter must have the following signature:

```
PROCEDURE user_procedure (
   parameter_name  IN  ANYDATA);
```

Here, *user_procedure* stands for the name of the procedure and *parameter_name* stands for the name of the parameter passed to the procedure. The parameter passed to the procedure is a ANYDATA encapsulation of a row LCR.

> **See Also:**
>
> Logical Change Record TYPEs for more information about LCRs

The procedure you create for error handling must have the following signature:

```
PROCEDURE user_procedure (
     message            IN  ANYDATA,
     error_stack_depth  IN  NUMBER,
     error_numbers      IN  DBMS_UTILITY.NUMBER_ARRAY,
     error_messages     IN  emsg_array);
```

If you want to retry the DML operation within the error handler, then have the error handler procedure run the `EXECUTE` member procedure for the LCR. The last error raised is on top of the error stack. To specify the error message at the top of the error stack, use `error_numbers(1)` and `error_messages(1)`.

> **Note:**
>
> - Each parameter is required and must have the specified datatype. However, you can change the names of the parameters.
> - The `emsg_array` value must be a user-defined array that is a table of type `VARCHAR2` with at least 76 characters.

Running an error handler results in one of the following outcomes:

- The error handler successfully resolves the error and returns control to the apply component.
- The error handler fails to resolve the error, and the error is raised. The raised error causes the transaction to be rolled back and placed in the error queue.

LOB Assembly

Do not modify `LONG`, `LONG RAW`, or nonassembled LOB column data in an LCR with procedure DML handlers, error handlers, or custom rule-based transformation functions. Procedure DML handlers and error handlers can modify LOB columns in row LCRs that have been constructed by LOB assembly.

The SET_DML_HANDLER Procedure and XStream Outbound Servers

Outbound servers ignore all apply handlers. This procedure has no effect on XStream outbound servers.

The SET_DML_HANDLER Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.16 SET_ENQUEUE_DESTINATION Procedure

This procedure sets the queue where the apply component automatically enqueues a message that satisfies the specified rule.

This procedure modifies the specified rule's action context to specify the queue. A rule action context is optional information associated with a rule that is interpreted by the client of the rules engine after the rule evaluates to `TRUE` for a message. In this case, the client of the rules engine is an apply component. The information in an action context is an object of type `SYS.RE$NV_LIST`, which consists of a list of name-value pairs.

A queue destination specified by this procedure always consists of the following name-value pair in an action context:

- The name is `APPLY$_ENQUEUE`.
- The value is an `ANYDATA` instance containing the queue name specified as a `VARCHAR2`.

Syntax

```
DBMS_APPLY_ADM.SET_ENQUEUE_DESTINATION(
  rule_name               IN  VARCHAR2,
  destination_queue_name  IN  VARCHAR2);
```

Parameters

Table 24-18 SET_ENQUEUE_DESTINATION Procedure Parameters

Parameter	Description
`rule_name`	The name of the rule, specified as `[`*`schema_name`*`.]`*`rule_name`*. For example, to specify a rule named `hr5` in the `hr` schema, enter `hr.hr5` for this parameter. If the schema is not specified, then the current user is the default.
`destination_queue_name`	The name of the queue into which the apply component enqueues the message. Specify the queue in the form `[`*`schema_name`*`.]`*`queue_name`*. Only local queues can be specified. For example, to specify a queue in the `hr` schema named `streams_queue`, enter `hr.streams_queue`. If the schema is not specified, then the current user is the default. If `NULL`, then an existing name-value pair with the name `APPLY$_ENQUEUE` is removed. If no name-value pair exists with the name `APPLY$_ENQUEUE` for the rule, then no action is taken. If non-`NULL` and a name-value pair exists for the rule with the name `APPLY$_ENQUEUE`, then it is removed, and a new name-value pair with the value specified by this parameter is added.

Usage Notes

The following usage notes apply to this procedure:

- The SET_ENQUEUE_DESTINATION Procedure and Apply Handlers
- Considerations for the SET_ENQUEUE_DESTINATION Procedure
- The SET_ENQUEUE_DESTINATION Procedure and XStream Outbound Servers
- The SET_ENQUEUE_DESTINATION Procedure and XStream Inbound Servers

The SET_ENQUEUE_DESTINATION Procedure and Apply Handlers

If an apply handler, such as a procedure DML handler, DDL handler, or message handler, processes a message that also is enqueued into a destination queue, then the apply handler processes the message before it is enqueued.

Considerations for the SET_ENQUEUE_DESTINATION Procedure

The following are considerations for using this procedure:

- This procedure does not verify that the specified queue exists. If the queue does not exist, then an error is raised when an apply component tries to enqueue a message into it.
- Oracle Replication capture processes, propagations, and messaging clients ignore the action context created by this procedure.

- The apply user of the apply component using the specified rule must have the necessary privileges to enqueue messages into the specified queue. If the queue is a secure queue, then the apply user must be a secure queue user of the queue.
- The specified rule must be in the positive rule set for an apply component. If the rule is in the negative rule set for an apply component, then the apply component does not enqueue the message into the destination queue.
- If the commit SCN for a message is less than or equal to the relevant instantiation SCN for the message, then the message is not enqueued into the destination queue, even if the message satisfies the apply component rule sets.

The SET_ENQUEUE_DESTINATION Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The SET_ENQUEUE_DESTINATION Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.17 SET_EXECUTE Procedure

This procedure specifies whether a message that satisfies the specified rule is executed by an apply component.

This procedure modifies the specified rule's action context to specify message execution. A rule action context is optional information associated with a rule that is interpreted by the client of the rules engine after the rule evaluates to `TRUE` for a message. In this case, the client of the rules engine is an apply component. The information in an action context is an object of type `SYS.RE$NV_LIST`, which consists of a list of name-value pairs.

A message execution directive specified by this procedure always consists of the following name-value pair in an action context:

- The name is `APPLY$_EXECUTE`.
- The value is an `ANYDATA` instance that contains `NO` as a `VARCHAR2`. When the value is `NO`, an apply component does not execute the message and does not send the message to any apply handler.

Syntax

```
DBMS_APPLY_ADM.SET_EXECUTE(
  rule_name  IN  VARCHAR2,
  execute    IN  BOOLEAN);
```

Parameters

Table 24-19 SET_EXECUTE Procedure Parameters

Parameter	Description
`rule_name`	The name of the rule, specified as `[`*`schema_name`*`.]`*`rule_name`*. For example, to specify a rule named `hr5` in the `hr` schema, enter `hr.hr5` for this parameter. If the schema is not specified, then the current user is the default.

Table 24-19 (Cont.) SET_EXECUTE Procedure Parameters

Parameter	Description
execute	If TRUE, then the procedure removes the name-value pair with the name APPLY$_EXECUTE for the specified rule. Removing the name-value pair means that the apply component executes messages that satisfy the rule. If no name-value pair with name APPLY$_EXECUTE exists for the rule, then no action is taken. If FALSE, then the procedure adds a name-value pair to the rule's action context. The name is APPLY$_EXECUTE and the value is NO. An apply component does not execute a message that satisfies the rule and does not send the message to any apply handler. If a name-value pair exists for the rule with the name APPLY$_EXECUTE, then it is removed, and a new one with the value NO is added. If NULL, then the procedure raises an error.

Usage Notes

The following usage notes apply to this procedure:

- Considerations for the SET_EXECUTE Procedure
- The SET_EXECUTE Procedure and XStream Outbound Servers
- The SET_EXECUTE Procedure and XStream Inbound Servers

Considerations for the SET_EXECUTE Procedure

The following are considerations for using this procedure:

- If the message is a logical change record (LCR) and the message is not executed, then the change encapsulated in the LCR is not made to the relevant local database object. Also, if the message is not executed, then it is not sent to any apply handler.
- Oracle Replication capture processes, propagations, and messaging clients ignore the action context created by this procedure.
- The specified rule must be in the positive rule set for an apply component for the apply component to follow the execution directive. If the rule is in the negative rule set for an apply component, then the apply component ignores the execution directive for the rule.

The SET_EXECUTE Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The SET_EXECUTE Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.18 SET_GLOBAL_INSTANTIATION_SCN Procedure

This procedure records the specified instantiation SCN for the specified source database and, optionally, for the schemas at the source database and the tables owned by these schemas. This procedure overwrites any existing instantiation SCN for the database, and, if it sets the

instantiation SCN for a schema or a table, then it overwrites any existing instantiation SCN for the schema or table.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This procedure gives you precise control over which DDL logical change records (DDL LCRs) from a source database are ignored and which DDL LCRs are applied by an apply component.

Syntax

```
DBMS_APPLY_ADM.SET_GLOBAL_INSTANTIATION_SCN(
  source_database_name  IN  VARCHAR2,
  instantiation_scn     IN  NUMBER,
  apply_database_link   IN  VARCHAR2  DEFAULT NULL,
  recursive             IN  BOOLEAN   DEFAULT FALSE,
  source_root_name      IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-20 SET_GLOBAL_INSTANTIATION_SCN Procedure Parameters

Parameter	Description
`source_database_name`	The global name of the source database. For example, `DBS1.NET`. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify `DBS1` and the domain is `NET`, then the procedure specifies `DBS1.NET` automatically.
`instantiation_scn`	The instantiation SCN. Specify `NULL` to remove the instantiation SCN metadata for the source database from the data dictionary.
`apply_database_link`	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database of a local apply component is a non-Oracle database.

Table 24-20 (Cont.) SET_GLOBAL_INSTANTIATION_SCN Procedure Parameters

Parameter	Description
recursive	If TRUE, then the procedure sets the instantiation SCN for the source database, all schemas in the source database, and all tables owned by the schemas in the source database. This procedure selects the schemas and tables from the ALL_USERS and ALL_TABLES data dictionary views, respectively, at the source database under the security context of the current user. If FALSE, then the procedure sets the global instantiation SCN for the source database, but does not set the instantiation SCN for any schemas or tables **Note:** If recursive is set to TRUE, then a database link from the destination database to the source database is required. This database link must have the same name as the global name of the source database and must be accessible to the current user. Also, a table must be accessible to the current user in either the ALL_TABLES or DBA_TABLES data dictionary view at the source database for this procedure to set the instantiation SCN for the table at the destination database.
source_root_name	The global name of the source root database. In a non-CDB, this parameter must be NULL. In a CDB, if you want to do the instantiation for a specific container than you must specify both source_database and source_root_name. If you want to do the instantiation for all the containers in the CDB, specify the source_root_name for the database and leave the source_database name as NULL.

Usage Notes

The following usage notes apply to this procedure:

- Instantiation SCNs and DDL LCRs
- The recursive Parameter
- Considerations for the SET_GLOBAL_INSTANTIATION_SCN Procedure
- The SET_GLOBAL_INSTANTIATION_SCN Procedure and XStream Outbound Servers
- The SET_GLOBAL_INSTANTIATION_SCN Procedure and XStream Inbound Servers
- The SET_GLOBAL_INSTANTIATION_SCN Procedure and CDBs

See Also:

- SET_SCHEMA_INSTANTIATION_SCN Procedure
- SET_TABLE_INSTANTIATION_SCN Procedure
- LCR$_DDL_RECORD Type for more information about DDL LCRs

Instantiation SCNs and DDL LCRs

If the commit SCN of a DDL LCR for a database object from a source database is less than or equal to the instantiation SCN for that source database at a destination database, then the apply component at the destination database disregards the DDL LCR. Otherwise, the apply component applies the DDL LCR.

The global instantiation SCN specified by this procedure is used for a DDL LCR only if the DDL LCR does not have `object_owner`, `base_table_owner`, `and base_table_name` specified. For example, the global instantiation SCN set by this procedure is used for DDL LCRs with a `command_type` of `CREATE USER`.

The recursive Parameter

If the `recursive` parameter is set to `TRUE`, then this procedure sets the instantiation SCN for each schema at a source database and for the tables owned by these schemas. This procedure uses the `SET_SCHEMA_INSTANTIATION_SCN` procedure to set the instantiation SCN for each schema, and it uses the `SET_TABLE_INSTANTIATION_SCN` procedure to set the instantiation SCN for each table. Each schema instantiation SCN is used for DDL LCRs on the schema, and each table instantiation SCN is used for DDL LCRs and row LCRs on the table.

If the `recursive` parameter is set to `FALSE`, then this procedure does not set the instantiation SCN for any schemas or tables.

Considerations for the SET_GLOBAL_INSTANTIATION_SCN Procedure

The following are considerations for using this procedure:

- Any instantiation SCN specified by this procedure is used only for LCRs captured by a capture process. It is not used for user-created LCRs.
- The instantiation SCN is not set for the `SYS` or `SYSTEM` schemas.

The SET_GLOBAL_INSTANTIATION_SCN Procedure and XStream Outbound Servers

Instantiation SCNs are not required for database objects processed by an outbound server. If an instantiation SCN is set for a database object, then the outbound server only sends the LCRs for the database object with SCN values that are greater than the instantiation SCN value. If a database object does not have an instantiation SCN set, then the outbound server skips the instantiation SCN check and sends all LCRs for that database object. In both cases, the outbound server only sends LCRs that satisfy its rule sets.

The `apply_database_link` parameter must be set to `NULL` or to the local database for this procedure to set an instantiation SCN for an outbound server.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about outbound servers and instantiation SCNs

The SET_GLOBAL_INSTANTIATION_SCN Procedure and XStream Inbound Servers

Inbound servers ignore instantiation SCNs. This procedure has no effect on XStream inbound servers.

The SET_GLOBAL_INSTANTIATION_SCN Procedure and CDBs

In a CDB, this procedure must be invoked from the same container as the apply process that uses the instantiation SCN information.

24.4.19 SET_KEY_COLUMNS Procedures

This procedure records the set of columns to be used as the substitute primary key for apply purposes and removes existing substitute primary key columns for the specified object if they exist.

This procedure is overloaded. The `column_list` and `column_table` parameters are mutually exclusive.

Syntax

```
DBMS_APPLY_ADM.SET_KEY_COLUMNS(
   object_name          IN  VARCHAR2,
   column_list          IN  VARCHAR2,
   apply_database_link  IN  VARCHAR2  DEFAULT NULL,
   apply_name           IN  VARCHAR2  DEFAULT NULL);

DBMS_APPLY_ADM.SET_KEY_COLUMNS(
   object_name          IN  VARCHAR2,
   column_table         IN  DBMS_UTILITY.NAME_ARRAY,
   apply_database_link  IN  VARCHAR2  DEFAULT NULL,
   apply_name           IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-21 SET_KEY_COLUMNS Procedure Parameters

Parameter	Description
`object_name`	The name of the table specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default. If the apply component is applying changes to a non-Oracle database in a heterogeneous environment, then the object name is not verified.
`column_list`	A comma-delimited list of the columns in the table to use as the substitute primary key, with no spaces between the column names. If the `column_list` parameter is empty or `NULL`, then the current set of key columns is removed.
`column_table`	A PL/SQL associative array of type `DBMS_UTILITY.NAME_ARRAY` of the columns in the table to use as the substitute primary key. The index for `column_table` must be 1-based, increasing, dense, and terminated by a `NULL`. If the `column_table` parameter is empty or `NULL`, then the current set of key columns is removed.
`apply_database_link`	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database is a non-Oracle database.
`apply_name`	The name of the apply component.

Usage Notes

The following usage notes apply to this procedure:

- Considerations for the SET_KEY_COLUMNS Procedure
- Duplicate Rows and Substitute Primary Key Columns

- The SET_KEY_COLUMNS Procedure and XStream Outbound Servers
- The SET_KEY_COLUMNS Procedure and XStream Inbound Servers
- The SET_KEY_COLUMNS Procedure and CDBs

Considerations for the SET_KEY_COLUMNS Procedure

The following are considerations for using this procedure:

- When not empty, the specified set of columns takes precedence over any primary key for the specified object. Do not specify substitute key columns if the object has primary key columns and you want to use those primary key columns as the key.
- Run this procedure at the destination database. At the source database, you must specify an unconditional supplemental log group for the substitute key columns.
- Unlike true primary keys, columns specified as substitute key column columns can contain `NULL`s. However, Oracle recommends that each column you specify as a substitute key column be a `NOT NULL` column. You also should create a single index that includes all of the columns in a substitute key. Following these guidelines improves performance for updates, deletes, and piecewise updates to LOBs because Oracle can locate the relevant row more efficiently.
- Do not permit applications to update the primary key or substitute key columns of a table. This ensures that Oracle can identify rows and preserve the integrity of the data.
- If there is neither a primary key, nor a unique index that has at least one `NOT NULL` column, nor a substitute key for a table, then the key consists of all of the table columns, except for columns of the following data types: LOB, `LONG`, `LONG RAW`, user-defined types (including object types, `REF`s, varrays, nested tables), and Oracle-supplied types (including `Any` types, XML types, spatial types, and media types).

Duplicate Rows and Substitute Primary Key Columns

A table has duplicate rows when all of the column values are identical for two or more rows in the table, excluding LOB, `LONG`, and `LONG RAW` columns. You can specify substitute primary key columns for a table at a destination database using by the `SET_KEY_COLUMNS` procedure. When substitute primary key columns are specified for a table with duplicate rows at a destination database, and the `allow_duplicate_rows` apply component parameter is set to `Y`, meet the following requirements to keep the table data synchronized at the source and destination databases:

- Ensure that supplemental logging is specified at source database for the columns specified as substitute key columns at the destination database. The substitute key columns must be in an unconditional log group at the source database.
- Ensure that the substitute key columns uniquely identify each row in the table at the destination database.

The rest of this section provides more details about these requirements.

When there is no key for a table and the `allow_duplicate_rows` apply component parameter is set to `Y`, a single row LCR with an `UPDATE` or `DELETE` command type only is applied to one of the duplicate rows. In this case, if the table at the source database and the table at the destination database have corresponding duplicate rows, then a change that changes all of the duplicate rows at the source database also changes all the duplicate rows at the destination database when the row LCRs resulting from the change are applied.

For example, suppose a table at a source database has two duplicate rows. An update is performed on the duplicate rows, resulting in two row LCRs. At the destination database, one row LCR is applied to one of the duplicate rows. At this point, the rows are no longer duplicate

at the destination database because one of the rows has changed. When the second row LCR is applied at the destination database, the rows are duplicate again. Similarly, if a delete is performed on these duplicate rows at the source database, then both rows are deleted at the destination database when the row LCRs resulting from the source change are applied.

When substitute primary key columns are specified for a table, row LCRs are identified with rows in the table during apply using the substitute primary key columns. If substitute primary key columns are specified for a table with duplicate rows at a destination database, and the `allow_duplicate_rows` apply component parameter is set to `Y`, then an update performed on duplicate rows at the source database can result in different changes when the row LCRs are applied at the destination database. Specifically, if the update does not change one of the columns specified as a substitute primary key column, then the same duplicate row can be updated multiple times at the destination database, while other duplicate rows might not be updated.

Also, if the substitute key columns do not identify each row in the table at the destination database uniquely, then a row LCR identified with multiple rows can update any one of the rows. In this case, the update in the row LCR might not be applied to the correct row in the table at the destination database.

An apply component ignores substitute primary key columns when it determines whether rows in a table are duplicates. An apply component determines that rows are duplicates only if all of the column values in the rows are identical (excluding LOB, `LONG`, and `LONG RAW` columns). Therefore, an apply component always raises an error if a single update or delete changes two or more nonduplicate rows in a table.

For example, consider a table with columns `c1`, `c2`, and `c3` on which the `SET_KEY_COLUMNS` procedure is used to designate column `c1` as the substitute primary key. If two rows have the same key value for the `c1` column, but different value for the `c2` or `c3` columns, then an apply component does not treat the rows as duplicates. If an update or delete modifies more than one row because the `c1` values in the rows are the same, then the apply component raises an error regardless of the setting for the `allow_duplicate_rows` apply component parameter.

> **See Also:**
>
> SET_PARAMETER Procedure for more information about the `allow_duplicate_rows` apply component parameter

The SET_KEY_COLUMNS Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The SET_KEY_COLUMNS Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

The SET_KEY_COLUMNS Procedure and CDBs

This procedure defines the columns that are used as a substitute primary key. You must perform the `SET_KEY_COLUMNS` procedure in the appropriate PDB.

24.4.20 SET_PARAMETER Procedure

This procedure sets an apply parameter to the specified value.

Syntax

```
DBMS_APPLY_ADM.SET_PARAMETER (
   apply_name  IN  VARCHAR2,
   parameter   IN  VARCHAR2,
   value       IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-22 SET_PARAMETER Procedure Parameters

Parameter	Description
apply_name	The apply component name. Do not specify an owner.
parameter	The name of the parameter you are setting.
value	The value to which the parameter is set. If NULL, then the parameter is set to its default value.

Apply Component Parameters

The following table lists the parameters for an apply component.

> **Note:**
>
> Starting from Oracle Database 21c release, OPTIMIZE_PROGRESS_TABLE is desupported.

Table 24-23 Apply Component Parameters

Parameter Name	Possible Values	Default	Description
allow_duplicate_rows	Y or N	N	If Y and more than one row is changed by a single row logical change record (row LCR) with an UPDATE or DELETE command type, then the apply component only updates or deletes one of the rows. If N, then the apply component raises an error when it encounters a single row LCR with an UPDATE or DELETE command type that changes more than one row in a table. **Note:** Regardless of the setting for this parameter, apply components do not allow changes to duplicate rows for tables with LOB, LONG, or LONG RAW columns. **See Also:** "Usage Notes" and "Duplicate Rows and Substitute Primary Key Columns"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
apply_sequence_nextval	Y or N	N for apply processes Y for XStream outbound servers and XStream inbound servers	Controls whether the apply component checks and adjusts sequence values. If Y, then the apply component checks and adjusts sequence values. For ascending sequences, setting this parameter to Y ensures that the destination sequence values are equal to or greater than the source sequence values. For descending sequences, setting this parameter to Y ensures that the destination sequence values are equal to or less than the source sequence values. If N, then the apply component does not check or adjust sequence values. **Note:** This parameter is intended for XStream. Do not set this parameter to Y for an apply process in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** SET_PARAMETER Procedure for information about the capture_sequence_nextval capture process parameter

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
batchsql_mode	DEPENDENT, DEPENDENT_EAGER, or SEQUENTIAL	DEPENDENT	Determines the batching method used to generate batch transactions for reordering. This parameter can be set to one of the following: • DEPENDENT - Batch transactions in a dependency-aware manner to minimize cross-batch dependencies and improve parallel processing performance in Oracle GoldenGate BATCHSQL mode. An executing batch has no unresolved dependencies. • DEPENDENT_EAGER - Batch transactions in a dependency-aware manner to minimize cross-batch dependencies and improve parallel processing performance in Oracle GoldenGate BATCHSQL mode. A batch can be executed when there are unresolved dependencies. The apply server waits for dependencies to be resolved before executing an LCR. • SEQUENTIAL - Batch transactions in a sequential manner. COMMIT_SERIALIZATION = FULL runs in this mode regardless of the Oracle GoldenGate BATCHSQL mode. **Note:** This parameter is intended for Oracle GoldenGate. Do not use this parameter in an Oracle Replication environment or in an XStream environment.
cdgranularity	ROW or COLGROUP	COLGROUP for XStream In ROW for Oracle GoldenGate	Specifies the conflict detection granularity. This parameter can be set to one of the following: • ROW - Conflict resolution is applied for all column groups if there is a conflict for any column group. • COLGROUP - Conflict resolution is applied for column groups that have a conflict. Conflict resolution is not applied for column groups that do not have a conflict. **Note:** This parameter is intended for XStream and Oracle GoldenGate. Do not set this parameter for an apply process in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations.

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
commit_serialization	DEPENDENT_TRANSACTIONS or FULL	DEPENDENT_TRANSACTIONS	The order in which applied transactions are committed. Apply servers can apply nondependent transactions at the destination database in an order that is different from the commit order at the source database. Dependent transactions are always applied at the destination database in the same order as they were committed at the source database. You control whether the apply servers can apply nondependent transactions in a different order at the destination database using the commit_serialization apply parameter. This parameter has the following settings: • DEPENDENT_TRANSACTIONS - The apply component can commit nondependent transactions in any order. Performance is best if you specify DEPENDENT_TRANSACTIONS. • FULL - The apply component commits applied transactions in the order in which they were committed at the source database. Regardless of the specification, applied transactions can execute in parallel subject to data dependencies and constraint dependencies. If you specify DEPENDENT_TRANSACTIONS, then a destination database might commit changes in a different order than the source database. For example, suppose two nondependent transactions are committed at the source database in the following order: 1. Transaction A 2. Transaction B At the destination database, these transactions might be committed in the opposite order: 1. Transaction B 2. Transaction A If you specify DEPENDENT_TRANSACTIONS and there are application constraints that are not enforced by the database, then use virtual dependency definitions or add RELY constraints to account for the application constraints. See *Oracle Database Data Warehousing Guide* for information about RELY constraints. **Note:** The NONE value is deprecated for this parameter. It is replaced by the DEPENDENT_TRANSACTIONS value.

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
			See Also: "Usage Notes"
compare_key_only	Y or N	N for apply processes Y for XStream inbound servers	If Y, then disables automatic conflict detection and only uses primary and unique key columns to identify the table row for a row LCR. If N, then enables automatic conflict detection and uses all of the old values in a row LCR to identify the table row for a row LCR. **Note:** The COMPARE_OLD_VALUES procedure in this package can disable comparison of old values for specified columns during apply. See COMPARE_OLD_VALUES Procedure. **See Also:** "Usage Notes"
compute_lcr_dep_on_arrival	Y or N	N	If Y, the dependencies are computed as the LCRs for the transaction are received. If N, the dependencies are computed only after all the LCRs for a transaction are received. If the target table has all of the same constraints as the source table, you can improve the performance by setting this parameter to Y. If the number of LCRs in transaction exceeds the value of the number of the eager_size parameter, then the dependencies for that transaction are calculated on arrival regardless of the setting of compute_lcr_dep_on_arrival. **Note:** This parameter is intended for XStream. Do not set this parameter to Y for an apply process in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations.
disable_on_error	Y or N	Y	If Y, then the apply component is disabled on the first unresolved error, even if the error is not irrecoverable. If N, then the apply component continues regardless of unresolved errors.

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
disable_on_limit	Y or N	N	If Y, then the apply component is disabled if the apply component terminates because it reached a value specified by the time_limit parameter or transaction_limit parameter. If N, then the apply component is restarted immediately after stopping because it reached a limit. When an apply component is restarted, it gets a new session identifier, and the processes associated with the apply component also get new session identifiers. However, the coordinator process number (APnn) remains the same.
eager_size	A positive integer	9500	The apply component usually waits until it receives a commit record before starting to apply changes of a transaction. If XStream is enabled and more than eager_size LCRs arrive for a given transaction, then apply starts processing the changes. If XStream is not enabled and more than eager_size LCRs arrive for a given transaction, then apply waits until the complete transaction is received before processing the changes. Since it is possible that all existing apply servers are handling complete transactions from the source, additional apply servers are automatically created to handle outstanding eager transactions. The apply parameter max_parallelism limits the maximum number of apply servers that can be used for an apply process. This apply parameter is relevant only if its value is less than the value of the txn_lcr_spill_threshold apply parameter. When the value of txn_lcr_spill_threshold is lower than the value of eager_size, transactions spill to disk before eager apply begins. **Note:** This parameter is intended for XStream. Do not set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations.

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
enable_xstream_table_stats	Y or N	Y	When this parameter is set to Y, statistics about the operations of applied transactions are collected and made available in the V$XSTREAM_TABLE_STATS view. When this parameter is set to N, no statistics are collected. **Note:** This parameter is intended for XStream. Do not set this parameter to Y for an apply process in an Oracle Replication environment.
excludetag	Comma-delimited list of Oracle Replication tags	NULL	Controls whether the capture process for an outbound server captures DML changes that are tagged with one of the specified Oracle Replication tags. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for an XStream Out environment in which multiple outbound servers use the same capture process. XStream inbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes" for the DBMS_CAPTURE_ADM.SET_PARAMETER procedure for more information about this parameter
excludetrans	Comma-delimited list of transaction names	NULL	Controls whether the capture process for an outbound server captures DML changes in the specified transaction names. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for an XStream Out environment in which multiple outbound servers use the same capture process. XStream inbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes" for the DBMS_CAPTURE_ADM.SET_PARAMETER procedure for more information about this parameter

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
excludeuser	Comma-delimited list of user names	NULL	Controls whether the capture process for an outbound server captures DML changes made by the specified users. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. Specify an exact pattern match for each user name. The pattern match is case sensitive. For example, specify HR for the hr user. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for an XStream Out environment in which multiple outbound servers use the same capture process. XStream inbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes" for the DBMS_CAPTURE_ADM.SET_PARAMETER procedure for more information about this parameter
excludeuserid	Comma-delimited list of user ID values	NULL	Controls whether the capture process for an outbound server captures data manipulation language (DML) changes made by the specified users. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. To view the user ID for a user, query the USER_ID column in the ALL_USERS data dictionary view. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for an XStream Out environment in which multiple outbound servers use the same capture process. XStream inbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes" for the DBMS_CAPTURE_ADM.SET_PARAMETER procedure for more information about this parameter

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
getapplops	Y or N	Y	If Y, then the capture process captures DML changes if the original user is not specified in the excludeuserid or excludeuser parameters and the transaction name is not specified in the excludetrans parameter. If N, then the capture process ignores DML changes if the original user is not specified in the excludeuserid or excludeuser parameters and the transaction name is not specified in the excludetrans parameter. In either case, the capture process captures a DML change only if it satisfies the capture process's rule sets. When N is set for both getapplops and getreplicates, no data is captured. **Note:** This parameter is intended for an XStream Out environment in which multiple outbound servers use the same capture process. XStream inbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes" for the DBMS_CAPTURE_ADM.SET_PARAMETER procedure for more information about this parameter
getreplicates	Y or N	N	If Y, then the capture process captures DML changes if the original user is specified in the excludeuserid or excludeuser parameters and the transaction name is specified in the excludetrans parameter. If N, then the capture process ignores DML changes if the original user is specified in the excludeuserid or excludeuser parameters and the transaction name is specified in the excludetrans parameter. In either case, the capture process captures a DML change only if it satisfies the capture process's rule sets. When N is set for both getapplops and getreplicates, no data is captured. **Note:** This parameter is intended for an XStream Out environment in which multiple outbound servers use the same capture process. XStream inbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes" for the DBMS_CAPTURE_ADM.SET_PARAMETER procedure for more information about this parameter

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
grouptransops	A positive integer from 1 to 10000	250 for apply processes and XStream inbound servers 10000 for XStream outbound servers	The minimum number of LCRs that can be grouped into a single transaction. The commit LCR for a transaction is not included in the LCR count for the transaction. This parameter enables an apply component to group LCRs from multiple transactions into a single transaction. The apply component groups only LCRs that are part of committed transactions. If a transaction has more LCRs than the setting for this parameter, then the transaction is applied as a single transaction. The apply component does not split a transaction into separate transactions. This parameter only takes effect if the parallelism parameter setting is 1. The grouptransops parameter is ignored if the parallelism parameter setting is greater than 1. **Note:** This parameter is intended for XStream outbound servers and inbound servers. An Oracle Apply process ignores this parameter unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
handlecollisions	Y or N	N	This parameter controls whether the apply component tries to resolve duplicate-record and missing-record errors when applying changes during data loading. This parameter should be set to N for normal replication activity. It should be set to Y only when data is being loaded (instantiated) and replication is enabled. If Y, then does the equivalent of OVERWRITE for INSERT, UPDATE and DELETE operations that get ROW_EXISTS errors, and ignores UPDATE and DELETE operations that get ROW_MISSING errors. Specifically, the apply component performs the following actions when this parameter is set to Y: • If the operation is an insert and the primary key or unique key exists, then the insert is converted to an update. • If the operation is an update that does not modify the primary key or unique key columns, and the row does not exist, then the change is ignored. • If the operation is an update that modifies the primary key or unique key columns, and the row does not exist, then the change is ignored. • If the operation is an update that modifies the primary key or unique key columns, and the a row with the new key values already exists, then delete the row with the old key values and replace the row with the new key values. • If the operation is a delete and the row does not exist, then the change is ignored. If N then it disables the above settings. **Note:** This parameter is intended for an XStream In environment with one or more inbound servers. Do not set this parameter in an Oracle Replication environment.

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
ignore_transaction	A valid source transaction ID or NULL	NULL	Instructs the apply component to ignore the specified transaction from the source database, effective immediately. Use caution when setting this parameter because ignoring a transaction might lead to data divergence between the source database and destination database. To ignore multiple transactions, specify each transaction in a separate call to the SET_PARAMETER procedure. The DBA_APPLY_PARAMETERS view displays a comma-delimited list of all transactions to be ignored. To clear the list of ignored transactions, run the SET_PARAMETER procedure and specify NULL for the ignore_transaction parameter. If NULL, then the apply component ignores this parameter. **Note:** An apply component ignores this parameter for transactions that were not captured by a capture process. **See Also:** "Usage Notes"
maximum_scn	A valid SCN or INFINITE	INFINITE	The apply component is disabled before applying a transaction with a commit SCN greater than or equal to the value specified. If INFINITE, then the apply component runs regardless of the SCN value. **Note:** An apply component ignores this parameter for transactions that were not captured by a capture process. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
max_parallelism	A positive integer	50	Limits the maximum number of apply servers that can be used for an apply component. When the apply parallelism parameter is set greater than one, the apply component adds apply servers when necessary to process transactions until it reaches the limit set by this parameter (max_parallelism). Transactions include both unassigned (eager) and assigned transactions. Apply servers that idle for more than 5 minutes are shut down until the configured parallelism is attained. Runtime statistics for servers that have been shut down are aggregated into apply server 0 so that accurate apply statistics for the entire apply process can be maintained. **Note:** This parameter is intended for XStream. Do not set this parameter to Y for an apply process in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations.
max_sga_size	A positive integer	INFINITE	Controls the amount of system global area (SGA) memory allocated specifically to the apply component, in megabytes. The memory is allocated for the duration of the apply component's session and is released when the apply component becomes disabled. **Note:** The sum of system global area (SGA) memory allocated for all components on a database must be less than the value set for the STREAMS_POOL_SIZE initialization parameter. If NULL, then the apply component uses the original default value. A NULL value has the same effect as resetting the parameter to its default value. **Note:** This parameter is intended for XStream. Do not use or attempt to set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
message_tracking_frequency	0 or a positive integer	2000000	The frequency at which messages applied by the inbound server are tracked automatically. For example, if this parameter is set to the default value of 2000000, then every two-millionth message is tracked automatically. The tracking label used for automatic message tracking is *inbound_server_name*:AUTOTRACK, where *inbound_server_name* is the name of the inbound server. Only the first 20 bytes of the inbound server name are used; the rest is truncated if it exceeds 20 bytes. If 0 (zero), then no messages are tracked automatically.
optimize_self_updates	Y or N	Y	This parameter affects conflict resolution when an update in the source database sets a column to its existing value. When this parameter is set to Y, a conflict between the value in the LCR and the corresponding column in the target database is considered resolved. When this parameter is set to N, the conflict is processed.
parallelism	A positive integer	4	The number of apply servers that can concurrently apply transactions. The reader server and the apply server process names are AS*nn*, where *nn* can include letters and numbers. The total number of AS*nn* processes is the value of the parallelism parameter plus one. For example, if parallelism is set to 4, then an apply component uses a total of five AS*nn* processes. In this case, there is one reader server and four apply servers. Setting the parallelism parameter to a number higher than the number of available operating system user processes can disable the apply component. Make sure the PROCESSES initialization parameter is set appropriately when you set the parallelism parameter. **Note:** When the value of this parameter is changed from 1 to a higher value for a running apply component, the apply component is stopped and restarted automatically. This can take some time depending on the size of the transactions currently being applied. When the value of this parameter is greater than 1, and the parameter value is decreased or increased, the apply component does not restart. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
parallelism_interval	0 or a positive integer	0	The parallelism interval is the interval in seconds at which the current workload activity is computed. The apply component calculates the mean throughput every 5 X parallelism_interval seconds. After each calculation, the apply component can increase or decrease the number of apply servers to try to improve throughput. If throughput is improved, then the apply component keeps the new number of apply servers. The parallelism interval is used only if the parallelism parameter is set to a value greater than one and the max_parallelism parameter value is greater than the parallelism parameter value. **Note:** This parameter is intended for an XStream In environment with one or more inbound servers. XStream outbound servers ignore this parameter. Do not set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
preserve_encryption	Y or N	Y	Whether to preserve encryption for columns encrypted using Transparent Data Encryption. If Y, then columns in tables at the destination database must be encrypted when corresponding columns in row LCRs are encrypted. If columns are encrypted in row LCRs but the corresponding columns are not encrypted in the tables at the destination database, then an error is raised when the apply component tries to apply the row LCRs. If N, then columns in tables at the destination database do not need to be encrypted when corresponding columns in row LCRs are encrypted. If columns are encrypted in row LCRs but the corresponding columns are not encrypted in the tables at the destination database, then the apply component applies the changes in the row LCRs. **Note:** When the value of this parameter is changed for a running apply component, the apply component is stopped and restarted automatically. This can take some time depending on the size of the transactions currently being applied. **See Also:** "Usage Notes"
rtrim_on_implicit_conversion	Y or N	Y	Whether to remove blank padding from the right end of a column when automatic data type conversion is performed during apply. If Y, then blank padding is removed when a CHAR or NCHAR source column in a row LCR is converted to a VARCHAR2, NVARCHAR2, or CLOB column in a table. If N, then blank padding is preserved in the column. **See Also:** "Usage Notes"
startup_seconds	0, a positive integer, or INFINITE	0	The maximum number of seconds to wait for another instantiation of the same apply component to finish. If the other instantiation of the same apply component does not finish within this time, then the apply component does not start. If INFINITE, then an apply component does not start until another instantiation of the same apply component finishes. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
suppresstriggers	Y or N	Y	This parameter controls whether triggers fire when a change is made by the apply component. If Y, triggers do not fire for changes made by the apply component. If N, triggers fire for changes made by the apply component. If a trigger's firing property is set to always fire, then the trigger always fires for changes made by the apply component, regardless of the value of the suppresstriggers parameter. A trigger's firing property is set to always fire by running the DBMS_DDL.SET_TRIGGER_FIRING_PROPERTY procedure with the fire_once parameter set to FALSE. **Note:** This parameter is intended for an XStream In environment with one or more inbound servers. Do not set this parameter in an Oracle Replication environment. **See Also:** "Usage Notes"
time_limit	A positive integer or INFINITE	INFINITE	The apply component stops as soon as possible after the specified number of seconds since it started. If INFINITE, then the apply component continues to run until it is stopped explicitly. **See Also:** "Usage Notes"
trace_level	0 or a positive integer	0	Set this parameter only under the guidance of Oracle Support Services. **See Also:** "Usage Notes"
transaction_limit	A positive integer or INFINITE	INFINITE	The apply component stops after applying the specified number of transactions. If INFINITE, then the apply component continues to run regardless of the number of transactions applied. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
txn_age_spill_threshold	A positive integer or INFINITE	900	The apply component begins to spill messages from memory to hard disk for a particular transaction when the amount of time that any message in the transaction has been in memory exceeds the specified number. The parameter specifies the age in seconds. When the reader server spills messages from memory, the messages are stored in a database table on the hard disk. These messages are not spilled from memory to a queue table. Message spilling occurs at the transaction level. For example, if this parameter is set to 900, and the reader server of an apply component detects that one message in a transaction has been in memory longer than 900 seconds, then all of the messages in the transaction spill from memory to hard disk. If INFINITE, then the apply component does not spill messages to the hard disk based on the age of the messages. Query the DBA_APPLY_SPILL_TXN data dictionary view for information about transactions spilled by an apply component. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
txn_lcr_spill_threshold	A positive integer or INFINITE	10000	The apply component begins to spill messages from memory to hard disk for a particular transaction when the number of messages in memory for the transaction exceeds the specified number. The number of messages in first chunk of messages spilled from memory equals the number specified for this parameter, and the number of messages spilled in future chunks is either 100 or the number specified for this parameter, whichever is less. If the reader server of an apply component has the specified number of messages in memory for a particular transaction, then when it detects the next message for this transaction, it spills the messages that are in memory to the hard disk. For example, if this parameter is set to 10000, and a transaction has 10,200 messages, then the reader server handles the transaction in the following way: 1. Reads the first 10,000 messages in the transaction into memory 2. Spills messages 1 - 10,000 to hard disk when it detects message 10,000 3. Reads the next 100 messages in the transaction into memory 4. Spills messages 10,001 - 10,100 to hard disk when it detects message 10,100 5. Reads the next 100 messages in the transaction into memory The apply component applies the first 10,100 messages from the hard disk and the last 100 messages from memory. When the reader server spills messages from memory, the messages are stored in a database table on the hard disk. These messages are not spilled from memory to a queue table. Message spilling occurs at the transaction level. For example, if this parameter is set to 10000, and the reader server of an apply component is assembling two transactions, one with 7,500 messages and another with 8,000 messages, then it does not spill any messages. If INFINITE, then the apply component does not spill messages to the hard disk based on the number of messages in a transaction. Query the DBA_APPLY_SPILL_TXN data dictionary view for information about transactions spilled by an apply component. **See Also:** "Usage Notes"

Table 24-23 (Cont.) Apply Component Parameters

Parameter Name	Possible Values	Default	Description
write_alert_log	Y or N	Y	If Y, then the apply component writes a message to the alert log on exit. If N, then the apply component does not write a message to the alert log on exit. The message specifies the reason why the apply component stopped.

Usage Notes

The following usage notes apply to this procedure:

- Delays Are Possible Before New Parameter Settings Take Effect
- Parameters Interpreted as Positive Integers
- Parameters with a System Change Number (SCN) Setting
- The SET_PARAMETER Procedure and Replication
- The SET_PARAMETER Procedure and XStream Outbound Servers
- The SET_PARAMETER Procedure and XStream Inbound Servers

Delays Are Possible Before New Parameter Settings Take Effect

When you alter a parameter value, a short amount of time might pass before the new value for the parameter takes effect.

Parameters Interpreted as Positive Integers

For all parameters that are interpreted as positive integers, the maximum possible value is 4,294,967,295. Where applicable, specify INFINITE for larger values.

Parameters with a System Change Number (SCN) Setting

For parameters that require an SCN setting, any valid SCN value can be specified.

The SET_PARAMETER Procedure and Replication

You can use the following parameters in Replication if you enable XStream performance optimizations for Oracle Replication using the procedure DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS:

- apply_sequence_nextval
- compute_lcr_dep_on_arrival
- eager_size
- grouptransops
- max_parallelism
- max_sga_size
- parallelism_interval

The SET_PARAMETER Procedure and XStream Outbound Servers

Outbound servers ignore the settings for the following apply parameters:

- allow_duplicate_rows
- commit_serialization
- compare_key_only
- compute_lcr_dep_on_arrival
- disable_on_error
- eager_size
- enable_xstream_table_stats
- grouptransops
- handlecollisions
- optimize_self_updates
- parallelism
- parallelism_interval
- preserve_encryption
- rtrim_on_implicit_conversion
- suppresstriggers

The commit_serialization parameter is always set to FULL for an outbound server, and the parallelism parameter is always set to 1 for an outbound server.

You can use the other apply parameters with outbound servers.

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

The SET_PARAMETER Procedure and XStream Inbound Servers

Inbound servers ignore the settings for the following apply parameters:

- excludetag
- excludetrans
- excludeuser
- excludeuserid
- getapplops
- getreplicates
- ignore_transaction
- maximum_scn

You can use all of the other apply component parameters with inbound servers.

The default setting for the compare_key_only parameter for an inbound server is Y.

The default setting for the `parallelism` parameter for an inbound server is `4`.

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

24.4.21 SET_REPERROR_HANDLER Procedure

This procedure specifies how a particular error is handled based on its error number.

You can choose between several predefined actions for a given error.

Syntax

```
DBMS_APPLY_ADM.SET_REPERROR_HANDLER(
  apply_name     IN  VARCHAR2,
  object         IN  VARCHAR2,
  error_number   IN  NUMBER,
  method         IN  VARCHAR2,
  source_object  IN  VARCHAR2  DEFAULT NULL,
  max_retries    IN  NUMBER    DEFAULT NULL,
  delay_csecs    IN  NUMBER    DEFAULT 6000);
```

Parameters

Table 24-24 SET_REPERROR_HANDLER Procedure Parameters

Parameter	Description
`apply_name`	The name of the apply process.
`object`	The schema and name of the target table, specified as `[`*`schema name`*`.]`*`table name`* for which an error handler is being added, modified, or removed. The table must exist. For example, if an update conflict handler is being added for table `employees` owned by user `hr`, then specify `hr.employees`. If the schema is not specified, then the current user is the default.
`error_number`	The error handling number. If `0`, then use the default for all error handling for `object`.
`method`	Specifies the action to take when the given `error_number` occurs. If `NULL`, remove the error handler for `error_number` The possible actions are: • `ABEND`: Stop the apply process when the error occurs. • `RECORD`: Move the LCR to the error queue when the error is encountered. • `IGNORE`: Silently ignore the error and do not apply the LCR. • `RETRY`: Retry the LCR `max_retries` times. • `RETRY_TRANSACTION`: Retry the transaction `max_retries` times. Wait `delay_csecs` centiseconds before each retry. `RECORD_TRANSACTION` : Move the entire transaction to the error queue if this error occurs.

Table 24-24 (Cont.) SET_REPERROR_HANDLER Procedure Parameters

Parameter	Description
source_object	The schema and object name of the source table, specified as [*schema_name*.]*table_name* for the table where the change originated. For example, if the change originated at the employees table owned by user hr, then specify hr.employees. If the schema is not specified, then the current user is the default.
max_retires	Maximum number of times to retry for RETRY and RETRY_TRANSACTION actions in method. Must be specified with either the RETRY or RETRY_TRANSACTION
delay_csecs	The number of centiseconds between retries for RETRY and RETRY_TRANSACTION action in method.

Usage Notes

The following usage notes apply to this procedure:

- Priority of Error Handlers

Priority of Error Handlers

Any conflict handling specified by SET_UPDATE_CONFLICT_HANDLER or SET_DML_CONFLICT_HANDLER is tried before the actions specified by SET_REPERROR_HANDLER. The PL/SQL procedure specified by SET_DML_HANDLER is called to handle the error if none of the previously mentioned methods resolve it.

24.4.22 SET_SCHEMA_INSTANTIATION_SCN Procedure

This procedure records the specified instantiation SCN for the specified schema in the specified source database and, optionally, for the tables owned by the schema at the source database. This procedure overwrites any existing instantiation SCN for the schema, and, if it sets the instantiation SCN for a table, it overwrites any existing instantiation SCN for the table.

This procedure gives you precise control over which DDL logical change records (LCRs) for a schema are ignored and which DDL LCRs are applied by an apply component.

Syntax

```
DBMS_APPLY_ADM.SET_SCHEMA_INSTANTIATION_SCN(
  source_schema_name    IN  VARCHAR2,
  source_database_name  IN  VARCHAR2,
  instantiation_scn     IN  NUMBER,
  apply_database_link   IN  VARCHAR2  DEFAULT NULL,
  recursive             IN  BOOLEAN   DEFAULT FALSE,
  source_root_name      IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-25 SET_SCHEMA_INSTANTIATION_SCN Procedure Parameters

Parameter	Description
source_schema_name	The name of the source schema. For example, hr. When setting an instantiation SCN for schema, always specify the name of the schema at the source database, even if a rule-based transformation or apply handler is configured to change the schema name.
source_database_name	The global name of the source database. For example, DBS1.NET. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is NET, then the procedure specifies DBS1.NET automatically.
instantiation_scn	The instantiation SCN. Specify NULL to remove the instantiation SCN metadata for the source schema from the data dictionary.
apply_database_link	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database of a local apply component is a non-Oracle database.
recursive	If TRUE, then the procedure sets the instantiation SCN for the specified schema and all tables owned by the schema in the source database. This procedure selects the tables owned by the specified schema from the ALL_TABLES data dictionary view at the source database under the security context of the current user. If FALSE, then the procedure sets the instantiation SCN for specified schema, but does not set the instantiation SCN for any tables **Note:** If recursive is set to TRUE, then a database link from the destination database to the source database is required. This database link must have the same name as the global name of the source database and must be accessible to the current user. Also, a table must be accessible to the current user in either the ALL_TABLES or DBA_TABLES data dictionary view at the source database for this procedure to set the instantiation SCN for the table at the destination database.
source_root_name	The global name of the source root database. In a non-CDB, this parameter must be NULL. In a CDB, both source_database and source_root_name must be specified to identify a specific container.

Usage Notes

The following usage notes apply to this procedure:

- The SET_SCHEMA_INSTANTIATION_SCN Procedure and LCRs
- Instantiation SCNs and DDL LCRs
- The recursive Parameter
- The SET_SCHEMA_INSTANTIATION_SCN Procedure and XStream Outbound Servers
- The SET_SCHEMA_INSTANTIATION_SCN Procedure and XStream Inbound Servers
- The SET_SCHEMA_INSTANTIATION_SCN Procedure and CDBs

See Also:

- SET_GLOBAL_INSTANTIATION_SCN Procedure
- SET_TABLE_INSTANTIATION_SCN Procedure
- LCR$_DDL_RECORD Type for more information about DDL LCRs

The SET_SCHEMA_INSTANTIATION_SCN Procedure and LCRs

Any instantiation SCN specified by this procedure is used only for LCRs captured by a capture process. It is not used for user-created LCRs.

Instantiation SCNs and DDL LCRs

If the commit SCN of a DDL LCR for a database object in a schema from a source database is less than or equal to the instantiation SCN for that database object at a destination database, then the apply component at the destination database disregards the DDL LCR. Otherwise, the apply component applies the DDL LCR.

The schema instantiation SCN specified by this procedure is used on the following types of DDL LCRs:

- DDL LCRs with a `command_type` of `CREATE TABLE`
- DDL LCRs with a non-`NULL` `object_owner` specified and neither `base_table_owner` nor `base_table_name` specified.

For example, the schema instantiation SCN set by this procedure is used for a DDL LCR with a `command_type` of `CREATE TABLE` and `ALTER USER`.

The schema instantiation SCN specified by this procedure is not used for DDL LCRs with a `command_type` of `CREATE USER`. A global instantiation SCN is needed for such DDL LCRs.

The recursive Parameter

If the `recursive` parameter is set to `TRUE`, then this procedure sets the table instantiation SCN for each table at the source database owned by the schema. This procedure uses the `SET_TABLE_INSTANTIATION_SCN` procedure to set the instantiation SCN for each table. Each table instantiation SCN is used for DDL LCRs and row LCRs on the table.

If the `recursive` parameter is set to `FALSE`, then this procedure does not set the instantiation SCN for any tables.

The SET_SCHEMA_INSTANTIATION_SCN Procedure and XStream Outbound Servers

Instantiation SCNs are not required for database objects processed by an outbound server. If an instantiation SCN is set for a database object, then the outbound server only sends the LCRs for the database object with SCN values that are greater than the instantiation SCN value. If a database object does not have an instantiation SCN set, then the outbound server skips the instantiation SCN check and sends all LCRs for that database object. In both cases, the outbound server only sends LCRs that satisfy its rule sets.

The `apply_database_link` parameter must be set to `NULL` or to the local database for this procedure to set an instantiation SCN for an outbound server.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about outbound servers and instantiation SCNs

The SET_SCHEMA_INSTANTIATION_SCN Procedure and XStream Inbound Servers

Inbound servers ignore instantiation SCNs. This procedure has no effect on XStream inbound servers.

The SET_SCHEMA_INSTANTIATION_SCN Procedure and CDBs

In a CDB, this procedure must be invoked from the same container as the apply process that uses the instantiation SCN information.

24.4.23 SET_TABLE_INSTANTIATION_SCN Procedure

This procedure records the specified instantiation SCN for the specified table in the specified source database. This procedure overwrites any existing instantiation SCN for the particular table.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This procedure gives you precise control over which logical change records (LCRs) for a table are ignored and which LCRs are applied by an apply component.

Syntax

```
DBMS_APPLY_ADM.SET_TABLE_INSTANTIATION_SCN(
  source_object_name    IN  VARCHAR2,
  source_database_name  IN  VARCHAR2,
  instantiation_scn     IN  NUMBER,
  apply_database_link   IN  VARCHAR2  DEFAULT NULL,
  source_root_name      IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-26 SET_TABLE_INSTANTIATION_SCN Procedure Parameters

Parameter	Description
source_object_name	The name of the source object specified as [*schema_name*.]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default. When setting an instantiation SCN for a database object, always specify the name of the schema and database object at the source database, even if a rule-based transformation or apply handler is configured to change the schema name or database object name.
source_database_name	The global name of the source database. For example, DBS1.NET. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is NET, then the procedure specifies DBS1.NET automatically.
instantiation_scn	The instantiation SCN. Specify NULL to remove the instantiation SCN metadata for the source table from the data dictionary.
apply_database_link	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database of a local apply component is a non-Oracle database. **Note:** This parameter must be NULL when the procedure is invoked from the root of a CDB.
source_root_name	The global name of the source root database. In a non-CDB, this parameter must be NULL. In a CDB, both source_database and source_root_name must be specified to identify a specific container.

Usage Notes

The following usage notes apply to this procedure:

- Instantiation SCNs and LCRs
- The SET_TABLE_INSTANTIATION_SCN Procedure and XStream Outbound Servers
- The SET_TABLE_INSTANTIATION_SCN Procedure and XStream Inbound Servers
- The SET_TABLE_INSTANTIATION_SCN Procedure and CDBs

Instantiation SCNs and LCRs

If the commit SCN of an LCR for a table from a source database is less than or equal to the instantiation SCN for that table at some destination database, then the apply component at the destination database disregards the LCR. Otherwise, the apply component applies the LCR.

The table instantiation SCN specified by this procedure is used on the following types of LCRs:

- Row LCRs for the table
- DDL LCRs that have a non-NULL base_table_owner and base_table_name specified, except for DDL LCRs with a command_type of CREATE TABLE

For example, the table instantiation SCN set by this procedure is used for DDL LCRs with a command_type of ALTER TABLE or CREATE TRIGGER.

> **Note:**
>
> The instantiation SCN specified by this procedure is used only for LCRs captured by a capture process. It is not used for user-created LCRs.

> **See Also:**
>
> - SET_GLOBAL_INSTANTIATION_SCN Procedure
> - SET_SCHEMA_INSTANTIATION_SCN Procedure
> - LCR$_ROW_RECORD Type for more information about row LCRs
> - LCR$_DDL_RECORD Type for more information about DDL LCRs

The SET_TABLE_INSTANTIATION_SCN Procedure and XStream Outbound Servers

Instantiation SCNs are not required for database objects processed by an outbound server. If an instantiation SCN is set for a database object, then the outbound server only sends the LCRs for the database object with SCN values that are greater than the instantiation SCN value. If a database object does not have an instantiation SCN set, then the outbound server skips the instantiation SCN check and sends all LCRs for that database object. In both cases, the outbound server only sends LCRs that satisfy its rule sets.

The `apply_database_link` parameter must be set to `NULL` or to the local database for this procedure to set an instantiation SCN for an outbound server.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about outbound servers and instantiation SCNs

The SET_TABLE_INSTANTIATION_SCN Procedure and XStream Inbound Servers

Inbound servers ignore instantiation SCNs. This procedure has no effect on XStream inbound servers.

The SET_TABLE_INSTANTIATION_SCN Procedure and CDBs

In a CDB, this procedure must be invoked from the same container as the apply process that uses the instantiation SCN information.

24.4.24 SET_UPDATE_CONFLICT_HANDLER Procedure

This procedure adds, modifies, or removes a prebuilt update conflict handler for the specified object.

Syntax

```
DBMS_APPLY_ADM.SET_UPDATE_CONFLICT_HANDLER(
   object_name          IN  VARCHAR2,
```

```
  method_name          IN  VARCHAR2,
  resolution_column    IN  VARCHAR2,
  column_list          IN  DBMS_UTILITY.NAME_ARRAY,
  apply_database_link  IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 24-27 SET_UPDATE_CONFLICT_HANDLER Procedure Parameters

Parameter	Description
`object_name`	The schema and name of the table, specified as `[schema_name.]object_name`, for which an update conflict handler is being added, modified, or removed. For example, if an update conflict handler is being added for table `employees` owned by user `hr`, then specify `hr.employees`. If the schema is not specified, then the current user is the default.
`method_name`	Type of update conflict handler to create. You can specify one of the prebuilt handlers, which determine whether the column list from the source database is applied for the row or whether the values in the row at the destination database are retained: • `MAXIMUM`: Applies the column list from the source database if it has the greater value for the resolution column. Otherwise, retains the values at the destination database. • `MINIMUM`: Applies the column list from the source database if it has the lesser value for the resolution column. Otherwise, retains the values at the destination database. • `OVERWRITE`: Applies the column list from the source database, overwriting the column values at the destination database. • `DISCARD`: Retains the column list from the destination database, discarding the column list from the source database. If `NULL`, then the procedure removes any existing update conflict handler with the same `object_name`, `resolution_column`, and `column_list`. If non-`NULL`, then the procedure replaces any existing update conflict handler with the same `object_name` and `resolution_column`.
`resolution_column`	Name of the column used to uniquely identify an update conflict handler. For the `MAXIMUM` and `MINIMUM` prebuilt methods, the resolution column is also used to resolve the conflict. The resolution column must be one of the columns listed in the `column_list` parameter. `NULL` is not allowed for this parameter. For the `OVERWRITE` and `DISCARD` prebuilt methods, you can specify any column in the column list.
`column_list`	List of columns for which the conflict handler is called. The same column cannot be in more than one column list. If a conflict occurs for one or more of the columns in the list when an apply component tries to apply a row logical change record (row LCR), then the conflict handler is called to resolve the conflict. The conflict handler is not called if a conflict occurs only for columns that are not in the list. **Note:** Prebuilt update conflict handlers do not support LOB, `LONG`, `LONG RAW`, user-defined type, and Oracle-suppled type columns. Therefore, you should not include these types of columns in the `column_list` parameter.

Table 24-27 (Cont.) SET_UPDATE_CONFLICT_HANDLER Procedure Parameters

Parameter	Description
`apply_database_link`	The name of the database link to a non-Oracle database. This parameter should be set only when the destination database is a non-Oracle database. **Note:** Currently, conflict handlers are not supported when applying changes to a non-Oracle database.

Usage Notes

The following usage notes apply to this procedure:

- Modifying an Existing Update Conflict Handler
- Removing an Existing Update Conflict Handler
- Series of Actions for Conflicts
- Procedure DML Handlers for Conflicts
- A Column Can Be in Only One Column List
- Update Conflict Handlers and Non-Oracle Databases
- The SET_UPDATE_CONFLICT_HANDLER Procedure and XStream Outbound Servers
- The SET_UPDATE_CONFLICT_HANDLER Procedure and XStream Inbound Servers

Modifying an Existing Update Conflict Handler

If you want to modify an existing update conflict handler, then you specify the table and resolution column of an the existing update conflict handler. You can modify the prebuilt method or the column list.

Removing an Existing Update Conflict Handler

If you want to remove an existing update conflict handler, then specify `NULL` for the prebuilt method and specify the table, column list, and resolution column of the existing update conflict handler.

Series of Actions for Conflicts

If an update conflict occurs, then Oracle completes the following series of actions:

1. Calls the appropriate update conflict handler to resolve the conflict
2. If no update conflict handler is specified or if the update conflict handler cannot resolve the conflict, then calls the appropriate error handler for the apply component, table, and operation to handle the error
3. If no error handler is specified or if the error handler cannot resolve the error, then raises an error and moves the transaction containing the row LCR that caused the error to the error queue

> **See Also:**
>
> "Signature of a DML Handler Procedure or Error Handler Procedure" for information about setting an error handler

Procedure DML Handlers for Conflicts

If you cannot use a prebuilt update conflict handler to meet your requirements, then you can create a PL/SQL procedure to use as a custom conflict handler. You use the `SET_DML_HANDLER` procedure to designate one or more custom conflict handlers for a particular table. In addition, a custom conflict handler can process LOB columns and use LOB assembly.

> **See Also:**
>
> SET_DML_HANDLER Procedure

A Column Can Be in Only One Column List

When a column is in a column list, and you try to add the same column to another column list, this procedure returns the following error:

```
ORA-00001: UNIQUE CONSTRAINT (SYS.APPLY$_CONF_HDLR_COLUMNS_UNQ1) VIOLATED
```

Update Conflict Handlers and Non-Oracle Databases

Setting an update conflict handler for an apply component that is applying to a non-Oracle database is not supported.

The SET_UPDATE_CONFLICT_HANDLER Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The SET_UPDATE_CONFLICT_HANDLER Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

Examples

The following is an example for setting an update conflict handler for the `employees` table in the `hr` schema:

```
DECLARE
  cols  DBMS_UTILITY.NAME_ARRAY;
BEGIN
  cols(1) := 'salary';
  cols(2) := 'commission_pct';
  DBMS_APPLY_ADM.SET_UPDATE_CONFLICT_HANDLER(
    object_name           =>  'hr.employees',
    method_name           =>  'MAXIMUM',
    resolution_column     =>  'salary',
    column_list           =>  cols);
END;
/
```

This example sets a conflict handler that is called if a conflict occurs for the `salary` or `commission_pct` column in the `hr.employees` table. If such a conflict occurs, then the `salary` column is evaluated to resolve the conflict. If a conflict occurs only for a column that is not in the column list, such as the `job_id` column, then this conflict handler is not called.

24.4.25 SET_VALUE_DEPENDENCY Procedure

This procedure sets or removes a value dependency. A value dependency is a virtual dependency definition that defines a relationship between the columns of two or more tables.

An apply component uses the name of a value dependencies to detect dependencies between row logical change records (row LCRs) that contain the columns defined in the value dependency. Value dependencies can define virtual foreign key relationships between tables, but, unlike foreign key relationships, value dependencies can involve more than two database objects.

This procedure is overloaded. The `attribute_list` and `attribute_table` parameters are mutually exclusive.

Syntax

```
DBMS_APPLY_ADM.SET_VALUE_DEPENDENCY(
   dependency_name IN VARCHAR2,
   object_name     IN VARCHAR2,
   attribute_list  IN VARCHAR2);

DBMS_APPLY_ADM.SET_VALUE_DEPENDENCY(
   dependency_name IN VARCHAR2,
   object_name     IN VARCHAR2,
   attribute_table IN DBMS_UTILITY.NAME_ARRAY);
```

Parameters

Table 24-28 SET_VALUE_DEPENDENCY Procedure Parameters

Parameter	Description
`dependency_name`	The name of the value dependency. If a dependency with the specified name does not exist, then it is created. If a dependency with the specified name exists, then the specified object and attributes are added to the dependency. If `NULL`, an error is raised.
`object_name`	The name of the table, specified as `[`*`schema_name`*`.]`*`table_name`*. For example, `hr.employees`. If the schema is not specified, then the current user is the default. If `NULL` and the specified dependency exists, then the dependency is removed. If `NULL` and the specified dependency does not exist, then an error is raised. If `NULL`, then `attribute_list` and `attribute_table` also must be `NULL`.
`attribute_list`	A comma-delimited list of column names in the table. There must be no spaces between entries.
`attribute_table`	A PL/SQL associative array of type `DBMS_UTILITY.NAME_ARRAY` that contains names of columns in the table. The first column name should be at position 1, the second at position 2, and so on. The table does not need to be `NULL` terminated.

Usage Notes

The following usage notes apply to this procedure:

- The SET_VALUE_DEPENDENCY Procedure and XStream Outbound Servers
- The SET_VALUE_DEPENDENCY Procedure and XStream Inbound Servers

The SET_VALUE_DEPENDENCY Procedure and XStream Outbound Servers

This procedure has no effect on XStream outbound servers.

The SET_VALUE_DEPENDENCY Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.26 START_APPLY Procedure

This procedure directs the apply component to start applying messages.

Syntax

```
DBMS_APPLY_ADM.START_APPLY(
   apply_name  IN  VARCHAR2);
```

Parameter

Table 24-29 START_APPLY Procedure Parameter

Parameter	Description
apply_name	The apply component name. A NULL setting is not allowed. Do not specify an owner.

Usage Notes

The following usage notes apply to this procedure:

- Apply Component Status
- The START_APPLY Procedure and XStream Outbound Servers
- The START_APPLY Procedure and XStream Inbound Servers

Apply Component Status

The apply component status is persistently recorded. Hence, if the status is ENABLED, then the apply component is started upon database instance startup. An apply component (*annn*) is an Oracle background process. The enqueue and dequeue state of DBMS_AQADM.START_QUEUE and DBMS_AQADM.STOP_QUEUE have no effect on the start status of an apply component.

The START_APPLY Procedure and XStream Outbound Servers

This procedure functions the same way for apply processes and outbound servers.

The START_APPLY Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

24.4.27 STOP_APPLY Procedure

This procedure stops the apply component from applying messages and rolls back any unfinished transactions being applied.

Syntax

```
DBMS_APPLY_ADM.STOP_APPLY(
   apply_name  IN  VARCHAR2,
   force       IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 24-30 STOP_APPLY Procedure Parameters

Parameter	Description
apply_name	The apply component name. A NULL setting is not allowed. Do not specify an owner.
force	If TRUE, then the procedure stops the apply component as soon as possible. If FALSE, then the procedure stops the apply component after ensuring that there are no gaps in the set of applied transactions. The behavior of the apply component depends on the setting specified for the force parameter and the setting specified for the commit_serialization apply component parameter. See "Usage Notes" for more information.

Usage Notes

The following usage notes apply to this procedure:

- Apply Component Status
- Queue Subprograms Have No Effect on Apply Component Status
- The STOP_APPLY force Parameter and the commit_serialization Apply Parameter
- The STOP_APPLY Procedure and XStream Outbound Servers
- The STOP_APPLY Procedure and XStream Inbound Servers

Apply Component Status

The apply component status is persistently recorded. Hence, if the status is DISABLED or ABORTED, then the apply component is not started upon database instance startup.

Queue Subprograms Have No Effect on Apply Component Status

The enqueue and dequeue state of DBMS_AQADM.START_QUEUE and DBMS_AQADM.STOP_QUEUE have no effect on the STOP status of an apply component.

The STOP_APPLY force Parameter and the commit_serialization Apply Parameter

The following table describes apply component behavior for each setting of the force parameter in the STOP_APPLY procedure and the commit_serialization apply component parameter. In all cases, the apply component rolls back any unfinished transactions when it stops.

force	commit_serialization	Apply Component Behavior
TRUE	FULL	The apply component stops immediately and does not apply any unfinished transactions.
TRUE	DEPENDENT_TRANSACTIONS	When the apply component stops, some transactions that have been applied locally might have committed at the source database at a later point in time than some transactions that have not been applied locally.
FALSE	FULL	The apply component stops after applying the next uncommitted transaction in the commit order, if any such transaction is in progress.
FALSE	DEPENDENT_TRANSACTIONS	Before stopping, the apply component applies all of the transactions that have a commit time that is earlier than the applied transaction with the most recent commit time.

For example, assume that the `commit_serialization` apply component parameter is set to `DEPENDENT_TRANSACTIONS` and there are three transactions: transaction 1 has the earliest commit time, transaction 2 is committed after transaction 1, and transaction 3 has the latest commit time. Also assume that an apply component has applied transaction 1 and transaction 3 and is in the process of applying transaction 2 when the `STOP_APPLY` procedure is run. Given this scenario, if the `force` parameter is set to `TRUE`, then transaction 2 is not applied, and the apply component stops (transaction 2 is rolled back). If, however, the `force` parameter is set to `FALSE`, then transaction 2 is applied before the apply component stops.

A different scenario would result if the `commit_serialization` apply component parameter is set to `FULL`. For example, assume that the `commit_serialization` apply component parameter is set to `FULL` and there are three transactions: transaction A has the earliest commit time, transaction B is committed after transaction A, and transaction C has the latest commit time. In this case, the apply component has applied transaction A and is in the process of applying transactions B and C when the `STOP_APPLY` procedure is run. Given this scenario, if the `force` parameter is set to `TRUE`, then transactions B and C are not applied, and the apply component stops (transactions B and C are rolled back). If, however, the `force` parameter is set to `FALSE`, then transaction B is applied before the apply component stops, and transaction C is rolled back.

See Also:

SET_PARAMETER Procedure for more information about the `commit_serialization` apply component parameter

The STOP_APPLY Procedure and XStream Outbound Servers

This procedure functions the same way for apply processes and outbound servers.

The STOP_APPLY Procedure and XStream Inbound Servers

This procedure functions the same way for apply processes and inbound servers.

25
DBMS_AQ

The DBMS_AQ package provides an interface to Oracle Advanced Queuing (AQ).

This chapter contains the following topics:

- Security Model
- Constants
- Data Structures
- Operational Notes
- Summary of DBMS_AQ Subprograms

> **See Also:**
>
> - *Oracle Database Advanced Queuing User's Guide*
> - Oracle Database Advanced Queuing (AQ) Types for information about TYPEs to use with DBMS_AQ.

25.1 DBMS_AQ Security Model

Initially, only SYS and SYSTEM have execution privilege for the procedures in DBMS_AQ and DBMS_AQADM.

Queue Security

To enqueue or dequeue, users need EXECUTE rights on DBMS_AQ and either ENQUEUE or DEQUEUE privileges on target queues, or ENQUEUE_ANY/DEQUEUE_ANY system privileges. Users who have been granted EXECUTE rights to DBMS_AQ and DBMS_AQADM are able to create, manage, and use queues in their own schemas. The MANAGE_ANY AQ system privilege is used to create and manage queues in other schemas.

As a database user, you do not need any explicit object-level or system-level privileges to enqueue or dequeue to queues in your own schema other than the EXECUTE right on DBMS_AQ.

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* for more information on queue privileges and access control.

OCI Applications and Queue Access

For an Oracle Call Interface (OCI) application to access a queue, the session user must be granted either the object privilege of the queue he intends to access or the ENQUEUE ANY QUEUE

or DEQUEUE ANY QUEUE system privileges. The EXECUTE right of DBMS_AQ is not checked against the session user's rights.

Security Required for Propagation

Propagation jobs are owned by SYS, but the propagation occurs in the security context of the queue table owner. Previously propagation jobs were owned by the user scheduling propagation, and propagation occurred in the security context of the user setting up the propagation schedule. The queue table owner must be granted EXECUTE privileges on the DBMS_AQADM package. Otherwise, the Oracle Database snapshot processes do not propagate and generate trace files with the error identifier SYS.DBMS_AQADM not defined. Private database links owned by the queue table owner can be used for propagation. The username specified in the connection string must have EXECUTE access on the DBMS_AQ and DBMS_AQADM packages on the remote database.

See Also:

- DBMS_AQADM
- *Oracle Database Advanced Queuing User's Guide* for more information on security required for propagation

25.2 DBMS_AQ Constants

This topic describes the constants used by DBMS_AQ.

The DBMS_AQ package uses the constants shown in the following table.

When using enumerated constants such as BROWSE, LOCKED, or REMOVE, the PL/SQL constants must be specified with the scope of the packages defining it. All types associated with the operational interfaces have to be prepended with DBMS_AQ. For example: DBMS_AQ.BROWSE.

Note:

The sequence_deviation attribute has no effect in releases prior to Oracle Streams AQ 10g Release 1 (10.1) if message_grouping parameter of DBMS_AQADM subprograms is set to TRANSACTIONAL. The sequence deviation feature is deprecated in Oracle Streams AQ 10g Release 2 (10.2).

Table 25-1 Enumerated Constants

Parameter	Options	Type	Description
VISIBILITY	IMMEDIATE	-	-
-	ON_COMMIT	-	-
DEQUEUE_MODE	BROWSE	-	-
-	LOCKED	-	-
-	REMOVE	-	-

Table 25-1 (Cont.) Enumerated Constants

Parameter	Options	Type	Description
-	REMOVE_NODATA	-	-
NAVIGATION	FIRST_MESSAGE	-	-
-	NEXT_MESSAGE	-	-
STATE	WAITING	-	-
-	READY	-	-
-	PROCESSED	-	-
-	EXPIRED	-	-
SEQUENCE_DEVIATION	BEFORE	-	-
-	TOP	-	-
WAIT	FOREVER	BINARY_INTEGER	-
-	NO_WAIT	BINARY_INTEGER	-
DELAY	NO_DELAY	-	-
EXPIRATION	NEVER	-	-
NAMESPACE	NAMESPACE_AQ	-	-
-	NAMESPACE_ANONYMOUS	-	-
NTFN_GROUPING_CLASS	NFTN_GROUPING_CLASS_TIME	NUMBER	-
NTFN_GROUPING_TYPE	NTFN_GROUPING_TYPE_SUMMARY	NUMBER	-
-	NTFN_GROUPING_TYPE_LAST	NUMBER	-
NTFN_GROUPING_REPEAT_COUNT	NTFN_GROUPING_FOREVER	NUMBER	-

25.3 DBMS_AQ Data Types

This topic lists and describes DBMS_AQ data types.

Table 25-2 DBMS_AQ Data Structures

Data Structures	Description
OBJECT_NAME	Names database objects
TYPE_NAME	Defines queue types
JSON	Starting from Oracle Database Release 21c, AQ also supports JSON payload type. A dedicated JSON data type permits us to store the JSON data in a post-parse binary format which allows much faster access to nested JSON values.
Oracle Database Advanced Queuing PL/SQL Callback	Specifies the user-defined PL/SQL procedure, defined in the database to be invoked on message notification

OBJECT_NAME

The `object_name` data structure names database objects. It applies to queues, queue tables, agent names, and object types.

Syntax

```
object_name := VARCHAR2;
object_name := [schema_name.]name;
```

Usage Notes

Names for objects are specified by an optional schema name and a name. If the schema name is not specified, the current schema is assumed. The name must follow object name guidelines in *Oracle Database SQL Language Reference* with regard to reserved characters. Schema names, agent names, and object type names can be up to 128 bytes long. Queue names and queue table names can be up to 122 bytes long. Maximum length of agent names and subscriber names can be 128 characters with datablock size greater than 2k. For 2k datablock size, maximum length for subscriber is 128 bytes.

TYPE_NAME

The `type_name` data structure defines queue types.

Syntax

```
type_name := VARCHAR2;
type_name := object_type | "RAW";
```

Attributes

Table 25-3 Type Name Attributes

Attribute	Description
object_type	Maximum number of attributes in the object type is limited to 900.
"RAW"	To store payload of type `RAW`, Oracle Database Advanced Queuing creates a queue table with a `LOB` column as the payload repository. The theoretical maximum size of the message payload is the maximum amount of data that can be stored in a `LOB` column. However, the maximum size of the payload is determined by which programmatic environment you use to access Oracle Database Advanced Queuing. For PL/SQL, Java and precompilers the limit is 32K; for the OCI the limit is 4G. Because the PL/SQL enqueue and dequeue interfaces accept `RAW` buffers as the payload parameters you will be limited to 32K bytes. In OCI, the maximum size of your `RAW` data will be limited to the maximum amount of contiguous memory (as an `OCIRaw` is simply an array of bytes) that the OCI Object Cache can allocate. Typically, this will be at least 32K bytes and much larger in many cases. Because `LOB` columns are used for storing `RAW` payload, the Oracle Database Advanced Queuing administrator can choose the `LOB` tablespace and configure the `LOB` storage by constructing a `LOB` storage string in the `storage_clause` parameter during queue table creation time.

Oracle Database Advanced Queuing PL/SQL Callback

The `plsqlcallback` data structure specifies the user-defined PL/SQL procedure, defined in the database to be invoked on message notification.

Syntax

If a notification message is expected for a `RAW` payload enqueue, then the PL/SQL callback must have the following signature:

```
procedure plsqlcallback(
  context  IN  RAW,
  reginfo  IN  SYS.AQ$_REG_INFO,
  descr    IN  SYS.AQ$_DESCRIPTOR,
  payload  IN  RAW,
  payloadl IN  NUMBER);
```

Attributes

Table 25-4 Oracle Database Advanced Queuing PL/SQL Callback Attributes

Attribute	Description
`context`	Specifies the context for the callback function that was passed by `dbms_aq.register`. See AQ$_REG_INFO Type.
`reginfo`	See AQ$_REG_INFO Type.
`descr`	See AQ$_DESCRIPTOR Type
`payload`	If a notification message is expected for a raw payload enqueue then this contains the raw payload that was enqueued into a non persistent queue. In case of a persistent queue with raw payload this parameter will be null.
`payloadl`	Specifies the length of `payload`. If `payload` is null, `payload1` = 0.

If the notification message is expected for an ADT payload enqueue, the PL/SQL callback must have the following signature:

```
procedure plsqlcallback(
  context  IN RAW,
  reginfo  IN SYS.AQ$_REG_INFO,
  descr    IN SYS.AQ$_DESCRIPTOR,
  payload  IN VARCHAR2,
  payloadl IN NUMBER);
```

25.4 DBMS_AQ Operational Notes

This topic lists various DBMS_AQ operational notes.

DBMS_AQ and DBMS_AQADM Java Classes

Java interfaces are available for `DBMS_AQ` and `DBMS_AQADM`. The Java interfaces are provided in the `$ORACLE_HOME/rdbms/jlib/aqapi.jar`. Users are required to have `EXECUTE` privileges on the `DBMS_AQIN` package to use these interfaces.

25.5 Summary of DBMS_AQ Subprograms

The `DBMS_AQ` package uses subprograms described in this table.

Table 25-5 DBMS_AQ Package Subprograms

Subprograms	Description
BIND_AGENT Procedure	Creates an entry for an Oracle Database Advanced Queuing agent in the LDAP directory
DEQUEUE Procedure	Dequeues a message from the specified queue
DEQUEUE_ARRAY Function	Dequeues an array of messages from the specified queue
ENQUEUE Procedure	Adds a message to the specified queue
ENQUEUE_ARRAY Function	Adds an array of messages to the specified queue
LISTEN Procedures	Listen to one or more queues on behalf of a list of agents
POST Procedure	Posts to a anonymous subscription which allows all clients who are registered for the subscription to get notifications
REGISTER Procedure	Registers for message notifications
SEEK Procedure	Enables a subscriber to seek any particular message of choice
UNBIND_AGENT Procedure	Removes an entry for an Oracle Database Advanced Queuing agent from the LDAP directory
UNREGISTER Procedure	Unregisters a subscription which turns off notification

> **Note:**
>
> `DBMS_AQ` does not have a purity level defined; therefore, you cannot call any procedure in this package from other procedures that have `RNDS`, `WNDS`, `RNPS` or `WNPS` constraints defined.

25.5.1 BIND_AGENT Procedure

This procedure creates an entry for an Oracle Database Advanced Queuing agent in the LDAP server.

Syntax

```
DBMS_AQ.BIND_AGENT(
   agent        IN SYS.AQ$_AGENT,
   certificate  IN VARCHAR2 default NULL);
```

Parameters

Table 25-6 BIND_AGENT Procedure Parameters

Parameter	Description
`agent`	Agent that is to be registered in LDAP server.
`certificate`	Location (LDAP distinguished name) of the "organizationalperson" entry in LDAP whose digital certificate (attribute `usercertificate`) is to be used for this agent. Example: "cn=OE, cn=ACME, cn=com" is a distinguished name for a OrganizationalPerson OE whose certificate will be used with the specified agent.

Usage Notes

In the LDAP server, digital certificates are stored as an attribute (`usercertificate`) of the `OrganizationalPerson` entity. The distinguished name for this `OrganizationalPerson` must be specified when binding the agent.

25.5.2 DEQUEUE Procedure

This procedure dequeues a message from the specified queue.

Syntax

```
DBMS_AQ.DEQUEUE (
   queue_name           IN      VARCHAR2,
   dequeue_options      IN      dequeue_options_t,
   message_properties   OUT     message_properties_t,
   payload              OUT     "<ADT_1>"
   msgid                OUT     RAW);
```

Parameters

Table 25-7 DEQUEUE Procedure Parameters

Parameter	Description
`queue_name`	Specifies the name of the queue.
`dequeue_options`	See DEQUEUE_OPTIONS_T Type.
`message_properties`	See MESSAGE_PROPERTIES_T Type.
`payload`	Not interpreted by Oracle Database Advanced Queuing. The payload must be specified according to the specification in the associated queue table. For the definition of *`type_name`* refer to `TYPE_NAME` in DBMS_AQ Data Types.
`msgid`	System generated identification of the message.

Usage Notes

The search criteria for messages to be dequeued is determined by the following parameters in `dequeue_options`:

- `consumer_name`
- `msgid`

  `Msgid` uniquely identifies the message to be dequeued. Only messages in the `READY` state are dequeued unless `msgid` is specified.

- `correlation`

  Correlation identifiers are application-defined identifiers that are not interpreted by Oracle Database Advanced Queuing.

- `deq_condition`

  Dequeue condition is an expression based on the message properties, the message data properties and PL/SQL functions. A `deq_condition` is specified as a Boolean expression using syntax similar to the `WHERE` clause of a SQL query. This Boolean expression can include conditions on message properties, user data properties (object payloads only), and

PL/SQL or SQL functions (as specified in the where clause of a SQL query). Message properties include `priority`, `corrid` and other columns in the queue table.

To specify dequeue conditions on a message payload (object payload), use attributes of the object type in clauses. You must prefix each attribute with `tab.user_data` as a qualifier to indicate the specific column of the queue table that stores the payload.

Example: `tab.user_data.orderstatus='EXPRESS'`

The dequeue order is determined by the values specified at the time the queue table is created unless overridden by the `msgid` and correlation ID in `dequeue_options`.

The database-consistent read mechanism is applicable for queue operations. For example, a `BROWSE` call may not see a message that is enqueued after the beginning of the browsing transaction.

The default `NAVIGATION` parameter during dequeue is `NEXT_MESSAGE`. This means that subsequent dequeues will retrieve the messages from the queue based on the snapshot obtained in the first dequeue. In particular, a message that is enqueued after the first dequeue command will be processed only after processing all the remaining messages in the queue. This is usually sufficient when all the messages have already been enqueued into the queue, or when the queue does not have a priority-based ordering. However, applications must use the `FIRST_MESSAGE` navigation option when the first message in the queue needs to be processed by every dequeue command. This usually becomes necessary when a higher priority message arrives in the queue while messages already-enqueued are being processed.

> **Note:**
>
> It may be more efficient to use the `FIRST_MESSAGE` navigation option when messages are concurrently enqueued. If the `FIRST_MESSAGE` option is not specified, Oracle Database Advanced Queuing continually generates the snapshot as of the first dequeue command, leading to poor performance. If the `FIRST_MESSAGE` option is specified, then Oracle Database Advanced Queuing uses a new snapshot for every dequeue command.

Messages enqueued in the same transaction into a queue that has been enabled for message grouping will form a group. If only one message is enqueued in the transaction, then this will effectively form a group of one message. There is no upper limit to the number of messages that can be grouped in a single transaction.

In queues that have not been enabled for message grouping, a dequeue in `LOCKED` or `REMOVE` mode locks only a single message. By contrast, a dequeue operation that seeks to dequeue a message that is part of a group will lock the entire group. This is useful when all the messages in a group need to be processed as an atomic unit.

When all the messages in a group have been dequeued, the dequeue returns an error indicating that all messages in the group have been processed. The application can then use the `NEXT_TRANSACTION` to start dequeuing messages from the next available group. In the event that no groups are available, the dequeue will time out after the specified `WAIT` period.

Using Secure Queues

For secure queues, you must specify `consumer_name` in the `dequeue_options` parameter. See DEQUEUE_OPTIONS_T Type for more information about `consumer_name`.

When you use secure queues, the following are required:

- You must have created a valid Oracle Database Advanced Queuing agent using `DBMS_AQADM.CREATE_AQ_AGENT`. See CREATE_AQ_AGENT Procedure .
- You must map the Oracle Database Advanced Queuing agent to a database user with dequeue privileges on the secure queue. Use `DBMS_AQADM.ENABLE_DB_ACCESS` to do this. See ENABLE_DB_ACCESS Procedure.

25.5.3 DEQUEUE_ARRAY Function

This function dequeues an array of messages and returns them in the form of an array of payloads, an array of message properties and an array of message IDs. This function returns the number of messages successfully dequeued.

Syntax

```
DBMS_AQ.DEQUEUE_ARRAY (
   queue_name                IN   VARCHAR2,
   dequeue_options           IN   dequeue_options_t,
   array_size                IN   pls_integer,
   message_properties_array  OUT  message_properties_array_t,
   payload_array             OUT  "<COLLECTION_1>",
   msgid_array               OUT  msgid_array_t,
   error_array               OUT  error_array_t)
RETURN pls_integer;
```

Parameters

Table 25-8 DEQUEUE_ARRAY Function Parameters

Parameter	Description
`queue_name`	The queue name from which messages are dequeued (same as single-row dequeue).
`dequeue_options`	The set of options which will be applied to all messages in the array (same as single-row dequeue).
`array_size`	The number of elements to dequeue. For buffered messages, `array_size` should be `1`.
`message_properties_array`	A record containing an array corresponding to each message property. Each payload element has a corresponding set of message properties. See MESSAGE_PROPERTIES_ARRAY_T Type.
`payload_array`	An array of dequeued payload data. "`<COLLECTION_1>`" can be an associative array, varray or nested table in its PL/SQL representation. Users can dequeue `JSON`, `RAW` and `ADT` payloads.
`msgid_array`	An array of message IDs of the dequeued messages. See MSGID_ARRAY_T Type.
`error_array`	Currently not implemented

Usage Notes

A nonzero wait time, as specified in `dequeue_options`, is recognized only when there are no messages in the queue. If the queue contains messages that are eligible for dequeue, then the `DEQUEUE_ARRAY` function will dequeue up to `array_size` messages and return immediately.

Dequeue by `message_id` is not supported. See DEQUEUE Procedure for more information on the `navigation` parameter. Existing `NAVIGATION` modes are supported. In addition, two new `NAVIGATION` modes are supported for queues enabled for message grouping:

- FIRST_MESSAGE_MULTI_GROUP
- NEXT_MESSAGE_MULTI_GROUP

> **See Also:**
>
> ENQUEUE_OPTIONS_T Type

For transaction grouped queues and ONE_GROUP navigation, messages are dequeued from a single transaction group only, subject to the array_size limit. In MULTI_GROUP navigation, messages are dequeued across multiple transaction groups, still subject to the array_size limit. ORA-25235 is returned to indicate the end of a transaction group.

DEQUEUE_ARRAY is not supported for buffered messages, but you can still use this procedure on individual buffered messages by setting array_size to one message.

25.5.4 ENQUEUE Procedure

This procedure adds a message to the specified queue.

Syntax

```
DBMS_AQ.ENQUEUE (
   queue_name          IN      VARCHAR2,
   enqueue_options     IN      enqueue_options_t,
   message_properties  IN      message_properties_t,
   payload             IN       "<ADT_1>",
   msgid               OUT     RAW);
```

Parameters

Table 25-9 ENQUEUE Procedure Parameters

Parameter	Description
queue_name	Specifies the name of the queue to which this message should be enqueued. The queue cannot be an exception queue.
enqueue_options	See ENQUEUE_OPTIONS_T Type.
message_properties	See MESSAGE_PROPERTIES_T Type.
payload	Not interpreted by Oracle Database Advanced Queuing. The payload must be specified according to the specification in the associated queue table. NULL is an acceptable parameter. For the definition of *type_name* refer to TYPE_NAME in DBMS_AQ Data Types.
msgid	System generated identification of the message. This is a globally unique identifier that can be used to identify the message at dequeue time.

Usage Notes

The sequence_deviation parameter in enqueue_options can be used to change the order of processing between two messages. The identity of the other message, if any, is specified by the enqueue_options parameter relative_msgid. The relationship is identified by the sequence_deviation parameter.

Specifying `sequence_deviation` for a message introduces some restrictions for the delay and priority values that can be specified for this message. The delay of this message must be less than or equal to the delay of the message before which this message is to be enqueued. The priority of this message must be greater than or equal to the priority of the message before which this message is to be enqueued.

> **Note:**
>
> The `sequence_deviation` attribute has no effect in releases prior to Oracle Streams AQ 10g Release 1 (10.1) if `message_grouping` parameter of `DBMS_AQADM` subprograms is set to `TRANSACTIONAL`. The sequence deviation feature is deprecated in Oracle Streams AQ 10g Release 2 (10.2).

If a message is enqueued to a multiconsumer queue with no recipient, and if the queue has no subscribers (or rule-based subscribers that match this message), then Oracle error `ORA_24033` is raised. This is a warning that the message will be discarded because there are no recipients or subscribers to whom it can be delivered.

Using Secure Queues

For secure queues, you must specify the `sender_id` in the `messages_properties` parameter. See MESSAGE_PROPERTIES_T Type for more information about `sender_id`.

When you use secure queues, the following are required:

- You must have created a valid Oracle Database Advanced Queuing agent using `DBMS_AQADM.CREATE_AQ_AGENT`. See CREATE_AQ_AGENT Procedure .
- You must map `sender_id` to a database user with enqueue privileges on the secure queue. Use `DBMS_AQADM.ENABLE_DB_ACCESS` to do this. See ENABLE_DB_ACCESS Procedure.

25.5.5 ENQUEUE_ARRAY Function

This function enqueues an array of payloads using a corresponding array of message properties. The output will be an array of message IDs of the enqueued messages.

Syntax

```
DBMS_AQ.ENQUEUE_ARRAY (
   queue_name                IN   VARCHAR2,
   enqueue_options           IN   enqueue_options_t,
   array_size                IN   pls_integer,
   message_properties_array  IN   message_properties_array_t,
   payload_array             IN   "<COLLECTION_1>",
   msgid_array               OUT  msgid_array_t,
   error_array               OUT  error_array_t)
RETURN pls_integer;
```

Parameters

Table 25-10 ENQUEUE_ARRAY Function Parameters

Parameter	Description
queue_name	The queue name in which messages are enqueued (same as single-row enqueue).
enqueue_options	See ENQUEUE_OPTIONS_T Type.
array_size	The number of elements to enqueue. For buffered messages, array_size should be 1.
message_properties_array	A record containing an array corresponding to each message property. For each property, the user must allocate array_size elements. See MESSAGE_PROPERTIES_ARRAY_T Type.
payload_array	An array of payload data. "<COLLECTION_1>" can be an associative array, VARRAY, or nested table in its PL/SQL representation. Users can enqueue JSON, RAW, and ADT payloads.
msgid_array	An array of message IDs for the enqueued messages. If an error occurs for a particular message, then its corresponding message ID is null. See MSGID_ARRAY_T Type.
error_array	Currently not implemented

Usage Notes

ENQUEUE_ARRAY is not supported for buffered messages, but you can still use this procedure on individual buffered messages by setting array_size to one message.

25.5.6 LISTEN Procedures

This procedure listens on one or more queues on behalf of a list of agents. The address field of the agent indicates the queue the agent wants to monitor. Only local queues are supported as addresses. Protocol is reserved for future use.

Syntax

```
DBMS_AQ.LISTEN (
   agent_list            IN    AQ$_AGENT_LIST_T,
   wait                  IN    BINARY_INTEGER DEFAULT DBMS_AQ.FOREVER,
   agent                 OUT   SYS.AQ$_AGENT);

DBMS_AQ.LISTEN (
   agent_list             IN   AQ$_AGENT_LIST_T,
   wait                   IN   BINARY_INTEGER DEFAULT FOREVER,
   listen_delivery_mode   IN   PLS_INTEGER DEFAULT DBMS_AQ.PERSISTENT,
   agent                  OUT  SYS.AQ$_AGENT,
   message_delivery_mode  OUT  PLS_INTEGER);

TYPE aq$_agent_list_t IS TABLE of aq$_agent INDEXED BY BINARY_INTEGER;
TYPE aq$_agent_list_t IS TABLE of aq$_agent INDEXED BY BINARY_INTEGER;
```

Parameters

Table 25-11 LISTEN Procedure Parameters

Parameter	Description
agent_list	List of agents to listen for
wait	Time out for the listen call in seconds. By default, the call will block forever.
listen_delivery_mode	The caller specifies whether it is interested in persistent, buffered messages or both types of messages, specifying a delivery mode of DBMS_AQ.PERSISTENT or DBMS_AQ.BUFFERED or DBMS_AQ.PERSISTENT_OR_BUFFERED
agent	Agent with a message available for consumption
message_delivery_mode	Returns the message type along with the queue and consumer for which there is a message

Usage Notes

If agent-address is a multiconsumer queue, then agent-name is mandatory. For single-consumer queues, agent-name must not be specified.

This procedure takes a list of agents as an argument. You specify the queue to be monitored in the address field of each agent listed. You also must specify the name of the agent when monitoring multiconsumer queues. For single-consumer queues, an agent name must not be specified. Only local queues are supported as addresses. Protocol is reserved for future use.

This is a blocking call that returns when there is a message ready for consumption for an agent in the list. If there are messages for more than one agent, only the first agent listed is returned. If there are no messages found when the wait time expires, an error is raised.

A successful return from the LISTEN call is only an indication that there is a message for one of the listed agents in one the specified queues. The interested agent must still dequeue the relevant message.

> **Note:**
>
> You cannot call LISTEN on nonpersistent queues.

25.5.7 POST Procedure

This procedure posts to a list of anonymous subscriptions that allows all clients who are registered for the subscriptions to get notifications.

Syntax

```
DBMS_AQ.POST (
  post_list       IN  SYS.AQ$_POST_INFO_LIST,
  post_count      IN  NUMBER);
```

Parameters

Table 25-12 POST Procedure Parameters

Parameter	Description
post_list	Specifies the list of anonymous subscriptions to which you want to post. It is a list of AQ$_POST_INFO_LIST Type.
post_count	Specifies the number of entries in the post_list.

Usage Notes

This procedure is used to post to anonymous subscriptions which allows all clients who are registered for the subscriptions to get notifications. Several subscriptions can be posted to at one time.

25.5.8 REGISTER Procedure

This procedure registers an e-mail address, user-defined PL/SQL procedure, or HTTP URL for message notification.

Syntax

```
DBMS_AQ.REGISTER (
   reg_list     IN  SYS.AQ$_REG_INFO_LIST,
   count        IN  NUMBER);
```

Parameters

Table 25-13 REGISTER Procedure Parameters

Parameter	Description
reg_list	Specifies the list of subscriptions to which you want to register for message notifications. It is a list of AQ$_REG_INFO Type.
count	Specifies the number of entries in the reg_list.

Usage Notes

- This procedure is used to register for notifications. You can specify an e-mail address to which message notifications are sent, register a procedure to be invoked on a notification, or register an HTTP URL to which the notification is posted. Interest in several subscriptions can be registered at one time.
- The procedure can also be used to register for grouping notifications using five grouping attributes:
  - Class – grouping criterion (currently only TIME criterion is supported)
  - Value – the value of the grouping criterion (currently only time in seconds for criterion TIME)
  - Type – summary or last, also contains count of notifications received in group (for AQ namespace only, not for DBCHANGE namespace)
  - Repeat count – how many times to perform grouping (Default is FOREVER)
  - Start time – when to start grouping (Default is current time)

- If you register for e-mail notifications, you should set the host name and port name for the SMTP server that will be used by the database to send e-mail notifications. If required, you should set the send-from e-mail address, which is set by the database as the `sent from` field. You need a Java-enabled database to use this feature.
- If you register for HTTP notifications, you may want to set the host name and port number for the proxy server and a list of no-proxy domains that will be used by the database to post HTTP notifications.

> **See Also:**
>
> DBMS_AQELM for more information on e-mail and HTTP notifications

25.5.9 SEEK Procedure

This procedure enables a subscriber to seek any particular message of choice. Once a seek operation is successful, the dequeues move from the new seek position onwards. A new call is added to facilitate the seek operation.

Syntax

```
DBMS_AQ.SEEK (
              queue_name         IN     VARCHAR2 not null,
              consumer_name      IN     VARCHAR2 default null,
              seek_input_array   IN     SEEK_INPUT_ARRAY_T,
              skip_option        IN     pls_integer DEFAULT DBMS_AQ.ERROR_IF_SKIPPED,
              seek_output_array OUT     SEEK_OUTPUT_ARRAY_T);
```

Parameters

Table 25-14 SEEK Procedure Parameters

Parameter	Description
`queue_name`	Specifies the name of the queue.
`consumer_name`	Specifies the name of the subscriber performing seek. `consumer_name` can be `NULL` for single-consumer queues.
`seek_input_array`	Array of type `seek_input_t`, where individual element in the array has per shard seek input. See SEEK_INPUT_T Type.
`skip_option`	Specifies the skip options. The following are the possible values: • `DBMS_AQ.ERROR_IF_SKIPPED` (default): Raises an error if seek operation can result in skipping some undequeued messsages when a seek forward is performed. • `DBMS_AQ.NO_DISCARD_SKIPPED`: Does not discard any skipped messages. Allows the seek operation to go through irrespective of the direction of the seek. Such undequeued messages before seek point will be shown as `READY_SKIPPED` in `AQ$<QUEUE_TABLE_NAME>`. • `DBMS_AQ.DISCARD_SKIPPED`: Discards any skipped messages. As soon as all the messages in a subshard are either dequeued or discarded by all its subscribers, the subshard is "retained" which will be truncated after retention time is over. Before retention time is over, the subscribers who have discarded the messages can seek back and dequeue the discarded messages if they wish to do so. The discarded messages will be shown as `READY_DISCARDED` in `AQ$<QUEUE_TABLE_NAME>`.

Table 25-14 (Cont.) SEEK Procedure Parameters

Parameter	Description
seek_output_array	Array of type seek_output_t, where individual element in the array has seeked from and seeked to fields per shard-priority. See SEEK_OUTPUT_T Type.

25.5.10 UNBIND_AGENT Procedure

This procedure removes the entry for an Oracle Database Advanced Queuing agent from the LDAP server.

Syntax

```
DBMS_AQ.UNBIND_AGENT(
   agent     IN SYS.AQ$_AGENT);
```

Parameters

Table 25-15 BIND_AGENT Procedure Parameters

Parameter	Description
agent	Agent that is to be removed from the LDAP server

25.5.11 UNREGISTER Procedure

This procedure unregisters a subscription which turns off notifications.

Syntax

```
DBMS_AQ.UNREGISTER (
   reg_list      IN  SYS.AQ$_REG_INFO_LIST,
   reg_count     IN  NUMBER);
```

Parameters

Table 25-16 UNREGISTER Procedure Parameters

Parameter	Description
reg_list	Specifies the list of subscriptions to which you want to register for message notifications. It is a list of AQ$_REG_INFO Type.
reg_count	Specifies the number of entries in the reg_list.

Usage Notes

This procedure is used to unregister a subscription which turns off notifications. Several subscriptions can be unregistered from at one time.

26
DBMS_AQADM

The DBMS_AQADM package provides procedures to manage Oracle Database Advanced Queuing (AQ) configuration and administration information.

> **See Also:**
>
> - *Oracle Database Advanced Queuing User's Guide*
> - Oracle Database Advanced Queuing (AQ) Types for information about the TYPEs to use with DBMS_AQADM

This chapter contains the following topics:

- Security Model
- Constants
- Subprogram Groups
- Summary of DBMS_AQADM Subprograms

26.1 DBMS_AQADM Security Model

Initially, only SYS and SYSTEM have execution privilege for the procedures in DBMS_AQADM and DBMS_AQ. Users who have been granted EXECUTE rights to DBMS_AQ and DBMS_AQADM are able to create, manage, and use queues in their own schemas. The MANAGE_ANY AQ system privilege is used to create and manage queues in other schemas and can be granted and revoked through DBMS_AQADM.GRANT_SYSTEM_PRIVILEGE and DBMS_AQADM.REVOKE_SYSTEM_PRIVILEGE respectively. Starting from Oracle Database 12*c* Release 2, MANAGE_ANY privilege will not allow access to SYS owned queues by users other than SYS.

User Roles

The database administrator has the option of granting the system privileges ENQUEUE_ANY and DEQUEUE_ANY, exercising DBMS_AQADM.GRANT_SYSTEM_PRIVILEGE and DBMS_AQADM.REVOKE_SYSTEM_PRIVILEGE directly to a database user, if you want the user to have this level of control.

The application developer gives rights to a queue by granting and revoking privileges at the object level by exercising DBMS_AQADM.GRANT_QUEUE_PRIVILEGE and DBMS_AQADM.REVOKE_QUEUE_PRIVILEGE. Starting from Oracle Database 12*c* Release 2, ENQUEUE_ANY and DEQUEUE_ANY privileges will not allow access to SYS owned queues by users other than SYS.

> **See Also:**
>
> - DBMS_AQ.
> - *Oracle Database Advanced Queuing User's Guide* for more information on queue privileges and access control.

Security Required for Propagation

Propagation jobs are owned by `SYS`, but the propagation occurs in the security context of the queue table owner. Previously propagation jobs were owned by the user scheduling propagation, and propagation occurred in the security context of the user setting up the propagation schedule. The queue table owner must be granted `EXECUTE` privileges on the `DBMS_AQADM` package. Otherwise, the Oracle Database snapshot processes do not propagate and generate trace files with the error identifier `SYS.DBMS_AQADM` not defined. Private database links owned by the queue table owner can be used for propagation. The username specified in the connection string must have `EXECUTE` access on the `DBMS_AQ` and `DBMS_AQADM` packages on the remote database.

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* for more information on security required for propagation.

Queue Table Migration

The `MIGRATE_QUEUE_TABLE` procedure requires that the `EXECUTE` privilege on `DBMS_AQADM` be granted to the queue table owner, who is probably an ordinary queue user. If you do not want ordinary queue users to be able to create and drop queues and queue tables, add and delete subscribers, and so forth, then you must revoke the `EXECUTE` privilege as soon as the migration is done.

> **See Also:**
>
> - "MIGRATE_QUEUE_TABLE Procedure."
> - *Oracle Database Advanced Queuing User's Guide* for more information on granting Oracle Database Advanced Queuing system privileges.

26.2 DBMS_AQADM Constants

When using enumerated constants, such as `INFINITE`, `TRANSACTIONAL`, or `NORMAL_QUEUE`, the symbol must be specified with the scope of the packages defining it. All types associated with the administrative interfaces must be prepended with `DBMS_AQADM`. For example: `DBMS_AQADM.NORMAL_QUEUE`.

Table 26-1 Enumerated Types in the Administrative Interface

Parameter	Options
retention	0, 1, 2...INFINITE
message_grouping	TRANSACTIONAL, NONE
queue_type	NORMAL_QUEUE, EXCEPTION_QUEUE, NON_PERSISTENT_QUEUE

> **See Also:**
>
> For more information on the Java classes and data structures used in both DBMS_AQ and DBMS_AQADM, see the DBMS_AQ package.

26.3 DBMS_AQADM Subprogram Groups

This section lists and describes the DBMS_AQADM subprogram groups.

This DBMS_AQADM package is made up of the following subprogram groups:

- Queue Table Subprograms
- Privilege Subprograms
- Queue Subprograms
- Subscriber Subprograms
- OKafka Subprograms
- Propagation Subprograms
- Miscellaneous Subprograms
- Oracle Database Advanced Queuing Agent Subprograms
- Alias Subprograms

26.3.1 DBMS_AQADM Queue Table Subprograms

This section lists and describes the DBMS_AQADM Queue Table subprograms.

Table 26-2 Queue Table Subprograms

Subprograms	Description
ALTER_QUEUE_TABLE Procedure	Alters the existing properties of a queue table
CREATE_QUEUE_TABLE Procedure	Creates a queue table for messages of a predefined type
DROP_QUEUE_TABLE Procedure	Drops an existing queue table
ENABLE_JMS_TYPES Procedure	A precondition for the enqueue of JMS types and XML types
MIGRATE_QUEUE_TABLE Procedure	Upgrades an 8.0-compatible queue table to an 8.1-compatible or higher queue table, or downgrades an 8.1-compatible or higher queue table to an 8.0-compatible queue table

Table 26-2 (Cont.) Queue Table Subprograms

Subprograms	Description
PURGE_QUEUE_TABLE Procedure	Purges messages from queue tables

26.3.2 DBMS_AQADM Privilege Subprograms

This sections lists and describes the DBMS_AQADM Privilege subprograms.

Table 26-3 Privilege Subprograms

Subprograms	Description
GRANT_QUEUE_PRIVILEGE Procedure	Grants privileges on a queue to users and roles
GRANT_SYSTEM_PRIVILEGE Procedure	Grants Oracle Database Advanced Queuing system privileges to users and roles
REVOKE_QUEUE_PRIVILEGE Procedure	Revokes privileges on a queue from users and roles
REVOKE_SYSTEM_PRIVILEGE Procedure	Revokes Oracle Database Advanced Queuing system privileges from users and roles

26.3.3 DBMS_AQADM Queue Subprograms

This sections lists and describes the DBMS_AQADM Queue subprograms.

Table 26-4 Queue Subprograms

Subprograms	Description
ALTER_QUEUE Procedure	Alters existing properties of a queue
CREATE_NP_QUEUE Procedure	Creates a nonpersistent `RAW` queue
CREATE_QUEUE Procedure	Creates a queue in the specified queue table
CREATE_SHARDED_QUEUE Procedure	Creates a queue and its queue table for a sharded queue all together.
DROP_SHARDED_QUEUE Procedure	Drops an existing sharded queue from the database queuing system
ALTER_SHARDED_QUEUE Procedure	Alters an sharded queue in the database queuing system
CREATE_EXCEPTION_QUEUE Procedure	Creates an exception queue for a sharded queue
DROP_QUEUE Procedure	Drops an existing queue
QUEUE_SUBSCRIBERS Function	Returns the subscribers to an 8.0-compatible multiconsumer queue in the PL/SQL index by table collection type `DBMS_AQADM.AQ$_subscriber_list_t`
START_QUEUE Procedure	Enables the specified queue for enqueuing or dequeuing
STOP_QUEUE Procedure	Disables enqueuing or dequeuing on the specified queue

26.3.4 DBMS_AQADM Subscriber Subprograms

This sections lists and describes the DBMS_AQADM Subscriber subprograms.

Table 26-5 Subscriber Subprograms

Subprograms	Description
ADD_SUBSCRIBER Procedure	Adds a default subscriber to a queue
ALTER_SUBSCRIBER Procedure	Alters existing properties of a subscriber to a specified queue
REMOVE_SUBSCRIBER Procedure	Removes a default subscriber from a queue

26.3.5 DBMS_AQADM OKafka Subprograms

This section lists and describes the DBMS_AQADM OKafka subprograms.

Table 26-6 OKafka Subprograms

Subprograms	Description
CREATE_DATABASE_KAFKA_TOPIC Procedure	Creates an OKafka Topic. OKafka topic is a multi-consumer transactional event queue specifically design to be used with OKafka application.
DROP_DATABASE_KAFKA_TOPIC Procedure	Drops an existing OKafka topic. OKafka topic is a multi-consumer transactional event queue specifically design to be used with OKafka application.

26.3.6 DBMS_AQADM Propagation Subprograms

This section lists and describes the DBMS_AQADM propagation subprograms.

Table 26-7 Propagation Subprograms

Subprograms	Description
ALTER_PROPAGATION_SCHEDULE Procedure	Alters parameters for a propagation schedule
DISABLE_PROPAGATION_SCHEDULE Procedure	Disables a propagation schedule
ENABLE_PROPAGATION_SCHEDULE Procedure	Enables a previously disabled propagation schedule
SCHEDULE_PROPAGATION Procedure	Schedules propagation of messages from a queue to a destination identified by a specific database link
UNSCHEDULE_PROPAGATION Procedure	Unschedules previously scheduled propagation of messages from a queue to a destination identified by a specific database link
VERIFY_QUEUE_TYPES Procedure	Verifies that the source and destination queues have identical types

26.3.7 DBMS_AQADM Miscellaneous Subprograms

This section lists and describes the DBMS_AQADM miscellaneous subprograms.

Table 26-8 Miscellaneous Subprograms

Subprograms	Description
GET_QUEUE_PARAMETER Procedure	Used to get different parameters for sharded queues at queue or database level.
GET_MAX_STREAMS_POOL Procedure	Retrieves the value of Oracle Database Advanced Queuing maximum streams pool memory limit
GET_MIN_STREAMS_POOL Procedure	Retrieves the value of Oracle Database Advanced Queuing minimum streams pool memory limit
GET_WATERMARK Procedure	Retrieves the value of watermark set by the `SET_WATERMARK` Procedure
SET_QUEUE_PARAMETER Procedure	Used to set different parameters for sharded queues at queue or database level.
SET_MAX_STREAMS_POOL Procedure	Used for Oracle Database Advanced Queuing to specify and limit maximum streams pool memory use
SET_MIN_STREAMS_POOL Procedure	Used for Oracle Database Advanced Queuing to specify and limit minimum streams pool memory use
SET_WATERMARK Procedure	Used for Oracle Database Advanced Queuing notification to specify and limit memory use
UNSET_QUEUE_PARAMETER Procedure	Used to unset different parameters for sharded queues at queue or database level.

26.3.8 DBMS_AQADM Agent Subprograms

This section lists and describes the DBMS_AQADM agent subprograms.

Table 26-9 Oracle Streams AQ Agent Subprograms

Subprograms	Description
ALTER_AQ_AGENT Procedure	Alters an agent registered for Oracle Database Advanced Queuing Internet access, and an Oracle Database Advanced Queuing agent that accesses secure queues
CREATE_AQ_AGENT Procedure	Registers an agent for Oracle Database Advanced Queuing Internet access using HTTP/SMTP protocols, and creates an Oracle Database Advanced Queuing agent to access secure queues
DISABLE_DB_ACCESS Procedure	Revokes the privileges of a specific database user from an Oracle Database Advanced Queuing Internet agent
DROP_AQ_AGENT Procedure	Drops an agent that was previously registered for Oracle Database Advanced Queuing Internet access
ENABLE_DB_ACCESS Procedure	Grants an Oracle Database Advanced Queuing Internet agent the privileges of a specific database user

26.3.9 DBMS_AQADM Alias Subprograms

This section lists and describes the DBMS_AQADM alias subprograms.

Table 26-10 Alias Subprograms

Subprograms	Description
ADD_ALIAS_TO_LDAP Procedure	Creates an alias for a queue, agent, or a JMS ConnectionFactory in LDAP
DEL_ALIAS_FROM_LDAP Procedure	Drops an alias for a queue, agent, or JMS ConnectionFactory in LDAP

26.4 Summary of DBMS_AQADM Subprograms

This section lists and describes the DBMS_AQADM package subprograms.

Table 26-11 DBMS_AQADM Package Subprograms

Subprograms	Description
ADD_ALIAS_TO_LDAP Procedure	Creates an alias for a queue, agent, or a JMS ConnectionFactory in LDAP
ADD_SUBSCRIBER Procedure	Adds a default subscriber to a queue
ALTER_AQ_AGENT Procedure	Alters an agent registered for Oracle Database Advanced Queuing Internet access, and an Oracle Database Advanced Queuing agent that accesses secure queues
ALTER_PROPAGATION_SCHEDULE Procedure	Alters parameters for a propagation schedule
ALTER_QUEUE Procedure	Alters existing properties of a queue
ALTER_QUEUE_TABLE Procedure	Alters the existing properties of a queue table
ALTER_SHARDED_QUEUE Procedure	Provides user the ability to alter the `cache_hint` and comment for the sharded queue *Starting with Oracle Database 20c, sharded queues are deprecated and will be desupported in a futur release. Use Transactional Event Queues(TEQ) instead.*
ALTER_SUBSCRIBER Procedure	Alters existing properties of a subscriber to a specified queue
ALTER_TRANSACTIONAL_EVENT_QUEUE Procedure	Provides user the ability to alter the `cache_hint` and comment for the TEQ queue
CREATE_AQ_AGENT Procedure	Registers an agent for Oracle Database Advanced Queuing Internet access using HTTP/SMTP protocols, and creates an Oracle Database Advanced Queuing agent to access secure queues
CREATE_NP_QUEUE Procedure	Creates a nonpersistent `RAW` queue
CREATE_QUEUE Procedure	Creates a queue in the specified queue table

Table 26-11 (Cont.) DBMS_AQADM Package Subprograms

Subprograms	Description
CREATE_SHARDED_QUEUE Procedure	Creates a queue and its queue table for a sharded queue all together. *Starting with Oracle Database 20c, sharded queues are deprecated and will be desupported in a futur release. Use Transactional Event Queues(TEQ) instead.*
CREATE_EXCEPTION_QUEUE Procedure	Creates an exception queue. *Starting with Oracle Database 20c, sharded queues are deprecated and will be desupported in a futur release. Use Transactional Event Queues(TEQ) instead.*
CREATE_EQ EXCEPTION_QUEUE Procedure	Creates an exception queue.
CREATE_TRANSACTIONAL_EVENT_QUEUE Procedure	Creates a queue and its queue table for a Transactional Event Queue (TEQ).
CREATE_QUEUE_TABLE Procedure	Creates a queue table for messages of a predefined type
DEL_ALIAS_FROM_LDAP Procedure	Drops an alias for a queue, agent, or JMS ConnectionFactory in LDAP
DISABLE_DB_ACCESS Procedure	Revokes the privileges of a specific database user from an Oracle Database Advanced Queuing Internet agent
DISABLE_PROPAGATION_SCHEDULE Procedure	Disables a propagation schedule
DROP_AQ_AGENT Procedure	Drops an agent that was previously registered for Oracle Database Advanced Queuing Internet access
DROP_QUEUE Procedure	Drops an existing queue
DROP_SHARDED_QUEUE Procedure	Drops an existing sharded queue from the database queuing system *Starting with Oracle Database 20c, sharded queues are deprecated and will be desupported in a futur release. Use Transactional Event Queues(TEQ) instead.*
DROP_TRANSACTIONAL_EVENT_QUEUE Procedure	Drops an existing TEQ queue from the database queuing system
DROP_QUEUE_TABLE Procedure	Drops an existing queue table
ENABLE_DB_ACCESS Procedure	Grants an Oracle Database Advanced Queuing Internet agent the privileges of a specific database user
ENABLE_JMS_TYPES Procedure	A precondition for the enqueue of JMS types and XML types
ENABLE_PROPAGATION_SCHEDULE Procedure	Enables a previously disabled propagation schedule
GET_MAX_STREAMS_POOL Procedure	Retrieves the value of Oracle Database Advanced Queuing maximum streams pool memory limit
GET_MIN_STREAMS_POOL Procedure	Retrieves the value of Oracle Database Advanced Queuing minimum streams pool memory limit

Table 26-11 (Cont.) DBMS_AQADM Package Subprograms

Subprograms	Description
GET_QUEUE_PARAMETER Procedure	Used to get different parameters for sharded queues at queue or database level.
GET_WATERMARK Procedure	Retrieves the value of watermark set by the `SET_WATERMARK` Procedure
GRANT_QUEUE_PRIVILEGE Procedure	Grants privileges on a queue to users and roles
GRANT_SYSTEM_PRIVILEGE Procedure	Grants Oracle Database Advanced Queuing system privileges to users and roles
MIGRATE_QUEUE_TABLE Procedure	Upgrades an 8.0-compatible queue table to an 8.1-compatible or higher queue table, or downgrades an 8.1-compatible or higher queue table to an 8.0-compatible queue table
PURGE_QUEUE_TABLE Procedure	Purges messages from queue tables
QUEUE_SUBSCRIBERS Function	Returns the subscribers to an 8.0-compatible multiconsumer queue in the PL/SQL index by table collection type `DBMS_AQADM.AQ$_subscriber_list_t`
REMOVE_SUBSCRIBER Procedure	Removes a default subscriber from a queue
REVOKE_QUEUE_PRIVILEGE Procedure	Revokes privileges on a queue from users and roles
REVOKE_SYSTEM_PRIVILEGE Procedure	Revokes Oracle Database Advanced Queuing system privileges from users and roles
SCHEDULE_PROPAGATION Procedure	Schedules propagation of messages from a queue to a destination identified by a specific database link
SET_QUEUE_PARAMETER Procedure	Used to set different parameters for sharded queues at queue or database level.
SET_MAX_STREAMS_POOL Procedure	Used for Oracle Database Advanced Queuing to specify and limit maximum streams pool memory use
SET_MIN_STREAMS_POOL Procedure	Used for Oracle Database Advanced Queuing to specify and limit minimum streams pool memory use
SET_WATERMARK Procedure	Used for Oracle Database Advanced Queuing notification to specify and limit memory use
START_QUEUE Procedure	Enables the specified queue for enqueuing or dequeuing
STOP_QUEUE Procedure	Disables enqueuing or dequeuing on the specified queue
UNSCHEDULE_PROPAGATION Procedure	Unschedules previously scheduled propagation of messages from a queue to a destination identified by a specific database link
UNSET_QUEUE_PARAMETER Procedure	Used to unset different parameters for sharded queues at queue or database level.
VERIFY_QUEUE_TYPES Procedure	Verifies that the source and destination queues have identical types

26.4.1 ADD_ALIAS_TO_LDAP Procedure

This procedure creates an alias for a queue, agent, or a JMS ConnectionFactory in LDAP. The alias will be placed directly under the database server's distinguished name in LDAP hierarchy.

Syntax

```
DBMS_AQADM.ADD_ALIAS_TO_LDAP(
   alias          IN VARCHAR2,
   obj_location   IN VARCHAR2);
```

Parameters

Table 26-12 ADD_ALIAS_TO_LDAP Procedure Parameters

Parameter	Description
`alias`	Name of the alias. Example: `west_shipping`.
`obj_location`	The distinguished name of the object (queue, agent or connection factory) to which `alias` refers.

Usage Notes

This method can be used to create aliases for queues, agents, and JMS ConnectionFactory objects. These object must exist before the alias is created. These aliases can be used for JNDI lookup in JMS and Oracle Database Advanced Queuing Internet access.

26.4.2 ADD_SUBSCRIBER Procedure

This procedure adds a default subscriber to a queue.

Syntax

```
DBMS_AQADM.ADD_SUBSCRIBER (
   queue_name      IN    VARCHAR2,
   subscriber      IN    sys.aq$_agent,
   rule            IN    VARCHAR2 DEFAULT NULL,
   transformation  IN    VARCHAR2 DEFAULT NULL
   queue_to_queue  IN    BOOLEAN DEFAULT FALSE,
   delivery_mode   IN    PLS_INTEGER DEFAULT DBMS_AQADM.PERSISTENT);
```

Parameters

Table 26-13 ADD_SUBSCRIBER Procedure Parameters

Parameter	Description
`queue_name`	Name of the queue.
`subscriber`	Agent on whose behalf the subscription is being defined.

Table 26-13 (Cont.) ADD_SUBSCRIBER Procedure Parameters

Parameter	Description
rule	A conditional expression based on the message properties, the message data properties and PL/SQL functions. A rule is specified as a Boolean expression using syntax similar to the WHERE clause of a SQL query. This Boolean expression can include conditions on message properties, user data properties (object payloads only), and PL/SQL or SQL functions (as specified in the where clause of a SQL query). Currently supported message properties are priority and corrid. To specify rules on a message payload (object payload), use attributes of the object type in clauses. You must prefix each attribute with tab.user_data as a qualifier to indicate the specific column of the queue table that stores the payload. The rule parameter cannot exceed 4000 characters.
transformation	Specifies a transformation that will be applied when this subscriber dequeues the message. The source type of the transformation must match the type of the queue. If the subscriber is remote, then the transformation is applied before propagation to the remote queue.
queue_to_queue	If TRUE, propagation is from queue-to-queue.
delivery_mode	The administrator may specify one of DBMS_AQADM.PERSISTENT, DBMS_AQADM.BUFFERED, or DBMS_AQADM.PERSISTENT_OR_BUFFERED for the delivery mode of the messages the subscriber is interested in. This parameter will not be modifiable by ALTER_SUBSCRIBER.

Usage Notes

A program can enqueue messages to a specific list of recipients or to the default list of subscribers. This operation only succeeds on queues that allow multiple consumers. This operation takes effect immediately, and the containing transaction is committed. Enqueue requests that are executed after the completion of this call will reflect the new behavior.

Any string within the rule must be quoted:

```
rule   => 'PRIORITY <= 3 AND CORRID =  ''FROM JAPAN'''
```

Note that these are all single quotation marks.

26.4.3 ALTER_AQ_AGENT Procedure

This procedure alters an agent registered for Oracle Database Advanced Queuing Internet access. It is also used to alter an Oracle Database Advanced Queuing agent that accesses secure queues.

Syntax

```
DBMS_AQADM.ALTER_AQ_AGENT (
  agent_name                IN VARCHAR2,
  certificate_location      IN VARCHAR2 DEFAULT NULL,
  enable_http               IN BOOLEAN DEFAULT FALSE,
  enable_smtp               IN BOOLEAN DEFAULT FALSE,
  enable_anyp               IN BOOLEAN DEFAULT FALSE )
```

Parameters

Table 26-14 ALTER_AQ_AGENT Procedure Parameters

Parameter	Description
agent_name	Specifies the username of the Oracle Database Advanced Queuing Internet agent.
certification_location	Agent's certificate location in LDAP (default is NULL). If the agent is allowed to access Oracle Database Advanced Queuing through SMTP, then its certificate must be registered in LDAP. For access through HTTP, the certificate location is not required.
enable_http	TRUE means the agent can access Oracle Database Advanced Queuing through HTTP. FALSE means the agent cannot access Oracle Database Advanced Queuing through HTTP.
enable_smtp	TRUE means the agent can access Oracle Database Advanced Queuing through SMTP (e-mail). FALSE means the agent cannot access Oracle Database Advanced Queuing through SMTP.
enable_anyp	TRUE means the agent can access Oracle Database Advanced Queuing through any protocol (HTTP or SMTP).

26.4.4 ALTER_PROPAGATION_SCHEDULE Procedure

This procedure alters parameters for a propagation schedule.

Syntax

```
DBMS_AQADM.ALTER_PROPAGATION_SCHEDULE (
   queue_name           IN    VARCHAR2,
   destination          IN    VARCHAR2 DEFAULT NULL,
   duration             IN    NUMBER   DEFAULT NULL,
   next_time            IN    VARCHAR2 DEFAULT NULL,
   latency              IN    NUMBER   DEFAULT 60,
   destination_queue    IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 26-15 ALTER_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
queue_name	Name of the source queue whose messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the user.
destination	Destination database link. Messages in the source queue for recipients at this destination are propagated. If it is NULL, then the destination is the local database and messages are propagated to other queues in the local database. The length of this field is currently limited to 128 bytes, and if the name is not fully qualified, then the default domain name is used.
duration	Duration of the propagation window in seconds. A NULL value means the propagation window is forever or until the propagation is unscheduled.

Table 26-15 (Cont.) ALTER_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
next_time	Date function to compute the start of the next propagation window from the end of the current window. If this value is NULL, then propagation is stopped at the end of the current window. For example, to start the window at the same time every day, next_time should be specified as SYSDATE + 1 - duration/86400.
latency	Maximum wait, in seconds, in the propagation window for a message to be propagated after it is enqueued. The default value is 60. Caution: if latency is not specified for this call, then latency will over-write any existing value with the default value. For example, if the latency is 60 seconds and there are no messages to be propagated during the propagation window, then messages from that queue for the destination are not propagated for at least 60 more seconds. It will be at least 60 seconds before the queue will be checked again for messages to be propagated for the specified destination. If the latency is 600, then the queue will not be checked for 10 minutes and if the latency is 0, then a job queue process will be waiting for messages to be enqueued for the destination and as soon as a message is enqueued it will be propagated.
destination_queue	Name of the target queue to which messages are to be propagated in the form of a dblink

26.4.5 ALTER_QUEUE Procedure

This procedure alters existing properties of a queue. The parameters max_retries, retention_time, and retry_delay are not supported for nonpersistent queues.

Syntax

```
DBMS_AQADM.ALTER_QUEUE (
   queue_name         IN    VARCHAR2,
   max_retries        IN    NUMBER   DEFAULT NULL,
   retry_delay        IN    NUMBER   DEFAULT NULL,
   retention_time     IN    NUMBER   DEFAULT NULL,
   auto_commit        IN    BOOLEAN  DEFAULT TRUE,
   comment            IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 26-16 ALTER_QUEUE Procedure Parameters

Parameter	Description
queue_name	Name of the queue that is to be altered
max_retries	Limits the number of times a dequeue with REMOVE mode can be attempted on a message. The maximum value of max_retries is 2**31 -1. A message is moved to an exception queue if RETRY_COUNT is greater than MAX_RETRIES. RETRY_COUNT is incremented when the application issues a rollback after executing the dequeue. If a dequeue transaction fails because the server process dies (including ALTER SYSTEM KILL SESSION) or SHUTDOWN ABORT on the instance, then RETRY_COUNT is not incremented. Note that max_retries is supported for all single consumer queues and 8.1-compatible or higher multiconsumer queues but not for 8.0-compatible multiconsumer queues.

Table 26-16 (Cont.) ALTER_QUEUE Procedure Parameters

Parameter	Description
retry_delay	Delay time in seconds before this message is scheduled for processing again after an application rollback. The default is NULL, which means that the value will not be altered. Note that retry_delay is supported for single consumer queues and 8.1-compatible or higher multiconsumer queues but not for 8.0-compatible multiconsumer queues.
retention_time	Retention time in seconds for which a message is retained in the queue table after being dequeued. The default is NULL, which means that the value will not be altered.
auto_commit	TRUE causes the current transaction, if any, to commit before the ALTER_QUEUE operation is carried out. The ALTER_QUEUE operation become persistent when the call returns. This is the default. FALSE means the operation is part of the current transaction and becomes persistent only when the caller enters a commit. Caution: This parameter has been deprecated.
comment	User-specified description of the queue. This user comment is added to the queue catalog. The default value is NULL, which means that the value will not be changed.

26.4.6 ALTER_QUEUE_TABLE Procedure

This procedure alters the existing properties of a queue table.

Syntax

```
DBMS_AQADM.ALTER_QUEUE_TABLE (
   queue_table          IN   VARCHAR2,
   comment              IN   VARCHAR2       DEFAULT NULL,
   primary_instance     IN   BINARY_INTEGER DEFAULT NULL,
   secondary_instance   IN   BINARY_INTEGER DEFAULT NULL,
   replication_mode     IN   BINARY_INTEGER DEFAULT NULL);
```

Parameters

Table 26-17 ALTER_QUEUE_TABLE Procedure Parameters

Parameter	Description
queue_table	Name of a queue table to be created.
comment	Modifies the user-specified description of the queue table. This user comment is added to the queue catalog. The default value is NULL which means that the value will not be changed.
primary_instance	This is the primary owner of the queue table. Queue monitor scheduling and propagation for the queues in the queue table will be done in this instance. The default value is NULL, which means that the current value will not be changed.
secondary_instance	The queue table fails over to the secondary instance if the primary instance is not available. The default value is NULL, which means that the current value will not be changed.

Table 26-17 (Cont.) ALTER_QUEUE_TABLE Procedure Parameters

Parameter	Description
replication_mode	DBMS_AQADM.REPLICATION_MODE if queue is being altered to be in the Replication Mode. DBMS_AQADM.NONE if queue is being altered to not be replicated. The default value is NULL.

26.4.7 ALTER_SHARDED_QUEUE Procedure

This procedure provides user the ability to alter a sharded queue.

> **See Also:**
>
> *Oracle® Database Advanced Queuing User's Guide* for information about sharded queues

Syntax

```
PROCEDURE ALTER_SHARDED_QUEUE(
    queue_name             IN VARCHAR2,
    max_retries            IN NUMBER         DEFAULT NULL,
    comment                IN VARCHAR2       DEFAULT NULL,
    queue_properties       IN QUEUE_PROPS_T  DEFAULT NULL,
    replication_mode       IN BINARY_INTEGER DEFAULT NULL);
```

Parameters

Table 26-18 ALTER_SHARDED_QUEUE Procedure Parameters

Parameter	Description
queue_name	This parameter specifies the name of the sharded queue. A maximum of 128 characters are allowed.
max_retries	The maximum number of retries allowed.
comment	The comment of the queue.
queue_properties	Properties such as Normal or Exception Queue, Retry delay, retention time, sort list and cache hint. Refer to QUEUE_PROPS_T Typefor more information about queue_properties.
replication_mode	Reserved for future use. DBMS_AQADM.REPLICATION_MODE if queue is being altered to be in the Replication Mode or else DBMS_AQADM.NONE. Default is NULL.

26.4.8 ALTER_SUBSCRIBER Procedure

This procedure alters existing properties of a subscriber to a specified queue. Only the rule can be altered.

Syntax

```
DBMS_AQADM.ALTER_SUBSCRIBER (
   queue_name     IN    VARCHAR2,
   subscriber     IN    sys.aq$_agent,
   rule           IN    VARCHAR2
   transformation IN    VARCHAR2);
```

Parameters

Table 26-19 ALTER_SUBSCRIBER Procedure Parameters

Parameter	Description
`queue_name`	Name of the queue.
`subscriber`	Agent on whose behalf the subscription is being altered. See "AQ__AGENT Type".
`rule`	A conditional expression based on the message properties, the message data properties and PL/SQL functions. The rule parameter cannot exceed 4000 characters. To eliminate the rule, set the rule parameter to `NULL`.
`transformation`	Specifies a transformation that will be applied when this subscriber dequeues the message. The source type of the transformation must match the type of the queue. If the subscriber is remote, then the transformation is applied before propagation to the remote queue.

Usage Notes

This procedure alters both the rule and the transformation for the subscriber. If you want to retain the existing value for either of them, you must specify its old value. The current values for rule and transformation for a subscriber can be obtained from the *schema*.AQ$*queue_table*_R and *schema*.AQ$*queue_table*_S views.

26.4.9 ALTER_TRANSACTIONAL_EVENT_QUEUE Procedure

This procedure provides user the ability to alter a TEQ queue.

> **See Also:**
>
> *Oracle® Database Advanced Queuing User's Guide* for information about TEQ queues

Syntax

```
PROCEDURE ALTER_TRANSACTIONAL_EVENT_QUEUE(
    queue_name              IN VARCHAR2,
    max_retries             IN NUMBER         DEFAULT NULL,
    comment                 IN VARCHAR2       DEFAULT NULL,
```

```
    queue_properties       IN QUEUE_PROPS_T  DEFAULT NULL,
    replication_mode       IN BINARY_INTEGER DEFAULT NULL);
```

Parameters

Table 26-20 ALTER_TRANSACTIONAL_EVENT_QUEUE Procedure Parameters

Parameter	Description
queue_name	This parameter specifies the name of the TEQ queue. A maximum of 128 characters are allowed.
max_retries	The maximum number of retries allowed.
comment	The comment of the queue.
queue_properties	Properties such as Normal or Exception Queue, Retry delay, retention time, sort list and cache hint. Refer to QUEUE_PROPS_T Typefor more information about queue_properties.
replication_mode	Reserved for future use. DBMS_AQADM.REPLICATION_MODE if queue is being altered to be in the Replication Mode or else DBMS_AQADM.NONE. Default is NULL.

26.4.10 CREATE_AQ_AGENT Procedure

This procedure registers an agent for Oracle Database Advanced Queuing Internet access using HTTP/SMTP protocols. It is also used to create an Oracle Database Advanced Queuing agent to access secure queues.

Syntax

```
DBMS_AQADM.CREATE_AQ_AGENT (
  agent_name                IN VARCHAR2,
  certificate_location      IN VARCHAR2 DEFAULT NULL,
  enable_http               IN BOOLEAN DEFAULT FALSE,
  enable_smtp               IN BOOLEAN DEFAULT FALSE,
  enable_anyp               IN BOOLEAN DEFAULT FALSE )
```

Parameters

Table 26-21 CREATE_AQ_AGENT Procedure Parameters

Parameter	Description
agent_name	Specifies the username of the Oracle Database Advanced Queuing Internet agent.
certification_location	Agent's certificate location in LDAP (default is NULL). If the agent is allowed to access Oracle Database Advanced Queuing through SMTP, then its certificate must be registered in LDAP. For access through HTTP, the certificate location is not required.
enable_http	TRUE means the agent can access Oracle Database Advanced Queuing through HTTP. FALSE means the agent cannot access Oracle Database Advanced Queuing through HTTP.
enable_smtp	TRUE means the agent can access Oracle Database Advanced Queuing through SMTP (e-mail). FALSE means the agent cannot access Oracle Database Advanced Queuing through SMTP.

Table 26-21 (Cont.) CREATE_AQ_AGENT Procedure Parameters

Parameter	Description
enable_anyp	TRUE means the agent can access Oracle Database Advanced Queuing through any protocol (HTTP or SMTP).

Usage Notes

The SYS.AQ$INTERNET_USERS view has a list of all Oracle Database Advanced Queuing Internet agents.

26.4.11 CREATE_DATABASE_KAFKA_TOPIC Procedure

This procedure create an OKafka Topic. OKafka Topic is a multi-consumer transactional event queue specifically design to be used with OKafka application.

This procedure will create a transactional event queue underneath which should only be access by OKafka programming interfaces documented at Kafka APIs for Oracle Transactional Event Queues.

Syntax

```
PROCEDURE CREATE_DATABASE_KAFKA_TOPIC(
            topicname IN VARCHAR2,
            partition_num IN NUMBER  Default 1,
            retentiontime IN NUMBER Default 7*24*3600,
            partition_assignment_mode IN NUMBER DEFAULT 2,
            replication_mode IN BINARY_INTEGER DEFAULT SYS.DBMS_AQADM.NONE)
```

Parameters

Table 26-22 CREATE_DATABASE_KAFKA_TOPIC Procedure Parameters

Parameter	Description
topicname	Specifies the name of the topic that is to be created. The name must be unique within a schema and must follow object name guidelines in Oracle SQL Reserved Words and Keywords with regard to reserved characters.
partition_num	Specifies the number of maximum partitions this topic can have
retentiontime	Specifies the retention time for this topic in seconds
partition_assignment_mode	Partition assignment strategy for this topic. Value for this parameter should be 2.
replication_mode	DBMS_AQADM.REPLICATION_MODE if queue is being created in the Replication Mode or else DBMS_AQADM.NONE. Default is DBMS_AQADM.NONE.

26.4.12 CREATE_NP_QUEUE Procedure

This procedure creates a nonpersistent `RAW` queue.

> **Note:**
>
> Nonpersistent queues are deprecated as of Release 10gR2. Oracle recommends using buffered messaging.

Syntax

```
DBMS_AQADM.CREATE_NP_QUEUE (
   queue_name              IN        VARCHAR2,
   multiple_consumers      IN        BOOLEAN  DEFAULT FALSE,
   comment                 IN        VARCHAR2 DEFAULT NULL);
```

Parameters

Table 26-23 CREATE_NP_QUEUE Procedure Parameters

Parameter	Description
`queue_name`	Name of the nonpersistent queue that is to be created. The name must be unique within a schema and must follow object name guidelines in *Oracle Database SQL Language Reference.*
`multiple_consumers`	`FALSE` means queues created in the table can only have one consumer for each message. This is the default. `TRUE` means queues created in the table can have multiple consumers for each message. Note that this parameter is distinguished at the queue level, because a nonpersistent queue does not inherit this characteristic from any user-created queue table.
`comment`	User-specified description of the queue. This user comment is added to the queue catalog.

Usage Notes

The queue may be either single-consumer or multiconsumer queue. All queue names must be unique within a schema. The queues are created in a 8.1-compatible or higher system-created queue table (`AQ$_MEM_SC` or `AQ$_MEM_MC`) in the same schema as that specified by the queue name.

If the queue name does not specify a schema name, the queue is created in the login user's schema. After a queue is created with `CREATE_NP_QUEUE`, it can be enabled by calling `START_QUEUE`. By default, the queue is created with both enqueue and dequeue disabled.

You cannot dequeue from a nonpersistent queue. The only way to retrieve a message from a nonpersistent queue is by using the OCI notification mechanism. You cannot invoke the `LISTEN` call on a nonpersistent queue.

26.4.13 CREATE_QUEUE Procedure

This procedure creates a queue in the specified queue table.

Syntax

```
DBMS_AQADM.CREATE_QUEUE (
   queue_name          IN       VARCHAR2,
   queue_table         IN       VARCHAR2,
   queue_type          IN       BINARY_INTEGER DEFAULT NORMAL_QUEUE,
   max_retries         IN       NUMBER         DEFAULT NULL,
   retry_delay         IN       NUMBER         DEFAULT 0,
   retention_time      IN       NUMBER         DEFAULT 0,
   dependency_tracking IN       BOOLEAN        DEFAULT FALSE,
   comment             IN       VARCHAR2       DEFAULT NULL,
   auto_commit         IN       BOOLEAN        DEFAULT TRUE);
```

Parameters

Table 26-24 CREATE_QUEUE Procedure Parameters

Parameter	Description
queue_name	Name of the queue that is to be created. The name must be unique within a schema and must follow object name guidelines in *Oracle Database SQL Language Reference* with regard to reserved characters.
queue_table	Name of the queue table that will contain the queue.
queue_type	Specifies whether the queue being created is an exception queue or a normal queue. NORMAL_QUEUE means the queue is a normal queue. This is the default. EXCEPTION_QUEUE means it is an exception queue. Only the dequeue operation is allowed on the exception queue.
max_retries	Limits the number of times a dequeue with the REMOVE mode can be attempted on a message. The maximum value of max_retries is 2**31 -1. A message is moved to an exception queue if RETRY_COUNT is greater than MAX_RETRIES. RETRY_COUNT is incremented when the application issues a rollback after executing the dequeue. If a dequeue transaction fails because the server process dies (including ALTER SYSTEM KILL SESSION) or SHUTDOWN ABORT on the instance, then RETRY_COUNT is not incremented. Note that max_retries is supported for all single consumer queues and 8.1-compatible or higher multiconsumer queues but not for 8.0-compatible multiconsumer queues.
retry_delay	Delay time, in seconds, before this message is scheduled for processing again after an application rollback. The default is 0, which means the message can be retried as soon as possible. This parameter has no effect if max_retries is set to 0. Note that retry_delay is supported for single consumer queues and 8.1-compatible or higher multiconsumer queues but not for 8.0-compatible multiconsumer queues.
retention_time	Number of seconds for which a message is retained in the queue table after being dequeued from the queue. INFINITE means the message is retained forever. NUMBER is the number of seconds for which to retain the messages. The default is 0, no retention.

Table 26-24 (Cont.) CREATE_QUEUE Procedure Parameters

Parameter	Description
dependency_tracking	Reserved for future use. FALSE is the default. TRUE is not permitted in this release.
comment	User-specified description of the queue. This user comment is added to the queue catalog.
auto_commit	TRUE causes the current transaction, if any, to commit before the CREATE_QUEUE operation is carried out. The CREATE_QUEUE operation becomes persistent when the call returns. This is the default. FALSE means the operation is part of the current transaction and becomes persistent only when the caller enters a commit. Caution: This parameter has been deprecated.

Usage Notes

All queue names must be unique within a schema. After a queue is created with CREATE_QUEUE, it can be enabled by calling START_QUEUE. By default, the queue is created with both enqueue and dequeue disabled.

26.4.14 CREATE_QUEUE_TABLE Procedure

This procedure creates a queue table for messages of a predefined type.

Syntax

```
DBMS_AQADM.CREATE_QUEUE_TABLE (
   queue_table          IN      VARCHAR2,
   queue_payload_type   IN      VARCHAR2,
   storage_clause       IN      VARCHAR2        DEFAULT NULL,
   sort_list            IN      VARCHAR2        DEFAULT NULL,
   multiple_consumers   IN      BOOLEAN         DEFAULT FALSE,
   message_grouping     IN      BINARY_INTEGER  DEFAULT NONE,
   comment              IN      VARCHAR2        DEFAULT NULL,
   auto_commit          IN      BOOLEAN         DEFAULT TRUE,
   primary_instance     IN      BINARY_INTEGER  DEFAULT 0,
   secondary_instance   IN      BINARY_INTEGER  DEFAULT 0,
   compatible           IN      VARCHAR2        DEFAULT NULL,
   secure               IN      BOOLEAN         DEFAULT FALSE
   replication_mode     IN      BINARY_INTEGER  DEFAULT NONE);
```

Parameters

Table 26-25 CREATE_QUEUE_TABLE Procedure Parameters

Parameter	Description
queue_table	Name of a queue table to be created
queue_payload_type	Type of the user data stored. See Type Name inDBMS_AQ Data Types for valid values for this parameter.

Table 26-25 (Cont.) CREATE_QUEUE_TABLE Procedure Parameters

Parameter	Description
storage_clause	Storage parameter. The storage parameter is included in the CREATE TABLE statement when the queue table is created. The storage_clause argument can take any text that can be used in a standard CREATE TABLE storage_clause argument.The storage parameter can be made up of any combinations of the following parameters: PCTFREE, PCTUSED, INITRANS, MAXTRANS, TABLESPACE, LOB, and a table storage clause. If a tablespace is not specified here, then the queue table and all its related objects are created in the default user tablespace. If a tablespace is specified here, then the queue table and all its related objects are created in the tablespace specified in the storage clause. See *Oracle Database SQL Language Reference* for the usage of these parameters.
sort_list	The columns to be used as the sort key in ascending order. This parameter has the following format: '*sort_column_1*,*sort_column_2*' The allowed column names are priority, enq_time, and commit_time.If both columns are specified, then *sort_column_1* defines the most significant order. After a queue table is created with a specific ordering mechanism, all queues in the queue table inherit the same defaults. The order of a queue table cannot be altered after the queue table has been created. If no sort list is specified, then all the queues in this queue table are sorted by the enqueue time in ascending order. This order is equivalent to FIFO order. Even with the default ordering defined, a dequeuer is allowed to choose a message to dequeue by specifying its msgid or correlation. msgid, correlation, and sequence_deviation take precedence over the default dequeueing order, if they are specified. When commit_time is specified for the sort_list parameter the resulting queue table uses commit-time ordering. See also "Priority and Ordering of Messages" in *Oracle Database Advanced Queuing User's Guide* for information about message ordering in Oracle Database Advanced Queuing.
multiple_consumers	FALSE means queues created in the table can only have one consumer for each message. This is the default. TRUE means queues created in the table can have multiple consumers for each message.
message_grouping	Message grouping behavior for queues created in the table. NONE means each message is treated individually. TRANSACTIONAL means messages enqueued as part of one transaction are considered part of the same group and can be dequeued as a group of related messages.
comment	User-specified description of the queue table. This user comment is added to the queue catalog.
auto_commit	TRUE causes the current transaction, if any, to commit before the CREATE_QUEUE_TABLE operation is carried out. The CREATE_QUEUE_TABLE operation becomes persistent when the call returns. This is the default. FALSE means the operation is part of the current transaction and becomes persistent only when the caller enters a commit. Note: This parameter has been deprecated.

Table 26-25 (Cont.) CREATE_QUEUE_TABLE Procedure Parameters

Parameter	Description
primary_instance	The primary owner of the queue table. Queue monitor scheduling and propagation for the queues in the queue table are done in this instance. The default value for primary instance is 0, which means queue monitor scheduling and propagation will be done in any available instance.
secondary_instance	The queue table fails over to the secondary instance if the primary instance is not available. The default value is 0, which means that the queue table will fail over to any available instance.
compatible	The lowest database version with which the queue is compatible. Currently the possible values are either 8.0, 8.1, or 10.0. If the database is in 10.1-compatible mode, the default value is 10.0. If the database is in 8.1-compatible or 9.2-compatible mode, the default value is 8.1. If the database is in 8.0 compatible mode, the default value is 8.0.
secure	This parameter must be set to TRUE if you want to use the queue table for secure queues. Secure queues are queues for which AQ agents must be associated explicitly with one or more database users who can perform queue operations, such as enqueue and dequeue. The owner of a secure queue can perform all queue operations on the queue, but other users cannot perform queue operations on a secure queue, unless they are configured as secure queue users.
replication_mode	DBMS_AQADM.REPLICATION_MODE if queue is being created in the Replication Mode or else DBMS_AQADM.NONE. Default is DBMS_AQADM.NONE.

Usage Notes

The sort keys for dequeue ordering, if any, must be defined at table creation time. The following objects are created at this time:

- aq$_*queue_table_name*_e, a default exception queue associated with the queue table
- aq$*queue_table_name*, a read-only view, which is used by Oracle Database Advanced Queuing applications for querying queue data
- aq$_*queue_table_name*_t, an index (or an index organized table (IOT) in the case of multiple consumer queues) for the queue monitor operations
- aq$_*queue_table_name*_i, an index (or an index organized table in the case of multiple consumer queues) for dequeue operations

For 8.1-compatible or higher queue tables, the following index-organized tables are created:

- aq$_*queue_table_name*_s, a table for storing information about the subscribers
- aq$_*queue_table_name*_r, a table for storing information about rules on subscriptions

aq$_*queue_table_name*_h, an index-organized table for storing the dequeue history data

CLOB, BLOB, and BFILE are valid attributes for Oracle Database Advanced Queuing object type payloads. However, only CLOB and BLOB can be propagated using Oracle Database Advanced Queuing propagation in Oracle8*i* release 8.1.5 or later. See the *Oracle Database Advanced Queuing User's Guide* for more information.

The default value of the compatible parameter depends on the database compatibility mode in the init.ora. If the database is in 10.1-compatible mode, the default value is 10.0. If the

database is in 8.1-compatible or 9.2-compatible mode, the default value is 8.1. If the database is in 8.0 compatible mode, the default value is 8.0

You can specify and modify the `primary_instance` and `secondary_instance` only in 8.1-compatible or higher mode. You cannot specify a secondary instance unless there is a primary instance.

26.4.15 CREATE_SHARDED_QUEUE Procedure

The `CREATE_SHARDED_QUEUE` API creates a queue and its queue table as appropriate for a sharded queue. This API cannot be used to create unsharded queues. Sharded queues must be created using this single integrated API that will automatically set AQ properties as needed

Starting with Oracle Database 20c, the `CREATE_SHARDED_QUEUE` procedure is deprecated and will be desupported in a future release. Use the CREATE_TRANSACTIONAL_EVENT_QUEUE Procedure procedure instead.

Sharded queues may be either a single consumer or a multi-consumer queue.

Syntax

```
PROCEDURE CREATE_SHARDED_QUEUE (
    queue_name              IN VARCHAR2,
    storage_clause          IN VARCHAR2       DEFAULT NULL,
    multiple_consumers      IN BOOLEAN        DEFAULT FALSE,
    max_retries             IN NUMBER         DEFAULT NULL,
    comment                 IN VARCHAR2       DEFAULT NULL,
    queue_payload_type      IN VARCHAR2       DEFAULT JMS_TYPE,
    queue_properties        IN QUEUE_PROPS_T  DEFAULT NULL,
    replication_mode        IN BINARY_INTEGER DEFAULT NONE,
    queue_type         IN BINARY_INTEGER DEFAULT SHARDED_QUEUE,
    squeue_ver              IN BINARY_INTEGER DEFAULT SQ);
```

Parameters

Table 26-26 CREATE_SHARDED_QUEUE Procedure Parameters

Parameter	Description
`queue_name`	This required parameter specifies the name of the new queue. Maximum of 128 characters allowed.
`storage_clause`	The storage parameter is included in the `CREATE TABLE` statement when the queue table is created. The `storage_clause` argument can take any text that can be used in a standard `CREATE TABLE storage_clause` argument. The storage parameter can be made up of any combinations of the following parameters: `PCTFREE`, `PCTUSED`, `INITRANS`, `MAXTRANS`, `TABLESPACE`, `LOB`, and a table storage clause. If a tablespace is not specified here, then the queue table and all its related objects are created in the default user tablespace. If a tablespace is specified here, then the queue table and all its related objects are created in the tablespace specified in the storage clause. See *Oracle Database SQL Language Reference* for the usage of these parameters.
`multiple_consumers`	`FALSE` means queues can only have one consumer for each message. This is the default. `TRUE` means queues created in the table can have multiple consumers for each message.

Table 26-26 (Cont.) CREATE_SHARDED_QUEUE Procedure Parameters

Parameter	Description
max_retries	This optional parameter limits the number of times that a dequeue can reattempted on a message after a failure. The maximum value of max_retries is 2**31 -1. After the retry limit has been exceeded, the message will be purged from the queue. RETRY_COUNT is incremented when the application issues a rollback after executing the dequeue. If a dequeue transaction fails because the server process dies (including ALTER SYSTEM KILL SESSION) or SHUTDOWN ABORT on the instance, then RETRY_COUNT is not incremented.
comment	This optional parameter is a user-specified description of the queue table. This user comment is added to the queue catalog.
queue_payload_type	Payload can be RAW, JSON, DBMS_AQADM.JMS_TYPE, or an object type. Default is DBMS_AQADM.JMS_TYPE. See DBMS_AQ Data Types.
queue_properties	Properties such as Normal or Exception Queue, Retry delay, retention time, sort list and cache hint. Refer to QUEUE_PROPS_T Type for more information about queue_properties.
replication_mode	Reserved for future use. DBMS_AQADM.REPLICATION_MODE if Queue is being created in the Replication Mode or else DBMS_AQADM.NONE. Default is DBMS_AQADM.NONE.
queue_kind	This parameter specifies the queue type. The following types are: CLASSIC_QUEUE CONSTANT BINARY_INTEGER := 0; SHARDED_QUEUE CONSTANT BINARY_INTEGER := 1; TRANSACTIONAL_EVENT_QUEUE CONSTANT BINARY_INTEGER := 2; The default value is SHARDED_QUEUE.
squeue_ver	The sharded queue phase/version. SQ CONSTANT BINARY_INTEGER := 0; TEQ CONSTANT BINARY_INTEGER := 1;

26.4.16 CREATE_EXCEPTION_QUEUE Procedure

The CREATE_EXCEPTION_QUEUE API creates an exception queue for a sharded queue.

Syntax

```
PROCEDURE CREATE_EXCEPTION_QUEUE(
    sharded_queue_name      IN VARCHAR2,
    exception_queue_name    IN VARCHAR2 DEFAULT NULL
    );
```

Parameters

Table 26-27 CREATE_EXCEPTION_QUEUE Procedure Parameters

Parameter	Description
sharded_queue_name	The name of the sharded queue.

Table 26-27 (Cont.) CREATE_EXCEPTION_QUEUE Procedure Parameters

Parameter	Description
exception_queue_name	The name of the exception queue.

26.4.17 CREATE_EQ EXCEPTION_QUEUE Procedure

The CREATE_EQ_EXCEPTION_QUEUE API creates an exception queue for a TEQ queue.

Syntax

```
PROCEDURE CREATE_EQ_EXCEPTION_QUEUE(
    teq_queue_name     IN VARCHAR2,
    exception_queue_name   IN VARCHAR2 DEFAULT NULL
    );
```

Parameters

Table 26-28 CREATE_EQ_EXCEPTION_QUEUE Procedure Parameters

Parameter	Description
teq_queue_name	The name of the TEQ queue.
exception_queue_name	The name of the exception queue.

26.4.18 CREATE_TRANSACTIONAL_EVENT_QUEUE Procedure

The CREATE_TRANSACTIONAL_EVENT_QUEUE API creates a queue and its queue table as appropriate for a Transactional Event Queue (TEQ). This API cannot be used to create AQ queues. TEQs must be created using this single integrated API that will automatically set AQ properties as needed.

TEQs may be either a single consumer or a multi-consumer queue.

Syntax

```
PROCEDURE CREATE_TRANSACTIONAL_EVENT_QUEUE (
    queue_name             IN VARCHAR2,
    storage_clause         IN VARCHAR2       DEFAULT NULL,
    multiple_consumers     IN BOOLEAN        DEFAULT FALSE,
    max_retries            IN NUMBER         DEFAULT NULL,
    comment                IN VARCHAR2       DEFAULT NULL,
    queue_payload_type     IN VARCHAR2       DEFAULT JMS_TYPE,
    queue_properties       IN QUEUE_PROPS_T  DEFAULT NULL,
    replication_mode       IN BINARY_INTEGER DEFAULT NONE);
```

Parameters

Table 26-29 CREATE_TRANSACTIONAL_EVENT_QUEUE Procedure Parameters

Parameter	Description
queue_name	This required parameter specifies the name of the new queue. Maximum of 128 characters allowed.

Table 26-29 (Cont.) CREATE_TRANSACTIONAL_EVENT_QUEUE Procedure Parameters

Parameter	Description
storage_clause	The storage parameter is included in the CREATE TABLE statement when the queue table is created. The storage_clause argument can take any text that can be used in a standard CREATE TABLE storage_clause argument. The storage parameter can be made up of any combinations of the following parameters: PCTFREE, PCTUSED, INITRANS, MAXTRANS, TABLESPACE, LOB, and a table storage clause. If a tablespace is not specified here, then the queue table and all its related objects are created in the default user tablespace. If a tablespace is specified here, then the queue table and all its related objects are created in the tablespace specified in the storage clause. See *Oracle Database SQL Language Reference* for the usage of these parameters.
multiple_consumers	FALSE means queues can only have one consumer for each message. This is the default. TRUE means queues created in the table can have multiple consumers for each message.
max_retries	This optional parameter limits the number of times that a dequeue can reattempted on a message after a failure. The maximum value of max_retries is 2**31 -1. After the retry limit has been exceeded, the message will be purged from the queue. RETRY_COUNT is incremented when the application issues a rollback after executing the dequeue. If a dequeue transaction fails because the server process dies (including ALTER SYSTEM KILL SESSION) or SHUTDOWN ABORT on the instance, then RETRY_COUNT is not incremented.
comment	This optional parameter is a user-specified description of the queue table. This user comment is added to the queue catalog.
queue_payload_type	Payload can be RAW, JSON, DBMS_AQADM.JMS_TYPE, or an object type. Default is DBMS_AQADM.JMS_TYPE. See DBMS_AQ Data Types.
queue_properties	Properties such as Normal or Exception Queue, Retry delay, retention time, sort list and cache hint. Refer to QUEUE_PROPS_T Type for more information about queue_properties.
replication_mode	Reserved for future use. DBMS_AQADM.REPLICATION_MODE if Queue is being created in the Replication Mode or else DBMS_AQADM.NONE. Default is DBMS_AQADM.NONE.

26.4.19 DEL_ALIAS_FROM_LDAP Procedure

This procedure drops an alias for a queue, agent, or JMS ConnectionFactory in LDAP.

Syntax

```
DBMS_AQADM.DEL_ALIAS_FROM_LDAP(
   alias IN VARCHAR2);
```

Parameters

Table 26-30 DEL_ALIAS_FROM_LDAP Procedure Parameters

Parameter	Description
alias	The alias to be removed.

26.4.20 DISABLE_DB_ACCESS Procedure

This procedure revokes the privileges of a specific database user from an Oracle Database Advanced Queuing Internet agent.

Syntax

```
DBMS_AQADM.DISABLE_DB_ACCESS (
  agent_name                IN VARCHAR2,
  db_username               IN VARCHAR2)
```

Parameters

Table 26-31 DISABLE_DB_ACCESS Procedure Parameters

Parameter	Description
agent_name	Specifies the username of the Oracle Database Advanced Queuing Internet agent.
db_username	Specifies the database user whose privileges are to be revoked from the Oracle Database Advanced Queuing Internet agent.

Usage Notes

The Oracle Database Advanced Queuing Internet agent should have been previously granted those privileges using the ENABLE_DB_ACCESS Procedure.

26.4.21 DISABLE_PROPAGATION_SCHEDULE Procedure

This procedure disables a propagation schedule.

Syntax

```
DBMS_AQADM.DISABLE_PROPAGATION_SCHEDULE (
   queue_name            IN   VARCHAR2,
   destination           IN   VARCHAR2 DEFAULT NULL,
   destination_queue  IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 26-32 DISABLE_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
queue_name	Name of the source queue whose messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the user.

Table 26-32 (Cont.) DISABLE_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
destination	Destination database link. Messages in the source queue for recipients at this destination are propagated. If it is NULL, then the destination is the local database and messages are propagated to other queues in the local database. The length of this field is currently limited to 128 bytes, and if the name is not fully qualified, then the default domain name is used.
destination_queue	Name of the target queue to which messages are to be propagated in the form of a dblink

26.4.22 DROP_AQ_AGENT Procedure

This procedure drops an agent that was previously registered for Oracle Database Advanced Queuing Internet access.

Syntax

```
DBMS_AQADM.DROP_AQ_AGENT (
  agent_name                IN VARCHAR2)
```

Parameters

Table 26-33 DROP_AQ_AGENT Procedure Parameters

Parameter	Description
agent_name	Specifies the username of the Oracle Database Advanced Queuing Internet agent

26.4.23 DROP_DATABASE_KAFKA_TOPIC Procedure

This procedure drops an existing OKafka topic. OKafka Topic is a multi-consumer transactional event queue specifically design to be used with OKafka application.

Syntax

```
PROCEDURE DROP_DATABASE_KAFKA_TOPIC(
            topicname IN VARCHAR2)
```

Parameters

Table 26-34 DROP_DATABASE_KAFKA_TOPIC Procedure Parameters

Parameter	Description
topicname	Specifies the name of the topic that is to be dropped. The name must be unique within a schema and must follow object name guidelines in Oracle SQL Reserved Words and Keywords with regard to reserved characters.

26.4.24 DROP_QUEUE Procedure

This procedure drops an existing queue.

Syntax

```
DBMS_AQADM.DROP_QUEUE (
   queue_name         IN    VARCHAR2,
   auto_commit        IN    BOOLEAN DEFAULT TRUE);
```

Parameters

Table 26-35 DROP_QUEUE Procedure Parameters

Parameter	Description
queue_name	Name of the queue that is to be dropped.
auto_commit	TRUE causes the current transaction, if any, to commit before the DROP_QUEUE operation is carried out. The DROP_QUEUE operation becomes persistent when the call returns. This is the default. FALSE means the operation is part of the current transaction and becomes persistent only when the caller enters a commit. Caution: This parameter has been deprecated.

Usage Notes

DROP_QUEUE is not allowed unless STOP_QUEUE has been called to disable the queue for both enqueuing and dequeuing. All the queue data is deleted as part of the drop operation.

26.4.25 DROP_QUEUE_TABLE Procedure

This procedure drops an existing queue table.

Syntax

```
DBMS_AQADM.DROP_QUEUE_TABLE (
   queue_table        IN    VARCHAR2,
   force              IN    BOOLEAN DEFAULT FALSE,
   auto_commit        IN    BOOLEAN DEFAULT TRUE);
```

Parameters

Table 26-36 DROP_QUEUE_TABLE Procedure Parameters

Parameter	Description
queue_table	Name of a queue table to be dropped.
force	FALSE means the operation does not succeed if there are any queues in the table. This is the default. TRUE means all queues in the table are stopped and dropped automatically.

Table 26-36 (Cont.) DROP_QUEUE_TABLE Procedure Parameters

Parameter	Description
auto_commit	TRUE causes the current transaction, if any, to commit before the DROP_QUEUE_TABLE operation is carried out. The DROP_QUEUE_TABLE operation becomes persistent when the call returns. This is the default. FALSE means the operation is part of the current transaction and becomes persistent only when the caller enters a commit. Caution: This parameter has been deprecated.

Usage Notes

All the queues in a queue table must be stopped and dropped before the queue table can be dropped. You must do this explicitly unless the force option is used, in which case this is done automatically.

26.4.26 DROP_SHARDED_QUEUE Procedure

This procedure drops an existing sharded queue from the database queuing system.

You must stop the queue before calling DROP_SHARDED_QUEUE. User must stop the queue explicitly if force is set to FALSE before calling DROP_SHARDED_QUEUE. If force is set to TRUE then queue will be stopped internally and then dropped.

Syntax

```
DBMS_AQADM.DROP_SHARDED_QUEUE(
       queue_name IN VARCHAR2,
       force      IN BOOLEAN DEFAULT FALSE )
```

Parameters

Table 26-37 DROP_SHARDED_QUEUE Procedure Parameters

Parameter	Description
queue_name	This required parameter specifies the name of the sharded queue.
force	The sharded queue is dropped even if the queue is not stopped.

26.4.27 DROP_TRANSACTIONAL_EVENT_QUEUE Procedure

This procedure drops an existing TEQ queue from the database queuing system.

You must stop the queue before calling DROP_TRANSACTIONAL_EVENT_QUEUE. User must stop the queue explicitly if force is set to FALSE before calling DROP_TRANSACTIONAL_EVENT_QUEUE. If force is set to TRUE then queue will be stopped internally and then dropped.

Syntax

```
DBMS_AQADM.DROP_TRANSACTIONAL_EVENT_QUEUE(
       queue_name IN VARCHAR2,
       force      IN BOOLEAN DEFAULT FALSE )
```

Parameters

Table 26-38 DROP_TRANSACTIONAL_EVENT_QUEUE Procedure Parameters

Parameter	Description
queue_name	This required parameter specifies the name of the TEQ queue.
force	The TEQ queue is dropped even if the queue is not stopped.

26.4.28 ENABLE_DB_ACCESS Procedure

This procedure grants an Oracle Database Advanced Queuing Internet agent the privileges of a specific database user.

Syntax

```
DBMS_AQADM.ENABLE_DB_ACCESS (
  agent_name                IN VARCHAR2,
  db_username               IN VARCHAR2)
```

Parameters

Table 26-39 ENABLE_DB_ACCESS Procedure Parameters

Parameter	Description
agent_name	Specifies the username of the Oracle Database Advanced Queuing Internet agent.
db_username	Specified the database user whose privileges are to be granted to the Oracle Database Advanced Queuing Internet agent.

Usage Notes

The Oracle Database Advanced Queuing Internet agent should have been previously created using the CREATE_AQ_AGENT Procedure.

For secure queues, the sender and receiver agent of the message must be mapped to the database user performing the enqueue or dequeue operation.

The SYS.AQ$INTERNET_USERS view has a list of all Oracle Database Advanced Queuing Internet agents and the names of the database users whose privileges are granted to them.

26.4.29 ENABLE_JMS_TYPES Procedure

Enqueue JMS types and XML types.

Syntax

```
DBMS_AQADM.ENABLE_JMS_TYPES (
   queue_table    IN    VARCHAR2);
```

Parameters

Table 26-40 ENABLE_JMS_TYPES Procedure Parameters

Parameter	Description
queue_table	Specifies name of the queue table to be enabled for JMS and XML types.

26.4.30 ENABLE_PROPAGATION_SCHEDULE Procedure

This procedure enables a previously disabled propagation schedule.

Syntax

```
DBMS_AQADM.ENABLE_PROPAGATION_SCHEDULE (
   queue_name          IN   VARCHAR2,
   destination         IN   VARCHAR2 DEFAULT NULL,
   destination_queue   IN   VARCHAR2  DEFAULT NULL);
```

Parameters

Table 26-41 ENABLE_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
queue_name	Name of the source queue whose messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the user.
destination	Destination database link. Messages in the source queue for recipients at this destination are propagated. If it is NULL, then the destination is the local database and messages are propagated to other queues in the local database. The length of this field is currently limited to 128 bytes, and if the name is not fully qualified, then the default domain name is used.
destination_queue	Name of the target queue to which messages are to be propagated in the form of a dblink

26.4.31 GET_MAX_STREAMS_POOL Procedure

This procedure retrieves the value of Oracle Database Advanced Queuing maximum streams pool memory limit.

Syntax

```
DBMS_AQADM.GET_MAX_STREAMS_POOL (
   value     OUT      NUMBER);
```

Parameters

Table 26-42 GET_MAX_STREAMS_POOL Procedure Parameter

Parameter	Description
value	Value in megabytes.

26.4.32 GET_MIN_STREAMS_POOL Procedure

This procedure retrieves the value of Oracle Database Advanced Queuing minimum streams pool memory limit.

Syntax

```
DBMS_AQADM.GET_MIN_STREAMS_POOL (
   value     OUT      NUMBER);
```

Parameters

Table 26-43 GET_MIN_STREAMS_POOL Procedure Parameter

Parameter	Description
value	Value in megabytes.

26.4.33 GET_QUEUE_PARAMETER Procedure

This procedure allows user to get different parameters for sharded queues at queue or database level.

For database level the queue_name should be NULL. Note that queue overrides database level parameter values.

> **See Also:**
>
> *Oracle® Database Advanced Queuing User's Guide* for information about sharded queues

Syntax

```
PROCEDURE  GET_QUEUE_PARAMETER(
    queue_name          IN VARCHAR2,
    param_name          IN VARCHAR2,
    param_value         OUT NUMBER);
```

Parameters

Table 26-44 GET_QUEUE_PARAMETER Procedure Parameters

Parameter	Description
queue_name	The name of the sharded queue.
param_name	The name of the parameter. Table 26-45 and Table 26-46 describe the valid parameter names.
param_value	The value of the parameter.

Table 26-45 Sharded queue parameters

Parameter Name	Scope	Allowed Values	Description
SHARD_NUM	Queue level	[1, UB4MAXVAL]	Number of shards created per instance for a queue, when KEY_BASED_ENQUEUE is not set. When KEY_BASED_ENQUEUE is set, the value is maximum number of shards created for the queue across all instances. The default value is 5.
KEY_BASED_ENQUEUE	Queue level	[0,1]	When set, the shard to which a message gets enqueued is determined by the key value specified in the message. Refer to key-based sharding (link) for more details. This parameter cannot be unset once set. When this parameter is not set (default), a session is bound to a shard at the time of first enqueue to the queue. All messages enqueued by the session will go to the same shard to which the session is bound.
STICKY_DEQUEUE	Queue level	[0,1]	When set, dequeue session sticks to a shard in the queue. A session is bound to a shard on first dequeue from the queue. All messages dequeued by the session come from the same shard to which it is bound. This parameter cannot be unset once set. When this parameter is not set, messages dequeued by a session can spread across multiple shards of the queue.

Table 26-46 Key-based Parameters

Parameter Name	Scope	Description
AQ$KEY_TO_SHARD_MAP	Queue level	Shard number to which a given key is mapped. When key-based sharding is enabled, this parameter is used to establish mapping between a key and a shard number or retrieve the shard number to which given key is mapped.
AQ$GET_KEY_SHARD_INST	Queue level	Instance number that owns the shard to which a given key is mapped. Applicable only when key-based sharding is enabled. It is a read-only parameter.

Example 26-1 Key to shard mapping

Users can explicitly map a new key to an existing shard or to a new shard.

To map a key value 'RED' to a shard with id 0 for queue named MY_SHQ1, submit the following:

```
execute
        dbms_aqadm.set_queue_paramter('MY_SHQ1',  'AQ$KEY_TO_SHARD_MAP=RED',
0);
```

If a key is being enqueued which has not been explicitly mapped to a shard, then one of the shards is chosen at random and mapped to that key permanently. Once a key is mapped to shard, the mapping cannot be changed.

To get the shard identifier for a key which is mapped already, submit the following:

```
declare
  pval number;
begin
  dbms_aqadm.get_queue_parameter('MY_SHQ1', 'AQ$GET_KEY_SHARD=RED', pval);
  dbms_output.put_line('The key RED is mapped to shard id ' || pval);
END;
/
```

Example 26-2 Key to instance mapping

User sharding performs best when cross instance enqueues are not involved. To know the instance where a key can be enqueued without any cross instance enqueue, submit the following:

```
declare
  pval number;
begin
  dbms_aqadm.get_queue_parameter('MY_SHQ1', 'AQ$GET_KEY_SHARD_INST=RED', pval);
  dbms_output.put_line('The key RED is owned by instance id ' || pval);
END;
/
```

26.4.34 GET_WATERMARK Procedure

This procedure retrieves the value of watermark set by `SET_WATERMARK`.

Syntax

```
DBMS_AQADM.GET_WATERMARK (
   wmvalue     IN      NUMBER);
```

Parameters

Table 26-47 GET_WATERMARK Procedure Parameter

Parameter	Description
`wmvalue`	Watermark value in megabytes.

26.4.35 GRANT_QUEUE_PRIVILEGE Procedure

This procedure grants privileges on a queue to users and roles. The privileges are `ENQUEUE` or `DEQUEUE`. Initially, only the queue table owner can use this procedure to grant privileges on the queues.

Syntax

```
DBMS_AQADM.GRANT_QUEUE_PRIVILEGE (
   privilege        IN    VARCHAR2,
   queue_name       IN    VARCHAR2,
   grantee          IN    VARCHAR2,
   grant_option     IN    BOOLEAN := FALSE);
```

Parameters

Table 26-48 GRANT_QUEUE_PRIVILEGE Procedure Parameters

Parameter	Description
`privilege`	The Oracle Database Advanced Queuing queue privilege to grant. The options are `ENQUEUE`, `DEQUEUE`, and `ALL`. `ALL` means both `ENQUEUE` and `DEQUEUE`.
`queue_name`	Name of the queue.
`grantee`	Grantee(s). The grantee(s) can be a user, a role, or the `PUBLIC` role.
`grant_option`	Specifies if the access privilege is granted with the `GRANT` option or not. If the privilege is granted with the `GRANT` option, then the grantee is allowed to use this procedure to grant the access privilege to other users or roles, regardless of the ownership of the queue table. The default is `FALSE`.

26.4.36 GRANT_SYSTEM_PRIVILEGE Procedure

This procedure grants Oracle Database Advanced Queuing system privileges to users and roles.

The privileges are `ENQUEUE_ANY`, `DEQUEUE_ANY`, and `MANAGE_ANY`. Initially, only `SYS` and `SYSTEM` can use this procedure successfully.

> **Note:**
>
> Starting from Oracle Database 12*c* Release 2, `MANAGE_ANY`, `ENQUEUE_ANY`, and `DEQUEUE_ANY` privileges will not allow access to `SYS` owned queues by users other than `SYS`.

Syntax

```
DBMS_AQADM.GRANT_SYSTEM_PRIVILEGE (
   privilege         IN    VARCHAR2,
   grantee           IN    VARCHAR2,
   admin_option      IN    BOOLEAN := FALSE);
```

Parameters

Table 26-49 GRANT_SYSTEM_PRIVILEGE Procedure Parameters

Parameter	Description
`privilege`	The Oracle Database Advanced Queuing system privilege to grant. The options are `ENQUEUE_ANY`, `DEQUEUE_ANY`, and `MANAGE_ANY`. `ENQUEUE_ANY` means users granted this privilege are allowed to enqueue messages to any queues in the database. `DEQUEUE_ANY` means users granted this privilege are allowed to dequeue messages from any queues in the database. `MANAGE_ANY` means users granted this privilege are allowed to run `DBMS_AQADM` calls on any schemas in the database. **Note:** Starting from Oracle Database 12*c* Release 2, `MANAGE_ANY`, `ENQUEUE_ANY`, and `DEQUEUE_ANY` privileges will not allow access to `SYS` owned queues by users other than `SYS`.
`grantee`	Grantee(s). The grantee(s) can be a user, a role, or the `PUBLIC` role.
`admin_option`	Specifies if the system privilege is granted with the `ADMIN` option or not. If the privilege is granted with the `ADMIN` option, then the grantee is allowed to use this procedure to grant the system privilege to other users or roles. The default is `FALSE`.

26.4.37 MIGRATE_QUEUE_TABLE Procedure

This procedure upgrades an 8.0-compatible queue table to an 8.1-compatible or higher queue table, or downgrades an 8.1-compatible or higher queue table to an 8.0-compatible queue table.

Syntax

```
DBMS_AQADM.MIGRATE_QUEUE_TABLE (
   queue_table    IN   VARCHAR2,
   compatible     IN   VARCHAR2);
```

Parameters

Table 26-50 MIGRATE_QUEUE_TABLE Procedure Parameters

Parameter	Description
queue_table	Specifies name of the queue table to be migrated.
compatible	Set this to 8.1 to upgrade an 8.0-compatible queue table, or set this to 8.0 to downgrade an 8.1-compatible queue table.

26.4.38 PURGE_QUEUE_TABLE Procedure

This procedure purges messages from queue tables. You can perform various purge operations on both single-consumer and multiconsumer queue tables for persistent and buffered messages.

Syntax

```
DBMS_AQADM.PURGE_QUEUE_TABLE(
   queue_table        IN   VARCHAR2,
   purge_condition    IN   VARCHAR2,
   purge_options      IN   aq$_purge_options_t);
```

where type aq$_purge_options_t is described in Oracle Database Advanced Queuing (AQ) Types.

Parameters

Table 26-51 PURGE_QUEUE_TABLE Procedure Parameters

Parameter	Description
queue_table	Specifies the name of the queue table to be purged.
purge_condition	Specifies the purge condition to use when purging the queue table. The purge condition must be in the format of a SQL WHERE clause, and it is case-sensitive. The condition is based on the columns of aq$queue_table_name view. When specifying the purge_condition, qualify the column names in aq$queue_table_name view with qtview. To purge all queues in a queue table, set purge_condition to either NULL (a bare null word, no quotes) or '' (two single quotes).

Table 26-51 (Cont.) PURGE_QUEUE_TABLE Procedure Parameters

Parameter	Description
purge_options	Type aq$_purge_options_t contains a block parameter and a delivery_mode parameter. • If block is TRUE, then an exclusive lock on all the queues in the queue table is held while purging the queue table. This will cause concurrent enqueuers and dequeuers to block while the queue table is purged. The purge call always succeeds if block is TRUE. The default for block is FALSE. This will not block enqueuers and dequeuers, but it can cause the purge to fail with an error during high concurrency times. • delivery_mode is used to specify whether DBMS_AQADM.PERSISTENT, DBMS_AQADM.BUFFERED or DBMS_AQADM.PERSISTENT_OR_BUFFERED types of messages are to be purged. You cannot implement arbitrary purge conditions if buffered messages have to be purged.

Usage Notes

- You an purge selected messages from the queue table by specifying a purge_condition. Table 26-51 describes these parameters. Messages can be enqueued to and dequeued from the queue table while the queue table is being purged.
- A trace file is generated in the udump destination when you run this procedure. It details what the procedure is doing.
- This procedure commits batches of messages in autonomous transactions. Several such autonomous transactions may get executed as a part of one purge_queue_table call depending on the number of messages in the queue table.

26.4.39 QUEUE_SUBSCRIBERS Function

This function returns the subscribers to an 8.0-compatible multiconsumer queue in the PL/SQL index by table collection type DBMS_AQADM.AQ$_subscriber_list_t.

Each element of the collection is of type sys.aq$_agent. This functionality is provided for 8.1-compatible queues by the AQ$*queue_table_name*_S view.

Syntax

```
DBMS_AQADM.QUEUE_SUBSCRIBERS (
   queue_name          IN          VARCHAR2);
RETURN aq$_subscriber_list_t IS
```

Parameters

Table 26-52 QUEUE_SUBSCRIBERS Function Parameters

Parameter	Description
queue_name	Specifies the queue whose subscribers are to be printed.

26.4.40 REMOVE_SUBSCRIBER Procedure

This procedure removes a default subscriber from a queue. This operation takes effect immediately, and the containing transaction is committed. All references to the subscriber in existing messages are removed as part of the operation.

Syntax

```
DBMS_AQADM.REMOVE_SUBSCRIBER (
   queue_name          IN         VARCHAR2,
   subscriber          IN         sys.aq$_agent);
```

Parameters

Table 26-53 REMOVE_SUBSCRIBER Procedure Parameters

Parameter	Description
queue_name	Name of the queue.
subscriber	Agent who is being removed. See AQ$_AGENT Type.

26.4.41 REVOKE_QUEUE_PRIVILEGE Procedure

This procedure revokes privileges on a queue from users and roles. The privileges are ENQUEUE or DEQUEUE.

Syntax

```
DBMS_AQADM.REVOKE_QUEUE_PRIVILEGE (
   privilege          IN      VARCHAR2,
   queue_name         IN      VARCHAR2,
   grantee            IN      VARCHAR2);
```

Parameters

Table 26-54 REVOKE_QUEUE_PRIVILEGE Procedure Parameters

Parameter	Description
privilege	The Oracle Database Advanced Queuing queue privilege to revoke. The options are ENQUEUE, DEQUEUE, and ALL. ALL means both ENQUEUE and DEQUEUE.
queue_name	Name of the queue.
grantee	Grantee(s). The grantee(s) can be a user, a role, or the PUBLIC role. If the privilege has been propagated by the grantee through the GRANT option, then the propagated privilege is also revoked.

Usage Notes

To revoke a privilege, the revoker must be the original grantor of the privilege. The privileges propagated through the GRANT option are revoked if the grantor's privileges are revoked.

26.4.42 REVOKE_SYSTEM_PRIVILEGE Procedure

This procedure revokes Oracle Database Advanced Queuing system privileges from users and roles. The privileges are `ENQUEUE_ANY`, `DEQUEUE_ANY` and `MANAGE_ANY`. The `ADMIN` option for a system privilege cannot be selectively revoked. Starting from Oracle Database 12*c* Release 2, `MANAGE_ANY`, `ENQUEUE_ANY`, and `DEQUEUE_ANY` privileges will not allow access to `SYS` owned queues by users other than `SYS`.

Syntax

```
DBMS_AQADM.REVOKE_SYSTEM_PRIVILEGE (
   privilege         IN   VARCHAR2,
   grantee           IN   VARCHAR2);
```

Parameters

Table 26-55 REVOKE_SYSTEM_PRIVILEGE Procedure Parameters

Parameter	Description
`privilege`	The Oracle Database Advanced Queuing system privilege to revoke. The options are `ENQUEUE_ANY`, `DEQUEUE_ANY`, and `MANAGE_ANY`. The `ADMIN` option for a system privilege cannot be selectively revoked. **Note:** Starting from Oracle Database 12*c* Release 2, `MANAGE_ANY`, `ENQUEUE_ANY`, and `DEQUEUE_ANY` privileges will not allow access to `SYS` owned queues by users other than `SYS`.
`grantee`	Grantee(s). The grantee(s) can be a user, a role, or the `PUBLIC` role.

26.4.43 SCHEDULE_PROPAGATION Procedure

This procedure schedules propagation of messages from a queue to a destination identified by a specific database link.

Syntax

```
DBMS_AQADM.SCHEDULE_PROPAGATION (
   queue_name          IN    VARCHAR2,
   destination         IN    VARCHAR2 DEFAULT NULL,
   start_time          IN    DATE     DEFAULT SYSDATE,
   duration            IN    NUMBER   DEFAULT NULL,
   next_time           IN    VARCHAR2 DEFAULT NULL,
   latency             IN    NUMBER   DEFAULT 60,
   destination_queue   IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 26-56 SCHEDULE_PROPAGATION Procedure Parameters

Parameter	Description
queue_name	Name of the source queue whose messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the administrative user.
destination	Destination database link. Messages in the source queue for recipients at this destination are propagated. If it is NULL, then the destination is the local database and messages are propagated to other queues in the local database. The length of this field is currently limited to 390 bytes, and if the name is not fully qualified, then the default domain name is used. The pattern *schema*.*queue*@dblink is used.
start_time	Initial start time for the propagation window for messages from the source queue to the destination.
duration	Duration of the propagation window in seconds. A NULL value means the propagation window is forever or until the propagation is unscheduled.
next_time	Date function to compute the start of the next propagation window from the end of the current window. If this value is NULL, then propagation is stopped at the end of the current window. For example, to start the window at the same time every day, next_time should be specified as SYSDATE + 1 - duration/86400.
latency	Maximum wait, in seconds, in the propagation window for a message to be propagated after it is enqueued. For example, if the latency is 60 seconds and there are no messages to be propagated during the propagation window, then messages from that queue for the destination are not propagated for at least 60 more seconds. It is at least 60 seconds before the queue is checked again for messages to be propagated for the specified destination. If the latency is 600, then the queue is not checked for 10 minutes, and if the latency is 0, then a job queue process will be waiting for messages to be enqueued for the destination. As soon as a message is enqueued, it is propagated.
destination_queue	Name of the target queue to which messages are to be propagated in the form of a dblink

Usage Notes

Messages may also be propagated to other queues in the same database by specifying a NULL destination. If a message has multiple recipients at the same destination in either the same or different queues, the message is propagated to all of them at the same time.

Oracle extensions for JMS such as JMS propagation and remote subscribers are not currently supported for sharded queues. Propagation between sharded and non-sharded queues is not supported

Related Topics

- *Oracle Database Advanced Queuing User's Guide*

26.4.44 SET_QUEUE_PARAMETER Procedure

This procedure allows user to set different parameters for sharded queues at queue or database level. For database level the `queue_name` should be `NULL`. Note that queue overrides database level parameter values.

> **See Also:**
>
> *Oracle® Database Advanced Queuing User's Guide* for information about sharded queues

Syntax

```
PROCEDURE  SET_QUEUE_PARAMETER(
  queue_name          IN VARCHAR2,
  param_name          IN VARCHAR2,
  param_value         IN NUMBER);
```

Parameters

Table 26-57 SET_QUEUE_PARAMETER Procedure Parameters

Parameter	Description
`queue_name`	The name of the sharded queue.
`param_name`	The name of the parameter. Table 26-58 and Table 26-59 describe the valid parameter names.
`param_value`	The value of the parameter.

Table 26-58 Sharded queue parameters

Parameter Name	Scope	Allowed Values	Description
`SHARD_NUM`	Queue level	`[1, UB4MAXVAL]`	Maximum number of shards allowed for the queue.

Table 26-58 (Cont.) Sharded queue parameters

Parameter Name	Scope	Allowed Values	Description
SUBSHARD_SIZE	Queue level	[1, UB2MAXVAL]	Default value is 20,000. This is the size of each partition for allocation of memory in a shard. Each shard consists of multiple subshards, which are allocated, used, and purged automatically based on message enqueue and dequeue rate for the shard.
KEY_BASED_ENQUEUE	Queue level	[0,1]	When set, the shard to which a message gets enqueued is determined by the key value specified in the message. Refer to key-based sharding (link) for more details. This parameter cannot be unset once set. When this parameter is not set (default), a session is bound to a shard at the time of first enqueue to the queue. All messages enqueued by the session will go to the same shard to which the session is bound.
STICKY_DEQUEUE	Queue level	[0,1]	When set, dequeue session sticks to a shard in the queue. A session is bound to a shard on first dequeue from the queue. All messages dequeued by the session come from the same shard to which it is bound. This parameter cannot be unset once set. When this parameter is not set, messages dequeued by a session can spread across multiple shards of the queue.

Table 26-59 Key-based Parameters

Parameter Name	Scope	Description
AQ$KEY_TO_SHARD_MAP	Queue level	Shard number to which a given key is mapped. When key-based sharding is enabled, this parameter is used to establish mapping between a key and a shard number or retrieve the shard number to which given key is mapped.
AQ$GET_KEY_SHARD_INST	Queue level	Instance number that owns the shard to which a given key is mapped. Applicable only when key-based sharding is enabled. It is a read-only parameter.

Example 26-3 Create a sharded queue with key based enqueues

To create a sharded queue with key based enqueues (user sharding), the PL/SQL procedure DBMS_AQADM.SET_QUEUE_PARAMETER is called after CREATE_SHARDED_QUEUE.

```
execute
        sys.dbms_aqadm.create_sharded_queue(queue_name => 'MY_SHQ1');


execute
        dbms_aqadm.set_queue_parameter('MY_SHQ1', 'KEY_BASED_ENQUEUE', 1);
```

Example 26-4 Create a sharded queue with sticky dequeues

To create a sharded queue with key based enqueues (user sharding), the PL/SQL procedure DBMS_AQADM.SET_QUEUE_PARAMETER is called after CREATE_SHARDED_QUEUE.

```
execute
        sys.dbms_aqadm.create_sharded_queue(queue_name => 'MY_SHQ1');


execute
        dbms_aqadm.set_queue_parameter('MY_SHQ1', 'STICKY_DEQUEUE', 1);
```

Example 26-5 Setting or changing the number of shards

Once a queue is created, it can have a maximum of 5 shards by default in non-Oracle RAC databases. In Oracle RAC databases, each queue can have a maximum of 5 shards per database instance by default.

```
execute
        dbms_aqadm.set_queue_parameter('MY_SHQ1', 'SHARD_NUM', 200);
```

> **Note:**
>
> Odd numbered shard identifiers are reserved for internal use. In the above example, the 200 shard identifiers used will be 0, 2, 4, …, 398.

26.4.45 SET_MAX_STREAMS_POOL Procedure

This procedure is used for Oracle Database Advanced Queuing to specify and limit maximum streams pool memory use.

Syntax

```
DBMS_AQADM.SET_MAX_STREAMS_POOL (
   value      IN      NUMBER);
```

Parameters

Table 26-60 SET_MAX_STREAMS_POOL Procedure Parameter

Parameter	Description
value	Value in megabytes.

26.4.46 SET_MIN_STREAMS_POOL Procedure

This procedure is used for Oracle Database AQ to specify and limit minimum streams pool memory use.

Syntax

```
DBMS_AQADM.SET_MIN_STREAMS_POOL (
   value      IN      NUMBER);
```

Parameters

Table 26-61 SET_MIN_STREAMS_POOL Procedure Parameter

Parameter	Description
value	Value in megabytes.

26.4.47 SET_WATERMARK Procedure

This procedure is used for Oracle Database Advanced Queuing notification to specify and limit memory use.

Syntax

```
DBMS_AQADM.SET_WATERMARK (
   wmvalue      IN      NUMBER);
```

Parameters

Table 26-62 SET_WATERMARK Procedure Parameter

Parameter	Description
wmvalue	Watermark value in megabytes.

26.4.48 START_QUEUE Procedure

This procedure enables the specified queue for enqueuing or dequeuing.

Syntax

```
DBMS_AQADM.START_QUEUE (
   queue_name      IN     VARCHAR2,
   enqueue         IN     BOOLEAN DEFAULT TRUE,
   dequeue         IN     BOOLEAN DEFAULT TRUE);
```

Parameters

Table 26-63 START_QUEUE Procedure Parameters

Parameter	Description
queue_name	Name of the queue to be enabled
enqueue	Specifies whether ENQUEUE should be enabled on this queue. TRUE means enable ENQUEUE. This is the default. FALSE means do not alter the current setting.
dequeue	Specifies whether DEQUEUE should be enabled on this queue. TRUE means enable DEQUEUE. This is the default. FALSE means do not alter the current setting.

Usage Notes

After creating a queue, the administrator must use START_QUEUE to enable the queue. The default is to enable it for both ENQUEUE and DEQUEUE. Only dequeue operations are allowed on an exception queue. This operation takes effect when the call completes and does not have any transactional characteristics.

26.4.49 STOP_QUEUE Procedure

This procedure disables enqueuing or dequeuing on the specified queue.

Syntax

```
DBMS_AQADM.STOP_QUEUE (
   queue_name      IN   VARCHAR2,
   enqueue         IN   BOOLEAN DEFAULT TRUE,
   dequeue         IN   BOOLEAN DEFAULT TRUE,
   wait            IN   BOOLEAN DEFAULT TRUE,
   free_memory     IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 26-64 STOP_QUEUE Procedure Parameters

Parameter	Description
queue_name	Name of the queue to be disabled
enqueue	Specifies whether ENQUEUE should be disabled on this queue. TRUE means disable ENQUEUE. This is the default. FALSE means do not alter the current setting.
dequeue	Specifies whether DEQUEUE should be disabled on this queue. TRUE means disable DEQUEUE. This is the default. FALSE means do not alter the current setting.
wait	Specifies whether to wait for the completion of outstanding transactions. TRUE means wait if there are any outstanding transactions. In this state no new transactions are allowed to enqueue to or dequeue from this queue. FALSE means return immediately either with a success or an error.
free_memory	Specifies whether the queue should be stopped.

Usage Notes

By default, this call disables both ENQUEUE and DEQUEUE. A queue cannot be stopped if there are outstanding transactions against the queue. This operation takes effect when the call completes and does not have any transactional characteristics.

26.4.50 UNSCHEDULE_PROPAGATION Procedure

This procedure unschedules previously scheduled propagation of messages from a queue to a destination identified by a specific database link.

Syntax

```
DBMS_AQADM.UNSCHEDULE_PROPAGATION (
   queue_name         IN  VARCHAR2,
   destination        IN  VARCHAR2 DEFAULT NULL
   destination_queue  IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 26-65 UNSCHEDULE_PROPAGATION Procedure Parameters

Parameter	Description
queue_name	Name of the source queue whose messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the administrative user.
destination	Destination database link. Messages in the source queue for recipients at this destination are propagated. If it is NULL, then the destination is the local database and messages are propagated to other queues in the local database. The length of this field is currently limited to 128 bytes, and if the name is not fully qualified, then the default domain name is used.
destination_queue	Name of the target queue to which messages are to be propagated in the form of a dblink

26.4.51 UNSET_QUEUE_PARAMETER Procedure

This procedure allows user to unset different parameters for sharded queues at queue or database level.

For database level the `queue_name` should be `NULL`. Note that queue overrides database level parameter values.

> **See Also:**
>
> *Oracle® Database Advanced Queuing User's Guide* for information about sharded queues

Syntax

```
PROCEDURE  UNSET_QUEUE_PARAMETER(
  queue_name          IN VARCHAR2,
  param_name          IN VARCHAR2);
```

Parameters

Table 26-66 UNSET_QUEUE_PARAMETER Procedure Parameters

Parameter	Description
`queue_name`	The name of the sharded queue.
`param_name`	The name of the parameter. Table 26-67 and Table 26-68 describe the valid parameter names.

Table 26-67 Sharded queue parameters

Parameter Name	Scope	Allowed Values	Description
`SHARD_NUM`	Queue level	`[1, UB4MAXVAL]`	Maximum number of shards allowed for the queue.
`SUBSHARD_SIZE`	Queue level	`[1, UB2MAXVAL]`	Default value is 20,000. This is the size of each partition for allocation of memory in a shard. Each shard consists of multiple subshards, which are allocated, used, and purged automatically based on message enqueue and dequeue rate for the shard.

Table 26-67 (Cont.) Sharded queue parameters

Parameter Name	Scope	Allowed Values	Description
KEY_BASED_ENQUEUE	Queue level	[0,1]	When set, the shard to which a message gets enqueued is determined by the key value specified in the message. Refer to key-based sharding (link) for more details. This parameter cannot be unset once set. When this parameter is not set (default), a session is bound to a shard at the time of first enqueue to the queue. All messages enqueued by the session will go to the same shard to which the session is bound.
STICKY_DEQUEUE	Queue level	[0,1]	When set, dequeue session sticks to a shard in the queue. A session is bound to a shard on first dequeue from the queue. All messages dequeued by the session come from the same shard to which it is bound. This parameter cannot be unset once set. When this parameter is not set, messages dequeued by a session can spread across multiple shards of the queue.

Table 26-68 Key-based Parameters

Parameter Name	Scope	Description
AQ$KEY_TO_SHARD_MAP	Queue level	Shard number to which a given key is mapped. When key-based sharding is enabled, this parameter is used to establish mapping between a key and a shard number or retrieve the shard number to which given key is mapped.

Table 26-68 (Cont.) Key-based Parameters

Parameter Name	Scope	Description
AQ$GET_KEY_SHARD_INST	Queue level	Instance number that owns the shard to which a given key is mapped. Applicable only when key-based sharding is enabled. It is a read-only parameter.

26.4.52 VERIFY_QUEUE_TYPES Procedure

This procedure verifies that the source and destination queues have identical types.

The result of the verification is stored in the table `sys.aq$_message_types`, overwriting all previous output of this command.

Syntax

```
DBMS_AQADM.VERIFY_QUEUE_TYPES (
   src_queue_name    IN    VARCHAR2,
   dest_queue_name   IN    VARCHAR2,
   destination       IN    VARCHAR2 DEFAULT NULL,
   rc                OUT   BINARY_INTEGER);
```

Parameters

Table 26-69 VERIFY_QUEUE_TYPES Procedure Parameters

Parameter	Description
src_queue_name	Name of the source queue whose messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the user.
dest_queue_name	Name of the destination queue where messages are to be propagated, including the schema name. If the schema name is not specified, then it defaults to the schema name of the user.
destination	Destination database link. Messages in the source queue for recipients at this destination are propagated. If it is `NULL`, then the destination is the local database and messages are propagated to other queues in the local database. The length of this field is currently limited to 128 bytes, and if the name is not fully qualified, then the default domain name is used.
rc	Return code for the result of the procedure. If there is no error, and if the source and destination queue types match, then the result is 1. If they do not match, then the result is 0. If an Oracle error is encountered, then it is returned in `rc`.

Note:

- `SYS.AQ$_MESSAGE_TYPES` can have multiple entries for the same source queue, destination queue, and database link, but with different transformations.
- `VERIFY_QUEUE_TYPES` check happens once per AQ propagation schedule and not for every propagated message send.
- In case the payload of the queue is modified then the existing propagation schedule between source and destination queue needs to be dropped and recreated.

27
DBMS_AQELM

The package provides subprograms to manage the configuration of Oracle Advanced Queuing (AQ) asynchronous notification by e-mail and HTTP.

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* for detailed information about `DBMS_AQELM`

This chapter contains the following topics:`DBMS_AQELM`

- Security Model
- Summary of DBMS_AQELM Subprograms

27.1 DBMS_AQELM Security Model

You need the administrator role `AQ_ADMINISTRATOR_ROLE` to run all procedures in `DBMS_AQELM`.

27.2 Summary of DBMS_AQELM Subprograms

This section lists and briefly describes the DBMS_AQELM subprograms.

Table 27-1 DBMS_AQELM Package Subprograms

Subprogram	Description
SET_MAILHOST Procedure	Sets the host name for the SMTP server that the database will uses send out e-mail notifications
SET_MAILPORT Procedure	Sets the port number for the SMTP server
SET_SENDFROM Procedure	Sets the sent-from e-mail address

27.2.1 SET_MAILHOST Procedure

This procedure sets the host name for the SMTP server. The database uses this SMTP server host name to send out e-mail notifications.

Syntax

```
DBMS_AQELM.SET_MAILHOST (
   mailhost  IN  VARCHAR2);
```

Parameters

Table 27-2 SET_MAILHOST Procedure Parameters

Parameter	Description
mailhost	SMTP server host name.

Usage Notes

As part of the configuration for e-mail notifications, a user with AQ_ADMINISTRATOR_ROLE or with EXECUTE permissions on the DBMS_AQELM package needs to set the host name before registering for e-mail notifications.

27.2.2 SET_MAILPORT Procedure

This procedure sets the port number for the SMTP server.

Syntax

```
DBMS_AQELM.SET_MAILPORT (
   mailport  IN NUMBER);
```

Parameters

Table 27-3 SET_MAILPORT Procedure Parameters

Parameter	Description
mailport	SMTP server port number.

Usage Notes

As part of the configuration for e-mail notifications, a user with AQ_ADMINISTRATOR_ROLE or with EXECUTE permissions on DBMS_AQELM package needs to set the port number before registering for e-mail notifications. The database uses this SMTP server port number to send out e-mail notifications. If not set, the SMTP mailport defaults to 25

27.2.3 SET_SENDFROM Procedure

This procedure sets the sent-from e-mail address. This e-mail address is used in the sent-from field in all the e-mail notifications sent out by the database to the registered e-mail addresses.

Syntax

```
DBMS_AQELM.SET_SENDFROM (
   sendfrom  IN  VARCHAR2);
```

Parameters

Table 27-4 SET_SENDFROM Procedure Parameters

Parameter	Description
sendfrom	The sent-from e-mail address.

Usage Notes

As part of the configuration for e-mail notifications, a user with `AQ_ADMINISTRATOR_ROLE` or with `EXECUTE` permissions on the `DBMS_AQELM` package should set the sent-from address before registering for e-mail notifications

28
DBMS_AQIN

The DBMS_AQIN package plays a part in providing secure access to the Oracle JMS interfaces.

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* for detailed information about DBMS_AQIN

This chapter contains the following topic:

- Security Model

28.1 DBMS_AQIN Security Model

While you should not call any subprograms in the DBMS_AQIN package directly, you must have the EXECUTE privilege on the DBMS_AQIN and DBMS_AQJMS packages to use the Oracle JMS interfaces.

Use the following syntax to accomplish this with regard to the DBMS_AQIN package:

```
GRANT EXECUTE ON DBMS_AQIN to user;
```

You must have EXECUTE privilege on the DBMS_AQIN and DBMS_AQJMS packages to use the Oracle JMS interfaces. You can also acquire these rights through the AQ_USER_ROLE or the AQ_ADMINSTRATOR_ROLE.

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* for more information on accessing standard and Oracle JMS applications.

29
DBMS_AQMIGTOOL

The DBMS_AQMIGTOOL package simplifies migration from Oracle Database Advanced Queuing (AQ) to Transactional Event Queue (TxEventQ) with orchestration automation, source and target compatibility diagnostics and remediation, and a unified user experience. Migration scenarios can be short or long-lived and be performed with or without AQ downtime, eliminating operational disruption.

> **See Also:**
>
> Migrating from AQ to TxEventQ in *Oracle Database Advanced Queuing User's Guide* for detailed information about DBMS_AQMIGTOOL

This chapter contains the following topics:

- Security Model
- DBMS_AQMIGTOOL Constants
- Summary of DBMS_AQMIGTOOL Subprograms

29.1 DBMS_AQMIGTOOL Security Model

All DBMS_AQMIGTOOL subprograms require the user to have EXECUTE privilege over the DBMS_AQMIGTOOL package. If the invoker of the package is the owner of the queue, then only EXECUTE privilege on the DBMS_AQMIGTOOL package is sufficient, But if the invoker is non-owner, then it also needs to have MANAGE_ANY queue system privilege, which can be granted through DBMS_AQADM.GRANT_SYSTEM_PRIVILEGE.

29.2 DBMS_AQMIGTOOL Constants

The DBMS_AQMIGTOOL package defines several constants that can be used for specifying parameter values.

When using enumerated constants, such as AUTOMATIC, INTERACTIVE, or ENABLE_EVALUATION, the symbol must be specified with the scope of the packages defining it. All types associated with the administrative interfaces must be prepended with DBMS_AQMIGTOOL. For example, DBMS_AQMIGTOOL.AUTOMATIC.

Table 29-1 DBMS_AQMIGTOOL Constants

Paramater	Options	Value
mig_mode	AUTOMATIC	1
	INTERACTIVE	2
	OFFLINE	3
	ONLY_DEFINITION	4
ordering	GLOBAL	1
	SESSION	2
checkmode	CURRENT	1
	ENABLE_EVALUATION	2
cancelmode	RESTORE	1
	NORESTORE	2
	Q_EMPTY	3
purge_option	ONLY_CQ	1
	ONLY_TXEVENTQ	2
	BOTH_Q	3

29.3 Summary of DBMS_AQMIGTOOL Subprograms

This section lists and briefly describes the DBMS_AQMIGTOOL subprograms.

Table 29-2 DBMS_AQMIGTOOL Package Subprograms

Subprogram	Description
CANCEL_MIGRATION Procedure	Cancels an ongoing migration by dropping the interim TxEventQ
CHECK_MIGRATED_MESSAGES Procedure	Provides a count of messages in the READY state within the AQ and the interim TxEventQ
CHECK_MIGRATION_TO_TXEVENTQ Procedure	Analyzes the AQ's definition and data and report any features unsupported in TxEventQ
CHECK_STATUS Procedure	Returns the current status of the ongoing migration process
CLEAR_UNSUPPORTED_FEATURE_TABLE Procedure	Clears entries from the USER_TXEVENTQ_MIGRATION_STATUS view
COMMIT_MIGRATION Procedure	Completes the migration process by dropping AQ, renaming the interim TxEventQ to match AQ's name, and enabling TxEventQ for all operations
DISABLE_MIGRATION_CHECK Procedure	Disables the internal AQ monitoring for detecting unsupported features and stops event insertion in the USER_TXEVENTQ_MIGRATION_STATUS view
INIT_MIGRATION Procedure	Analyzes the AQ's definition and data for unsupported features and then starts the migration process by creating an interim TxEventQ copying the AQ's configuration
PURGE_QUEUE_MESSAGES Procedure	Purges messages from the queue

Table 29-2 (Cont.) DBMS_AQMIGTOOL Package Subprograms

Subprogram	Description
RECOVER_MIGRATION Procedure	Helps recovery from any failures during execution of migration procedures such as `DBMS_AQMIGTOOL.INIT_MIGRATION`, `DBMS_AQMIGTOOL.COMMIT_MIGRATION`, or `DBMS_AQMIGTOOL.CANCEL_MIGRATION`
RENAME_QUEUE Procedure	Renames the TxEventQ along with its default exception queue if present

29.3.1 CANCEL_MIGRATION Procedure

This procedure serves the purpose of canceling an ongoing migration. It involves dropping the interim TxEventQ, which was created during the execution of `DBMS_AQMIGTOOL.INIT_MIGRATION`.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.CANCEL_MIGRATION (
    cqschema     IN VARCHAR2,
    cqname       IN VARCHAR2,
    cancelmode   IN NUMBER DEFAULT DBMS_AQMIGTOOL.RESTORE );
```

Parameters

Table 29-3 CANCEL_MIGRATION Procedure Parameters

Parameter	Description
`cqschema`	Specifies the schema name where the queue exists
`cqname`	Specifies the queue name for which the migration needs to be canceled
`cancelmode`	Specifies the mode in which the user wants to cancel the migration. The following are the possible values: `DBMS_AQMIGTOOL.RESTORE` (default): This option restores the messages from the interim TxEventQ into the AQ, including their message state. New MSGID will get populated for restored messages. **Note:** The priority behavior of restored messages may change as AQ and TxEventQ default values are different. `DBMS_AQMIGTOOL.NORESTORE`: Messages within interim TxEventQ will be discarded. `DBMS_AQMIGTOOL.EMPTY`: If the interim TxEventQ is not empty, selecting this option will trigger an exception, prompting the user to dequeue all messages from the interim TxEventQ. This mode is useful if the user wishes to prevent message migration while transitioning to the AQ.

Usage Notes

A prerequisite for this procedure is that the migration must already be started on the queue, that is, `DBMS_AQMIGTOOL.INIT_MIGRATION` should be invoked before executing this procedure. For `DBMS_AQMIGTOOL.RESTORE` mode, the TxEventQ's exception queue messages are not restored to AQ or its exception queue.

29.3.2 CHECK_MIGRATED_MESSAGES Procedure

This procedure calculates the count of messages in the `READY` state within both the AQ and the interim TxEventQ. This count provides valuable insight before using `DBMS_AQMIGTOOL.COMMIT_MIGRATION` or `DBMS_AQMIGTOOL.CANCEL_MIGRATION`. The calculated count is independent of the number of subscribers.

Syntax

```
PROCEDURE SYS.DBMS_AQMIGTOOL.CHECK_MIGRATED_MESSAGES (
    cqschema                          IN VARCHAR2,
    cqname                            IN VARCHAR2,
    txeventq_migrated_message         IN OUT NUMBER,
    cq_pending_messages               IN OUT NUMBER);
```

Parameters

Table 29-4 CHECK_MIGRATED_MESSAGES Procedure Parameters

Parameter	Description
`cqschema`	Specifies the schema name where the queue exists
`cqname`	Specifies the name of the queue on which the migration process has started
`txeventq_migrated_message`	Represents the count of messages in the `READY` state within the interim TxEventQ. The count helps estimate the potential fallback time if the user opts to execute `DBMS_AQMIGTOOL.CANCEL_MIGRATION`.
`cq_pending_messages`	Represents the count of messages in the `READY` state within the AQ. The count helps determine the remaining number of `READY` state messages until the AQ is empty, which is a prerequisite for using the `DBMS_AQMIGTOOL.COMMIT_MIGRATION` procedure.

Usage Notes

A prerequisite for this procedure is that the migration must already be started on the queue, meaning `DBMS_AQMIGTOOL.INIT_MIGRATION` should be invoked before executing this procedure.

29.3.3 CHECK_MIGRATION_TO_TXEVENTQ Procedure

This procedure examines the AQ's definition and data and reports any features that are unsupported in TxEventQ. If no unsupported features are detected, then `migration_report` will be empty.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.CHECK_MIGRATION_TO_TXEVENTQ (
    cqschema           IN VARCHAR2,
    cqname             IN VARCHAR2,
```

```
    migration_report    IN OUT TXEVENTQ_MIGREPORT_ARRAY,
    checkmode           IN NUMBER DEFAULT DBMS_AQMIGTOOL.ENABLE_EVALUATION);
```

Parameters

Table 29-5 CHECK_MIGRATION_TO_TXEVENTQ Procedure Parameters

Parameter	Description
cqschema	Specifies the schema name where the queue exists
cqname	Specifies the queue name for which unsupported features need to be checked
migration_report	A Varray containing details of unsupported events and their corresponding descriptions. It holds the most recent 20 unsupported features.
checkmode	Specifies the mode in which the user wants to check. The following are the possible values: DBMS_AQMIGTOOL.CURRENT: This mode generates a report using the current definition of the AQ and its data. DBMS_AQMIGTOOL.ENABLE_EVALUATION (Default): This mode, along with the CURRENT option report, will enable the monitoring of the AQ. It helps the capturing runtime-specific unsupported features. As more workload is applied to the AQ, the unsupported features, if found any, are recorded in an internal table accessible through the USER_TXEVENTQ_MIGRATION_STATUS view. Users can disable monitoring using DBMS_AQMIGTOOL.DISABLE_MIGRATION_CHECK.

Usage Notes

Several features like relative message identifier, sequence deviation, and transformation are not supported in TxEventQ. If the queue uses any of them, they will be recorded in the migration_report.

It is recommended to use the DBMS_AQMIGTOOL.CHECK_MIGRATION_TO_TXEVENTQ procedure to detect unsupported features before beginning the migration process.

> **See Also:**
>
> Limitations and Workarounds in *Oracle Database Transactional Event Queues and Advanced Queuing User's Guide*

29.3.4 CHECK_STATUS Procedure

This procedure returns the status of the migration process. In case any unsupported features are detected, the procedure will return details about the most recent unsupported feature, including its description. On the other hand, if no unsupported features are detected, it will return a status of 'No Compatibility Error'.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.CHECK_STATUS (
    cqschema                    IN VARCHAR2,
    cqname                      IN VARCHAR2,
    status                      IN OUT VARCHAR2,
    migration_comment           IN OUT VARCHAR2);
```

Parameters

Table 29-6 CHECK_STATUS Procedure Parameters

Parameter	Description
cqschema	Specifies the schema name where the queue exists
cqname	Specifies the name of the queue for which the migration process status needs to be checked
status	Return the compatibility status. In case of incompatibility, that is, detection of unsupported features; the most recent unsupported event will be returned, and the status return format will be: "Compatibility Error: <feature_name> Unsupported Feature".
migration_comment	If the status is incompatible, the description of the unsupported event will be provided

Usage Notes

A prerequisite for this procedure is that the migration must already be started on the queue, meaning DBMS_AQMIGTOOL.INIT_MIGRATION should be invoked before executing this procedure.

The following table will store the information regarding events during pre-init or post-init migration. Users can access this information through security views: DBA_TXEVENTQ_MIGRATION_STATUS, USER_TXEVENTQ_MIGRATION_STATUS, and ALL_TXEVENTQ_MIGRATION_STATUS.

```
sys.aq$_migration_status(
        migration_id            RAW(16);
        source_schema           VARCHAR2(128) NOT NULL,
        source_queue            VARCHAR2(128) NOT NULL,
        source_queue_table      VARCHAR2(128),
        target_schema           VARCHAR2(128) NOT NULL,
        target_queue            VARCHAR2(128) NOT NULL,
        status                  NUMBER,
        event                   VARCHAR(128),
        event_timestamp         TIMESTAMP WITH TIME ZONE,
        event_id                NUMBER,
        event_error             VARCHAR2(1024),
        ordering                VARCHAR(30),
        suffix                  VARCHAR2(2),
        mig_mode                NUMBER,
        spare1                  NUMBER,
        spare2                  VARCHAR2(30),
        spare3                  TIMESTAMP WITH TIME ZONE
)
```

A unique migration_id will be assigned to each initiated migration.

> **See Also:**
>
> DBA_TXEVENTQ_MIGRATION_STATUS in Oracle Database Reference
>
> USER_TXEVENTQ_MIGRATION_STATUS in Oracle Database Reference
>
> ALL_TXEVENTQ_MIGRATION_STATUS in Oracle Database Reference

29.3.5 CLEAR_UNSUPPORTED_FEATURE_TABLE Procedure

This procedure removes entries from the underlying table of the `USER_TXEVENTQ_MIGRATION_STATUS` view. This view stores records related to unsupported features detected by the `DBMS_AQMIGTOOL.CHECK_MIGRATION_TO_TXEVENTQ` procedure and details of other migration procedure calls (`INIT_MIGRATION`/`COMMIT_MIGRATION`/ `CANCEL_MIGRATION`) used for internal purposes.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.CLEAR_UNSUPPORTED_FEATURE_TABLE (
    cqschema             IN   VARCHAR2,
    cqname               IN   VARCHAR2 DEFAULT NULL,
    eraseall            IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 29-7 CLEAR_UNSUPPORTED_FEATURE_TABLE Procedure Parameters

Parameter	Description
`cqschema`	Specifies the schema name where the queue exists
`cqname`	Specifies the name of the queue for which records need to be cleared
`eraseall`	`TRUE` erases all the records for the specified queue. **Note:** The `TRUE` value is intended solely for internal purposes and should not be used without consulting Oracle support. `FALSE` erases records related to unsupported features only.

Usage Notes

Users can use this procedure to erase the records generated by the `DBMS_AQMIGTOOL.CHECK_MIGRATION_TO_TXEVENTQ` procedure. Dropping the queue or the user will also erase records for that queue. However, executing `DBMS_AQMIGTOOL.COMMIT_MIGRATION` or `DBMS_AQMIGTOOL.CANCEL_MIGRATION` will not clear the records for the queue. Therefore, this procedure offers the flexibility to erase records for a queue at any point.

29.3.6 COMMIT_MIGRATION Procedure

This procedure completes the migration process. It drops AQ and renames the interim TxEventQ to the AQ's name, and enables the TxEventQ for all operations. It is important to note that an empty AQ (that is, with no messages in the `READY` state) is required to execute the procedure successfully; otherwise, an exception will be raised.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.COMMIT_MIGRATION (
    cqschema         IN VARCHAR,
    cqname           IN VARCHAR,
    ignore_warning   IN BOOLEAN DEFAULT FALSE));
```

Parameters

Table 29-8 COMMIT_MIGRATION Procedure Parameters

Parameter	Description
`cqschema`	Specifies the schema name where the queue exists
`cqname`	Specifies the queue name for which the migration needs to be completed
`ignore_warning`	After `DBMS_AQMIGTOOL.INIT_MIGRATION`, all unsupported events are recorded as warnings. `TRUE` means the procedure will ignore the warnings and complete the migration. `FALSE` means the procedure will throw an error if there are any warnings.

Usage Notes

A prerequisite for this procedure is that the migration must already be started on the queue. In other words, the `DBMS_AQMIGTOOL.INIT_MIGRATION` procedure should be invoked before executing this procedure. The messages from the AQ's exception queue will not be copied to the TxEventQ's exception queue.

29.3.7 DISABLE_MIGRATION_CHECK Procedure

This procedure disables the internal monitoring of the AQ aimed at detecting unsupported features. It also stops the recording of events for the `USER_TXEVENTQ_MIGRATION_STATUS` view.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.DISABLE_MIGRATION_CHECK (
    cqschema       IN VARCHAR2)
    cqname         IN VARCHAR2);
```

Parameters

Table 29-9 DISABLE_MIGRATION_CHECK Procedure Parameter

Parameter	Description
`cqschema`	Specifies the schema name where the queue exists
`cqname`	Specifies the name of the queue on which migration monitoring needs to be disabled

Usage Notes

As a pre-migration step, the user can start recording unsupported features by invoking DBMS_AQMIGTOOL.CHECK_MIGRATION_TO_TXEVENTQ with DBMS_AQMIGTOOL.ENABLE_EVALUATION option. This can be followed by running a comprehensive workload on the AQ to detect potential issues and then stop recording unsupported features by calling DBMS_AQMIGTOOL.DISABLE_MIGRATION_CHECK. The user can modify the workload if any migration issues are found before repeating the process. If no migration issues were found, migration of the AQ can be attempted.

29.3.8 INIT_MIGRATION Procedure

This procedure examines the definition and data of the AQ to detect any features that are not supported in TxEventQ. If any unsupported features are detected, an exception is raised. Otherwise, the procedure starts the migration process by creating an interim TxEventQ that copies the configuration of the AQ, which includes payload type, rules, subscribers, privileges, PL/SQL notifications, and more.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.INIT_MIGRATION (
    cqschema            IN VARCHAR2,
    cqname              IN VARCHAR2 DEFAULT NULL,
    txeventqschema      IN VARCHAR2 DEFAULT NULL,
    txeventqname        IN VARCHAR2 DEFAULT NULL,
    mig_mode            IN NUMBER DEFAULT DBMS_AQMIGTOOL.INTERACTIVE,
    ordering            IN NUMBER DEFAULT DBMS_AQMIGTOOL.GLOBAL,
    suffix              IN VARCHAR2 DEFAULT 'M');
```

Parameters

Table 29-10 INIT_MIGRATION Procedure Parameters

Parameter	Description
cqschema	Specifies the schema name where the queue (AQ) exists
cqname	Specifies the name of the queue (AQ) for which the migration process is to be started. If cqname is NULL, then migration initiates for all AQs in cqschema, ignoring queues with unsupported features. Execution displays the count of queues where DBMS_AQMIGTOOL.INIT_MIGRATION succeeded.
txeventqschema	Specifies the schema name where the target TxEventQ is intended to be created. Only provide this value for DBMS_AQMIGTOOL.ONLY_DEFINITION mode. If the invoker is on a different schema than txeventqschema, it needs MANAGE_ANY queue system privilege to execute this procedure successfully.
txeventqname	Specifies the name of the queue for the target TxEventQ to be created. Only provide this value for DBMS_AQMIGTOOL.ONLY_DEFINITION mode.

Table 29-10 (Cont.) INIT_MIGRATION Procedure Parameters

Parameter	Description
mig_mode	Specifies the migration mode. The following are the possible values: DBMS_AQMIGTOOL.AUTOMATIC: Enqueue and dequeue operations are allowed in this mode, but a background job will attempt to execute DBMS_AQMIGTOOL.COMMIT_MIGRATION once no messages are left in the READY state in AQ and no unsupported features are detected. DBMS_AQMIGTOOL.INTERACTIVE (Default): In this mode, both enqueue and dequeue operations are allowed. DBMS_AQMIGTOOL.OFFLINE: Only dequeue operations are allowed in this mode, which helps in reducing the workload by restricting the new enqueue operations. DBMS_AQMIGTOOL.ONLY_DEFINITION: This mode creates a separate TxEventQ with the same configuration as the AQ instead of an interim TxEventQ setup. This completes the migration process, with AQ and TxEventQ remaining in the system. The messages present in AQ will not be copied to the newly created TxEventQ. **Note:** For the DBMS_AQMIGTOOL.ONLY_DEFINITION mode, there is no need to call DBMS_AQMIGTOOL.COMMIT_MIGRATION or DBMS_AQMIGTOOL.CANCEL_MIGRATION to complete or cancel the migration; this DBMS_AQMIGTOOL.INIT_MIGRATION call is sufficient. However, for other modes, the user must explicitly call other procedures in the DBMS_AQMIGTOOL package to proceed further.
ordering	Specifies the message level ordering the user wants to follow. The following are the possible values: DBMS_AQMIGTOOL.GLOBAL (Default): This option implements global-level message ordering by setting the number of event streams in the TxEventQ to one. DBMS_AQMIGTOOL.SESSION: This option imposes a message order only for a session. The number of event streams in the TxEventQ will be set based on the initialization parameter _aq_init_shards. Oracle recommends using this option to achieve the full performance benefits of TxEventQ.
suffix	Specifies a single character suffix for naming the interim TxEventQ. The interim TxEventQ name will be in the format <cqname>_<suffix>. The default value for the suffix is M.

Usage Notes

- This procedure will also create the default exception queue on the TxEventQ, following the naming format AQ$_<TxEventQ_name>_E.
- The following points are not relevant for DBMS_AQMIGTOOL.ONLY_DEFINITION mode but apply to other modes:

  - It restricts AQ from administrative operations to maintain interim TxEventQ configuration integrity until the migration is completed or canceled.
  - The Enqueue - Dequeue operations are allowed on AQ:
    * Enqueue requests for new messages are directed to the interim TxEventQ.
    * Messages are first dequeued from AQ. If no messages are in the `READY` state, then messages are dequeued from the interim TxEventQ.
  - Users are restricted from performing all direct operations on the interim TxEventQ. Enqueue and dequeue operations on interim TxEventQ will always be performed through AQ.

- If the procedure triggers an exception due to the detection of unsupported features, it is recommended to use the `DBMS_AQMIGTOOL.CHECK_MIGRATION_TO_TXEVENTQ` procedure to obtain a detailed list of the detected unsupported features.
- If there is a name conflict, such as when a queue (AQ) being migrated shares the same name as the queue table in which it resides, then only the `DBMS_AQMIGTOOL.ONLY_DEFINITION` mode is supported. For any other modes, attempting to proceed will result in an exception being raised.

29.3.9 PURGE_QUEUE_MESSAGES Procedure

This procedure purges messages from the queue. It can perform message purging from the AQ, the interim TxEventQ, or both, depending on user input. Specifically, one of the prerequisites for executing `DBMS_AQMIGTOOL.COMMIT_MIGRATION` is to ensure an empty AQ, meaning no messages in the `READY` state. This procedure allows users to efficiently purge all messages from the AQ to fulfill this requirement.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.PURGE_QUEUE_MESSAGES (
    cqschema       IN VARCHAR2,
    cqname         IN VARCHAR2,
    purge_option   IN NUMBER DEFAULT DBMS_AQMIGTOOL.ONLY_CQ);
```

Parameters

Table 29-11 PURGE_QUEUE_MESSAGES Procedure Parameters

Parameter	Description
`cqschema`	Specifies the schema name where the queue exists
`cqname`	Specifies the name of the queue where the migration process started
`purge_option`	Specifies the option from which queue messages need to be purged. The following options are available: `DBMS_AQMIGTOOL.ONLY_CQ` (Default): Purge messages only from the AQ. `DBMS_AQMIGTOOL.ONLY_TXEVENTQ`: Purge messages only from the interim TxEventQ. `DBMS_AQMIGTOOL.BOTH`: Purge messages from both the AQ and the interim TxEventQ.

Usage Notes

Suppose the count of messages in the `READY` state within the AQ is obtained from the `DBMS_AQMIGTOOL.CHECK_MIGRATED_MESSAGE` procedure is large, and the user wishes to speed

up the DBMS_AQMIGTOOL.COMMIT_MIGRATION process without waiting for dequeues to consume the messages, the DBMS_AQMIGTOOL.PURGE_QUEUE_MESSAGES procedure can be used.

A prerequisite for this procedure is that the migration must already be started on the queue, meaning DBMS_AQMIGTOOL.INIT_MIGRATION should be invoked before executing this procedure.

29.3.10 RECOVER_MIGRATION Procedure

This procedure restores the migration state to the nearest feasible and consistent point, either before or after the execution of DBMS_AQMIGTOOL.CANCEL_MIGRATION, DBMS_AQMIGTOOL.COMMIT_MIGRATION, or DBMS_AQMIGTOOL.INIT_MIGRATION. The recovered state is then displayed to the user through the output parameter recovery_message, providing guidance for further action. If migration procedures experience unexpected failures, such as instance crashes, then this procedure can be used to recover the migration to the nearest consistent state like before INIT_MIGRATION, after INIT_MIGRATION, after COMMIT_MIRGATION, and after CANCEL_MIGRATION.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.RECOVER_MIGRATION (
    cqschema              IN VARCHAR2,
    cqname                IN VARCHAR2,
    recovery_message      OUT VARCHAR2);
```

Parameters

Table 29-12 RECOVER_MIGRATION Procedure Parameters

Parameter	Description
cqschema	Specifies the schema name where the queue exists
cqname	Specifies the name of the queue where the migration procedure was attempted
recovery_message	Returns a descriptive message indicating the restored migration state's nearest feasible and consistent point

Usage Notes

The following table shows all the possible recovery_message:

Table 29-13 Recommended action table according to recovery_message

Error while executing migration procedure:	recovery_message	Recommended Action
DBMS_AQMIGTOOL.INIT_MIGRATION	State is restored to before INIT_MIGRATION call execution.	To start the migration, the user must explicitly call DBMS_AQMIGTOOL.INIT_MIGRATION again.
DBMS_AQMIGTOOL.INIT_MIGRATION	State is restored to after INIT_MIGRATION call execution.	No further action is needed to start the migration, as this procedure has successfully started.

Table 29-13 (Cont.) Recommended action table according to recovery_message

Error while executing migration procedure:	recovery_message	Recommended Action
DBMS_AQMIGTOOL.COMMIT_MIGRATION	State is restored to after COMMIT_MIGRATION call execution.	No further action is needed to complete the migration, as this procedure has successfully completed it.
DBMS_AQMIGTOOL.CANCEL_MIGRATION	State is restored to before CANCEL_MIGRATION call execution.	To proceed with canceling the migration, the user must explicitly call DBMS_AQMIGTOOL.CANCEL_MIGRATION again.
DBMS_AQMIGTOOL.CANCEL_MIGRATION	State is restored to after CANCEL_MIGRATION call execution.	No further action is needed to cancel the migration, as this procedure has successfully cancelled it.
No migration procedure	No need for recovery call.	Since no migration procedure execution is detected, restoring it to the nearest feasible and consistent point is not required.

29.3.11 RENAME_QUEUE Procedure

This procedure renames the TxEventQ along with its default exception queue if present.

Syntax

```
PROCEDURE DBMS_AQMIGTOOL.RENAME_QUEUE (
    schema             IN VARCHAR2,
    qname              IN VARCHAR2,
    new_qname          IN VARCHAR2);
```

Parameters

Table 29-14 RENAME_QUEUE Procedure Parameters

Parameter	Description
schema	Specifies the schema name where the queue exists
qname	Specifies the current name of the queue to be renamed
new_qname	Specify the new name to be given to the existing queue. The new name must be unique within a schema and must follow object name guidelines in Oracle Database SQL Language Reference with regard to reserved characters.

Usage Notes

This procedure requires a prerequisite step: the user must perform DBMS_AQADM.STOP_QUEUE to ensure there are no concurrent enqueue and dequeue transactions.

If the default exception queue is present, it will be renamed from <schema>.AQ$_<qname>_E to <schema>.AQ$_<new_qname>_E.

30
DBMS_ASSERT

The DBMS_ASSERT package provides an interface to validate properties of the input value.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information about "Avoiding SQL Injection in PL/SQL"

This chapter contains the following topics:

- Operational Notes
- Summary of DBMS_ASSERT Subprograms

30.1 DBMS_ASSERT Operational Notes

If the condition which determines the property asserted in a function is not met then a value error is raised. Otherwise the input value is returned through the return value. Most functions return the value unchanged, however, several functions modify the value.

30.2 Summary of DBMS_ASSERT Subprograms

This section describes the subprograms of the DBMS_ASSERT package.

Table 30-1 DBMS_ASSERT Package Subprograms

Subprogram	Description
ENQUOTE_LITERAL Function	Enquotes a string literal
ENQUOTE_NAME Function	Ensures that a string is enclosed by quotation marks, then checks that the result is a valid SQL identifier.
NOOP Functions	Returns the value without any checking
QUALIFIED_SQL_NAME Function	Verifies that the input string is a qualified SQL name
SCHEMA_NAME Function	Verifies that the input string is an existing schema name
SIMPLE_SQL_NAME Function	Verifies that the input string is a simple SQL name
SQL_OBJECT_NAME Function	Verifies that the input parameter string is a qualified SQL identifier of an existing SQL object

30.2.1 ENQUOTE_LITERAL Function

This function adds leading and trailing single quotes to a string literal.

Syntax

```
DBMS_ASSERT.ENQUOTE_LITERAL (
   str            VARCHAR2)
RETURN VARCHAR2;
```

Parameters

Table 30-2 *ENQUOTE_LITERAL Function Parameters*

Parameter	Description
`str`	String to enquote

Usage Notes

- Verify that all single quotes except leading and trailing characters are paired with adjacent single quotes.
- No additional quotes are added if the name was already in quotes.

30.2.2 ENQUOTE_NAME Function

This function encloses the provided string in double quotes (quotation marks). No additional quotes are added if the string was already in quotes (quotation marks). The quoted string is then checked to see if it is a valid (quoted) simple SQL name.

For more information on Database object names and qualifiers, see Oracle Database SQL Language Reference.

Syntax

```
DBMS_ASSERT.ENQUOTE_NAME (
   str            VARCHAR2,
   capitalize     BOOLEAN DEFAULT TRUE)
RETURN VARCHAR2;
```

Parameters

Table 30-3 *ENQUOTE_NAME Function Parameters*

Parameter	Description
`str`	String to enquote
`capitalize`	If `TRUE` or defaulted, alphabetic characters of `str` which was not in quotes are translated to upper case

Usage Notes

- No additional quotes are added if the name was already in quotes.
- Verify that all other double quotes in the string are adjacent pairs of double quotes.

Examples

```
-- This procedure creates a single column table in the createOneColumnTable's schema.
create or replace procedure createOneColumnTable(proposedTableNamevarchar2) is
BEGIN
  IF
    (proposedTableName is NULL)
  THEN
    raise value_error;
  END IF;
  -- The use of ENQUOTE_NAME ensures that the table will be created in the
  -- definer's schema and not in some other schema even if the definer has
  -- privileges to create tables in other schemas.
  EXECUTE IMMEDIATE 'create table ' ||
  DBMS_ASSERT.ENQUOTE_NAME(proposedTableName) || ' (c1 number)';
  EXCEPTION
    WHEN
      others
    THEN
      dbms_output.put_line('Table creation failed due to: ' || SQLERRM);
END;
/

-- Examples of ENQUOTE_NAME showing input/output relationships
BEGIN
  -- 'eMp' becomes '"EMP"' since it is unquoted
  dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('eMp'));
END;
/

BEGIN
  -- For quoted strings, the case is preserved
  dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('"EmP"'));
END;
/

-- Invalid identifier example
BEGIN
  dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('SCOTT."EMP"'));
END;
/

-- CHR(0) examples
-- The following examples illustrates that CHR(0), the NULL character, cannot appear
-- in the string; such a string poses a SQL injection risk.
BEGIN
  dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('BAD' || CHR(0) || 'IDENTIFIER'));
END;
/

BEGIN
   dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('"SCOTT' || CHR(0) || '.EMP"'));
END;
/

-- Oracle allows a period (.) to be a part of a quoted string
BEGIN
  dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('"SCOTT.EMP"'));
END;
/
```

```
-- The single quotation mark ('), as opposed to a double quotation mark, can appear in
the string
-- Note: In Oracle, a single quotation mark is specified in a literal using two single
-- quotes. The first quotation mark escapes the second quotation mark in the same way
that
-- backslash (\) in POSIX is an escape character.
BEGIN
   dbms_output.put_line(DBMS_ASSERT.ENQUOTE_NAME('"O''LEARY"'));
END;
/
```

30.2.3 NOOP Functions

This function returns the value without any checking.

Syntax

```
DBMS_ASSERT.NOOP (
   str      VARCHAR2 CHARACTER SET ANY_CS)
 RETURN     VARCHAR2 CHARACTER SET str%CHARSET;

DBMS_ASSERT.NOOP (
   str      CLOB CHARACTER SET ANY_CS)
 RETURN     CLOB CHARACTER SET str%CHARSET;
```

Parameters

Table 30-4 *NOOP Function Parameters*

Parameter	Description
str	Input value

30.2.4 QUALIFIED_SQL_NAME Function

This function verifies that the input string is a qualified SQL name.

Syntax

```
DBMS_ASSERT.QUALIFIED_SQL_NAME (
   str      VARCHAR2 CHARACTER SET ANY_CS)
 RETURN     VARCHAR2 CHARACTER SET str%CHARSET;
```

Parameters

Table 30-5 *QUALIFIED_SQL_NAME Function Parameters*

Parameter	Description
str	Input value

Exceptions

ORA44004: string is not a qualified SQL name

Usage Notes

A qualified SQL name <qualified name> can be expressed by the following grammar:

```
<local qualified name> ::= <simple name> {'.' <simple name>}
<database link name> ::= <local qualified name> ['@' <connection string>]
<connection string> ::= <simple name>
<qualified name> ::= <local qualified name> ['@' <database link name>]
```

30.2.5 SCHEMA_NAME Function

This function verifies that the input string is an existing schema name.

Syntax

```
DBMS_ASSERT.SCHEMA_NAME (
   str      VARCHAR2 CHARACTER SET ANY_CS)
 RETURN     VARCHAR2 CHARACTER SET str%CHARSET;
```

Parameters

Table 30-6 *SCHEMA_NAME Function Parameters*

Parameter	Description
str	Input value

Exceptions

ORA44001: Invalid schema name

Usage Notes

By definition, a schema name need not be just a simple SQL name. For example, "FIRST LAST" is a valid schema name. As a consequence, care must be taken to quote the output of schema name before concatenating it with SQL text.

30.2.6 SIMPLE_SQL_NAME Function

This function verifies that the input string is a simple SQL name.

Syntax

```
DBMS_ASSERT.SIMPLE_SQL_NAME (
   str      VARCHAR2 CHARACTER SET ANY_CS)
 RETURN     VARCHAR2 CHARACTER SET str%CHARSET;
```

Parameters

Table 30-7 *SIMPLE_SQL_NAME Function Parameters*

Parameter	Description
str	Input value

Exceptions

ORA44003: string is not a simple SQL name

Usage Notes

- The input value must be meet the following conditions:
  - The name must begin with an alphabetic character. It may contain alphanumeric characters as well as the characters _, $, and # in the second and subsequent character positions.
  - Quoted SQL names are also allowed.
  - Quoted names must be enclosed in double quotes.
  - Quoted names allow any characters between the quotes.
  - Quotes inside the name are represented by two quote characters in a row, for example, "a name with "" inside" is a valid quoted name.
  - The input parameter may have any number of leading and/or trailing white space characters.
- The length of the name is not checked.

30.2.7 SQL_OBJECT_NAME Function

This function verifies that the input parameter string is a qualified SQL identifier of an existing SQL object.

Syntax

```
DBMS_ASSERT.SQL_OBJECT_NAME (
   str      VARCHAR2 CHARACTER SET ANY_CS)
 RETURN     VARCHAR2 CHARACTER SET str%CHARSET;
```

Parameters

Table 30-8 *SQL_OBJECT_NAME Function Parameters*

Parameter	Description
`str`	Input value

Exceptions

`ORA44002`: Invalid object name

Usage Notes

The use of synonyms requires that the base object exists.

31
DBMS_AUDIT_MGMT

The DBMS_AUDIT_MGMT package provides subprograms to manage audit trail records. These subprograms enable audit administrators to manage the audit trail. In a mixed-mode environment, these audit trails comprise the database, operating system (OS), and XML audit trails. In a unified auditing environment, this comprises the unified audit trail.

This chapter contains the following topics:

- DBMS_AUDIT_MGMT Overview
- DBMS_AUDIT_MGMT Deprecated and Desupported Subprograms
- DBMS_AUDIT_MGMT Security Model
- DBMS_AUDIT_MGMT Constants
- DBMS_AUDIT_MGMT Views
- Subprogram Groups
- Summary of DBMS_AUDIT_MGMT Subprograms

> **See Also:**
>
> - *Oracle Database Security Guide* regarding verifying security access with auditing
> - *Oracle Database Upgrade Guide* regarding migrating to unified auditing

31.1 DBMS_AUDIT_MGMT Overview

Database auditing helps meet your database security and compliance requirements. In a mixed mode environment, audit records are written to database tables, operating system (OS) files, or XML files depending on the AUDIT_TRAIL initialization parameter setting. If you have upgraded to unified auditing, then the audit records are written to the unified audit trail.

> **Note:**
>
> Traditional auditing is desupported in Oracle Database 23ai. Oracle recommends that you use unified auditing instead. Any current traditional audit settings that you have will still be honored, but you cannot create new traditional audit settings. You can delete existing traditional audit settings. See *Oracle Database Security Guide* for more information.

In a mixed mode environment, when AUDIT_TRAIL is set to DB, database records are written to the AUD$ table. In a unified auditing environment, audit records are written to a read-only table in the AUDSYS schema. The contents of this table are available from the UNIFIED_AUDIT_TRAIL data dictionary view.When AUDIT_TRAIL is set to OS, audit records are written to operating

system files. When AUDIT_TRAIL is set to XML, audit records are written to operating system files in XML format.

With Unified Auditing facility, all audit records are written to the unified audit trail in a uniform format and are made available through the UNIFIED_AUDIT_TRAIL views.

It is important to manage your audit records properly in order to ensure efficient performance and disk space management. The DBMS_AUDIT_MGMT subprograms enable you to efficiently manage your audit trail records.

If you have not yet migrated to unified auditing, the DBMS_AUDIT_MGMT package provides a subprogram that allows you to move the database audit trail tables out of the SYSTEM tablespace. This improves overall database performance by reducing the load on the SYSTEM tablespace. It also enables you to dedicate an optimized tablespace for audit records.

For a mixed mode environment, the DBMS_AUDIT_MGMT subprograms also enable you to manage your operating system and XML audit files. You can define properties like the maximum size and age of an audit file. This enables you to keep the file sizes of OS and XML audit files in check.

The DBMS_AUDIT_MGMT subprograms enable you to perform cleanup operations on all audit trail types. Audit trails can be cleaned based on the last archive timestamp value. The last archive timestamp represents the timestamp of the most recent audit record that was securely archived.

The DBMS_AUDIT_MGMT package provides a subprogram that enables audit administrators to set the last archive timestamp for archived audit records. This subprogram can also be used by external archival systems to set the last archive timestamp.

The DBMS_AUDIT_MGMT subprograms also enable you to configure jobs that periodically delete audit trail records. The frequency with which these jobs should run can be controlled by the audit administrator.

> **See Also:**
>
> *Oracle Database Security Guide* for a detailed description of unified auditing

31.2 DBMS_AUDIT_MGMT Deprecated and Desupported Subprograms

Oracle recommends that you do not use deprecated subprograms in new applications. Support for deprecated features is for backward compatibility only.

The following have been deprecated from the DBMS_AUDIT_MGMT package starting in Oracle Database 12c release 12.2:

- DBMS_AUDIT_MGMT.FLUSH_UNIFIED_AUDIT_TRAIL procedure (desupported starting in Oracle Database 23.4)
- AUDIT_TRAIL_WRITE mode of the AUDIT_TRAIL_PROPERTY parameter of the DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_PROPERTY procedure

These are no longer necessary because audit records now bypass the common logging infrastructure queues and are directly written to a new internal relational table.

31.3 DBMS_AUDIT_MGMT Security Model

All DBMS_AUDIT_MGMT subprograms require the user to have EXECUTE privilege over the DBMS_AUDIT_MGMT package. The SYSDBA and AUDIT_ADMIN roles have EXECUTE privileges on the package by default.

Oracle strongly recommends that only audit administrators should have the EXECUTE privilege on the DBMS_AUDIT_MGMT package and be granted the AUDIT_ADMIN role.

Executions of the DBMS_AUDIT_MGMT subprograms are always audited.

31.4 DBMS_AUDIT_MGMT Constants

The DBMS_AUDIT_MGMT package defines several constants that can be used for specifying parameter values.

These constants shown in the following tables:

- Table 31-1
- Table 31-2
- Table 31-3

Audit trails can be classified based on whether audit records are written to database tables, operating system files, or XML files. The following table lists the audit trail type constants.

Table 31-1 DBMS_AUDIT_MGMT Constants - Audit Trail Types

Constant	Type	Description
AUDIT_TRAIL_ALL	PLS_INTEGER	All audit trail types. This includes the standard database audit trail (SYS.AUD$, SYS.FGA_LOG$ and unified audit trail tables), operating system (OS) audit trail, and XML audit trail.
AUDIT_TRAIL_AUD_STD	PLS_INTEGER	Standard database audit records in the SYS.AUD$ table
AUDIT_TRAIL_DB_STD	PLS_INTEGER	Both standard audit (SYS.AUD$) and FGA audit(SYS.FGA_LOG$) records
AUDIT_TRAIL_FGA_STD	PLS_INTEGER	Standard database fine-grained auditing (FGA) records in the SYS.FGA_LOG$ table
AUDIT_TRAIL_FILES	PLS_INTEGER	Both operating system (OS) and XML audit trails
AUDIT_TRAIL_OS	PLS_INTEGER	Operating system audit trail. This refers to the audit records stored in operating system files.
AUDIT_TRAIL_UNIFIED	PLS_INTEGER	Unified audit trail. In unified auditing, all audit records are written to the unified audit trail and are made available through the unified audit trail views, such as UNIFIED_AUDIT_TRAIL.
AUDIT_TRAIL_UNIFIED_FILES	PLS_INTEGER	Operating system spillover files in each database (primary or standby); used in archive or purge operations
AUDIT_TRAIL_UNIFIED_TABLE	PLS_INTEGER	Records from the AUDSYS.AUD$UNIFIED table; used in archive or purge operations

Table 31-1 (Cont.) DBMS_AUDIT_MGMT Constants - Audit Trail Types

Constant	Type	Description
AUDIT_TRAIL_XML	PLS_INTEGER	XML audit trail. This refers to the audit records stored in XML files.

Audit trail properties determine the audit configuration settings. The following table lists the constants related to audit trail properties.

Table 31-2 DBMS_AUDIT_MGMT Constants - Audit Trail Properties

Constant	Type	Description
AUDIT_TRAIL_WRITE_MODE	PLS_INTEGER	**Note:** AUDIT_TRAIL_WRITE_MODE has been deprecated starting in Oracle Database 12*c* release 12.2. It is retained only for backwards compatibility, and has no effect on Oracle Database versions 12.2 and higher. A value of AUDIT_TRAIL_IMMEDIATE_WRITE indicates that the audit record must be immediately persisted and not to be queued. By contrast, AUDIT_TRAIL_QUEUED_WRITE indicates that the audit record can be queued and persisting can be done according the database's flushing strategy. See Also *Oracle Database Security Guide*
CLEAN_UP_INTERVAL	PLS_INTEGER	Interval, in hours, after which the cleanup procedure is called to clear audit records in the specified audit trail
DB_DELETE_BATCH_SIZE	PLS_INTEGER	Specifies the batch size to be used for deleting audit records in database audit tables. The audit records are deleted in batches of size equal to DB_DELETE_BATCH_SIZE.
FILE_DELETE_BATCH_SIZE	PLS_INTEGER	Specifies the batch size to be used for deleting audit files in the audit directory. The audit files are deleted in batches of size equal to FILE_DELETE_BATCH_SIZE.
OS_FILE_MAX_AGE	PLS_INTEGER	Specifies the maximum number of days for which an operating system (OS) or XML audit file can be kept open before a new audit file gets created
OS_FILE_MAX_SIZE	PLS_INTEGER	Specifies the maximum size, in kilobytes (KB), to which an operating system (OS) or XML audit file can grow before a new file is opened

The audit trail purge job cleans the audit trail. The following table lists the constants related to purge job status values.

Table 31-3 DBMS_AUDIT_MGMT Constants - Purge Job Status

Constant	Type	Description
PURGE_JOB_DISABLE	PLS_INTEGER	Disables a purge job
PURGE_JOB_ENABLE	PLS_INTEGER	Enables a purge job

31.5 DBMS_AUDIT_MGMT Views

DBMS_AUDIT_MGMT views are used to display `DBMS_AUDIT_MGMT` configuration and cleanup events.

These views listed in the following table.

Table 31-4 Views used by DBMS_AUDIT_MGMT

View	Description
`DBA_AUDIT_MGMT_CLEAN_EVENTS`	Displays the cleanup event history
`DBA_AUDIT_MGMT_CLEANUP_JOBS`	Displays the currently configured audit trail purge jobs
`DBA_AUDIT_MGMT_CONFIG_PARAMS`	Displays the currently configured audit trail properties
`DBA_AUDIT_MGMT_LAST_ARCH_TS`	Displays the last archive timestamps set for the audit trails

> **See Also:**
>
> *Oracle Database Reference* for more information on these views

31.6 DBMS_AUDIT_MGMT Subprogram Groups

The `DBMS_AUDIT_MGMT` package subprograms can be grouped into two categories: Audit Trail Management Subprograms and Audit Trail Cleanup Subprograms.

- Audit Trail Management Subprograms
- Audit Trail Cleanup Subprograms

31.6.1 DBMS_AUDIT_MGMT Audit Trail Management Subprograms

Audit trail management subprograms enable you to manage audit trail properties.

Table 31-5 Audit Trail Management Subprograms

Subprogram	Description
ALTER_PARTITION_INTERVAL Procedure	Changes the unified audit internal relational table's partition interval
CLEAR_AUDIT_TRAIL_PROPERTY Procedure	Clears the value for the audit trail property that you specify
GET_ AUDIT_TRAIL_PROPERTY_VALUE Function	Returns the property value set by the SET_AUDIT_TRAIL_PROPERTY Procedure
GET_LAST_ARCHIVE_TIMESTAMP Function	Returns the timestamp set by the SET_LAST_ARCHIVE_TIMESTAMP Procedure in that database instance
LOAD_UNIFIED_AUDIT_FILES Procedure	Loads the data from the spillover OS audit files in a unified audit trail into the designated unified audit trail tablespace

Table 31-5 (Cont.) Audit Trail Management Subprograms

Subprogram	Description
SET_AUDIT_TRAIL_LOCATION Procedure	Moves the audit trail tables from their current tablespace to a user-specified tablespace
SET_AUDIT_TRAIL_PROPERTY Procedure	Sets an audit trail property for the audit trail type that you specify
SET_LAST_ARCHIVE_TIMESTAMP Procedure	Sets a timestamp indicating when the audit records or files were last archived
TRANSFER_UNIFIED_AUDIT_RECORDS Procedure	Transfers audit records from the common logging infrastructure (CLI) swap table to the `AUDSYS.AUD$UNIFIED` relational table

The Summary of DBMS_AUDIT_MGMT Subprograms contains a complete listing of all subprograms in the package.

31.6.2 DBMS_AUDIT_MGMT Audit Trail Cleanup Subprograms

Audit trail cleanup subprograms help you perform cleanup related operations on the audit trail records.

Table 31-6 Audit Trail Cleanup Subprograms

Subprogram	Description
CLEAN_AUDIT_TRAIL Procedure	Deletes audit trail records or files that have been archived
CLEAR_LAST_ARCHIVE_TIMESTAMP Procedure	Clears the timestamp set by the SET_LAST_ARCHIVE_TIMESTAMP Procedure
CREATE_PURGE_JOB Procedure	Creates a purge job for periodically deleting the audit trail records or files
DEINIT_CLEANUP Procedure	Undoes the setup and initialization performed by the INIT_CLEANUP Procedure
DROP_OLD_UNIFIED_AUDIT_TABLES Procedure	Drops old unified audit tables following the cloning of a pluggable database (PDB)
DROP_PURGE_JOB Procedure	Drops the purge job created using the CREATE_PURGE_JOB Procedure
INIT_CLEANUP Procedure	Sets up the audit management infrastructure and sets a default cleanup interval for audit trail records or files
IS_CLEANUP_INITIALIZED Function	Checks to see if the INIT_CLEANUP Procedure has been run for an audit trail type
SET_PURGE_JOB_INTERVAL Procedure	Sets the interval at which the CLEAN_AUDIT_TRAIL Procedure is called for the purge job that you specify
SET_PURGE_JOB_STATUS Procedure	Enables or disables the purge job that you specify

The Summary of DBMS_AUDIT_MGMT Subprograms contains a complete listing of all subprograms in the package.

31.7 Summary of DBMS_AUDIT_MGMT Subprograms

This table lists and describes the subprograms of the `DBMS_AUDIT_MGMT` package

Table 31-7 DBMS_AUDIT_MGMT Package Subprograms

Subprogram	Description
ALTER_PARTITION_INTERVAL Procedure	Changes the unified audit internal relational table's partition interval
CLEAN_AUDIT_TRAIL Procedure	Deletes audit trail records that have been archived
CLEAR_AUDIT_TRAIL_PROPERTY Procedure	Clears the value for the audit trail property that you specify
CLEAR_LAST_ARCHIVE_TIMESTAMP Procedure	Clears the timestamp set by the SET_LAST_ARCHIVE_TIMESTAMP Procedure
CREATE_PURGE_JOB Procedure	Creates a purge job for periodically deleting the audit trail records
DEINIT_CLEANUP Procedure	Undoes the setup and initialization performed by the INIT_CLEANUP Procedure
DROP_OLD_UNIFIED_AUDIT_TABLES Procedure	Drops old unified audit tables following the cloning of a pluggable database (PDB)
DROP_PURGE_JOB Procedure	Drops the purge job created using the CREATE_PURGE_JOB Procedure
GET_AUDIT_COMMIT_DELAY Function	Returns the audit commit delay time as the number of seconds. This is the maximum time that it takes to `COMMIT` an audit record to the database audit trail.
GET_ AUDIT_TRAIL_PROPERTY_VALUE Function	Returns the property value set by the SET_AUDIT_TRAIL_PROPERTY Procedure
GET_LAST_ARCHIVE_TIMESTAMP Function	Returns the timestamp set by the SET_LAST_ARCHIVE_TIMESTAMP Procedure in that database instance
INIT_CLEANUP Procedure	Sets up the audit management infrastructure and sets a default cleanup interval for audit trail records
IS_CLEANUP_INITIALIZED Function	Checks to see if the INIT_CLEANUP Procedure has been run for an audit trail type
LOAD_UNIFIED_AUDIT_FILES Procedure	Loads the data from the spillover OS audit files in a unified audit trail into the designated unified audit trail tablespace
SET_AUDIT_TRAIL_LOCATION Procedure	Moves the audit trail tables from their current tablespace to a user-specified tablespace
SET_AUDIT_TRAIL_PROPERTY Procedure	Sets the audit trail properties for the audit trail type that you specify
SET_LAST_ARCHIVE_TIMESTAMP Procedure	Sets a timestamp indicating when the audit records were last archived
SET_PURGE_JOB_INTERVAL Procedure	Sets the interval at which the CLEAN_AUDIT_TRAIL Procedure is called for the purge job that you specify
SET_PURGE_JOB_STATUS Procedure	Enables or disables the purge job that you specify
TRANSFER_UNIFIED_AUDIT_RECORDS Procedure	Transfers audit records from the common logging infrastructure (CLI) swap table to the `AUDSYS.AUD$UNIFIED` relational table

31.7.1 ALTER_PARTITION_INTERVAL Procedure

This procedure changes the unified audit internal relational table's partition interval.

Syntax

```
DBMS_AUDIT_MGMT.ALTER_PARTITION_INTERVAL(
   interval_number          IN BINARY_INTEGER,
   interval_frequency       IN VARCHAR2);
```

Parameters

Table 31-8 ALTER_PARTITION_INTERVAL Procedure Parameters

Parameter	Description
`interval_number`	Sets how often the database creates partitions for the unified audit internal relational table. For example, to specify that the partition is created every two days, you must enter `2`.
`interval_frequency`	Sets the frequency for the value that was set in the `interval_number` setting. For example, for a partition to be created every two days, with `interval_number` set to `2`, you must set `interval_frequency` to `DAY`. Supported values are `YEAR`, `MONTH`, and `DAY`. The default is `DAY`.

Usage Notes

- The interval frequency that you choose depends on the rate of audit records that are generated in your database. The default setting is for one month. If you have a high audit record rate and are using the default, then too many audit records may be generated in the same partition. In this case, you should change the interval frequency to a more frequent interval, such as one month or one day. If the audit record rate generation is not so high, then you may want to keep it at the default of one month.

Example

The following example sets the partition interval to occur every two months.

```
BEGIN
 DBMS_AUDIT_MGMT.ALTER_PARTITION_INTERVAL(
   interval_number       => 2,
   interval_frequency    => 'MONTH');
 END;
```

31.7.2 CLEAN_AUDIT_TRAIL Procedure

This procedure deletes audit trail records.

The `CLEAN_AUDIT_TRAIL` procedure is usually called after the SET_LAST_ARCHIVE_TIMESTAMP Procedure has been used to set the last archived timestamp for the audit records.

Syntax

```
DBMS_AUDIT_MGMT.CLEAN_AUDIT_TRAIL(
   audit_trail_type         IN PLS_INTEGER,
   use_last_arch_timestamp  IN BOOLEAN DEFAULT TRUE,
```

```
container               IN PLS_INTEGER DEFAULT CONTAINER_CURRENT,
database_id             IN NUMBER DEFAULT NULL,
container_guid          IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 31-9 CLEAN_AUDIT_TRAIL Procedure Parameters

Parameter	Description
audit_trail_type	The audit trail type for which the cleanup operation needs to be performed. Audit trail types are listed in Table 31-1 .
use_last_arch_timestamp	Specifies whether the last archived timestamp should be used for deciding on the records that should be deleted. A value of TRUE indicates that only audit records created before the last archive timestamp should be deleted. A value of FALSE indicates that all audit records should be deleted. The default value is TRUE. Oracle recommends using this value, as this helps guard against inadvertent deletion of records.
container	Values: CONTAINER_CURRENT for the connected pluggable database (PDB) or CONTAINER_ALL for all pluggable databases (PDBs). When CONTAINER is set to CONTAINER_ALL, this purges the audit trail in all the PDBs, otherwise it only purges from the connected PDB.
database_id	Database ID (DBID) of the audit records to cleanup
container_guid	Container GUID of the audit records to cleanup **Note:** This parameter has been deprecated but is currently retained for backward compatibility.

Usage Notes

The following usage notes apply:

- When cleaning up operating system (OS) or XML audit files, only files in the current audit directory, specified by the AUDIT_FILE_DEST parameter, are cleaned up.
- For Windows platforms, no cleanup is performed when the audit_trail_type parameter is set to DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS. This is because operating system (OS) audit records on Windows are written to the Windows Event Viewer.
- For Unix platforms, no cleanup is performed for cases where the operating system (OS) audit records are written to the syslog. When the audit_trail_type parameter is set to DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS, it removes only the *.aud files under the directory specified by the AUDIT_FILE_DEST initialization parameter.

> **See Also:**
>
> "AUDIT_SYSLOG_LEVEL" in the *Oracle Database Reference*

- When the audit_trail_type parameter is set to DBMS_AUDIT_MGMT.AUDIT_TRAIL_XML, this procedure only removes XML audit files (*.xml) from the current audit directory.

  Oracle database maintains a book-keeping file (adx_*$ORACLE_SID*.txt) for the XML audit files. This file is not removed by the cleanup procedure.

- If the cleanup of the unified audit trail is performed when the `use_last_arch_timestamp` parameter is set to `TRUE`:
  - If you set the `database_id` value for the cleanup operation, then this value is used with the last archive timestamp while `CLEAN_AUDIT_TRAIL` runs. While cleaning up the unified audit records that are present in the database tables, if the `HIGH_VALUE` of an audit table partition is less than last archive timestamp, then those partitions will be dropped directly and for remaining audit table partitions, records will be deleted based on the `database_id` value along with last archive timestamp. However, for the unified audit records that are present during the cleanup of spillover operating system audit files, the `database_id` value is ignored. Cleanup for operating system audit files is based on the last archive timestamp only. If you want to have the `database_id` value used for the cleanup operation of unified audit trail records that are present in the spillover operating system audit files, then load the contents of these files into database tables by using the `DBMS_AUDIT_MGMT.LOAD_UNIFIED_AUDIT_FILES` procedure before you run `DBMS_AUDIT_MGMT.CLEAN_AUDIT_TRAIL`.
  - If you do not set the `database_id` value when you invoke the `CLEAN_AUDIT_TRAIL` procedure, then Oracle Database purges unified audit records based on the last archive timestamp value irrespective of `database_id` values and irrespective of the location (that is, database tables or spillover operating system audit files) where the unified audit records reside. During cleanup, last archive timestamp value of current database ID of the container will be used.
- If the cleanup of the unified audit trail is performed when the `use_last_arch_timestamp` parameter is set to `FALSE`:
  - If you set the `database_id` value for the cleanup operation, then this value is used while `CLEAN_AUDIT_TRAIL` cleans up the unified audit trail records that are present in database tables. However, the `database_id` value is not used for the cleanup of unified audit trail records that are present in spillover operating system files. If want the `database_id` value to be used in the `CLEAN_AUDIT_TRAIL` operation of unified audit records that are present in spillover operating system audit files, then load the contents of these files to the database tables by using the `DBMS_AUDIT_MGMT.LOAD_UNIFIED_AUDIT_FILES` procedure before you run `DBMS_AUDIT_MGMT.CLEAN_AUDIT_TRAIL`.
  - If you do not set the `database_id` value when you invoke the `CLEAN_AUDIT_TRAIL` procedure, then Oracle Database purges all unified audit records irrespective of `database_id` values and irrespective of the location (that is, database tables or spillover operating system audit files) where the unified audit records reside.
- `CLEAN_AUDIT_TRAIL` procedure expects that the last archive timestamp set via SET_LAST_ARCHIVE_TIMESTAMP has been set usring SYS_EXTRACT_UTC if it want to refer to SYSTIMESTAMP. Reference to data in Unified Audit trail is done in UTC timeformat.
- If the PDB database is read only then clean up of that database audit record does not take place and returns without performing a clean up job on `AUDSYS.AUD$UNIFIED` table.
- If timestamp of database is ahead of current system timestamp then it will return an error.
- In a multitenant setup, if one of the PDB has timestamp ahead of the current system timestamp (in UTC format) then records for that PDB are not cleaned where as the other database do clean up their audit trail records leaving the last activity of "clean up" in the audit trail.

- The SET_LAST_ARCHIVE_TIMESTAMP procedure and CLEAN_AUDIT_TRAIL procedure should be not be executed in the same transaction block. Otherwise, the results are usually unpredictable.
- Make sure that the transaction for SET_LAST_ARCHIVE_TIMESTAMP is completed before you call the CLEAN_AUDIT_TRAIL procedure. To have predictable results in terms of cleaning records until the time provided in SET_LAST_ARCHIVE_TIMESTAMP(UTC format), make sure that SET_LAST_ARCHIVE_TIMESTAMP and CLEAN_AUDIT_TRAIL procedures are not called in the same transaction. It is recommended to commit before you call the CLEAN_AUDIT_TRAIL procedure.
- Run the CLEAN_AUDIT_TRAIL procedure when USE_LAST_ARCH_TIMESTAMP = TRUE uses the last archived timestamp that was set for the current DB_UNIQUE_NAME column value.

Examples

The following example calls the CLEAN_AUDIT_TRAIL procedure to clean up the operating system (OS) audit trail records that were updated before the last archive timestamp.

```
BEGIN
DBMS_AUDIT_MGMT.CLEAN_AUDIT_TRAIL(
   audit_trail_type         =>  DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
   use_last_arch_timestamp  =>  TRUE);
END;
```

31.7.3 CLEAR_AUDIT_TRAIL_PROPERTY Procedure

This procedure clears the value for the specified audit trail property.

Audit trail properties are set using the SET_AUDIT_TRAIL_PROPERTY Procedure.

The CLEAR_AUDIT_TRAIL_PROPERTY procedure can optionally reset the property value to it's default value through the use_default_values parameter.

Syntax

```
DBMS_AUDIT_MGMT.CLEAR_AUDIT_TRAIL_PROPERTY(
   audit_trail_type        IN PLS_INTEGER,
   audit_trail_property    IN PLS_INTEGER,
   use_default_values      IN BOOLEAN DEFAULT FALSE) ;
```

Parameters

Table 31-10 CLEAR_AUDIT_TRAIL_PROPERTY Procedure Parameters

Parameter	Description
audit_trail_type	The audit trail type for which the property needs to be cleared. Audit trail types are listed in Table 31-1
audit_trail_property	The audit trail property whose value needs to be cleared. You cannot clear the value for the CLEANUP_INTERVAL property. Audit trail properties are listed in Table 31-2
use_default_values	Specifies whether the default value of the audit_trail_property should be used in place of the cleared value. A value of TRUE causes the default value of the parameter to be used. A value of FALSE causes the audit_trail_property to have no value. The default value for this parameter is FALSE.

Usage Notes

The following usage notes apply:

- You can use this procedure to clear the value for an audit trail property that you do not wish to use. For example, if you do not want a restriction on the operating system audit file size, then you can use this procedure to reset the `OS_FILE_MAX_SIZE` property.

  You can also use this procedure to reset an audit trail property to it's default value. You need to set `use_default_values` to `TRUE` when invoking the procedure.
- The `DB_DELETE_BATCH_SIZE` property needs to be individually cleared for the `AUDIT_TRAIL_AUD_STD` and `AUDIT_TRAIL_FGA_STD` audit trail types. You cannot clear this property collectively using the `AUDIT_TRAIL_DB_STD` and `AUDIT_TRAIL_ALL` audit trail types.
- If you clear the value of the `DB_DELETE_BATCH_SIZE` property with `use_default_value` set to `FALSE`, the default value of `DB_DELETE_BATCH_SIZE` is still assumed. This is because audit records are always deleted in batches.
- The `FILE_DELETE_BATCH_SIZE` property needs to be individually cleared for the `AUDIT_TRAIL_OS` and `AUDIT_TRAIL_XML` audit trail types. You cannot clear this property collectively using the `AUDIT_TRAIL_FILES` and `AUDIT_TRAIL_ALL` audit trail types.
- If you clear the value of the `FILE_DELETE_BATCH_SIZE` property with `use_default_value` set to `FALSE`, the default value of `FILE_DELETE_BATCH_SIZE` is still assumed. This is because audit files are always deleted in batches.
- You cannot clear the value for the `CLEANUP_INTERVAL` property.
- You cannot clear the value for the `AUDIT_TRAIL_WRITE_MODE` property.

Examples

The following example calls the `CLEAR_AUDIT_TRAIL_PROPERTY` procedure to clear the value for the audit trail property, `OS_FILE_MAX_SIZE`. The procedure uses a value of `FALSE` for the `USE_DEFAULT_VALUES` parameter. This means that there will be no maximum size threshold for operating system (OS) audit files.

```
BEGIN
DBMS_AUDIT_MGMT.CLEAR_AUDIT_TRAIL_PROPERTY(
   AUDIT_TRAIL_TYPE        =>  DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
   AUDIT_TRAIL_PROPERTY    =>  DBMS_AUDIT_MGMT.OS_FILE_MAX_SIZE,
   USE_DEFAULT_VALUES      =>  FALSE );
END;
```

31.7.4 CLEAR_LAST_ARCHIVE_TIMESTAMP Procedure

This procedure clears the timestamp set by the SET_LAST_ARCHIVE_TIMESTAMP Procedure.

Syntax

```
DBMS_AUDIT_MGMT.CLEAR_LAST_ARCHIVE_TIMESTAMP(
   audit_trail_type         IN PLS_INTEGER,
   rac_instance_number      IN PLS_INTEGER DEFAULT NULL,
   container                IN PLS_INTEGER DEFAULT CONTAINER_CURRENT,
   database_id              IN NUMBER DEFAULT NULL,
   container_guid           IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 31-11 CLEAR_LAST_ARCHIVE_TIMESTAMP Procedure Parameters

Parameter	Description
audit_trail_type	The audit trail type for which the timestamp needs to be cleared. Audit trail types are listed in Table 31-1 .
rac_instance_number	The instance number for the Oracle Real Application Clusters (Oracle RAC) instance. The default value is NULL. The rac_instance_number is not relevant for single instance databases. You can find the instance number by issuing the SHOW PARAMETER INSTANCE_NUMBER command in SQL*Plus.
container	Values: CONTAINER_CURRENT for the connected pluggable database (PDB) or CONTAINER_ALL for all pluggable databases (PDBs). When CONTAINER is set to CONTAINER_ALL, this clears the last archive timestamp from all the PDBs, otherwise it clears from only the connected PDB.
database_id	Database ID (DBID) of the audit records to cleanup
container_guid	Container GUID of the audit records to cleanup **Note:** This parameter has been deprecated but is currently retained for backward compatibility.

Usage Notes

The following usage notes apply:

- The timestamp for only one audit_trail_type can be cleared at a time.
- The following are invalid audit_trail_type values for this procedure and cannot be used:
  - AUDIT_TRAIL_ALL
  - AUDIT_TRAIL_DB_STD
  - AUDIT_TRAIL_FILES

Examples

The following example calls the CLEAR_LAST_ARCHIVE_TIMESTAMP procedure to clear the timestamp value for the operating system (OS) audit trail type.

```
BEGIN
DBMS_AUDIT_MGMT.CLEAR_LAST_ARCHIVE_TIMESTAMP(
   audit_trail_type     =>  DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
   rac_instance_number  =>  1);
END;
```

Related Topics

- SET_LAST_ARCHIVE_TIMESTAMP Procedure
  This procedure sets a timestamp indicating when the audit records were last archived. The audit administrator provides the timestamp to be attached to the audit records.

31.7.5 CREATE_PURGE_JOB Procedure

This procedure creates a purge job for periodically deleting the audit trail records.

This procedure carries out the cleanup operation at intervals specified by the user. It calls the CLEAN_AUDIT_TRAIL Procedure to perform the cleanup operation.

The SET_PURGE_JOB_INTERVAL Procedure is used to modify the frequency of the purge job.

The SET_PURGE_JOB_STATUS Procedure is used to enable or disable the purge job.

The DROP_PURGE_JOB Procedure is used to drop a purge job created with the `CREATE_PURGE_JOB` procedure.

Syntax

```
DBMS_AUDIT_MGMT.CREATE_PURGE_JOB(
   audit_trail_type            IN PLS_INTEGER,
   audit_trail_purge_interval  IN PLS_INTEGER,
   audit_trail_purge_name      IN VARCHAR2,
   use_last_arch_timestamp     IN BOOLEAN DEFAULT TRUE,
   container                   IN PLS_INTEGER DEFAULT CONTAINER_CURRENT) ;
```

Parameters

Table 31-12 CREATE_PURGE_JOB Procedure Parameters

Parameter	Description
`audit_trail_type`	The audit trail type for which the purge job needs to be created. Audit trail types are listed in Table 31-1 .
`audit_trail_purge_interval`	The interval, in hours, at which the clean up procedure is called. A lower value means that the cleanup is performed more often.
`audit_trail_purge_name`	A name to identify the purge job.
`use_last_arch_timestamp`	Specifies whether the last archived timestamp should be used for deciding on the records that should be deleted. A value of `TRUE` indicates that only audit records created before the last archive timestamp should be deleted. A value of `FALSE` indicates that all audit records should be deleted. The default value is `TRUE`.
`container`	Values: `CONTAINER_CURRENT` for the connected pluggable database (PDB) or `CONTAINER_ALL` for all pluggable databases (PDBs). When `CONTAINER` is set to `CONTAINER_ALL`, it creates one job in the Root PDB and the invocation of this job will invoke cleanup in all the PDBs.

Usage Notes

Use this procedure to schedule the CLEAN_AUDIT_TRAIL Procedure for your audit trail records.

Examples

The following example calls the `CREATE_PURGE_JOB` procedure to create a cleanup job called `CLEANUP`, for all audit trail types. It sets the `audit_trail_purge_interval` parameter to 100. This means that the cleanup job is invoked every 100 hours. It also sets the `use_last_arch_timestamp` parameter value to `TRUE`. This means that all audit records older than the last archive timestamp are deleted.

```
BEGIN
DBMS_AUDIT_MGMT.CREATE_PURGE_JOB(
  audit_trail_type             =>  DBMS_AUDIT_MGMT.AUDIT_TRAIL_ALL,
  audit_trail_purge_interval   =>  100 /* hours */,
  audit_trail_purge_name       =>  'CLEANUP',
  use_last_arch_timestamp      =>  TRUE);
END;
```

31.7.6 DEINIT_CLEANUP Procedure

This procedure undoes the setup and initialization performed by the `INIT_CLEANUP` procedure. This procedure is deprecated starting in Oracle Database 23ai.Traditional auditing packages and functions are deprecated in Oracle Database 23ai. The `DEINIT_CLEANUP` procedure clears the value of the `default_cleanup_interval` parameter. However, when used for audit tables, it does not move the audit trail tables back to their original tablespace.

Syntax

```
DBMS_AUDIT_MGMT.DEINIT_CLEANUP(
   audit_trail_type  IN PLS_INTEGER,
   container         IN PLS_INTEGER DEFAULT CONTAINER_CURRENT);
```

Parameters

Table 31-13 DEINIT_CLEANUP Procedure Parameters

Parameter	Description
`audit_trail_type`	The audit trail type for which the procedure needs to be called. Audit trail types are listed in Table 31-1
`container`	Values: `CONTAINER_CURRENT` for the connected pluggable database (PDB) or `CONTAINER_ALL` for all pluggable databases (PDBs). When `CONTAINER` is set to `CONTAINER_ALL`, this de-initializes the audit trail from cleanup in all the pluggable databases, otherwise it de-initializes the audit trail from cleanup in the connected PDB only.

Usage Notes

With the desupport of traditional auditing, the PL/SQL packages and functions associated with traditional auditing are deprecated, This deprecation includes the packages and functions `INIT_CLEANUP`, `DEINIT_CLEANUP`, and `IS_CLEANUP_INITIALIZED`. While these packages or functions continue to operate in Oracle Database 23ai, you can neither add to or modify traditional auditing configurations.

You cannot run this procedure for `AUDIT_TRAIL_UNIFIED`. Doing so it will raise `ORA-46250 : Invalid value for argument 'AUDIT_TRAIL_TYPE'`

Examples

The following example clears the `default_cleanup_interval` parameter setting for the standard database audit trail:

```
BEGIN
DBMS_AUDIT_MGMT.DEINIT_CLEANUP(
  AUDIT_TRAIL_TYPE  => DBMS_AUDIT_MGMT.AUDIT_TRAIL_AUD_STD);
END;
```

Related Topics

- INIT_CLEANUP Procedure

31.7.7 DROP_OLD_UNIFIED_AUDIT_TABLES Procedure

This procedure drops old unified audit tables following the cloning of a pluggable database (PDB).

Syntax

```
DBMS_AUDIT_MGMT.DROP_OLD_UNIFIED_AUDIT_TABLES(
   container_guid    IN VARCHAR2) ;
```

Parameters

Table 31-14 DROP_OLD_UNIFIED_AUDIT_TABLES Procedure Parameters

Parameter	Description
`container_guid`	Container GUID of the old unified audit tables

Usage Notes

When a pluggable database gets cloned, the unified audit tables get newly created in the new pluggable database. To drop the old unified audit tables, use the `DROP_OLD_UNIFIED_AUDIT_TABLES` by specifying the old `GUID` of the PDB from which the clone was created. You can query the historical `GUIDs` from the `DBA_PDB_HISTORY` view for the given PDB.

Only use the `DBMS_AUDIT_MGMT.DROP_OLD_UNIFIED_AUDIT_TABLES` procedure if the database was upgraded from Oracle Database release 12.1 or earlier. If the database was upgraded from a later release (including release 12.2), then an `ORA-55906: Secure file log [id: 0 name: ORA$AUDIT_NEXTGEN_LOG] does not exist` error will appear. This is because starting with release 12.2, the common logging infrastructure tables that this procedure searches are no longer being created. Instead, the `AUD$UNIFIED` relational table stores the unified audit records.

Examples

```
BEGIN
    DBMS_AUDIT_MGMT.DROP_OLD_UNIFIED_AUDIT_TABLES (
      container_guid => 'E4721865A9321CB5E043EFA9E80A2D77');
END;
```

31.7.8 DROP_PURGE_JOB Procedure

This procedure drops the purge job created using the CREATE_PURGE_JOB Procedure. The name of the purge job is passed as an argument.

Syntax

```
DBMS_AUDIT_MGMT.DROP_PURGE_JOB(
   audit_trail_purge_name     IN VARCHAR2) ;
```

Parameters

Table 31-15 DROP_PURGE_JOB Procedure Parameters

Parameter	Description
audit_trail_purge_name	The name of the purge job which is being deleted. This is the purge job name that you specified with the CREATE_PURGE_JOB Procedure.

Examples

The following example calls the DROP_PURGE_JOB procedure to drop the purge job called CLEANUP.

```
BEGIN
DBMS_AUDIT_MGMT.DROP_PURGE_JOB(
  AUDIT_TRAIL_PURGE_NAME  => 'CLEANUP');
END;
```

Related Topics

- CREATE_PURGE_JOB Procedure
  This procedure creates a purge job for periodically deleting the audit trail records.

31.7.9 GET_AUDIT_COMMIT_DELAY Function

This function returns the audit commit delay time as the number of seconds. audit commit delay time is the maximum time that it takes to COMMIT an audit record to the database audit trail. If it takes more time to COMMIT an audit record than defined by the audit commit delay time, then a copy of the audit record is written to the operating system (OS) audit trail.

The audit commit delay time value is useful when determining the last archive timestamp for database audit records.

Syntax

```
DBMS_AUDIT_MGMT.GET_AUDIT_COMMIT_DELAY
  RETURN NUMBER;
```

31.7.10 GET_ AUDIT_TRAIL_PROPERTY_VALUE Function

This procedure returns the property value set by the SET_AUDIT_TRAIL_PROPERTY Procedure.

Syntax

```
DBMS_AUDIT_MGMT.GET_AUDIT_TRAIL_PROPERTY_VALUE(
   audit_trail_type         IN PLS_INTEGER,
   audit_trail_property     IN PLS_INTEGER)
 RETURN NUMBER;
```

Parameters

Table 31-16 GET_ AUDIT_TRAIL_PROPERTY_VALUE Function Parameters

Parameter	Description
audit_trail_type	The audit trail type for the timestamp to be retrieved. Audit trail types are listed in Table 31-1 .
audit_trail_property	The audit trail property that is being queried. Audit trail properties are listed in Table 31-2 .

Return Values

If the property value is cached in SGA memory, this function will return the value set by the SET_AUDIT_TRAIL_PROPERTY Procedure. Else it will return NULL.

The GET_AUDIT_TRAIL_PROPERTY_VALUE function may return an ORA-46250 error if the audit trail property value has been set to DBMS_AUDIT_MGMT.CLEAN_UP_INTERVAL. To find the cleanup interval of the purge job, query SYS.DAM_CLEANUP_JOBS$.

Examples

The following example prints the property value of OS_FILE_MAX_AGE set by the SET_AUDIT_TRAIL_PROPERTY Procedure.

```
SET_AUDIT_TRAIL_PROPERTY.
SET SERVEROUTPUT ON
DECLARE
 OS_MAX_AGE_VAL NUMBER;
BEGIN
 OS_MAX_AGE_VAL := DBMS_AUDIT_MGMT.GET_AUDIT_TRAIL_PROPERTY_VALUE(
      audit_trail_type          => DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
      audit_trail_property      => DBMS_AUDIT_MGMT. OS_FILE_MAX_AGE);
 IF OS_MAX_AGE_VAL is not NULL THEN
  DBMS_OUTPUT.PUT_LINE('The Maximum Age configured for OS Audit files is: ' ||
                       OS_MAX_AGE_VAL);
 END IF;
END;
```

Related Topics

- SET_AUDIT_TRAIL_PROPERTY Procedure
  This procedure sets an audit trail property for the audit trail type that is specified.

31.7.11 GET_LAST_ARCHIVE_TIMESTAMP Function

This procedure returns the timestamp set by the `SET_LAST_ARCHIVE_TIMESTAMP` procedure in that database instance.

Syntax

```
DBMS_AUDIT_MGMT.GET_LAST_ARCHIVE_TIMESTAMP(
   audit_trail_type     IN PLS_INTEGER)
 RETURN TIMESTAMP;
```

Parameters

Table 31-17 GET_LAST_ARCHIVE_TIMESTAMP Function Parameters

Parameter	Description
`audit_trail_type`	The audit trail type for the timestamp to be retrieved. Supported Audit trail types for this procedure are `AUDIT_TRAIL_OS` , `AUDIT_TRAIL_XML` and `AUDIT_TRAIL_UNIFIED`. All Audit trail types are listed in Table 31-1 .

Return Values

The `GET_LAST_ARCHIVE_TIMESTAMP` function returns last archive timestamp.

Usage Notes

- Alternatively, you can query the `DBA_AUDIT_MGMT_LAST_ARCH_TS` data dictionary view to check the last archived timestamp for a database that is in read-write mode.
- The `DBA_AUDIT_MGMT_LAST_ARCH_TS` data dictionary view can contain multiple last archive timestamps with different `DB_UNIQUE_NAME` column values. `GET_LAST_ARCHIVE_TIMESTAMP` will return the last archive timestamp where the `DB_UNIQUE_NAME` column value is equal to the current `DB_UNIQUE_NAME` value.

Examples

The following example prints the timestamp set by the `SET_LAST_ARCHIVE_TIMESTAMP` procedure on a `READ ONLY` database.

```
SET SERVEROUTPUT ON
DECLARE
 LAT_TS TIMESTAMP;
BEGIN
 LAT_TS := DBMS_AUDIT_MGMT.GET_LAST_ARCHIVE_TIMESTAMP(
      audit_trail_type          => DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS);
 IF LAT_TS is not NULL THEN
  DBMS_OUTPUT.PUT_LINE('The Last Archive Timestamp is: ' || to_char(LAT_TS));
 END IF;
END;
```

Related Topics

- SET_LAST_ARCHIVE_TIMESTAMP Procedure
  This procedure sets a timestamp indicating when the audit records were last archived. The audit administrator provides the timestamp to be attached to the audit records.

31.7.12 INIT_CLEANUP Procedure

This procedure sets up the audit management infrastructure and a default cleanup interval for the audit trail records. This procedure is deprecated starting in Oracle Database 23ai. Traditional auditing packages and functions are deprecated in Oracle Database 23ai.
If the audit trail tables are in the `SYSTEM` tablespace, then the procedure moves them to the `SYSAUX` tablespace. If you are using unified auditing, you do not need to run this procedure because the unified audit trail tables are in the `SYSAUX` tablespace by default. If you are not using unified auditing, refer to *Oracle Database Upgrade Guide* for documentation which references an environment without unified auditing.

Moving the audit trail tables out of the `SYSTEM` tablespace enhances overall database performance. The `INIT_CLEANUP` procedure moves the audit trail tables to the `SYSAUX` tablespace. If the SET_AUDIT_TRAIL_LOCATION Procedure has already moved the audit tables elsewhere, then no tables are moved.

The SET_AUDIT_TRAIL_LOCATION Procedure enables you to specify an alternate target tablespace for the database audit tables.

The `INIT_CLEANUP` procedure is currently not relevant for the `AUDIT_TRAIL_OS`, `AUDIT_TRAIL_XML`, and `AUDIT_TRAIL_FILES` audit trail types. No preliminary set up is required for these audit trail types.

> **See Also:**
>
> Table 31-1 for a list of all audit trail types

This procedure also sets a default cleanup interval for the audit trail records.

Syntax

```
DBMS_AUDIT_MGMT.INIT_CLEANUP(
   audit_trail_type          IN PLS_INTEGER,
   default_cleanup_interval  IN PLS_INTEGER
   container                 IN PLS_INTEGER DEFAULT CONTAINER_CURRENT);
```

Parameters

Table 31-18 INIT_CLEANUP Procedure Parameters

Parameter	Description
`audit_trail_type`	The audit trail type for which the clean up operation needs to be initialized. Audit trail types are listed in Table 31-1 except `AUDIT_TRAIL_UNIFIED`
`default_cleanup_interval`	The default time interval, in hours, after which the cleanup procedure should be called. The minimum value is 1 and the maximum is 999.

Table 31-18 (Cont.) INIT_CLEANUP Procedure Parameters

Parameter	Description
container	Values: CONTAINER_CURRENT for the connected pluggable database (PDB) or CONTAINER_ALL for all Open and Available pluggable databases (PDBs). When CONTAINER is set to CONTAINER_ALL, this initializes the audit trails for clean up in all the Open and Available pluggable databases, otherwise this initializes the audit trail in the connected PDB only. When you add a new PDB you need to initialize the audit trails for clean up in the new PDB using the CONTAINER_CURRENT option.

Usage Notes

With the desupport of traditional auditing, the PL/SQL packages and functions associated with traditional auditing are deprecated, This deprecation includes the packages and functions INIT_CLEANUP, DEINIT_CLEANUP, and IS_CLEANUP_INITIALIZED. While these packages or functions continue to operate in Oracle Database 23ai, you can neither add to or modify traditional auditing configurations.

The following usage notes apply:

- This procedure may involve data movement across tablespaces. This can be a resource intensive operation especially if your database audit trail tables are already populated. Oracle recommends that you invoke the procedure during non-peak hours.
- You should ensure that the SYSAUX tablespace, into which the audit trail tables are being moved, has sufficient space to accommodate the audit trail tables. You should also optimize the SYSAUX tablespace for frequent write operations.
- You can change the default_cleanup_interval later using the SET_AUDIT_TRAIL_PROPERTY Procedure.
- If you do not wish to move the audit trail tables to the SYSAUX tablespace, then you should use the DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_LOCATION procedure to move the audit trail tables to another tablespace before calling the INIT_CLEANUP procedure.
- Invoking this procedure with AUDIT_TRAIL_UNIFIED results in ORA-46250. It requires no initializations for cleanup since it is cleanup-ready by default.

> **See Also:**
>
> "SET_AUDIT_TRAIL_LOCATION Procedure"

Examples

The following example calls the INIT_CLEANUP procedure to set a default_cleanup_interval of 12 hours for all audit trail types:

```
BEGIN
DBMS_AUDIT_MGMT.INIT_CLEANUP(
               audit_trail_type   => DBMS_AUDIT_MGMT.AUDIT_TRAIL_ALL,
      default_cleanup_interval    => 12 /* hours */);
END;
```

> **See Also:**
>
> Table 31-1 for a list of all audit trail types

31.7.13 IS_CLEANUP_INITIALIZED Function

This function checks to see if the `INIT_CLEANUP` procedure has been run for an audit trail type. This procedure is deprecated starting in Oracle Database 23ai. Traditional auditing packages and functions are deprecated in Oracle Database 23ai.
With the desupport of traditional auditing, the PL/SQL packages and functions associated with traditional auditing are deprecated, This deprecation includes the packages and functions `INIT_CLEANUP`, `DEINIT_CLEANUP`, and `IS_CLEANUP_INITIALIZED`. While these packages or functions continue to operate in Oracle Database 23ai, you can neither add to or modify traditional auditing configurations.

The `IS_CLEANUP_INITIALIZED` function returns `TRUE` if the procedure has already been run for the audit trail type. It returns `FALSE` if the procedure has not been run for the audit trail type.

This function is currently not relevant for the `AUDIT_TRAIL_OS`, `AUDIT_TRAIL_XML`, and `AUDIT_TRAIL_FILES` audit trail types. The function always returns TRUE for these audit trail types. No preliminary set up is required for these audit trail types.

> **See Also:**
>
> Table 31-1 for a list of all audit trail types

Syntax

```
DBMS_AUDIT_MGMT.IS_CLEANUP_INITIALIZED(
   audit_trail_type  IN PLS_INTEGER
   container         IN PLS_INTEGER DEFAULT CONTAINER_CURRENT)
 RETURN BOOLEAN;
```

Parameters

Table 31-19 IS_CLEANUP_INITIALIZED Function Parameters

Parameter	Description
`audit_trail_type`	The audit trail type for which the function needs to be called. Note that this does not apply to `AUDIT_TRAIL_UNIFIED`. Audit trail types are listed in Table 31-1
`container`	Values: `CONTAINER_CURRENT` for the connected pluggable database (PDB) or `CONTAINER_ALL` for all pluggable databases (PDBs). • When `CONTAINER` is set to `CONTAINER_ALL`, this function returns the initialization status of all the pluggable databases. The function returns `FALSE` even if one of the PDBs is not initialized. • When `CONTAINER` is set to `CONTAINER_CURRENT`, this returns the initialization status of the connected PDB.

Examples

The following example checks to see if the standard database audit trail type has been initialized for cleanup operation. If the audit trail type has not been initialized, then it calls the INIT_CLEANUP Procedure to initialize the audit trail type.

```
BEGIN
 IF
   NOT DBMS_AUDIT_MGMT.IS_CLEANUP_INITIALIZED(DBMS_AUDIT_MGMT.AUDIT_TRAIL_AUD_STD)
 THEN
   DBMS_AUDIT_MGMT.INIT_CLEANUP(
      audit_trail_type          => DBMS_AUDIT_MGMT.AUDIT_TRAIL_AUD_STD,
      default_cleanup_interval  => 12 /* hours */);
 END IF;
END;
```

Related Topics

- INIT_CLEANUP Procedure

31.7.14 LOAD_UNIFIED_AUDIT_FILES Procedure

This procedure loads the data from the spillover OS audit files in a unified audit trail into the designated unified audit trail tablespace.

> **See Also:**
>
> *Oracle Database Security Guide* for information about moving the OS audit trail records into the unified audit trail

Syntax

```
DBMS_AUDIT_MGMT.LOAD_UNIFIED_AUDIT_FILES (
   container       IN BINARY_INTEGER,
   load_batch_size IN PLS_INTEGER);
```

Parameters

Table 31-20 LOAD_UNIFIED_AUDIT_FILES Procedure Parameters

Parameter	Description
container	Values: CONTAINER_CURRENT for the connected pluggable database (PDB) or CONTAINER_ALL for all pluggable databases (PDBs). • CONTAINER_CURRENT - loads the unified audit files from $ORACLE_BASE/audit/$ORACLE_SID OS directory to the tables in only current PDB • CONTAINER_ALL - loads the unified audit files from $ORACLE_BASE/audit/$ORACLE_SID OS directory to the tables in the respective PDBs, but for all the active PDBs

Table 31-20 (Cont.) LOAD_UNIFIED_AUDIT_FILES Procedure Parameters

Parameter	Description
load_batch_size	Specifies the number of spillover OS audit files to be loaded into the designated unified audit trail tablespace. The load_batch_size parameter can have a minimum value of 1 and maximum value of 32767. The default value is equivalent to DEFAULT_FILE_LOAD_BATCH_SIZE.

Usage Notes

- Ensure that you set the audit data designated tablespace to be online before you run the LOAD_UNIFIED_AUDIT_FILES procedure.

31.7.15 SET_AUDIT_TRAIL_LOCATION Procedure

This procedure moves the audit trail tables from their current tablespace to a user-specified tablespace.

The SET_AUDIT_TRAIL_LOCATION procedure is not relevant for the AUDIT_TRAIL_OS, AUDIT_TRAIL_XML, and AUDIT_TRAIL_FILES audit trail types. The AUDIT_FILE_DEST initialization parameter is the only way you can specify the destination directory for these audit trail types.

> **See Also:**
>
> - Table 31-1 for a list of all audit trail types
> - AUDIT_FILE_DEST in the *Oracle Database Reference*

Syntax

```
DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_LOCATION(
   audit_trail_type            IN PLS_INTEGER,
   audit_trail_location_value  IN VARCHAR2) ;
```

Parameters

Table 31-21 SET_AUDIT_TRAIL_LOCATION Procedure Parameters

Parameter	Description
audit_trail_type	The audit trail type for which the audit trail location needs to be set. Audit trail types are listed in Table 31-1
audit_trail_location_value	Target location or tablespace for the audit trail records

Usage Notes

The following usage notes apply:

- Before changing the audit trail location, check the block size of the new location that you plan to use. The block size must be the same size as that of the current audit trail location.

Otherwise, an `ORA-14520: Tablespace string block size [string] does not match existing object block size [string]` error appears. To check the block size of a tablespace, query the `USER_TABLESPACES` or `DBA_TABLESPACES` data dictionary view. See *Oracle Database VLDB and Partitioning Guide* for more information about partition restrictions for multiple block sizes.

- By default, audit trail records are written to the `SYSAUX` tablespace. You can change the tablespace location to a user tablespace using this procedure. However the designated tablespace must be an automatic storage space management (ASSM) tablespace.
- This procedure is valid for the following `audit_trail_type` values:
  - `AUDIT_TRAIL_AUD_STD`
  - `AUDIT_TRAIL_FGA_STD`
  - `AUDIT_TRAIL_DB_STD`
  - `AUDIT_TRAIL_UNIFIED`
- For the `audit_trail_type` values `AUDIT_TRAIL_AUD_STD`, `AUDIT_TRAIL_FGA_STD`, and `AUDIT_TRAIL_DB_STD`, you should ensure that the target tablespace, into which the audit trail tables are being moved, has sufficient space to accommodate the audit trail tables. You should also optimize the target tablespace for frequent write operations.
- This procedure involves data movement across tablespaces. For the `audit_trail_type` values `AUDIT_TRAIL_AUD_STD`, `AUDIT_TRAIL_FGA_STD`, and `AUDIT_TRAIL_DB_STD`, this can be a resource intensive operation especially if your database audit trail tables are already populated. Oracle recommends that you invoke the procedure during non-peak hours.
- You optionally can specify an encrypted tablespace for the audit trail location.
- When `AUDIT_TRAIL_TYPE` is `AUDIT_TRAIL_UNIFIED`, this procedure sets the tablespace for newer audit records in the unified audit trail but does not move the older audit records. Thus, it is not resource intensive for the unified audit trail.
- The `UNIFIED_AUDIT_TRAIL` data dictionary view is built on top of an internal relational table. This table is an interval partitioned table (irrespective of database editions) with a default interval of 1 day. This setting means that when you execute the `DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_LOCATION` procedure, only newly created partitions of the internal table are created in the new tablespace that is set as part of this procedure. Existing partitions of this table remain in the earlier tablespace (`SYSAUX` is the default tablespace for this internal table). If you want to change this table's partition interval, then use the `DBMS_AUDIT_MGMT.ALTER_PARTITION_INTERVAL` procedure.
- As a best practice to improve performance and the availability of audit records, enable the `AUTOEXTEND` option on the audit tablespace.
- If the tablespace that is specified for the `audit_trail_location_value` parameter has a block size that is different from the current tablespace, then this procedure for `AUDIT_TRAIL_UNIFIED` `audit_trail_type` fails with an `ORA-14520: Tablespace string block size [string] does not match existing object block size [string]` error. This is due to the restrictions of multiple block sizes for partition tables.

Examples

The following example moves the database audit trail tables, `AUD$` and `FGA_LOG$`, from the current tablespace to a user-created tablespace called `RECORDS`:

```
BEGIN
DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_LOCATION(
       audit_trail_type            => DBMS_AUDIT_MGMT.AUDIT_TRAIL_DB_STD,
```

```
      audit_trail_location_value  =>  'RECORDS');
END;
```

31.7.16 SET_AUDIT_TRAIL_PROPERTY Procedure

This procedure sets an audit trail property for the audit trail type that is specified.

The procedure sets the properties `OS_FILE_MAX_SIZE`, `OS_FILE_MAX_AGE`, and `FILE_DELETE_BATCH_SIZE` for operating system (OS), XML, and unified audit trail types. The `OS_FILE_MAX_SIZE` and `OS_FILE_MAX_AGE` properties determine the maximum size and age of an audit trail file before a new audit trail file is created. The `OS_FILE_MAX_SIZE` and `OS_FILE_MAX_AGE` properties determine the maximum size and age of an audit trail file before a new audit trail file gets created. The `FILE_DELETE_BATCH_SIZE` property specifies the number of audit trail files that are deleted in one batch.

The procedure sets the properties `DB_DELETE_BATCH_SIZE` and `CLEANUP_INTERVAL` for the database audit trail type. `DB_DELETE_BATCH_SIZE` specifies the batch size in which records get deleted from audit trail tables. This ensures that if a cleanup operation gets interrupted midway, the process does not need to start afresh the next time it is invoked. This is because all batches before the last processed batch are already committed.

The `CLEANUP_INTERVAL` specifies the frequency, in hours, with which the cleanup procedure is called.

Syntax

```
DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_PROPERTY(
   audit_trail_type            IN PLS_INTEGER,
   audit_trail_property        IN PLS_INTEGER,
   audit_trail_property_value  IN PLS_INTEGER) ;
```

Parameters

Table 31-22 SET_AUDIT_TRAIL_PROPERTY Procedure Parameters

Parameter	Description
`audit_trail_type`	The audit trail type for which the property needs to be set. Audit trail types are listed in Table 31-1
`audit_trail_property`	The audit trail property that is being set. Audit trail properties are listed in Table 31-2

Table 31-22 (Cont.) SET_AUDIT_TRAIL_PROPERTY Procedure Parameters

Parameter	Description
audit_trail_property_value	The value of the property specified using audit_trail_property. The following are valid values for audit trail properties: • OS_FILE_MAX_SIZE can have a minimum value of 1 and maximum value of 2000000. The default value is 10000. OS_FILE_MAX_SIZE is measured in kilobytes (KB). • OS_FILE_MAX_AGE can have a minimum value of 1 and a maximum value of 497. The default value is 5. OS_FILE_MAX_AGE is measured in days. • DB_DELETE_BATCH_SIZE can have a minimum value of 100 and a maximum value of 1000000. The default value is 10000. DB_DELETE_BATCH_SIZE is measured as the number of audit records that are deleted in one batch. • FILE_DELETE_BATCH_SIZE can have a minimum value of 100 and a maximum value of 1000000. The default value is 1000. FILE_DELETE_BATCH_SIZE is measured as the number of audit files that are deleted in one batch. • CLEANUP_INTERVAL can have a minimum value of 1 and a maximum value of 999. The default value is set using the INIT_CLEANUP Procedure. CLEANUP_INTERVAL is measured in hours.

Usage Notes

The following usage notes apply:

- The audit trail properties for which you do not explicitly set values use their default values.
- If you have set both the OS_FILE_MAX_SIZE and OS_FILE_MAX_AGE properties for an operating system (OS), XML or unified audit trail type, then a new audit trail file is created depending on which of these two limits is reached first

  For example, let us take a scenario where OS_FILE_MAX_SIZE is 10000 and OS_FILE_MAX_AGE is 5. If the operating system audit file is already more than 5 days old and has a size of 9000 KB, then a new audit file is opened. This is because one of the limits has been reached.
- You can set the OS_FILE_MAX_SIZE and OS_FILE_MAX_AGE properties for the unified audit trail type by using AUDIT_TRAIL_UNIFIED or AUDIT_TRAIL_UNIFIED_FILES. Note that AUDIT_TRAIL_UNIFIED is deprecated
- The DB_DELETE_BATCH_SIZE property needs to be individually set for the AUDIT_TRAIL_AUD_STD and AUDIT_TRAIL_FGA_STD audit trail types. You cannot set this property collectively using the AUDIT_TRAIL_DB_STD and AUDIT_TRAIL_ALL audit trail types.
- The DB_DELETE_BATCH_SIZE property enables you to control the number of audit records that are deleted in one batch. Setting a large value for this parameter requires increased allocation for the undo log space.
- The FILE_DELETE_BATCH_SIZE property needs to be individually set for the AUDIT_TRAIL_OS and AUDIT_TRAIL_XML audit trail types. You cannot set this property collectively using the AUDIT_TRAIL_FILES and AUDIT_TRAIL_ALL audit trail types.

- The FILE_DELETE_BATCH_SIZE property enables you to control the number of audit files that are deleted in one batch. Setting a very large value may engage the GEN0 background process for a long time.
- In Oracle Database Standard Edition, you can only associate the tablespace for unified auditing once. You should perform this association before you generate any audit records for the unified audit trail. The default tablespace is SYSAUX. After you have associated the tablespace, you cannot modify it on the Standard Edition because the partitioning feature is not supported in the Standard Edition.

Examples

The following example calls the SET_AUDIT_TRAIL_PROPERTY procedure to set the OS_FILE_MAX_SIZE property for the operating system (OS) audit trail. It sets this property value to 102400. This means that a new audit file gets created every time the current audit file size reaches 100 MB.

```
BEGIN
DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_PROPERTY(
       audit_trail_type            => DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
       audit_trail_property        =>  DBMS_AUDIT_MGMT.OS_FILE_MAX_SIZE,
       audit_trail_property_value  =>  102400 /* 100MB*/ );
END;
```

The following example calls the SET_AUDIT_TRAIL_PROPERTY procedure to set the OS_FILE_MAX_AGE property for the operating system (OS) audit trail. It sets this property value to 5. This means that a new audit file gets created every sixth day.

```
BEGIN
DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_PROPERTY(
       audit_trail_type            => DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
       audit_trail_property        =>  DBMS_AUDIT_MGMT.OS_FILE_MAX_AGE,
       audit_trail_property_value  =>  5 /* days */);
END;
```

The following example calls the SET_AUDIT_TRAIL_PROPERTY procedure to set the DB_DELETE_BATCH_SIZE property for the AUDIT_TRAIL_AUD_STD audit trail. It sets this property value to 100000. This means that during a cleanup operation, audit records are deleted from the SYS.AUD$ table in batches of size 100000.

```
BEGIN
DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL_PROPERTY(
       audit_trail_type              => DBMS_AUDIT_MGMT.AUDIT_TRAIL_AUD_STD,
       audit_trail_property          => DBMS_AUDIT_MGMT.DB_DELETE_BATCH_SIZE,
       audit_trail_property_value    => 100000 /* delete batch size */);
END;
```

31.7.17 SET_LAST_ARCHIVE_TIMESTAMP Procedure

This procedure sets a timestamp indicating when the audit records were last archived. The audit administrator provides the timestamp to be attached to the audit records.

The CLEAN_AUDIT_TRAIL Procedure uses this timestamp to decide on the audit records to be deleted.

Syntax

```
DBMS_AUDIT_MGMT.SET_LAST_ARCHIVE_TIMESTAMP(
   audit_trail_type         IN PLS_INTEGER,
   last_archive_time        IN TIMESTAMP,
```

```
rac_instance_number      IN PLS_INTEGER DEFAULT NULL,
container                IN PLS_INTEGER DEFAULT CONTAINER_CURRENT,
database_id              IN NUMBER DEFAULT NULL,
container_guid           IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 31-23 SET_LAST_ARCHIVE_TIMESTAMP Procedure Parameters

Parameter	Description
audit_trail_type	The audit trail type for which the timestamp needs to be set. Audit trail types are listed in Table 31-1 .
last_archive_time	The TIMESTAMP value based on which the audit records or files should be deleted. This indicates the last time when the audit records or files were archived. **Note:** ORA-46250 error is thrown if any future date or TIMESTAMP value is given.
rac_instance_number	The instance number for the Oracle Real Application Clusters (Oracle RAC) instance. The default value is NULL. The rac_instance_number parameter is not accepted when AUDIT_TRAIL_TYPE has values: • audit_trail_aud_std • audit_trail_db_std • audit_trail_fga_std • audit_trail_unified Set rac_instance_number to 1 for a single-instance database.
container	Values: CONTAINER_CURRENT for the connected pluggable database (PDB) or CONTAINER_ALL for all pluggable databases (PDBs). When CONTAINER is set to CONTAINER_ALL, this sets the value for last archive timestamp in all the pluggable databases, otherwise it sets the value in the connected PDB only.
database_id	Database ID (DBID) of the audit records to cleanup
container_guid	Container GUID of the audit records to cleanup **Note:** This parameter has been deprecated but is currently retained for backward compatibility.

Usage Notes

The following usage notes apply:

- The last_archive_time must be specified in Coordinated Universal Time (UTC) when the audit trail types are AUDIT_TRAIL_AUD_STD, AUDIT_TRAIL_FGA_STD, or AUDIT_TRAIL_UNIFIED. This is because the database audit trails store the timestamps in UTC. UTC is also known as Greenwich Mean Time (GMT).
- The last_archive_time must be specified as the local time zone time when the audit trail types are AUDIT_TRAIL_OS or AUDIT_TRAIL_XML. The time zone must be the time zone of the machine where the OS or XML audit files were created. This is because the operating

system audit files are cleaned based on the audit file's Last Modification Timestamp property. The Last Modification Timestamp property value is stored in the local time zone of the machine.

- When you use an Oracle Real Application Clusters (Oracle RAC) database, Oracle recommends that you use the Network Time Protocol (NTP) to synchronize individual Oracle RAC nodes.
- The `SET_LAST_ARCHIVE_TIMESTAMP` procedure and `CLEAN_AUDIT_TRAIL` procedure should be not be executed in the same transaction block. Otherwise, the results are usually unpredictable.
- If this timestamp set to a future date, an error is returned.
- Make sure that the transaction for `SET_LAST_ARCHIVE_TIMESTAMP` is completed before you call the `CLEAN_AUDIT_TRAIL` procedure. To have predictable results in terms of cleaning records until the time provided in `SET_LAST_ARCHIVE_TIMESTAMP`(UTC format), make sure that `SET_LAST_ARCHIVE_TIMESTAMP` and `CLEAN_AUDIT_TRAIL` procedures are not called in the same transaction. It is recommended to `commit` before you call the `CLEAN_AUDIT_TRAIL` procedure.
- When you invoke `SET_LAST_ARCHIVE_TIMESTAMP` using database links, be cognizant that there could be time difference between the source and the target databases.
- When you set the last archive timestamp on a read-write database, the `DB_UNIQUE_NAME` column of the `DBA_AUDIT_MGMT_LAST_ARCH_TS` data dictionary view is also populated. You cannot configure the `DB_UNIQUE_NAME` value, so you cannot set it by using the `SET_LAST_ARCHIVE_TIMESTAMP` procedure.
- For every configuration of the last archive timestamp, Oracle Database cleans entries from the `DBA_AUDIT_MGMT_LAST_ARCH_TS` data dictionary view where the `DB_UNIQUE_NAME` column value is not equal to the current `DB_UNIQUE_NAME` value. This avoids having to reuse the old last archive timestamp when multiple switch-over operations take place and keeps in check the growing rows of the `DBA_AUDIT_MGMT_LAST_ARCH_TS` view.

Examples

The following example calls the `SET_LAST_ARCHIVE_TIMESTAMP` procedure to set the last archive timestamp for the operating system (OS) audit trail type on Oracle RAC instance 1. It uses the `TO_TIMESTAMP` function to convert a character string into a timestamp value.

A subsequent call to the CLEAN_AUDIT_TRAIL Procedure, with `use_last_arch_timestamp` set to `TRUE`, will delete all those OS audit files from the current `AUDIT_FILE_DEST` directory that were modified before 10-Sep-2012 14:10:10.0.

```
BEGIN
DBMS_AUDIT_MGMT.SET_LAST_ARCHIVE_TIMESTAMP(
   audit_trail_type     => DBMS_AUDIT_MGMT.AUDIT_TRAIL_OS,
   last_archive_time    => TO_TIMESTAMP('12-SEP-0714:10:10.0','DD-MON-RRHH24:MI:SS.FF'),
   rac_instance_number  => 1);
END;
```

31.7.18 SET_PURGE_JOB_INTERVAL Procedure

This procedure sets the interval at which the CLEAN_AUDIT_TRAIL Procedure is called for the purge job specified.

The purge job must have already been created using the CREATE_PURGE_JOB Procedure.

Syntax

```
DBMS_AUDIT_MGMT.SET_PURGE_JOB_INTERVAL(
   audit_trail_purge_name      IN VARCHAR2,
   audit_trail_interval_value  IN PLS_INTEGER) ;
```

Parameters

Table 31-24 SET_PURGE_JOB_INTERVAL Procedure Parameters

Parameter	Description
audit_trail_purge_name	The name of the purge job for which the interval is being set. This is the purge job name that you specified with the CREATE_PURGE_JOB Procedure.
audit_trail_interval_value	The interval, in hours, at which the clean up procedure should be called. This value modifies the audit_trail_purge_interval parameter set using the CREATE_PURGE_JOB Procedure

Usage Notes

Use this procedure to modify the audit_trail_purge_interval parameter set using the CREATE_PURGE_JOB Procedure.

Examples

The following example calls the SET_PURGE_JOB_INTERVAL procedure to change the frequency at which the purge job called CLEANUP is invoked. The new interval is set to 24 hours.

```
BEGIN
DBMS_AUDIT_MGMT.SET_PURGE_JOB_INTERVAL(
  AUDIT_TRAIL_PURGE_NAME       => 'CLEANUP',
  AUDIT_TRAIL_INTERVAL_VALUE   => 24 );
END;
```

Related Topics

- CLEAN_AUDIT_TRAIL Procedure
  This procedure deletes audit trail records.

31.7.19 SET_PURGE_JOB_STATUS Procedure

This procedure enables or disables the specified purge job.

The purge job must have already been created using the CREATE_PURGE_JOB Procedure.

Syntax

```
DBMS_AUDIT_MGMT.SET_PURGE_JOB_STATUS(
   audit_trail_purge_name     IN VARCHAR2,
   audit_trail_status_value   IN PLS_INTEGER) ;
```

Parameters

Table 31-25 SET_PURGE_JOB_STATUS Procedure Parameters

Parameter	Description
audit_trail_purge_name	The name of the purge job for which the status is being set. This is the purge job name that you specified with the CREATE_PURGE_JOB Procedure.
audit_trail_status_value	One of the values specified in Table 31-3. The value PURGE_JOB_ENABLE enables the specified purge job. The value PURGE_JOB_DISABLE disables the specified purge job.

Examples

The following example calls the SET_PURGE_JOB_STATUS procedure to enable the CLEANUP purge job.

```
BEGIN
DBMS_AUDIT_MGMT.SET_PURGE_JOB_STATUS(
  audit_trail_purge_name      => 'CLEANUP',
  audit_trail_status_value    => DBMS_AUDIT_MGMT.PURGE_JOB_ENABLE);
END;
```

31.7.20 TRANSFER_UNIFIED_AUDIT_RECORDS Procedure

This procedure transfers unified audit records that were in a pre-upgraded Oracle database to an internal relational table that is designed to improve read performance.

In the pre-upgraded Oracle database, these records resided in the common logging infrastructure (CLI) SGA back-end tables.

Syntax

```
DBMS_AUDIT_MGMT.TRANSFER_UNIFIED_AUDIT_RECORDS (
   container_guid       IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 31-26 TRANSFER_UNIFIED_AUDIT_RECORDS Procedure Parameters

Parameter	Description
container_guid	The GUID of the container of the associated CLI back-end table. This back-end table contains the audit records from the pre-upgraded Oracle database. If you omit this setting, then the GUID of the current container is used.

Usage Notes

- It is not mandatory to run DBMS_AUDIT_MGMT.TRANSFER_UNIFIED_AUDIT_RECORDS after an upgrade, but for better read performance of the unified audit trail, Oracle highly recommends that you run this procedure.
- The DBMS_AUDIT_MGMT.TRANSFER_UNIFIED_AUDIT_RECORDS is designed to be a one-time operation, to be performed after you upgrade from Oracle Database 12*c* release 12.1.

- There is one CLI back-end table per GUID of the container. You can find the GUIDs for containers by querying the `PDB_GUID` column of the `DBA_PDB_HISTORY` data dictionary view. Execute the `DBMS_AUDIT_MGMT.TRANSFER_UNIFIED_AUDIT_RECORDS` procedure by passing each of these GUIDs one by one to ensure that you move the unified audit records from these CLI back-end tables to the `AUDSYS.AUD$UNIFIED` table.
- In a multitenant environment, you must run the `DBMS_AUDIT_MGMT.TRANSFER_UNIFIED_AUDIT_RECORDS` procedure only in the container to which the transfer operation applies, whether it is the root or an individual PDB. You cannot run this procedure in the root, for example, to transfer audit records in a PDB.
- If you have a high rate of audit record generation and your database supports partitioning, then you may want to use the `DBMS_AUDIT_MGMT.ALTER_PARTITION_INTERVAL` procedure to alter the partition interval setting for the internal relational table. See ALTER_PARTITION_INTERVAL Procedure for more information.

32
DBMS_AUDIT_UTIL

The DBMS_AUDIT_UTIL package provides functions that enable you to format the output of queries to the DBA_FGA_AUDIT_TRAIL, DBA_AUDIT_TRAIL, UNIFIED_AUDIT_TRAIL, and V$XML_AUDIT_TRAIL views.

This chapter contains the following topics:

- DBMS_AUDIT_UTIL Overview
- DBMS_AUDIT_UTIL Security Model
- DBMS_AUDIT_UTIL Views
- Summary of DBMS_AUDIT_UTIL Subprograms

32.1 DBMS_AUDIT_UTIL Overview

The functions in the DBMS_AUDIT_UTIL package enable you to format the output of queries to the RLS_INFO column of several audit trail views so that the output appear in separate rows.

These functions use a cursor to find and format each row of the corresponding view. To use the functions in this package, include the function in a query to one of the following views:

- DBA_FGA_AUDIT_TRAIL data dictionary view, for the DECODE_RLS_INFO_ATRAIL_FGA function
- DBA_AUDIT_TRAIL data dictionary view, for the DECODE_RLS_INFO_ATRAIL_STD function
- UNIFIED_AUDIT_TRAIL data dictionary view, for the DECODE_RLS_INFO_ATRAIL_UNI function
- V$XML_AUDIT_TRAIL dynamic view, for the DECODE_RLS_INFO_ATRAIL_XML function

32.2 DBMS_AUDIT_UTIL Security Model

All DBMS_AUDIT_UTIL subprograms require the user to have EXECUTE privilege on the DBMS_AUDIT_UTIL package.

The SYSDBA administrative privilege and AUDIT_ADMIN and AUDIT_VIEWER roles have the EXECUTE privilege on the DBMS_AUDIT_UTIL package by default. An auditor can view audit data after being granted the AUDIT_VIEWER role.

Oracle strongly recommends that only audit administrators have the EXECUTE privilege on the DBMS_AUDIT_UTIL package and be granted the AUDIT_VIEWER role.

32.3 DBMS_AUDIT_UTIL Views

The views in this section display the audit information used by the DBMS_AUDIT_UTIL package function.

Table 32-1 displays the DBMS_AUDIT_UTIL views.

Table 32-1 Views Used by DBMS_AUDIT_UTIL

View	Description
DBA_FGA_AUDIT_TRAIL	Displays fine-grained audit record information; used with the DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_FGA function
DBA_AUDIT_TRAIL	Displays standard audit record information; used with the DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_STD function
UNIFIED_AUDIT_TRAIL	Displays unified audit trail information; used with the DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_XML function
V$XML_AUDIT_TRAIL	Displays XML audit record information; used with the DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_UNI function

32.4 Summary of DBMS_AUDIT_UTIL Subprograms

This table lists the DBMS_AUDIT_UTIL subprograms and their descriptions.

Table 32-2 DBMS_AUDIT_UTIL Package Subprograms

Subprogram	Description
DECODE_RLS_INFO_ATRAIL_FGA Function	Reformats the output for queries to the RLS_INFO column of the DBA_FGA_AUDIT_TRAIL data dictionary view
DECODE_RLS_INFO_ATRAIL_STD Function	Reformats the output for queries to the RLS_INFO column of the DBA_AUDIT_TRAIL data dictionary view
DECODE_RLS_INFO_ATRAIL_UNI Function	Reformats the output for queries to the RLS_INFO column of the UNIFIED_AUDIT_TRAIL data dictionary view
DECODE_RLS_INFO_ATRAIL_XML Function	Reformats the output for queries to the RLS_INFO column of the V$XML_AUDIT_TRAIL dynamic view

32.4.1 DECODE_RLS_INFO_ATRAIL_FGA Function

This function reformats the output for queries to the RLS_INFO column of the DBA_FGA_AUDIT_TRAIL data dictionary view so that the output is easily readable. It is used for the concatenated Oracle Virtual Private Database predicates for multiple fine-grained audit policies in an environment that has not been enabled for unified auditing. It returns the output in separate rows.

> **See Also:**
>
> *Oracle Database Security Guide* regarding fine-grained auditing

Syntax

```
DECODE_RLS_INFO_ATRAIL_FGA(
  IN_CURSOR                 REF CURSOR              IN)
RETURN PIPELINED ROW;
```

Parameters

Except for the `IN_CURSOR` parameter, the parameters for the `DECODE_RLS_INFO_ATRAIL_FGA` function are the same as the columns in the `DBA_FGA_AUDIT_TRAIL` data dictionary view. See *Oracle Database Reference* for more information about this view.

Usage Notes

- To use this function, include it in a query to the `DBA_FGA_AUDIT_TRAIL` data dictionary view, using a cursor similar to the example shown in the following section.
- See *Oracle Database Reference* for more information about the `DBA_FGA_AUDIT_TRAIL` data dictionary view.

Example

```
SELECT DB_USER, OBJECT_NAME, SQL_TEXT
  RLS_PREDICATE, RLS_POLICY_TYPE, RLS_POLICY_OWNER, RLS_POLICY_NAME
  FROM TABLE (DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_FGA
  (CURSOR (SELECT * FROM DBA_FGA_AUDIT_TRAIL)));
```

Return Values

A piped row with decoded values of `DBA_FGA_AUDIT_TRAIL.RLS_INFO` column.

32.4.2 DECODE_RLS_INFO_ATRAIL_STD Function

This function reformats the output for queries to the `RLS_INFO` column of the `DBA_AUDIT_TRAIL` data dictionary view so that the output is easily readable. It is used for the concatenated Oracle Virtual Private Database predicates for multiple standard audit records in an environment that has not been enabled for unified auditing. It returns the output in separate rows.

> **See Also:**
>
> *Oracle Database Security Guide* regarding auditing

Syntax

```
DECODE_RLS_INFO_ATRAIL_STD(
  IN_CURSOR                 REF CURSOR              IN)
RETURN PIPELINED ROW;
```

Parameters

Except for the `IN_CURSOR` parameter, the parameters for the `DECODE_RLS_INFO_ATRAIL_STD` function are the same as the columns in the `DBA_AUDIT_TRAIL` data dictionary view. See *Oracle Database Reference* for more information about this view.

Usage Notes

- To use this function, include it in a query to the `DBA_AUDIT_TRAIL` data dictionary view, using a cursor similar to the example shown in the following section.
- See *Oracle Database Reference* for more information about the `DBA_AUDIT_TRAIL` data dictionary view.

Example

```
SELECT USERNAME, USERHOST, ACTION, OBJ_NAME, OBJ_PRIVILEGE
  RLS_PREDICATE, RLS_POLICY_TYPE, RLS_POLICY_OWNER, RLS_POLICY_NAME
  FROM TABLE (DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_STD
  (CURSOR (SELECT * FROM DBA_AUDIT_TRAIL)));
```

Return Values

A piped row with decoded values of `DBA_AUDIT_TRAIL.RLS_INFO` column

32.4.3 DECODE_RLS_INFO_ATRAIL_UNI Function

This function reformats the output for queries to the `RLS_INFO` column of the `V$XML_AUDIT_TRAIL` dynamic view so that the output is easily readable. It is used for the concatenated Oracle Virtual Private Database predicates for multiple XML audit records in an environment that has not been enabled for unified auditing. It returns the output in separate rows.

> **See Also:**
>
> *Oracle Database Security Guide* regarding unified auditing

Syntax

```
DECODE_RLS_INFO_ATRAIL_UNI(
  IN_CURSOR                 REF CURSOR              IN)
RETURN PIPELINED ROW;
```

Parameters

Except for the `IN_CURSOR` parameter, the parameters for the `DECODE_RLS_INFO_ATRAIL_UNI` function are the same as the columns in the `V$XML_AUDIT_TRAIL` data dictionary view. See *Oracle Database Reference* for more information about this view.

Usage Notes

- To use this function, include it in a query to the `V$XML_AUDIT_TRAIL` dynamic view, using a cursor similar to the example shown in the following section.
- See *Oracle Database Reference* for more information about the `V$XML_AUDIT_TRAIL` dynamic view.

Example

```
SELECT OBJECT_NAME, SQL_TEXT
  RLS_PREDICATE, RLS_POLICY_TYPE, RLS_POLICY_OWNER, RLS_POLICY_NAME
```

```
  FROM TABLE (DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_UNI
  (CURSOR (SELECT * FROM V$XML_AUDIT_TRAIL)));
```

Return Values

A piped row with decoded values of `V$XML_AUDIT_TRAIL.RLS_INFO` column

32.4.4 DECODE_RLS_INFO_ATRAIL_XML Function

This function reformats the output for queries to the `RLS_INFO` column of the `UNIFIED_AUDIT_TRAIL` data dictionary view so that the output is easily readable. It is used for the concatenated Oracle Virtual Private Database predicates for multiple audit records from unified audit policies. It returns the output in separate rows.

> **See Also:**
>
> *Oracle Database Security Guide* regarding auditing

Syntax

```
DECODE_RLS_INFO_ATRAIL_XML(
  IN_CURSOR                 REF CURSOR              IN)
RETURN PIPELINED ROW;
```

Parameters

Except for the `IN_CURSOR` parameter, the parameters for the `DECODE_RLS_INFO_ATRAIL_XML` function are the same as the columns in the `UNIFIED_AUDIT_TRAIL` data dictionary view. See *Oracle Database Reference* for more information about this view.

Usage Notes

- To use this function, include it in a query to the `UNIFIED_AUDIT_TRAIL` data dictionary view, using a cursor similar to the example shown in the following section.
- See *Oracle Database Reference* for more information about the `UNIFIED_AUDIT_TRAIL` data dictionary view.

Example

```
SELECT DBUSERNAME, ACTION_NAME, OBJECT_NAME, SQL_TEXT,
  RLS_PREDICATE, RLS_POLICY_TYPE, RLS_POLICY_OWNER, RLS_POLICY_NAME
  FROM TABLE (DBMS_AUDIT_UTIL.DECODE_RLS_INFO_ATRAIL_XML
  (CURSOR (SELECT * FROM UNIFIED_AUDIT_TRAIL)));
```

Return Values

A piped row with decoded values of `UNIFIED_AUDIT_TRAIL.RLS_INFO` column

33
DBMS_AUTO_CLUSTERING

The `DBMS_AUTO_CLUSTERING` package contains a collection of subprograms that recommend clustering and zone map schemes for improved query performance.

This chapter contains the following topics:

- DBMS_AUTO_CLUSTERING Overview
- Summary of DBMS_AUTO_CLUSTERING Subprograms

33.1 DBMS_AUTO_CLUSTERING Overview

The `DBMS_AUTO_CLUSTERING` package generates clustering recommendations, such as tables to cluster and clustering schemes. It also verifies the impact of the clustering recommendation and implements the clustering recommendation.

33.2 Summary of DBMS_AUTO_CLUSTERING Subprograms

This table lists the `DBMS_AUTO_CLUSTERING` subprograms and their descriptions.

Table 33-1 DBMS_AUTO_CLUSTERING Package Subprograms

Subprogram	Description
CONFIGURE Procedure	Configures settings related to automatic clustering.
RECOMMEND_CLUSTERING_METHOD Function	Generates a recommendation. No recommendation will be returned if the recommendation criteria are not met. The generated recommendation is stored in the `DBA_AUTO_CLUSTERING_RECOMMENDATION` view. The function will return the recommendation ID when a recommendation is generated and `NULL` when no recommendation is generated.
VERIFY_RECOMMENDATION Procedure	Verifies the performance of the clustering recommendation. It retrieves the information used for generating the recommendation based on input arguments. This procedure finds the original table and the SQL tuning set used for the recommendation task. It creates a table and applies the clustering recommendation to the table. It runs a workload against the clustered table and the original table and returns a detailed report on performance benefits. The verification result can be `ACCEPTED` or `REJECTED`.
APPLY_RECOMMENDATION Procedure	Clusters an input table using the recommendation (identified by `recommendation_id`). The value of the incremental column specifies whether to use full clustering or incremental clustering.
GET_RECOMMENDATION Function	Returns the latest recommendation ID for a given input.

Table 33-1 (Cont.) DBMS_AUTO_CLUSTERING Package Subprograms

Subprogram	Description
REPORT_ACTIVITY Function	Returns a report of the auto-clustering operations executed during a specific period in a database.
REPORT_LAST_ACTIVITY Function	Returns a report of the last automatic indexing operation executed in a database.

33.2.1 CONFIGURE Procedure

This procedure configures settings related to automatic clustering and zone map.

Syntax

```
DBMS_AUTO_CLUSTERING.CONFIGURE (
     parameter_name      IN VARCHAR2,
     parameter_value     IN VARCHAR2,
     allow               IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 33-2 CONFIGURE Procedure Parameters

Parameter	Description
parameter_name	Automatic clustering configuration setting. It can have one of the following values: • AUTO_CLUSTERING_SCHEMA: Schemas to include or exclude from the automatic clustering recommendation. Its behavior is controlled by allow parameter. The automatic clustering process manages two schema lists – the *inclusion list* and the *exclusion list*. The inclusion list contains the schemas that can use cluster tables. The exclusion list contains the schemas that cannot use cluster tables. When automatic clustering is enabled for a database, both lists are initially empty, and all schemas in the database can use auto-clustering in all modules in the database. If the inclusion list contains at least one module, then only the modules listed in the inclusion list can use auto-clustering. If the inclusion list is empty and the exclusion list contains at least one module, then all the modules can use auto-clustering, except the module listed in the exclusion list. If both the lists (the inclusion list and the exclusion list) contain at least one module, then all the modules can use auto-clustering, except the modules listed in the exclusion list. • AUTO_CLUSTERING_REPORT_RETENTION: Number of days for which automatic clustering logs are retained in the database before they are deleted. As automatic clustering report is generated based on these logs, it cannot be generated for a period beyond the value specified for AUTO_CLUSTERING_REPORT_RETENTION. Default value is 31 days[1]. • AUTO_CLUSTERING_WORKLOAD_WINDOW: Maximum number of hours for which auto-clustering will investigate queries from the latest snapshots from AutoSTS to make recommendations. For example, if the parameter is 24, then auto-clustering will look at statements in latest 24 hours from the AutoSTS. Valid value is an integer value between 1 and 8760 (24*36). Default value is 24. • AUTO_CLUSTERING_WORKLOAD_MIN_TIME: Minimum database time, in seconds, for a query to be considered for auto-clustering recommendation. Queries below this threshold will not be considered for recommendations.

Table 33-2 (Cont.) CONFIGURE Procedure Parameters

Parameter	Description
	Valid value is an integer value between 0 and 3600 (1 hour). Default value is 120 (2 minutes).
parameter_value	Value for the configuration setting specified in the parameter_name. When it is set to NULL, the configuration setting is assigned the default value.
allow	This parameter is applicable only for the AUTO_CLUSTERING_SCHEMA and it can have one of the following values: • TRUE: Add the specified modules to the inclusion list. • FALSE: Add the specified modules to the exclusion list. • NULL: Remove the specified modules from the list to which it is currently added.

[1] Default value for auto partitioning is 90 days and 373 days for auto indexing.

33.2.2 RECOMMEND_CLUSTERING_METHOD Function

This function generates a recommendation. No recommendation will be returned if the recommendation criteria are not met. The generated recommendation can be retrieved from the DBA_AUTO_CLUSTERING_RECOMMENDATION view. The function will return the recommendation ID when a recommendation is generated and NULL when no recommendation is generated.

Syntax

```
DBMS_AUTO_CLUSTERING.RECOMMEND_CLUSTERING_METHOD (
     sqlset_owner       IN  VARCHAR2 DEFAULT 'SYS',
     sqlset_name        IN  VARCHAR2 DEFAULT 'SYS_AUTO_STS',
     table_owner        IN  VARCHAR2 DEFAULT  NULL,
     table_name         IN  VARCHAR2 DEFAULT  NULL,
     report_type        IN  VARCHAR2 DEFAULT 'TEXT',
     report_section     IN  VARCHAR2 DEFAULT 'SUMMARY',
     report_level       IN  VARCHAR2 DEFAULT 'TYPICAL',
     verification       IN  BOOLEAN  DEFAULT  TRUE)
RETURN RAW;
```

Parameters

Table 33-3 RECOMMEND_CLUSTERING_METHOD Procedure Parameters

Parameter	Description
sqlset_owner	Owner of the SQL Tuning set representing the workload to be considered. Default is SYS.

Table 33-3 (Cont.) RECOMMEND_CLUSTERING_METHOD Procedure Parameters

Parameter	Description
sqlset_name	Name of the SQL Tuning set representing the workload to be considered. Default is SYS_AUTO_STS.
table_owner, table_name	Name of a table for generating recommendation.
report_type	Type parameter for generating report for the recommendation activity. Default is TEXT.
report_section	Section parameter used to generate persistent report for the recommended clustering method. Default is SUMMARY.
report_level	Level parameter used to generate report for the recommended clustering method. Default is TYPICAL.
verification	Verification parameter indicates whether the generated recommendation is verified as part of RECOMMEND_CLUSTERING_METHOD. Default is TRUE, which means, the generated recommendation is verified to see how it performs based on the workload. If it is set to FALSE, only recommendations are generated, and their performance against workload and schema is not verified.

33.2.3 VERIFY_RECOMMENDATION Procedure

This procedure verifies the performance of the clustering recommendation. It retrieves the information used for generating the recommendation based on the input arguments. This procedure finds the original table and the SQL tuning set used for the recommendation task identified by the recommendation ID. It creates a table and applies the clustering recommendation to the table. It runs a workload against the clustered table and the original table and returns a detailed report on performance benefits. The verification result can be ACCEPTED or REJECTED.

Syntax

```
DBMS_AUTO_CLUSTERING.VERIFY_RECOMMENDATION (
     recommendation_id      IN  RAW,
     table_owner               IN  VARCHAR2 DEFAULT NULL,
     table_name             IN  VARCHAR2 DEFAULT NULL)
RETURN CLOB;
```

Parameters

Table 33-4 VERIFY_RECOMMENDATION Procedure Parameters

Parameter	Description
recommendation_id	Recommendation ID returned from RECOMMEND_CLUSTERING_METHOD function that generates recommendation for several tables in a given workload.
table_owner	Name of a owner to validate as a candidate for automatic clustering within the given recommendation id.
table_name	Name of a table to validate as a candidate for automatic clustering within the given recommendation id.

33.2.4 APPLY_RECOMMENDATION Procedure

This procedure clusters an input table using the recommendation (identified by recommendation_id). The value of the incremental column specifies whether to use full clustering or incremental clustering. Full clustering clusters an input table after applying the recommended clustering clause and performs an alter table move online to cluster data in the table. Incremental clustering applies the recommended clustering clause, and Automatic Data Compression (ADO) background will cluster data in the background.

Syntax

```
DBMS_AUTO_CLUSTERING.APPLY_RECOMMENDATION(
    recommendation_id IN RAW DEFAULT NULL,
    table_name        IN VARCHAR2 DEFAULT NULL,
    table_owner       IN VARCHAR2 DEFAULT NULL,
    apply_mode        IN VARCHAR2 DEFAULT 'FULL',
    zonemap_creation  IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 33-5 APPLY_RECOMMENDATION Procedure Parameters

Parameter	Description
recommendation_id	The task ID that recommends clustering for the given table. If not specified, the clustering recommendation made in the latest task will be used.
table_name	Name of the input table.
table_owner	Owner of the input table.
apply_mode	Mode to specify the clustering to be performed: • FULL: full clustering using online redefinition • INCREMENTAL: background incremental clustering

Table 33-5 (Cont.) APPLY_RECOMMENDATION Procedure Parameters

Parameter	Description
zonemap_creation	If this value is set to FALSE, it will not create a zonemap. It will only apply the clustering clause. Default is TRUE.

33.2.5 GET_RECOMMENDATION Function

This function returns the latest recommendation ID for a given input.

Syntax

```
DBMS_AUTO_CLUSTERING.GET_RECOMMENDATION (
     table_owner        IN  VARCHAR2,
     table_name         IN  VARCHAR2,
RETURN RAW;
```

Parameters

Table 33-6 GET_RECOMMENDATION Function Parameters

Parameter	Description
table_owner	Owner of the table to cluster
table_name	Name of the table to cluster

33.2.6 REPORT ACTIVITY Function

This function returns a report of the auto-clustering operations executed during a specific period in a database.

Syntax

```
DBMS_AUTO_CLUSTERING.REPORT_ACTIVITY (
     activity_start IN  TIMESTAMP WITH TIME ZONE DEFAULT
SYSTIMESTAMP-1,
     activity_end   IN  TIMESTAMP WITH TIME ZONE DEFAULT SYSTIMESTAMP,
     type           IN  VARCHAR2 DEFAULT 'TEXT',
     section        IN  VARCHAR2 DEFAULT 'ALL',
     level          IN  VARCHAR2 DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 33-7 REPORT ACTIVITY Function Parameters

Parameter	Description
activity_start	Time from when automatic clustering operations are considered for the report. If no value is specified or NULL is specified, then the report is generated for the last automatic clustering operation that was executed.
activity_end	Time until the automatic clustering operations are considered for the report. If no value is specified or NULL is specified, then the report is generated for the last automatic clustering operation that was executed.
type	Format of the report. It can have one of the following values: • TEXT • HTML • XML The default value is TEXT
section	Sections to include in the report. It can have a combination of the following values: • SUMMARY: Include only the summary details section in the report. • ALL: Include all the sections in the report. This is the default value.
level	Level of automatic clustering information to include in the report. It can have one of the following values: • BASIC: Include basic automatic clustering information in the report. • TYPICAL: Include typical automatic clustering information in the report. This is the default value. • ALL: Include all the automatic clustering information in the report.

33.2.7 REPORT_LAST_ACTIVITY Function

This function returns a report of the last automatic clustering operation executed in a database.

Syntax

```
DBMS_AUTO_CLUSTERING.REPORT_LAST_ACTIVITY (
     type             IN  VARCHAR2 DEFAULT 'TEXT',
     section          IN  VARCHAR2 DEFAULT 'ALL',
     level            IN  VARCHAR2 DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 33-8 REPORT_LAST_ACTIVITY Function Parameters

Parameter	Description
`type`	Format of the report. It can have one of the following values: • `TEXT` • `HTML` • `XML` The default value is `TEXT`.
`section`	Sections to include in the report. It can have a combination of the following values: • `SUMMARY`: Include only the summary details section in the report. • `ALL`: Include all the sections in the report. This is the default value.
`level`	Level of automatic clustering information to include in the report. It can have one of the following values: • `BASIC`: Include basic automatic clustering information in the report. • `TYPICAL`: Include typical automatic clustering information in the report. This is the default value. • `ALL`: Include all the automatic clustering information in the report.

34
DBMS_AUTO_SQLTUNE

The DBMS_AUTO_SQLTUNE package is the interface for managing the Automatic SQL Tuning task. Unlike DBMS_SQLTUNE, the DBMS_AUTO_SQLTUNE package requires the DBA role.

The chapter contains the following topics:

- DBMS_AUTO_SQLTUNE Overview
- DBMS_AUTO_SQLTUNE Security Model
- Summary of DBMS_AUTO_SQLTUNE Subprograms

34.1 DBMS_AUTO_SQLTUNE Overview

The DBMS_AUTO_SQLTUNE package is the interface to SQL Tuning Advisor (DBMS_SQLTUNE) when run within the AutoTask framework.

The database creates the automated system task SYS_AUTO_SQL_TUNING_TASK as part of the catalog scripts. This task automatically chooses a set of high-load SQL from AWR and runs SQL Tuning Advisor on this SQL. The automated task performs the same comprehensive analysis as any other SQL Tuning task.
The automated task tests any SQL profiles it finds by executing both the old and new query plans. Automatic SQL Tuning differs from manual SQL tuning in one important way. If automatic implementation of SQL profiles is enabled (the default is disabled), then the database implements any SQL profiles that promise a great performance benefit. The implementation occurs at tuning time so that the database can immediately benefit from the new plan. You can enable or disable automatic implementation by using the SET_AUTO_TUNING_TASK_PARAMETER API to set the ACCEPT_SQL_PROFILES parameter.

In each maintenance window, the automated tuning task stores its results as a new execution. Each execution result has the same task name but a different execution name. Query the DBA_ADVISOR_EXECUTIONS view for information about task executions. To view reports that span multiple executions, use the REPORT_AUTO_TUNING_TASK Function.

34.2 DBMS_AUTO_SQLTUNE Security Model

This package is available to users with the DBA role. For other users, you must grant the EXECUTE privilege on the package explicitly. Note that the EXECUTE_AUTO_TUNING_TASK procedure is an exception: only SYS can invoke it.

Users can call APIs in this package to control how the automatic tuning task behaves when it runs, such as enabling automatic SQL profile creation and configuring the total and per-SQL time limits under which the task runs. Because these settings affect the overall performance of the database, it may not be appropriate for all users with the ADVISOR privilege to have access to this package.

34.3 Summary of DBMS_AUTO_SQLTUNE Subprograms

The DBMS_AUTO_SQLTUNE package contains EXECUTE, REPORT, and SET subprograms.

Table 34-1 DBMS_AUTO_SQLTUNE Package Subprograms

Subprogram	Description
EXECUTE_AUTO_TUNING_TASK Function and Procedure	Executes the Automatic SQL Tuning task immediately (`SYS` only)
REPORT_AUTO_TUNING_TASK Function	Displays a text report of the automatic tuning task's history
SET_AUTO_TUNING_TASK_PARAMETER Procedures	Changes a task parameter value for the daily automatic runs

34.3.1 EXECUTE_AUTO_TUNING_TASK Function and Procedure

This function and procedure executes the Automatic SQL Tuning task (`SYS_AUTO_SQL_TUNING_TASK`).

Both the function and the procedure run in the context of a new task execution. The difference is that the function returns the name of the new execution.

Syntax

```
DBMS_AUTO_SQLTUNE.EXECUTE_AUTO_TUNING_TASK(
   execution_name    IN VARCHAR2               := NULL,
   execution_params  IN DBMS_ADVISOR.argList   := NULL,
   execution_desc    IN VARCHAR2               := NULL)
 RETURN VARCHAR2;

DBMS_AUTO_SQLTUNE.EXECUTE_AUTO_TUNING_TASK(
   execution_name    IN VARCHAR2               := NULL,
   execution_params  IN DBMS_ADVISOR.argList   := NULL,
   execution_desc    IN VARCHAR2               := NULL);
```

Parameters

Table 34-2 EXECUTE_AUTO_TUNING_TASK Function and Procedure Parameters

Parameter	Description
`execution_name`	A name to qualify and identify an execution. If not specified, it is generated by the advisor and returned by function.
`execution_params`	List of parameters (name, value) for the specified execution. The execution parameters have effect only on the execution for which they are specified. They override the values for the parameters stored in the task (set through the SET_AUTO_TUNING_TASK_PARAMETER Procedures).
`execution_desc`	A 256-length string describing the execution

Usage Notes

Only `SYS` can invoke this subprogram. A tuning task can be executed multiple times without having to reset it.

Examples

```
EXEC DBMS_AUTO_SQLTUNE.EXECUTE_AUTO_TUNING_TASK('SYS_AUTO_SQL_TUNING_TASK');
```

34.3.2 REPORT_AUTO_TUNING_TASK Function

This procedure displays the results of an Automatic SQL Tuning task.

Syntax

```
DBMS_AUTO_SQLTUNE.REPORT_AUTO_TUNING_TASK(
   begin_exec       IN   VARCHAR2   := NULL,
   end_exec         IN   VARCHAR2   := NULL,
   type             IN   VARCHAR2   := 'TEXT',
   level            IN   VARCHAR2   := 'TYPICAL',
   section          IN   VARCHAR2   := ALL,
   object_id        IN   NUMBER     := NULL,
   result_limit     IN   NUMBER     := NULL)
RETURN CLOB;
```

Parameters

Table 34-3 REPORT_AUTO_TUNING_TASK Function Parameters

Parameter	Description
`begin_exec`	Name of the beginning task execution to use. If `NULL`, the report is generated for the most recent task execution.
`end_exec`	Name of the ending task execution to use. If `NULL`, the report is generated for the most recent task execution.
`type`	Type of the report to produce. Possible values are `TEXT` which produces a text report.
`level`	Level of detail in the report: • `BASIC`: simple version of the report. Just show info about the actions taken by the advisor. • `TYPICAL`: show information about every statement analyzed, including requests not implemented. • `ALL`: highly detailed report level, also provides annotations about statements skipped over.
`section`	Section of the report to include: • `SUMMARY`: summary information • `FINDINGS`: tuning findings • `PLAN`: explain plans • `INFORMATION`: general information • `ERROR`: statements with errors • `ALL`: all sections
`object_id`	Advisor framework object id that represents a single statement to restrict reporting to. `NULL` for all statements. Only valid for reports that target a single execution.
`result_limit`	Maximum number of SQL statements to show in the report

Return Values

A `CLOB` containing the desired report.

Examples

```
-- Get the whole report for the most recent execution
SELECT DBMS_AUTO_SQLTUNE.REPORT_AUTO_TUNING_TASK
FROM   DUAL;

-- Show the summary for a range of executions
SELECT DBMS_AUTO_SQLTUNE.REPORT_AUTO_TUNING_TASK(:begin_exec, :end_exec,
'TEXT',
       'TYPICAL', 'SUMMARY')
FROM   DUAL;

-- Show the findings for the statement of interest
SELECT DBMS_AUTO_SQLTUNE.REPORT_AUTO_TUNING_TASK(:exec, :exec, 'TEXT',
       'TYPICAL', 'FINDINGS', 5)
FROM   DUAL;
```

34.3.3 SET_AUTO_TUNING_TASK_PARAMETER Procedures

This procedure updates the value of a SQL tuning parameter of type `VARCHAR2` or `NUMBER` for `SYS_AUTO_SQL_TUNING_TASK`.

Syntax

```
DBMS_AUTO_SQLTUNE.SET_AUTO_TUNING_TASK_PARAMETER(
   parameter    IN  VARCHAR2,
   value        IN  VARCHAR2);

DBMS_AUTO_SQLTUNE.SET_AUTO_TUNING_TASK_PARAMETER(
   parameter    IN  VARCHAR2,
   value        IN  NUMBER);
```

Parameters

Table 34-4 SET_AUTO_TUNING_TASK_PARAMETER Procedure Parameters

Parameter	Description
parameter	Name of the parameter to set. The possible tuning parameters that can be set by this procedure using the parameter in the form VARCHAR2: • MODE: tuning scope (comprehensive, limited) • USERNAME: user name under which the statement is parsed • DAYS_TO_EXPIRE: number of days until the task is deleted • EXECUTION_DAYS_TO_EXPIRE: number of days until the task execution is deleted • DEFAULT_EXECUTION_TYPE: default execution when none is specified by the EXECUTE_AUTO_TUNING_TASK Function and Procedure • TIME_LIMIT: global timeout in seconds • LOCAL_TIME_LIMIT: per-statement timeout in seconds • TEST_EXECUTE: – FULL - test-execute for as much time as necessary, up to the local time limit for the SQL (or the global task time limit if no SQL time limit is set) – AUTO - test-execute for an automatically chosen time proportional to the tuning time – OFF - do not test-execute • BASIC_FILTER: basic filter for SQL tuning set • OBJECT_FILTER: object filter for SQL tuning set • PLAN_FILTER: plan filter for SQL tuning set (see SELECT_SQLSET for possible values) • RANK_MEASURE1: first ranking measure for SQL tuning set • RANK_MEASURE2: second ranking measure for SQL tuning set • RANK_MEASURE3: third ranking measure for SQL tuning set • RESUME_FILTER: extra filter for SQL tuning sets besides BASIC_FILTER • SQL_LIMIT: maximum number of SQL statements to tune • SQL_PERCENTAGE: percentage filter of SQL tuning set statements The following parameters are supported for the automatic tuning task only: • ACCEPT_SQL_PROFILES: whether the task should accept SQL profiles automatically (TRUE or FALSE) • MAX_AUTO_SQL_PROFILES: maximum number of automatic SQL profiles allowed on the system, in sum • MAX_SQL_PROFILES_PER_EXEC: maximum number of SQL profiles that can be automatically implemented per execution of the task.
value	New value of the specified parameter

35
DBMS_AUTO_INDEX

The DBMS_AUTO_INDEX package provides the interface for managing auto indexes in an Oracle database.

This chapter contains the following topics:

- DBMS_AUTO_INDEX Overview
- Summary of DBMS_AUTO_INDEX Subprograms

35.1 DBMS_AUTO_INDEX Overview

The DBMS_AUTO_INDEX package is the interface for configuring auto indexes and generating reports of auto indexing operations in an Oracle database.

35.2 Summary of DBMS_AUTO_INDEX Subprograms

This table lists the DBMS_AUTO_INDEX package subprograms and briefly describes them.

Table 35-1 DBMS_AUTO_INDEX Package Subprograms

Procedure	Description
CONFIGURE Procedure	Configures settings related to automatic indexing.
DROP_AUTO_INDEXES Procedure	This procedure can be used to manually drop the automatically created indexes that overrides the retention parameter setting.
DROP_SECONDARY_INDEXES Procedure	Deletes all the indexes, except the ones used for constraints, from a schema or a table.
REPORT_ACTIVITY Function	Returns a report of the automatic indexing operations executed during a specific period in a database.
REPORT_LAST_ACTIVITY Function	Returns a report of the last automatic indexing operation executed in a database.
RECOMMEND Function	Run auto index on demand. This function analyzes all statements within a workload, creates invisible auto indexes and evaluates the performance of the statements in the workload with and without the candidate indexes.

35.2.1 CONFIGURE Procedure

This procedure configures settings related to automatic indexing.

Syntax

```
DBMS_AUTO_INDEX.CONFIGURE (
   parameter_name         IN VARCHAR2,
```

```
parameter_value         IN VARCHAR2,
allow                   IN BOOLEAN  DEFAULT TRUE);
```

Parameters

Table 35-2 CONFIGURE Procedure Parameters

Parameter	Description
parameter_name	Automatic indexing configuration setting. It can have one of the following values: • AUTO_INDEX_MODE: Modes of operation of auto indexes. It can have one of the following values: – IMPLEMENT: In this mode, new auto indexes are created as *visible* indexes and any existing *invisible* auto indexes are also set to *visible* indexes. In this mode, auto indexes are available to be used in SQL statements. – REPORT ONLY: In this mode, new auto indexes are created as *invisible* indexes and are not available to be used in SQL statements. – OFF: Setting the mode to OFF prevents new auto indexes from being considered and created. However, it does not disable existing auto indexes. • AUTO_INDEX_SCHEMA: Schemas to include or exclude from using auto indexes. Its value is case-sensitive and can include wildcards. Its behavior is controlled by the allow parameter. The automatic indexing process manages two schema lists – the *inclusion list* and the *exclusion list*. The inclusion list contains the schemas that can use auto indexes. The exclusion list contains the schemas that cannot use auto indexes. Initially, both these lists are empty and all the schemas in the database can use auto indexes when automatic indexing is enabled for a database. If the inclusion list contains at least one schema, then only the schemas listed in the inclusion list can use auto indexes. If the inclusion list is empty and the exclusion list contains at least one schema, then all the schemas can use auto indexes, except the schemas listed in the exclusion list. If both the lists (the inclusion list and the exclusion list) contain at least one schema, then all the schemas can use auto indexes, except the schemas listed in the exclusion list. • AUTO_INDEX_RETENTION_FOR_AUTO: Number of days for which the unused auto indexes are retained in the database, after which they are deleted. Default value is 373 days. • AUTO_INDEX_RETENTION_FOR_MANUAL: Number of days for which the unused manually created indexes (non-auto indexes) are retained in the database, after which they are

Table 35-2 (Cont.) CONFIGURE Procedure Parameters

Parameter	Description
	deleted. When it is set to `NULL`, the manually created indexes are not deleted by the automatic indexing process. Default value is `NULL`. • `AUTO_INDEX_REPORT_RETENTION`: Number of days for which automatic indexing logs are retained in the database before they are deleted. As automatic indexing report is generated based on these logs, automatic indexing report cannot be generated for a period beyond the value specified for `AUTO_INDEX_REPORT_RETENTION`. Default value is 373 days. • `AUTO_INDEX_DEFAULT_TABLESPACE`: Tablespace to use to store auto indexes. Default is `NULL`, which means the default permanent tablespace specified during the database creation is used to store auto indexes. Note that you cannot specify an Oracle-owned tablespace (such as `SYSAUX`) as the default tablespace. • `AUTO_INDEX_SPACE_BUDGET`: Percentage of tablespace size to use for auto indexes. This configuration setting can be used only when the default tablespace specified during the database creation is used for storing auto indexes. • `AUTO_INDEX_COMPRESSION`: Values to enable and disable advanced index compression for auto indexes. The supported values are: – `ON`: to enable advanced index compression for auto indexes – `OFF`: to disable advanced index compression for auto indexes The default value is `OFF`. • `AUTO_INDEX_TABLE`: You can use the `AUTO_INDEX_TABLE` configuration setting to specify tables that can use auto indexes. When you enable automatic indexing for a schema, all the tables in that schema can use auto indexes. However, if there is a conflict between the schema level and table level setting, the table level setting takes precedence. The parameter value is `<schema_name>.<table_name>`. You can then specify `TRUE` or `FALSE` to enable or disable auto indexes on the table respectively. To remove all tables from inclusion/exclusion list run the statement: `EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_TABLE', NULL);`

Table 35-2 (Cont.) CONFIGURE Procedure Parameters

Parameter	Description
parameter_value	Value for the configuration setting specified in parameter_name. When it is set to NULL, the configuration setting is assigned the default value.
allow	This parameter is applicable only for the AUTO_INDEX_SCHEMA configuration setting and it can have one of the following values: • TRUE: Add the specified schema to the inclusion list. • FALSE: Add the specified schema to the exclusion list. • NULL: Remove the specified schema from the list to which it is currently added. Refer to the description of the AUTO_INDEX_SCHEMA configuration setting for more information about the inclusion list and the exclusion list.

Examples

These examples are based on the assumption that the inclusion list and the exclusion list are initially empty.

The following example adds the SH and HR schemas to the exclusion list, so that only the SH and HR schemas cannot use auto indexes.

```
begin
    dbms_auto_index.configure(
         parameter_name  => 'AUTO_INDEX_SCHEMA',
         parameter_value => 'SH',
         allow           => FALSE);

    dbms_auto_index.configure(
         parameter_name  => 'AUTO_INDEX_SCHEMA',
         parameter_value => 'HR',
         allow           => FALSE);
end;
```

The following example removes the HR schema from the exclusion list, so that it can also use auto indexes. Now, only the SH schema cannot use auto indexes, because it is the only schema added to the exclusion list.

```
begin
    dbms_auto_index.configure(
         parameter_name  => 'AUTO_INDEX_SCHEMA',
         parameter_value => 'HR',
         allow           => NULL);
end;
```

The following example removes all the schemas from the exclusion list, so that all the schemas can use auto indexes.

```
begin
    dbms_auto_index.configure(
```

```
        parameter_name  => 'AUTO_INDEX_SCHEMA',
        parameter_value => NULL,
        allow           => TRUE);
end;
```

The following example adds the `HR` schema to the inclusion list, so that only the `HR` schema can use auto indexes.

```
begin
    dbms_auto_index.configure(
        parameter_name  => 'AUTO_INDEX_SCHEMA',
        parameter_value => 'HR',
        allow           => TRUE);
end;
```

The following example sets the retention period for auto indexes to 90 days.

```
begin
    dbms_auto_index.configure(
        parameter_name  => 'AUTO_INDEX_RETENTION_FOR_AUTO',
        parameter_value => '90');
end;
```

The following example sets the retention period for auto indexes to the default value of 373 days.

```
begin
    dbms_auto_index.configure(
        parameter_name  => 'AUTO_INDEX_RETENTION_FOR_AUTO',
        parameter_value => NULL);
end;
```

The following example enables a table:

```
EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_TABLE', 'SH.SALES', TRUE);
```

To remove all tables from inclusion/exclusion list:

```
EXEC DBMS_AUTO_INDEX.CONFIGURE('AUTO_INDEX_TABLE', NULL);
```

35.2.2 DROP_AUTO_INDEXES Procedure

This procedure can be used to manually drop the automatically created indexes that overrides the retention parameter setting.

Syntax

```
DBMS_AUTO_INDEX.DROP_AUTO_INDEXES (
   owner           IN  VARCHAR2 DEFAULT NULL,
   index_name      IN  VARCHAR2 DEFAULT NULL,
   allow_recreate  IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 35-3 DROP_AUTO_INDEXES Procedure Parameters

Parameter	Description
`owner`	The name of the index owner.

Table 35-3 (Cont.) DROP_AUTO_INDEXES Procedure Parameters

Parameter	Description
index_name	The name of the index.
allow_recreate	Set this parameter to allow or disallow automatic creation of the dropped index again.

Examples

Drop a single index and allow recreate:

```
exec dbms_auto_index.drop_auto_indexes('SH','"SYS_AI_612ud3j5ngf0c"',TRUE);
```

Drop all indexes owned by SH and allow recreate:

```
exec dbms_auto_index.drop_auto_indexes('SH',NULL,TRUE);
```

Drop all indexes owned by HR, disallowing recreate and then change the recreation status back to allow:

```
exec dbms_auto_index.drop_auto_indexes('HR',NULL);
exec dbms_auto_index.drop_auto_indexes('HR', NULL, TRUE);
```

Usage Notes

- If the values of the parameters owner and index_name are explicitly set to `NULL`, all auto indexes which the user has privileges on will be dropped, except the ones used for constraints.
- If owner is explicitly specified and `index_name` is set to `NULL`, all auto indexes within the given schema will be dropped, except the ones used for constraints. The dropped indexes are not recreated automatically by the system by default. Set `allow_recreate` parameter to `TRUE` to change this behavior.
- This procedure updates the `allow_recreate` status associated with the dropped indexes from `FALSE` to `TRUE` and vice-versa.

35.2.3 DROP_SECONDARY_INDEXES Procedure

This procedure deletes all the indexes, except the ones used for constraints, from a schema or a table.

Syntax

```
DBMS_AUTO_INDEX.DROP_SECONDARY_INDEXES (
   ownname     IN  VARCHAR2 DEFAULT NULL,
   tabname     IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 35-4 DROP_SECONDARY_INDEXES Procedure Parameters

Parameter	Description
ownname	(Optional) Name of the schema from which all the indexes need to be deleted. **Note:** The indexes used for constraints are not deleted.
tabname	(Optional) Name of the table from which all the indexes need to be deleted. **Note:** The indexes used for constraints are not deleted.

Examples

The following example deletes all the indexes, except the ones used for constraints, from the SH schema.

```
begin
    dbms_auto_index.drop_secondary_indexes('SH');
end;
```

The following example deletes all the indexes, except the ones used for constraints, from the EMP table in the HR schema.

```
begin
    dbms_auto_index.drop_secondary_indexes('HR', 'EMP');
end;
```

The following example deletes all the indexes, except the ones used for constraints, for which the user has the delete privileges from all the schemas in a database.

```
begin
    dbms_auto_index.drop_secondary_indexes;
end;
```

35.2.4 REPORT_ACTIVITY Function

This function returns a report of the automatic indexing operations executed during a specific period in a database.

Syntax

```
DBMS_AUTO_INDEX.REPORT_ACTIVITY (
   activity_start  IN  TIMESTAMP WITH TIME ZONE DEFAULT SYSTIMESTAMP - 1,
   activity_end    IN  TIMESTAMP WITH TIME ZONE DEFAULT SYSTIMESTAMP,
   type            IN  VARCHAR2 DEFAULT 'TEXT',
   section         IN  VARCHAR2 DEFAULT 'ALL',
   level           IN  VARCHAR2 DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 35-5 REPORT_ACTIVITY Function Parameters

Parameter	Description
activity_start	Time starting from which the executed automatic indexing operations are considered for the report. If NULL is specified, the last executed automatic indexing operation is considered for the report. If no value is specified for this parameter, then the current time minus one day (24 hours) is considered at the start time.
activity_end	Time till which the executed automatic indexing operations are considered for the report. If no value is specified, then the current time is considered as the end time.
type	Format of the report. It can have one of the following values: • TEXT • HTML • XML The default value is TEXT.
section	Sections to include in the report. It can have a combination of the following values: • SUMMARY: Include only the summary details section in the report. • INDEX_DETAILS: Include only the auto index details section in the report. • VERIFICATION_DETAILS: Include only the auto index verification details section in the report. • ERRORS: Include only the error details section in the report. • ALL: Include all the sections (summary details, auto index details, auto index verification details, and error details) in the report. This is the default value. A combination of these values can be specified using the + or – operators as shown in the following examples: • SUMMARY +INDEX_DETAILS +ERRORS: Include summary details, auto index details, and error details sections in the report. • ALL -ERRORS: Include all the sections in the report, except the error details section.

Table 35-5 (Cont.) REPORT_ACTIVITY Function Parameters

Parameter	Description
level	Level of automatic indexing information to include in the report. It can have one of the following values: • BASIC: Include basic automatic indexing information in the report. • TYPICAL: Include typical automatic indexing information in the report. This is the default value. • ALL: Include all the automatic indexing information in the report.

Return Value

A report of the automatic indexing operations executed during the specified period in a database.

Examples

The following example generates a typical report of the automatic indexing operations executed in the last 24 hours. The report is generated in the text format and contains all the sections (summary details, auto index details, auto index verification details, and error details).

```
declare
  report clob := null;
begin
  report := dbms_auto_index.report_activity();
end;
```

35.2.5 REPORT_LAST_ACTIVITY Function

This function returns a report of the last automatic indexing operation executed in a database.

Syntax

```
DBMS_AUTO_INDEX.REPORT_LAST_ACTIVITY (
   type            IN  VARCHAR2 DEFAULT 'TEXT',
   section         IN  VARCHAR2 DEFAULT 'ALL',
   level           IN  VARCHAR2 DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 35-6 REPORT_LAST_ACTIVITY Function Parameters

Parameter	Description
`type`	Format of the report. It can have one of the following values: • `TEXT` • `HTML` • `XML` The default value is `TEXT`.
`section`	Sections to include in the report. It can have a combination of the following values: • `SUMMARY`: Include only the summary details section in the report. • `INDEX_DETAILS`: Include only the auto index details section in the report. • `VERIFICATION_DETAILS`: Include only the auto index verification details section in the report. • `ERRORS`: Include only the error details section in the report. • `ALL`: Include all the sections (summary details, auto index details, auto index verification details, and error details) in the report. This is the default value. A combination of these values can be specified using the + or – operators as shown in the following examples: • `SUMMARY +INDEX_DETAILS +ERRORS`: Include summary details, auto index details, and error details sections in the report. • `ALL -ERRORS`: Include all the sections in the report, except the error details section.
`level`	Level of automatic indexing information to include in the report. It can have one of the following values: • `BASIC`: Include basic automatic indexing information in the report. • `TYPICAL`: Include typical automatic indexing information in the report. This is the default value. • `ALL`: Include all the automatic indexing information in the report.

Return Value

A report of the last automatic indexing operation executed in a database.

Examples

The following example generates a typical report of the last automatic indexing operation executed in a database. The report is generated in the text format and contains all the sections (summary details, auto index details, auto index verification details, and error details).

```
declare
  report clob := null;
begin
  report := dbms_auto_index.report_last_activity();
end;
```

35.2.6 RECOMMEND Function

Run auto indexing on demand.

This function analyzes all statements within a workload, creates invisible auto indexes and evaluates the performance of the statements in the workload both with and without the candidate indexes. If called in `IMPLEMENT` mode, this function marks as visible any indexes that improve performance. If called in `REPORT ONLY` mode, all indexes created in this task are dropped at the end of the task. In that case you can run `DBMS_AUTO_INDEX.REPORT_ACTIVITY()` to view the results.

Syntax

```
DBMS_AUTO_INDEX.RECOMMEND (
   WORKLOAD_START_TIME IN TIMESTAMP DEFAULT NULL,
   WORKLOAD_END_TIME   IN TIMESTAMP DEFAULT NULL,
   AUTO_INDEX_MODE     IN VARCHAR2 )
RETURN VARCHAR2;
```

Parameters

Table 35-7 RECOMMEND Function Parameters

Parameter	Description
`WORKLOAD_START_TIME`	The starting point for analysis within the ASTS (Automatic SQL tuning set) workload. The default NULL means all statements from the start of the workload to `WORKLOAD_END_TIME` are analyzed.
`WORKLOAD_END_TIME`	The end point for analysis within the ASTS workload. The default NULL means that all statements from `WORKLOAD_START_TIME` to the end of the workload are analyzed.
`AUTO_INDEX_MODE`	• `REPORT ONLY` Analysis is performed, but there is no change. Created indexes are dropped. See `DBMS_AUTO_INDEX.REPORT_ACTIVITY()` to view verification details. • `IMPLEMENT` Marks performance-improving indexes as visible. The default is `REPORT ONLY`.

Return Value

The execution name of the task is returned.

Example 35-1 Using the RECOMMEND Function

Execute the function and then call `DBMS_AUTO_INDEX.REPORT_LAST_ACTIVITY` to return the report.

```
var tname varchar2(100)
EXEC :tname := dbms_auto_index.recommend()
SELECT :tname task_name FROM dual;
```

Get the report:

```
set linesize 250
set trims on
set pagesize 1000
set long 10000000
column report format a120
spool report.txt
SELECT dbms_auto_index.report_last_activity('text','all','all') report FROM
dual;
spool off
```

36
DBMS_AUTO_MV

DBMS_AUTO_MV contains subprograms for configuring automatic materialized views.

This chapter contains the following topics:

- Using DBMS_AUTO_MV
- Summary of DBMS_AUTO_MV Subprograms

36.1 Using DBMS_AUTO_MV

The DBMS_AUTO_MV package contains functions and procedures to manage automatic materialized views.

This package is owned by SYS, so EXECUTE package privilege is required by non-SYS users. Grant EXECUTE privilege on the package to the users.

36.2 Summary of DBMS_AUTO_MV Subprograms

This table lists the DBMS_AUTO_MV package subprograms and briefly describes them.

Table 36-1 DBMS_AUTO_MV Package Subprograms

Subprogram	Description
CONFIGURE Procedure	Enables, disables, and configures the various parameters of the automatic materialized view feature.
DROP_AUTO_MVS Procedure	Drops automatic materialized views that were created. This routine can only be executed by DBA.
RECOMMEND Function	This function allows a user to manually generate automatic materialized view recommendations for SQL statements in a given SQL tuning set
REFRESH Procedure	This procedure allows a user to manually refresh all stale automatic materialized views in the system unconditionally.
REPORT_ACTIVITY Function	This function generates a report on the automatic materialized view activities and usage for a specified time duration. The report can be generated in text, HTML, or XML formats as specified by the argument type.
REPORT_LAST_ACTIVITY Function	This function generates a report on the most recent automatic materialized view activities and usage. The report can be generated in text, HTML, or XML formats as specified by the argument type.

36.2.1 CONFIGURE Procedure

This procedure enables, disables, and configures the various parameters of the automatic materialized view feature.

Syntax

```
DBMS_AUTO_MV.CONFIGURE (
   parameter   IN VARCHAR2,
   value       IN VARCHAR2,
   allow       IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 36-2 CONFIGURE Procedure Parameters

Parameter	Description
`parameter`	The name of the parameter to be modified. Parameter names are not case sensitive.
`value`	The value of the specified parameter.
`allow`	This parameter allows or disallows various schemas or app modules. It is applicable only for the `AUTO_MV_SCHEMA` and `AUTO_MV_APP_MODULE` parameters.

Parameters

Table 36-3 CONFIGURE Parameters Names

Parameter	Description
`AUTO_MV_MODE`	Enables, disables, or engages report-only mode. Possible values are: • `OFF`: No recommendations are made. This is the default value. • `REPORT ONLY`: This mode generates recommendations and stores them in the internal repository so they are available to the DBA through `DBA_AUTO_MV_ANALYSIS_RECOMMENDATIONS`. • `IMPLEMENT`: This mode generates, verifies, and publishes recommendations, or drops them.
`AUTO_MV_MAINT_TASK`	Activates and deactivated automatic maintenance of materialized views. • `ENABLE`: Activates automatic maintenance of materialized views. • `DISABLE`: Deactivates the automatic maintenance of materialized views. If automatic maintenance of materialized views is in progress, it finishes the maintenance. This is the default value. • `CLEANUP_AND_DISABLE`: Drops all automatic materialized views, and deactivates automatic maintenance of materialized views. If automatic materialized views maintenance is in progress, it finishes the maintenance before the task is deactivated.

Table 36-3 (Cont.) CONFIGURE Parameters Names

Parameter	Description
AUTO_MV_SPACE_BUDGET	Specifies the amount of space budget available for implementing automatic materialized views. The total space value is the sum of currently space used by all user tables (i.e. not system tables). The calculation does not include user access structures (like indexes or materialized views). Possible values are: • Budget in percent: A positive number ending with % symbol designating the percentage of currently utilized space for all user tables. • Budget in GB: A positive integer ending with GB that indicates the absolute space limit for automatic materialized views. For example, 10GB indicates 10 Gigabytes. The minimum value is 1GB but no maximum value. The default budget is 10% of the total size of user tables.
AUTO_MV_DEFAULT_TABLESPACE	Specifies the tablespace to place automatic materialized views. Possible values are: • Tablespace name: A valid Oracle tablespace name to be used when creating new automatic materialized views. Quoted identifiers are supported. • NULL: A new automatic materialized view is created in the default tablespace of the owner of parent object. If automatic materialized view has more than one parent object, such as materialized views defined on multiple base tables, the default tablespace of the owner of largest base table is selected. This is the default value. If the value is changed dynamically, it takes effect the next time automatic materialized views recommendations are implemented.
AUTO_MV_TEMP_TABLESPACE	Specifies the temporary tablespace while creating or refreshing automatic materialized views. Possible values are: • Tablespace name: A valid Oracle temp tablespace name to be used when creating new automatic materialized views and the data needs to be spilled to temp. • NULL: The temp table space assigned to the owner of the largest parent object of the automatic materialized views. This is the default value. If the value is changed dynamically, it takes effect the next time recommendations are implemented.
AUTO_MV_RETENTION	Specifies the number of days automatic materialized views exists without being utilized by a query. When the expiry period is reached, the materialized view is dropped. Positive integer: An integer between 1 and 373. The default value is 33 days.
AUTO_MV_ANALYZE_REPORT_RETENTION	Specifies the maximum number of days to retain analysis and recommendation history. Positive integer: An integer value between 0 and 90. Value 0 implies that history is not maintained. The default value is 31. The history of analysis and verification is retained in the DBA_AUTO_MV_* dictionary tables.
AUTO_MV_ANALYZE_WORKLOAD_WINDOW	Specifies the maximum number of hours to make recommendations. Positive integer: An integer value between 1 and 8760. The default value is 24.
AUTO_MV_ANALYZE_WORKLOAD_MIN_TIME	Specifies the minimum time in seconds for a query to be considered for automatic materialized views recommendation. Queries below this threshold are not considered for recommendations. Positive value: An integer value between 0 and 3600 (1 hour). The default value is 120 (2 minutes).

Table 36-3 (Cont.) CONFIGURE Parameters Names

Parameter	Description
AUTO_MV_SCHEMA	Specifies the schemas to include or exclude from creation of automatic materialized views. Possible values are: • TRUE: Adds the specified schema to the inclusion list. • FALSE: Adds the specified schema to the exclusion list. • NULL: Removes the specified schema from the list to which it is currently added. If both the lists (the inclusion list and the exclusion list) contain at least one schema, then all the schemas can use automatic materialized views, except the schemas listed in the exclusion list.
AUTO_MV_APP_MODULE	Specifies application modules to include or exclude from creation of automatic materialized views. Possible values are: • TRUE: Adds the specified application module to the inclusion list. • FALSE: Adds the specified application module to the exclusion list. • NULL: Removes the specified application module from the list to which it is currently added. Initially, the inclusion list and the exclusion list are empty and we can create automatic materialized views under all application modules when automatic automatic materialized views are enabled for a database.
AUTO_MV_VERIFY_REPORT_RETENTION	Specifies the maximum number of days to retain the verification history. Positive value: An integer value between 0 and 90 for the number of days the history of analysis and verification is retained in the DBA_AUTO_MV_* dictionary tables . Value 0 implies that history is not maintained. The default value is 31.
AUTO_MV_MAINT_REPORT_RETENTION	Specifies the maximum number of days to retain history of automatic materialized view maintenance. Positive integer: An integer value between 0 and 90 for the number of days the history of automatic materialized view refreshes is retained in the DBA_AUTO_MV_REFRESH_* dictionary tables. Value 0 means no history is maintained. The default value is 31.

Example

```
begin
   dbms_auto_mv.configure ('AUTO_MV_SPACE_BUDGET',  '50%');
end;
begin
   dbms_auto_mv.configure ('AUTO_MV_SCHEMA',  'SH', FALSE);
   dbms_auto_mv.configure ('AUTO_MV_SCHEMA',  'SCOTT');
end;
/
```

36.2.2 DROP_AUTO_MVS Procedure

This procedure drops automatic materialized views that were created. This routine can only be executed by DBA.

Syntax

```
DBMS_AUTO_MV.DROP_AUTO_MVS (
   owner            IN VARCHAR2,
   mv_name          IN VARCHAR2,
   allow_recreate   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 36-4 DROP_AUTO_MVS Procedure Parameters

Parameter	Description
OWNER	Specifies the name of the owner of the automatic materialized views. If OWNER is explicity specified and MV_NAME is set to null, all automatic materialized views which the user has privileges are dropped. If OWNER is explicity specified and MV_NAME is set to null, all automatic materialized views with the given OWNER are dropped. Dropped automatic materialized views are not recreated automatically by the system as default.
MV_NAME	The name of the automatic materialized views.
ALLOW_RECREATE	Enables or disables the automatic creation of dropped automatic materialized views. FALSE disables the automatic creation of dropped automatic materialized views. This is default. TRUE enables the automatic creation of dropped automatic materialized views.

Examples

```
begin
   dbms_auto_mv.DROP_AUTO_MVS ('SH');
end;
/

begin
   dbms_auto_mv.DROP_AUTO_MVS ('SH', 'AUTO_MV$$_G2MKPB9SA1FB7');
end;
/
```

36.2.3 RECOMMEND Function

This function allows a user to manually generate automatic materialized view recommendations for SQL statements in a given SQL tuning set

Syntax

```
DBMS_AUTO_MV.RECOMMEND (
   sts_owner              IN  VARCHAR2  DEFAULT 'SYS',
   sts_name               IN  VARCHAR2  DEFAULT 'SYS_AUTO_STS',
   workload_start_time    IN  TIMESTAMP DEFAULT NULL,
   workload_end_time      IN  TIMESTAMP DEFAULT NULL,
   automv_mode            IN  VARCHAR2  DEFAULT 'REPORT ONLY')
RETURN VARCHAR2;
```

Parameters

Table 36-5 RECOMMEND Function Parameters

Parameter	Description
sts_owner	The name of the owner of the SQL tuning set. The default value is SYS.
sts_name	The name of the SQL tuning set. The default value is SYS_AUTO_STS.
workload_start_time	The start time of the workload window. The value NULL means that the default is chosen. The default is SYSDATE minus the number of hours defined by AUTO_MV_ANALYZE_WORKLOAD_WINDOW.
workload_end_time	The end time of the workload window.
automv_mode	When AUTOMV_MODE is set to REPORT ONLY, which is the default mode, the function will only output the recommendations. If this parameter is set to IMPLEMENT, then the recommended automatic materialized views will be verified and implemented.

Return Value

The execution name (execution_name) to be used in DBA_AUTO_MV% catalog views.

Examples

In the following examples, the default SQL tuning set, SYS_AUTO_STS is used. Make sure that SYS_AUTO_STS contains the required workload to generate the automatic materialized views.

Example 1: Generate and report recommendations using SYS_AUTO_STS for the past 24 hours. Note that the default behavior of this function is REPORT ONLY, so no automatic materialized view will be implemented.

```
var exec_name varchar2(200);
begin
    :exec_name := dbms_auto_mv.recommend();
end;
/

SELECT * FROM DBA_AUTO_MV_ANALYSIS_RECOMMENDATIONS
        WHERE exec_name = :exec_name;
```

Example 2: Generate and publish recommendations using SYS_AUTO_STS for the past 24 hours.

```
var exec_name varchar2(200);
begin
    :exec_name := dbms_auto_mv.recommend(automv_mode=>'IMPLEMENT');
end;
/
```

36.2.4 REFRESH Procedure

This procedure allows a user to manually refresh all stale automatic materialized views in the system unconditionally.

Syntax

```
DBMS_AUTO_MV.REFRESH ();
```

Example

```
begin
    dbms_auto_mv.refresh();
end;
/
```

36.2.5 REPORT_ACTIVITY Function

This function generates a report on the automatic materialized view activities and usage for a specified time duration. The report can be generated in text, HTML, or XML formats as specified by the argument type.

Syntax

```
DBMS_AUTO_MV.REPORT_ACTIVITY (
   activity_start        IN  TIMESTAMP WITH TIME ZONE DEFAULT SYSTIMESTAMP -1,
   activity_end          IN  TIMESTAMP WITH TIME ZONE DEFAULT SYSTIMESTAMP,
   type                  IN  VARCHAR2  DEFAULT 'TEXT',
   section               IN  VARCHAR2  DEFAULT 'ALL',
   level                 IN  VARCHAR2  DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 36-6 REPORT_ACTIVITY Function Parameters

Parameter	Description
activity_start	The start time for report generation.
activity_end	The end time for report generation.
type	The format type in which the report needs to be generated. The possible values are: • TEXT • HTML • XML The default value is TEXT.

Table 36-6 (Cont.) REPORT_ACTIVITY Function Parameters

Parameter	Description
section	The section can be a combination of the following: • SUMMARY • MV_DETAILS • QUERY_DETAILS • VERIFICATION_DETAILS • ALL The default value is ALL. You can generate a specific combination of report by using + or – operators. For example, when section is specified as SUMMARY+MV_DETAILS, the generated report will contain only the summary and the details about the automatic materialized view.
level	The level can be either BASIC, TYPICAL, or ALL. When the level is set to BASIC, a minimum set of information regarding the most recent automatic materialized view activity is reported. On the other hand, when the level is set to ALL, a detailed report is generated. The default value is TYPICAL.

Return Value

This functions returns the report as a CLOB.

Examples

Example 1: The following call to REPORT_ACTIVITY() generates an HTML output for all the automatic materialized view activities:

```
select dbms_auto_mv.report_activity(type => 'HTML') from dual;
```

36.2.6 REPORT_LAST_ACTIVITY Function

This function generates a report on the most recent automatic materialized view activities and usage. The report can be generated in text, HTML, or XML formats as specified by the argument type.

Syntax

```
DBMS_AUTO_MV.REPORT_LAST_ACTIVITY (
   type                 IN  VARCHAR2  DEFAULT 'TEXT',
   section              IN  VARCHAR2  DEFAULT 'ALL',
   level                IN  VARCHAR2  DEFAULT 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 36-7 REPORT_LAST_ACTIVITY Function Parameters

Parameter	Description
type	The format type in which the report needs to be generated. The possible values are: • TEXT • HTML • XML The default value is TEXT.
section	The section can be a combination of the following: • SUMMARY • MV_DETAILS • QUERY_DETAILS • VERIFICATION_DETAILS • ALL The default value is ALL. You can generate a specific combination of report by using + or – operators. For example, when section is specified as SUMMARY+MV_DETAILS, the generated report will contain only the summary and the details about the automatic materialized view.
level	The level can be either BASIC, TYPICAL, or ALL. When the level is set to BASIC, a minimum set of information regarding the most recent automatic materialized view activity is reported. On the other hand, when the level is set to ALL, a detailed report is generated. The default value is TYPICAL.

Return Value

This functions returns the report as a CLOB.

Examples

Example 1: The following call to REPORT_LAST_ACTIVITY() generates a TEXT output:

```
select dbms_auto_mv.report_last_activity('TEXT', 'ALL', 'TYPICAL') from dual;
```

Example 2: The following call to REPORT_LAST_ACTIVITY() generates an XML output:

```
select dbms_auto_mv.report_last_activity('XML', 'ALL', 'TYPICAL') from dual;
```

Example 3: The following call to REPORT_LAST_ACTIVITY() generates an HTML output:

```
select dbms_auto_mv.report_last_activity('HTML', 'ALL', 'TYPICAL') from dual;
```

37
DBMS_AUTO_REPORT

The DBMS_AUTO_REPORT package provides an interface to view SQL Monitoring and Real-time Automatic Database Diagnostic Monitor (ADDM) data that has been captured into Automatic Workload Repository (AWR). It also provides subprograms to control the behavior of how these data are captured to AWR.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* for more information about reporting database operations

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_AUTO_REPORT Subprograms

37.1 DBMS_AUTO_REPORT Overview

This package provides an interface to view SQL Monitoring and Real-time ADDM data that has been captured into AWR. It also provides subprograms to control the behavior of how these data are captured to AWR.Captured data are stored in AWR and exposed via 2 views: DBA_HIST_REPORTS and DBA_HIST_REPORTS_DETAILS.

37.2 DBMS_AUTO_REPORT Security Model

This package is available to PUBLIC and performs its own security checking.

37.3 Summary of DBMS_AUTO_REPORT Subprograms

This table describes the parameters of the DBMS_AUTO_REPORT package subprograms.

Table 37-1 DBMS_AUTO_REPORT Package Subprograms

Subprogram	Description
FINISH_REPORT_CAPTURE Procedure	Ends the complete capture of SQL monitor data that was started with the START_REPORT_CAPTURE Procedure.
REPORT_REPOSITORY_DETAIL Function	Obtains the stored report for a given report ID
REPORT_REPOSITORY_DETAIL_XML Function	Obtains the stored XML report for a given report ID

Table 37-1 (Cont.) DBMS_AUTO_REPORT Package Subprograms

Subprogram	Description
REPORT_REPOSITORY_LIST_XML Function	Obtains an XML report of the list of SQL Monitor and Real-time ADDM data captured in AWR
START_REPORT_CAPTURE Procedure	Captures SQL monitor data of any newly monitored SQLs every minute since the last run of the capture cycle, and stores it in AWR.
PURGE_ALL_REPORTS Function	Purges all SQL Monitor content from WRP$_ tables associated with a single DBID (if specified), or from all non-local DBIDs if no DBID is specified.
PURGE_REPORTS_BY_RETENTION Function	Purges SQL Monitor reports older than the retention period from WRP$_ tables associated with a specified container DBID, or if no DBID is specified, then it purges old reports from all non-local containers.

37.3.1 FINISH_REPORT_CAPTURE Procedure

This procedure ends the complete capture of SQL monitor data that was started with the `START_REPORT_CAPTURE` procedure.

After calling this subprogram, capture of data continues every minute except that it is not captured for all active SQLs but only for those deemed important, namely the top 5 SQLs (by elapsed time, or elapsed time*DOP in case of PQ) whose monitoring has completed.

Syntax

```
DBMS_AUTO_REPORT.FINISH_REPORT_CAPTURE;
```

Related Topics

- START_REPORT_CAPTURE Procedure
  This procedure captures SQL monitor data of any newly monitored SQLs every minute since the last run of the capture cycle, and stores it in AWR.

37.3.2 REPORT_REPOSITORY_DETAIL Function

This procedure obtains the stored report for a given report ID in the specified format such as XML or HTML.

Syntax

```
DBMS_AUTO_REPORT.REPORT_REPOSITORY_DETAIL (
   rid              IN NUMBER    DEFAULT NULL,
   type             IN VARCHAR2  DEFAULT 'XML',
   base_path        IN VARCHAR2  DEFAULT NULL)
 RETURNS CLOB
```

Parameters

Table 37-2 REPORT_REPOSITORY_DETAIL Function Parameters

Parameter	Description
rid	ID of the stored report which returned by the function

Table 37-2 (Cont.) REPORT_REPOSITORY_DETAIL Function Parameters

Parameter	Description
type	Desired format of the report. Values can be 'XML', 'TEXT', 'HTML', 'EM' or 'ACTIVE'. The last two options generate a report in the same format called active HTML. Default value is 'XML'.
base_path	Unused/Non-operative

Return Values

The persisted report for the given record ID

37.3.3 REPORT_REPOSITORY_DETAIL_XML Function

This procedure obtains the stored XML report for a given report ID.

Syntax

```
DBMS_AUTO_REPORT.REPORT_REPOSITORY_DETAIL_XML (
   rid              IN NUMBER    DEFAULT NULL,
   base_path        IN VARCHAR2  DEFAULT NULL)
 RETURNS XMLTYPE
```

Parameters

Table 37-3 *REPORT_REPOSITORY_DETAIL_XML Function Parameters*

Parameter	Description
rid	ID of the stored report which returned by the function
base_path	Unused/Non-operative

Return Values

The persisted XML report for the given record ID

37.3.4 REPORT_REPOSITORY_LIST_XML Function

This procedure obtains an XML report of the list of SQL Monitor and Real-time ADDM data captured in AWR.

The input parameters can be used to select and restrict which captured data will be included in the list report. All parameters are optional.

Syntax

```
DBMS_AUTO_REPORT.REPORT_REPOSITORY_LIST_XML (
    active_since              IN DATE     DEFAULT NULL,
    active_upto               IN DATE     DEFAULT NULL,
    snapshot_id               IN NUMBER   DEFAULT NULL,
    dbid                      IN NUMBER   DEFAULT NULL,
    inst_id                   IN NUMBER   DEFAULT NULL,
    con_dbid                  IN NUMBER   DEFAULT NULL,
    session_id                IN NUMBER   DEFAULT NULL,
    session_serial            IN NUMBER   DEFAULT NULL,
```

```
    component_name           IN VARCHAR2 DEFAULT NULL,
    key1                     IN VARCHAR2 DEFAULT NULL,
    key2                     IN VARCHAR2 DEFAULT NULL,
    key3                     IN VARCHAR2 DEFAULT NULL,
    report_level             IN VARCHAR2 DEFAULT 'TYPICAL',
    base_path                IN VARCHAR2 DEFAULT NULL)
RETURNS XMLTYPE
```

Parameters

Table 37-4 *REPORT_REPOSITORY_LIST_XML Function Parameters*

Parameter	Description
active_since	Start of a time range used to select data. When a time range is specified, only those data are included in the list that were active during the time range. When no value is specified the time range is chosen as the last 24 hours ending at the current system time.
active_upto	Same as active_since except that it is the end of the time range
snapshot_id	If a value is specified, only those data captured during the specified snapshot ID are included in the list report. If no value is specified, no filtering is performed on snapshot ID.
dbid	If a value is specified, only those data captured for the specified database ID are included in the list report. If no value is specified, no filtering is performed on database ID
inst_id	If a value is specified, only those data captured on the specified instance number are included in the list report. If no value is specified, no filtering is performed on the instance ID.
con_dbid	If a value is specified, only those data captured on the specified container DBID are included in the list report. If no value is specified, no filtering is performed on the container DBID.
session_id	If a value is specified, only those data captured for the specified session ID are included in the list report. If no value is specified, no filtering is performed on session ID.
session_serial	If a value is specified, only those data captured for the specified session are included in the list report. If no value is specified, no filtering is performed on session serial number. This parameter should be used in conjunction with the session_id parameter.
component_name	Can be 'sqlmonitor' for SQL Monitor data or 'rtaddm' for Real-time ADDM data. If a value is specified then data pertaining only to the specified component will be included in the list report. If no value is specified, no filtering is performed.
key1	Key value relevant to a component. For SQL Monitor, key1 is the SQL ID of the captured SQL statement. If a value is specified, only those data having specified value for key1 are included, else no filtering is performed on key1.
key2	Key value relevant to a component. For SQL Monitor, key2 is the SQL execution ID of the captured SQL statement. If a value is specified, only those data having specified value for key2 are included, else no filtering is performed on key2.
key3	Key value relevant to a component. For SQL Monitor, key3 is the SQL execution start time of the captured SQL statement. If a value is specified, then only those data having specified value for key3 are included, else no filtering is performed on key3.
report_level	Currently only 'TYPICAL' is used

Table 37-4 (Cont.) *REPORT_REPOSITORY_LIST_XML Function Parameters*

Parameter	Description
base_path	Unused/Non-operative

37.3.5 START_REPORT_CAPTURE Procedure

This procedure captures SQL monitor data of any newly monitored SQLs every minute since the last run of the capture cycle, and stores it in AWR.

Every capture cycle attempts to capture data for SQLs that are not currently executing or queued. This is a complete capture since data of all newly monitored SQLs is captured. It continues to run every minute until it is explicitly ended with the FINISH_REPORT_CAPTURE Procedure. In the case of a RAC system, the capture will start on each node of the cluster.

Syntax

```
DBMS_AUTO_REPORT.START_REPORT_CAPTURE;
```

37.3.6 PURGE_ALL_REPORTS Function

This function purges all SQL Monitor content from WRP$_ tables associated with a single DBID (if specified), or from all non-local DBIDs if no DBID is specified.

Syntax

```
DBMS_AUTO_REPORT.PURGE_ALL_REPORTS (
   dbid  IN NUMBER DEFAULT NULL,
);
```

This function does not return a value.

37.3.7 PURGE_REPORTS_BY_RETENTION Function

This function purges SQL Monitor reports older than the retention period from WRP$_ tables associated with a specified container DBID, or if no DBID is specified, then it purges old reports from all non-local containers.

Only the reports that are older than the retention period are purged.

Syntax

```
DBMS_AUTO_REPORT.PURGE_REPORTS_BY_RETENTION (
   dbid  IN NUMBER DEFAULT NULL,
);
```

This function does not return a value.

38
DBMS_AUTO_TASK_ADMIN

The DBMS_AUTO_TASK_ADMIN package provides an interface to AUTOTASK functionality. It is used by the DBA as well as Enterprise Manager to access the AUTOTASK controls. Enterprise Manager also uses the AUTOTASK Advisor.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more information about "Configuring Automated Maintenance Task"

This chapter contains the following sections:

- Deprecated Subprograms
- Security Model
- Constants
- Summary of DBMS_AUTO_TASK_ADMIN Subprograms

38.1 DBMS_AUTO_TASK_ADMIN Deprecated Subprograms

The DBMS_AUTO_TASK_ADMIN OVERRIDE_PRIORITY subprogram has been deprecated.

> **Note:**
>
> Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

- OVERRIDE_PRIORITY Procedures

38.2 DBMS_AUTO_TASK_ADMIN Security Model

DBMS_AUTO_TASK_ADMIN is a definer's rights package, and EXECUTE is automatically granted to DBA, IMP_FULL_DATABASE and DATAPUMP_IMP_FULL_DATABASE.

38.3 DBMS_AUTO_TASK_ADMIN Constants

The DBMS_AUTO_TASK_ADMIN package defines several constants that can be used for specifying parameter values.

These constants shown in the following table:

Table 38-1 DBMS_AUTO_TASK_ADMIN Constants

Name	Type	Value	Description
PRIORITY_MEDIUM	VARCHAR2	'MEDIUM'	Task with this priority should be executed as time permits
PRIORITY_HIGH	VARCHAR2	'HIGH'	Task with this priority should be executed within the current Maintenance Window
PRIORITY_URGENT	VARCHAR2	'URGENT'	Task with this priority is to be executed at the earliest opportunity

38.4 Summary of DBMS_AUTO_TASK_ADMIN Subprograms

This table lists the DBMS_AUTO_TASK_ADMIN subprograms and briefly describes them.

Table 38-2 DBMS_AUTO_TASK_ADMIN Package Subprograms

Method	Description
DISABLE Procedures	Prevents AUTOTASK from executing any requests from a specified client or operation.
ENABLE Procedures	Allows a previously disabled client, operation, target type, or individual target to be enabled under AUTOTASK control
GET_CLIENT_ATTRIBUTES Procedure	Returns values of select client attributes
GET_P1_RESOURCES Procedure	Returns percent of resources allocated to each AUTOTASK High Priority Consumer Groups
OVERRIDE_PRIORITY Procedures	Manually overrides task priority.
SET_CLIENT_SERVICE Procedure	Associates an AUTOTASK Client with a specified Service
SET_P1_RESOURCES Procedure	Sets percentage-based resource allocation for each High Priority Consumer Group used by AUTOTASK Clients

38.4.1 DISABLE Procedures

This procedure prevents AUTOTASK from executing any requests from a specified client or operation.

Syntax

Disables all AUTOTASK functionality.

```
DBMS_AUTO_TASK_ADMIN.DISABLE;
```

Disables all tasks for the client or operation.

```
DBMS_AUTO_TASK_ADMIN.DISABLE (
   client_name       IN    VARCHAR2,
   operation         IN    VARCHAR2,
   window_name       IN    VARCHAR2);
```

Parameters

Table 38-3 DISABLE Procedure Parameters

Parameter	Description
client_name	Name of the client, as found in DBA_AUTOTASK_CLIENT View
operation	Name of the operation as specified in DBA_AUTOTASK_OPERATION View
window_name	Optional name of the window in which client is to be disabled

Usage Notes

- If operation and window_name are both NULL, the client is disabled.
- If operation is not NULL, window_name is ignored and the operation is disabled
- If operation is NULL and window_name is not NULL, the client is disabled in the specified window.
- Auto STS Capture Task is accepted.

38.4.2 ENABLE Procedures

This procedure allows a previously disabled client, operation, target type, or individual target to be enabled under AUTOTASK control.

Specifying the DEFERRED option postpones the effect of the call until the start of the next maintenance window. If IMMEDIATE option is specified the effect of this call is immediate – as long as there is a currently open maintenance window.

Syntax

Re-enabling AUTOTASK. This version enables the specified client. Note that any explicitly disabled tasks or operations must be re-enabled individually.

```
DBMS_AUTO_TASK_ADMIN.ENABLE;
```

Re-enabling a client or operation.Note that any explicitly disabled tasks or operations must be re-enabled individually.

```
DBMS_AUTO_TASK_ADMIN.ENABLE (
   client_name       IN    VARCHAR2,
   operation         IN    VARCHAR2,
   window_name       IN    VARCHAR2);
```

Parameters

Table 38-4 ENABLE Procedure Parameters

Parameter	Description
client_name	Name of the client, as found in DBA_AUTOTASK_CLIENT View
operation	Name of the operation as specified in DBA_AUTOTASK_OPERATION View
window_name	Optional name of the window in which client is to be enabled

Usage Notes

- If `operation` and `window_name` are both `NULL`, the client is enabled.
- If `operation` is not `NULL`, `window_name` is ignored and the specified operation is enabled
- If `operation` is `NULL` and `window_name` is not `NULL`, the client is enabled in the specified window.
- `Auto STS Capture Task` is accepted.

38.4.3 GET_CLIENT_ATTRIBUTES Procedure

This procedure returns values of select client attributes.

Syntax

```
DBMS_AUTO_TASK_ADMIN.GET_CLIENT_ATTRIBUTES(
   client_name        IN    VARCHAR2,   service_name       OUT   VARCHAR2,
window_group       OUT   VARCHAR2);
```

Parameters

Table 38-5 GET_CLIENT_ATTRIBUTES Procedure Parameters

Parameter	Description
`client_name`	Name of the client, as found in `DBA_AUTOTASK_CLIENT` View
`service_name`	Service name for client, may be `NULL`
`window_group`	Name of the window group in which the client is active

38.4.4 GET_P1_RESOURCES Procedure

This procedure returns percent of resources allocated to each `AUTOTASK` High Priority Consumer Group.

Syntax

```
DBMS_AUTO_TASK_ADMIN.GET_P1_RESOURCES(
   stats_group_pct     OUT   NUMBER,
   seg_group_pct       OUT   NUMBER,
   tune_group_pct      OUT   NUMBER,
   health_group_pct    OUT   NUMBER);
```

Parameters

Table 38-6 GET_P1_RESOURCES Procedure Parameters

Parameter	Description
`stats_group_pct`	%resources for Statistics Gathering
`seq_group_pct`	%resources for Space Management
`tune_group_pct`	%resources for SQL Tuning
`health_group_pct`	%resources for Health Checks

Usage Notes

Values will add up to 100%.

38.4.5 OVERRIDE_PRIORITY Procedures

This deprecated procedure is used to manually override task priority.

> **Note:**
>
> This subprogram is deprecated and becomes nonoperative with Oracle Database 12*c*.

This can be done at the client, operation or individual task level. This priority assignment is honored during the next maintenance window in which the named client is active. Specifically, setting the priority to `URGENT` causes a high priority job to be generated at the start of the maintenance window. Setting `priority` to `CLEAR` removes the override.

Syntax

Override Priority for a Client.

```
DBMS_AUTO_TASK_ADMIN.OVERRIDE_PRIORITY (
   client_name       IN    VARCHAR2,
   priority          IN    VARCHAR2);
```

Override Priority for an Operation.

```
DBMS_AUTO_TASK_ADMIN.OVERRIDE_PRIORITY (
   client_name       IN    VARCHAR2,
   operation         IN    VARCHAR2,
   priority          IN    VARCHAR2);
```

Override Priority for a Task.

```
DBMS_AUTO_TASK_ADMIN.OVERRIDE_PRIORITY (
   client_name       IN    VARCHAR2,
   operation         IN    VARCHAR2,
   task_target_type  IN    VARCHAR2,
   task_target_name  IN    VARCHAR2,
   priority          IN    VARCHAR2);
```

Parameters

Table 38-7 OVERRIDE_PRIORITY Procedure Parameters

Parameter	Description
`client_name`	Name of the client, as found in `DBA_AUTOTASK_CLIENT` View
`priority`	`URGENT`, `HIGH`, `MEDIUM` or `LOW`
`operation`	Name of the operation as specified in `DBA_AUTOTASK_OPERATION` View
`task_target_type`	Type of target to be affected, as found in V$`AUTOTASK_TARGET_TYPE` View

Table 38-7 (Cont.) OVERRIDE_PRIORITY Procedure Parameters

Parameter	Description
task_target_name	Name of the specific target to be affected

38.4.6 SET_CLIENT_SERVICE Procedure

This procedure associates an AUTOTASK Client with a specified Service.

Syntax

```
DBMS_AUTO_TASK_ADMIN.SET_CLIENT_SERVICE(
   client_name         IN    VARCHAR2,
   service_name        IN   VARCHAR2);
```

Parameters

Table 38-8 SET_CLIENT_SERVICE Procedure Parameters

Parameter	Description
client_name	Name of the client, as found in DBA_AUTOTASK_CLIENT View
service_name	Service name for client, may be NULL

Usage Notes

All work performed on behalf of the Client takes place only on instances where the service is enabled.

38.4.7 SET_P1_RESOURCES Procedure

This procedure sets percentage-based resource allocation for each High Priority Consumer Group used by AUTOTASK Clients.

Syntax

```
DBMS_AUTO_TASK_ADMIN.SET_P1_RESOURCES(
   stats_group_pct     IN   NUMBER,
   seg_group_pct       IN   NUMBER,
   tune_group_pct      IN   NUMBER,
   health_group_pct    IN   NUMBER);
```

Parameters

Table 38-9 SET_P1_RESOURCES Procedure Parameters

Parameter	Description
stats_group_pct	%resources for Statistics Gathering
seq_group_pct	%resources for Space Management
tune_group_pct	%resources for SQL Tuning
health_group_pct	%resources for Health Checks

Usage Notes

Values must be integers in the range `0` to `100`, and must add up to `100` (percent), otherwise, an exception is raised.

39
DBMS_AUTO_ZONEMAP

The DBMS_AUTO_ZONEMAP package provides autonomous maintenance and creation of zone map.

This chapter contains the following topics:

- DBMS_AUTO_TASK_ADMIN Security Model
- Summary of DBMS_AUTO_ZONEMAP Subprograms

39.1 DBMS_AUTO_TASK_ADMIN Security Model

The subprograms in DBMS_AUTO_TASK_ADMIN package can be executed by users with DBA privileges.

39.2 Summary of DBMS_AUTO_ZONEMAP Subprograms

This table lists the DBMS_AUTO_ZONEMAP package subprograms and briefly describes them.

Table 39-1 DBMS_AUTO_ZONEMAP Package Subprograms

Procedure	Description
ACTIVITY_REPORT Function	This function reports auto zone map activity for a given time window.
CONFIGURE Procedure	This procedure sets configuration options for auto zone map.

39.2.1 ACTIVITY_REPORT Function

This function reports auto zone map activity for a given time window.

Syntax

```
DBMS_AUTO_ZONEMAP.ACTIVITY_REPORT (
   start_time       IN TIMESTAMP WITH TIME ZONE  DEFAULT NULL,
   end_time         IN TIMESTAMP WITH TIME ZONE  DEFAULT NULL,
   type             IN VARCHAR2  DEFAULT NULL,
   section          IN VARCHAR2  DEFAULT NULL,
   level            IN VARCHAR2  DEFAULT NULL);
 RETURNS CLOB
```

Parameters

Table 39-2 ACTIVITY_REPORT Function Parameters

Parameter	Description
`start_time`	Timestamp from which auto zone map executions are observed for the report. If the value is `NULL` , the function reports everything from the beginning of auto zone map maintenance subject to purging. The default value is `NULL`.
`end_time`	Timestamp until which auto zone map executions are observed for the report. If the value is `NULL`, the function reports everything to the end of auto zone map maintenance subject to purging. The default value is `NULL`. **Note:** If you specify `NULL` for both `start_time` and `end_time`, the function will display the report from the last run.
`type`	Output type of the report. The possible values are `TEXT`, `XML`, and `HTML`. The default value is `TEXT`.
`section`	Particular section in the report. The possible values are: • `SUMMARY`: Very high level numbers summary on new zone map created and maintained for the given time window. • `DETAILS`: Detailed summary report on names and other details of new zone map created and maintained for the given time window. It also includes findings details. • `ALL`: In addition to summary and details, it includes time series based execution / action logs. The default value is `ALL`.

Table 39-2 (Cont.) ACTIVITY_REPORT Function Parameters

Parameter	Description
`level`	Format of the report. It represents the level of details with in each section. Possible values are: • `BASIC`: Represents very high level details in executive summary. You will only see numbers on zone map that were created, complete rebuilt and fast rebuilt. In new zone map details section, only new zone map name, date created and base table name re displayed. Under maintenance details section, only zone map name, previous state and current state are displayed. Under findings section, only object name and blacklist reason are displayed. Under action logs section, only important time series based log messages pertaining to zone map creation and maintenance are displayed. • `TYPICAL`: Full overview on executive summary section. Under `new zone map details` section, schema name, column list, and date created are displyed. Under `zone map maintenance details` section, refresh type and date are displayed. In `findings` section, timestamp and exception message are displayed. In `action logs` section more comprehensive logs than basic, which will have information about candidate column list, findings information, and creation DDLs are displayed. This is the default value. • `ALL`: On top of typical level, DOP used for each operation for creating or maintaining zone map, time took to process each DDL and other details in action logs are displayed. All log messages with details on clustering ratios of columns, exception messages and other details are also displayed.

Return Values

Report in `CLOB` format. The default report format is `TEXT`. Possible formats are `TEXT`, `XML`, and `HTML`.

Usage Notes

- DBMS_AUTO_ZONEMAP.ACTIVITY_REPORT(): returns report for all the execution history of last execution in `TEXT` format and all the sections will be displayed with typical level.
- DBMS_AUTO_ZONEMAP.ACTIVITY_REPORT(systimestamp-2, NULL): returns report for all the execution history for last two days in `TEXT` format and all the sections will be displayed with typical level.
- DBMS_AUTO_ZONEMAP.ACTIVITY_REPORT(NULL, systimestamp-1): returns report for all the execution history from beginning to yesterday in `TEXT` format and all the sections will be displayed with typical level.
- DBMS_AUTO_ZONEMAP.ACTIVITY_REPORT(NULL, NULL, ‘HTML’, ‘DETAILS’, ‘BASIC’): returns last execution’s report in `HTML` format and only details section will be displayed with only basic details in each section
- DBMS_AUTO_ZONEMAP.ACTIVITY_REPORT(systimestamp – 2, systimestmap, ‘TEXT’, ‘ALL’,’TYPICAL’): returns report for last 48 hours in text format and all the sections will be displayed with typical details.

39.2.2 CONFIGURE Procedure

This procedure sets configuration options for auto zone map; specifically to enable or disable feature and to control foreground or background mode of the feature.

Syntax

```
DBMS_AUTO_ZONEMAP.CONFIGURE (
   parameter_name       IN    VARCHAR2,
   parameter_value      IN    VARCHAR2);
```

Parameters

Table 39-3 CONFIGURE Procedure Parameters

Parameter	Description
`parameter_name`	`AUTO_ZONEMAP_MODE` is the only configure parameter name that is currently allowed. If you specify any other name, an invalid argument error message is displayed.
`parameter_value`	The four values are allowed for this parameter: • `ON`: Turns on auto zone map feature completely. Both for foreground and background zone map creation and maintenance. • `OFF`: Turns off auto zone map feature completely. Both for foreground and background zone map creation and maintenance. • `FOREGROUND`: Turns on only for foreground zone map creation and maintenance. • `BACKGROUND`: Turns on only for background zone map creation and maintenance.

40
DBMS_AVTUNE

The DBMS_AVTUNE package analyzes query access levels and aggregation workloads for each auto-cache enabled analytic view (AV) and creates or drops AV auto-caches to improve the overall performance of the SQL queries on that AV.

This chapter contains the following topics:

- DBMS_AVTUNE Overview
- DBMS_AVTUNE Security Model
- Summary of DBMS_AVTUNE Subprograms

40.1 DBMS_AVTUNE Overview

This package allows you to create a cache on an existing AV. It also enables the AV for auto caching. You can also create, refresh, and remove a star cache.

Usage Notes

- It is intended that the owner of the AV is also the user calling this package. Calling this package on an AV in another schema is not supported.
- Before you use any other procedure in this package, you need to call the AUTO_CACHE_ENABLE procedure. After the AV is enabled, auto-tuning will continue till the AV is disabled. You can call any other procedures on an enabled AV to manually tune the AV for more fine-grained control of the performance.

40.2 DBMS_AVTUNE Security Model

The DBMS_AVTUNE package runs under AUTHID CURRENT_USER. The SQL generated through use of the package is executed as the current user, and EXECUTE privilege on the package is granted to PUBLIC.

40.3 Summary of DBMS_AVTUNE Subprograms

This table lists the DBMS_AVTUNE package subprograms and briefly describes them.

Table 40-1 DBMS_AVTUNE Package Subprograms

Procedure	Description
AUTO_CACHE_CREATE Procedure	This procedure creates a cache on an existing analytic view.
AUTO_CACHE_DISABLE Procedure	This procedure disables an analytic view for auto caching and removes all auto-caches.
AUTO_CACHE_ENABLE Procedure	This procedure enables analytic views for auto caching and sets auto-cache settings for analytic views. This procedure throws an error if the AV is already enabled.

Table 40-1 (Cont.) DBMS_AVTUNE Package Subprograms

Procedure	Description
AUTO_CACHE_MODIFY Procedure	This procedure modifies tuning parameters originally supplied in the enable call.
AUTO_CACHE_REFRESH Procedure	This procedure forces a refresh immediately of all auto-caches for the analytic view.
AUTO_CACHE_REMOVE Procedure	This procedure manually removes the auto-cache defined for the specified levels.
AUTO_CACHE_TUNE Procedure	This procedure forces the tuning process to run immediately for the analytic view and implement changes.
AUTO_CACHE_STAR_DISABLE Procedure	This procedure removes the star cache on the attribute dimension `dim_name`.
AUTO_CACHE_STAR_ENABLE Procedure	This procedure creates a star cache on the atttribute dimension `dim_name`.
AUTO_CACHE_STAR_MODIFY Procedure	This procedure modifies the parameters specified in the enable call.
AUTO_CACHE_STAR_REFRESH Procedure	This procedure refreshes the star cache for the attribute dimension dim_name.

40.3.1 AUTO_CACHE_CREATE Procedure

This procedure creates a cache on an existing analytic view.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_CREATE (
    av_name         IN  VARCHAR2,
    level_group     IN  LEVEL_LIST,
    av_owner        IN  VARCHAR2);
```

Parameters

Table 40-2 AUTO_CACHE_CREATE Procedure Parameters

Parameter	Description
`av_name`	The name of the analytic view.
`level_group`	Represents the levels to cache.
`av_owner`	The owner of the AV, defaulting to the current schema.

Example

```
execute dbms_avtune.auto_cache_create('UNITS_HCUBE_AVTUNE',
  dbms_avtune.level_list(dbms_avtune.level('PRODUCT', 'ROLLUP', 'CLASS'),
  dbms_avtune.level('TIME', 'CALENDAR', 'YEAR')));
```

40.3.2 AUTO_CACHE_DISABLE Procedure

This procedure disables an analytic view for auto caching and removes all auto-caches.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_DISABLE (
    av_name         IN  VARCHAR2,
    av_owner        IN  VARCHAR2);
```

Parameters

Table 40-3 AUTO_CACHE_DISABLE Procedure Parameters

Parameter	Description
av_name	The name of the analytic view to disable caching on.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.3 AUTO_CACHE_ENABLE Procedure

This procedure enables analytic views for auto caching and sets auto-cache settings for analytic views. This procedure throws an error if the AV is already enabled. It is not possible to update an already enabled AV using this procedure.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_ENABLE (
    av_name           IN   VARCHAR2,
    refresh_intvl     IN   NUMBER,
    num_queries       IN   NUMBER,
    avg_query_time    IN   NUMBER,
    total_cache_pct   IN   NUMBER,
    init_max_pct      IN   NUMBER,
    init_numhier      IN   NUMBER,
    init_fixed_lvls   IN   LEVEL_LIST,
    run_mode          IN   VARCHAR2,
    av_owner          IN   VARCHAR2);
```

Parameters

Table 40-4 AUTO_CACHE_ENABLE Procedure Parameters

Parameter	Description
av_name	The name of the analytic view to enable caching on.
refresh_intvl	The number of minutes to wait between each cache refresh. The default value is 10 minutes.
num_queries	The minimum number of queries to consider for caches. The default value is 1.
avg_query_time	The average number of seconds each query should take before being considered for caches. The default value is 3.
total_cache_pct	The maximum percentage of the fact table you wish to cache. The default value is 50.

Table 40-4 (Cont.) AUTO_CACHE_ENABLE Procedure Parameters

Parameter	Description
init_max_pct	The maximum percentage of the fact table for the initial cache. The default value is 1%.
init_numhier	The maximum number of hierarchies to use for each dimension in the initial cache.
init_fixed_lvls	The levels you definitely want in the initial cache.
run_mode	This parameter is not used. Always use the default value.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.4 AUTO_CACHE_MODIFY Procedure

This procedure modifies the tuning parameters originally supplied in the enable call. The parameters remain unchanged if NULL is specified.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_MODIFY (
    av_name           IN   VARCHAR2,
    refresh_intvl     IN   NUMBER,
    num_queries       IN   NUMBER,
    avg_query_time    IN   NUMBER,
    total_cache_pct   IN   NUMBER,
    av_owner          IN   VARCHAR2);
```

Parameters

Table 40-5 AUTO_CACHE_MODIFY Procedure Parameters

Parameter	Description
av_name	The name of the analytic view to enable caching on.
refresh_intvl	The number of minutes to wait between each cache refresh. The default value is 10 minutes.
num_queries	The minimum number of queries to consider for caches. The default value is 1.
avg_query_time	The average number of seconds each query should take before being considered for caches. The default value is 3.
total_cache_pct	The maximum percentage of the fact table you wish to cache. The default value is 50.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.5 AUTO_CACHE_REFRESH Procedure

This procedure forces a refresh immediately of all auto-caches for the analytic view.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_REFRESH (
    av_name           IN   VARCHAR2,
    av_owner          IN   VARCHAR2);
```

Parameters

Table 40-6 AUTO_CACHE_REFRESH Procedure Parameters

Parameter	Description
av_name	The name of the analytic view.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.6 AUTO_CACHE_REMOVE Procedure

This procedure manually removes the auto-cache defined for the specified levels.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_REMOVE (
    av_name         IN  VARCHAR2,
    level_group     IN  LEVEL_LIST,
    av_owner        IN  VARCHAR2);
```

Parameters

Table 40-7 AUTO_CACHE_REMOVE Procedure Parameters

Parameter	Description
av_name	The name of the analytic view.
level_group	Represents the levels to cache.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.7 AUTO_CACHE_TUNE Procedure

This procedure forces the tuning process to run immediately for the analytic view and implement changes.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_TUNE (
    av_name           IN   VARCHAR2,
    num_queries       IN   NUMBER,
    avg_query_time    IN   NUMBER,
    run_mode          IN   VARCHAR2,
    av_owner          IN   VARCHAR2);
```

Parameters

Table 40-8 AUTO_CACHE_TUNE Procedure Parameters

Parameter	Description
av_name	The name of the analytic view.
num_queries	The minimum number of queries to consider for caches. The default value is 1.

Table 40-8 (Cont.) AUTO_CACHE_TUNE Procedure Parameters

Parameter	Description
avg_query_time	The average number of seconds each query should take before being considered for caches. The default value is 3.
run_mode	The run_mode parameter is not used. Always use the default value.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.8 AUTO_CACHE_STAR_DISABLE Procedure

This procedure removes the star cache on the attribute dimension dim_name.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_STAR_DISABLE (
    dim_name       IN   VARCHAR2,
    dim_owner      IN   VARCHAR2);
```

Parameters

Table 40-9 AUTO_CACHE_STAR_DISABLE Procedure Parameters

Parameter	Description
dim_name	The name of the attribute dimension.
dim_owner	The owner of the attribute dimension.

40.3.9 AUTO_CACHE_STAR_ENABLE Procedure

This procedure creates a star cache on the atttribute dimension dim_name. The av_name can optionally be passed to qualify the attribute dimension.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_STAR_ENABLE (
    dim_name          IN   VARCHAR2,
    av_name           IN   VARCHAR2,
    refresh_intvl     IN   NUMBER,
    run_mode          IN   VARCHAR2,
    dim_owner         IN   VARCHAR2,
    av_owner          IN   VARCHAR2);
```

Parameters

Table 40-10 AUTO_CACHE_STAR_ENABLE Procedure Parameters

Parameter	Description
dim_name	The name of the attribute dimension.
av_name	The name of the AV to associate the star cache with, if any. The av_name can optionally be specified to associate the star cache with the given AV. When this is done then the star cache will be refreshed and maintained with the AV instead of independently.

Table 40-10 (Cont.) AUTO_CACHE_STAR_ENABLE Procedure Parameters

Parameter	Description
refresh_intvl	The number of minutes to wait between each cache refresh. The default value is 10 minutes.
run_mode	This parameter is not used. Always use the default value.
dim_owner	The owner of the attribute dimension.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.10 AUTO_CACHE_STAR_MODIFY Procedure

This procedure modifies the parameters specified in the enable call. The parameters remain unchanged if NULL is specified.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_STAR_MODIFY (
    dim_name          IN   VARCHAR2,
    av_name           IN   VARCHAR2,
    refresh_intvl     IN   NUMBER,
    dim_owner         IN   VARCHAR2,
    av_owner          IN   VARCHAR2);
```

Parameters

Table 40-11 AUTO_CACHE_STAR_MODIFY Procedure Parameters

Parameter	Description
dim_name	The name of the attribute dimension.
av_name	The name of the associated analytic view.
refresh_intvl	The number of minutes to wait for each cache refresh. The default value is 10 minutes.
dim_owner	The owner of the attribute dimension.
av_owner	The owner of the AV, defaulting to the current schema.

40.3.11 AUTO_CACHE_STAR_REFRESH Procedure

This procedure refreshes the star cache for the attribute dimension dim_name.

Syntax

```
DBMS_AVTUNE.AUTO_CACHE_STAR_REFRESH (
    dim_name       IN   VARCHAR2,
    dim_owner      IN   VARCHAR2);
```

Parameters

Table 40-12 AUTO_CACHE_STAR_REFRESH Procedure Parameters

Parameter	Description
dim_name	The name of the attribute dimension.

Table 40-12 (Cont.) AUTO_CACHE_STAR_REFRESH Procedure Parameters

Parameter	Description
dim_owner	The owner of the attribute dimension.

41
DBMS_AW_STATS

DBMS_AW_STATS contains subprograms for managing optimizer statistics for cubes and dimensions. Generating the statistics does not have a significant performance cost.

> **See Also:**
>
> *Oracle OLAP User's Guide* regarding use of the OLAP option to support business intelligence and analytical applications

This chapter contains the following topics:

- Using DBMS_AW_STATS
- Summary of DBMS_AW_STATS Subprograms

41.1 Using DBMS_AW_STATS

Cubes and dimensions are first class data objects that support multidimensional analytics. They are stored in a container called an analytic workspace. Multidimensional objects and analytics are available with the OLAP option to Oracle Database.

Optimizer statistics are used to create execution plans for queries that join two cube views or join a cube view to a table or a view of a table. They are also used for query rewrite to cube materialized views. You need to generate the statistics only for these types of queries.

Queries against a single cube do not use optimizer statistics. These queries are automatically optimized within the analytic workspace.

41.2 Summary of DBMS_AW_STATS Subprograms

DBMS_AW_STATS uses the ANALYZE and CLEAR procedures.

Table 41-1 DBMS_AW_STATS Package Subprograms

Subprogram	Description
ANALYZE Procedure	Generates optimizer statistics on cubes and cube dimensions.
CLEAR Procedure	Clears optimizer statistics from cubes and cube dimensions.

41.2.1 ANALYZE Procedure

This procedure generates optimizer statistics on a cube or a cube dimension.

These statistics are used to generate some execution plans, as described in "Using DBMS_AW_STATS".

For a cube, the statistics are for all of the measures and calculated measures associated with the cube. These statistics include:

- The average length of data values
- The length of the largest data value
- The minimum value
- The number of distinct values
- The number of null values

For a dimension, the statistics are for the dimension and its attributes, levels, and hierarchies. These statistics include:

- The average length of a value
- The length of the largest value
- The minimum value
- The maximum value

Syntax

```
DBMS_AW_STATS.ANALYZE
     (inname       IN VARCHAR2);
```

Parameters

Table 41-2 ANALYZE Procedure Parameters

Parameter	Description
inname	The qualified name of a cube or a dimension. For a cube, the format of a qualified name is *owner.cube_name*. For a dimension, the format is *owner.dimension_name*.

Usage Notes

Always analyze the dimensions first, then the cube.

After analyzing a dimension, analyze all cubes that use that dimension.

Example

This sample script generates optimizer statistics on UNITS_CUBE and its dimensions.

```
BEGIN
     DBMS_AW_STATS.ANALYZE('time');
     DBMS_AW_STATS.ANALYZE('customer');
     DBMS_AW_STATS.ANALYZE('product');
     DBMS_AW_STATS.ANALYZE('channel');
     DBMS_AW_STATS.ANALYZE('units_cube');
END;
/
```

The following statements create and display an execution plan for aSELECT statement that joins columns from UNITS_CUBE_VIEW, CUSTOMER_PRIMARY_VIEW, and the ACCOUNTS table:

```
EXPLAIN PLAN FOR SELECT
     cu.long_description customer,
```

```
                    a.city city,
                    a.zip_pc zip,
                    cu.level_name "LEVEL",
                    round(f.sales) sales
                /* From dimension views and cube view */
                FROM time_calendar_view t,
                    product_primary_view p,
                    customer_view cu,
                    channel_view ch,
                    units_cube_view f,
                    account a
                /* Create level filters instead of GROUP BY */
                WHERE t.long_description = '2004'
                    AND p.level_name ='TOTAL'
                    AND cu.customer_account_id like 'COMP%'
                    AND ch.level_name = 'TOTAL'
                /* Join dimension views to cube view */
                    AND t.dim_key = f.TIME
                    AND p.dim_key = f.product
                    AND cu.dim_key = f.customer
                    AND ch.dim_key = f.channel
                    AND a.account_id = cu.customer_account_id
                ORDER BY zip;

SQL>  SELECT plan_table_output FROM TABLE(dbms_xplan.display());

PLAN_TABLE_OUTPUT
-------------------------------------------------------------------------------------------------
Plan hash value: 3890178023

-------------------------------------------------------------------------------------------------
| Id  | Operation                          | Name       | Rows  | Bytes | Cost (%CPU)| Time     |
-------------------------------------------------------------------------------------------------
|   0 | SELECT STATEMENT                   |            |     1 |    89 |     6  (34)| 00:00:01 |
|   1 |  SORT ORDER BY                     |            |     1 |    89 |     6  (34)| 00:00:01 |
|*  2 |   HASH JOIN                        |            |     1 |    89 |     5  (20)| 00:00:01 |
|   3 |    JOINED CUBE SCAN PARTIAL OUTER|            |       |       |            |          |
|   4 |     CUBE ACCESS                    | UNITS_CUBE |       |       |            |          |
|   5 |     CUBE ACCESS                    | CHANNEL    |       |       |            |          |
|   6 |     CUBE ACCESS                    | CUSTOMER   |       |       |            |          |
|   7 |     CUBE ACCESS                    | PRODUCT    |       |       |            |          |
|*  8 |     CUBE ACCESS                    | TIME       |     1 |    55 |     2   (0)| 00:00:01 |
|*  9 |    TABLE ACCESS FULL               | ACCOUNT    |     3 |   102 |     2   (0)| 00:00:01 |
-------------------------------------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------------

   2 - access("A"."ACCOUNT_ID"=SYS_OP_ATG(VALUE(KOKBF$),39,40,2))
   8 - filter(SYS_OP_ATG(VALUE(KOKBF$),16,17,2)='2004' AND
              SYS_OP_ATG(VALUE(KOKBF$),39,40,2) LIKE 'COMP%' AND
              SYS_OP_ATG(VALUE(KOKBF$),47,48,2)='TOTAL' AND
              SYS_OP_ATG(VALUE(KOKBF$),25,26,2)='TOTAL')
   9 - filter("A"."ACCOUNT_ID" LIKE 'COMP%')

Note
-----
   - dynamic statistics used for this statement

30 rows selected.
```

41.2.2 CLEAR Procedure

This procedure clears the statistics generated by the ANALYZE procedure.

Syntax

```
DBMS_AW_STATS.CLEAR (
          inname        IN  VARCHAR2;
```

Parameters

Table 41-3 CLEAR Procedure Parameters

Parameter	Description
inname	The qualified name of a cube or a dimension. For a cube, the format of a qualified name is *owner.cube_name*. For a dimension, the format is *owner.dimension_name*.

Examples

The following scripts clears the statistics from UNITS_CUBE and its dimensions.

```
BEGIN
     DBMS_AW_STATS.clear('units_cube');
     DBMS_AW_STATS.clear('time');
     DBMS_AW_STATS.clear('customer');
     DBMS_AW_STATS.clear('product');
     DBMS_AW_STATS.clear('channel');
END;
/
```

42
DBMS_BLOCKCHAIN_TABLE

A blockchain table is an append-only table designed for centralized blockchain applications. The `DBMS_BLOCKCHAIN_TABLE` package allows you do operations like the following: delete rows in a blockchain table that are beyond the row retention defined for the blockchain table; get the bytes that are input to the cryptographic hash for a row so you can verify the hash in the row; sign a row you inserted into a blockchain table after the row is added to a chain in the blockchain table; and have the database verify the hashes and digital signatures on some or all rows in a blockchain table. V2 blockchain tables support schema evolution, delegate signatures, and countersignatures in addition to the functionality found in V1 blockchain tables. Blockchain tables support only `DER` encoding for `X.509` certificates, not `PEM` encoding.

This chapter contains the following topics:

- DBMS_BLOCKCHAIN_TABLE Overview
- DBMS_BLOCKCHAIN_TABLE Security Model
- Summary of DBMS_BLOCKCHAIN_TABLE Subprograms

See Also:

- *Oracle Database Administrator's Guide*
- *Oracle Database Concepts*
- *Oracle Database SQL Language Reference*
- *Oracle Database Reference*
- For information on hidden columns in blockchain tables, see Hidden Columns in Blockchain Tables

42.1 DBMS_BLOCKCHAIN_TABLE Overview

In Oracle Blockchain Table, peers are database users who trust the database to maintain a tamper-resistant ledger.

The ledger is implemented as a blockchain table, which is defined and managed by the application. Existing applications can protect against fraud without requiring a new infrastructure or programming model. Although transaction throughput is lower than for a standard table, performance for a blockchain table is better than for a decentralized blockchain.

The `DBMS_BLOCKCHAIN_TABLE` package lets you do the following:

- delete rows in a blockchain table that are beyond the row retention defined for the blockchain table
- get the bytes that are input to the signature algorithm so you can sign a row you inserted into the blockchain table

- get the bytes that are input to the cryptographic hash for a row so you can verify the hash in the row
- sign a row you inserted into a blockchain table after the row is added to a chain in the blockchain table
- have the database verify the hashes and signatures on some or all rows in a blockchain table
- enable a delegate with sufficient privileges to sign a row inserted by another user
- procure a countersignature for the row that is being signed by the end-user or delegate

42.2 DBMS_BLOCKCHAIN_TABLE Security Model

The `DBMS_BLOCKCHAIN_TABLE` package is owned by `SYS` and is installed as part of database installation. The routines in the package are run with invoker's rights (run with the privileges of the current user). Thus any user with select privileges on the blockchain table should be able to validate the row contents of that table.

Any user with delete privileges on the blockchain table can delete rows beyond the retention period defined for the blockchain table.

A user that inserted a row into the blockchain table can add a digital signature to the row after the row is added to a chain in the blockchain table. This user can secure a countersignature on this row either when signing the row or at a later time after the row has been signed either by the user or by a delegate.

A delegate signer for a row in a blockchain table needs the `SIGN` privilege on the table. In addition, the ID number of the delegate user may be specified in the hidden column `ORABCTAB_DELEGATE_USER_NUMBER$` when the row is inserted.

A user with the `SIGN` privilege on a blockchain table can secure a countersignature on any row in the blockchain table that has a user signature or a delegate signature. Similarly, the owner of a blockchain table can secure a countersignature on any row in the blockchain table that has a user signature or a delegate signature.

42.3 Summary of DBMS_BLOCKCHAIN_TABLE Subprograms

The `DBMS_BLOCKCHAIN_TABLE` package uses `ADD_INTERVAL_PARTITIONING`, `COUNTERSIGN_ROW`, `COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS`, `DELETE_EXPIRED_ROWS`, `GET_BLOCKCHAIN_DIGEST`, `GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS`, `GET_BYTES_FOR_ROW_HASH`, `GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_COLUMNS`, `GET_BYTES_FOR_ROW_SIGNATURE`, `GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_KEY_COLUMNS`, `GET_SIGNED_BLOCKCHAIN_DIGEST`, `GET_SIGNED_BLOCKCHAIN_DIGEST_FOR SELECTED_ROWS`, `SIGN_ROW`, `SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS`, `SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE`, `SIGN_ROW_WITH_COUNTERSIGNATURE`, `VERIFY_ROWS`, `VERIFY_TABLE_BLOCKCHAIN`, and `VERIFY_USER_BLOCKCHAIN_ROWS` subprograms to perform various functions.

Table 42-1 DBMS_BLOCKCHAIN_TABLE Package Subprograms

Subprogram	Description
ADD_INTERVAL_PARTITIONING Procedure	This procedure adds interval partitioning to an existing, non-partitioned, V1 or V2 blockchain table.

Table 42-1 (Cont.) DBMS_BLOCKCHAIN_TABLE Package Subprograms

Subprogram	Description
COUNTERSIGN_ROW Procedure	This procedure procures a countersignature on a specified row in a blockchain table. The countersignature will be produced by signing the row data content using the table owner's private key stored in the database wallet.
COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure	It is an extension of `COUNTERSIGN_ROW` that uses at most three user columns to identify exactly one row in a blockchain table and procure a countersignature on that row.
DELETE_EXPIRED_ROWS Procedure	This procedure deletes rows outside the retention window created before `before_timestamp` if the timestamp is specified; otherwise, this procedure deletes all rows outside the retention window. This procedure commits before deleting any expired rows and commits after deleting any expired rows.
GET_BLOCKCHAIN_DIGEST Function	This function generates and returns a cryptographic hash of the digest for a specified blockchain table.
GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function	This function generates and returns a cryptographic hash of the digest for user-specified rows in a blockchain table.
GET_BYTES_FOR_ROW_HASH Procedure	This procedure returns in `row_data` the bytes for the particular row identified (a series of meta-data-value, column-data-value pairs in column position order) followed by the hash for the previous row in the chain in the data format specified.
GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_COLUMNS Procedure	It is an extension of `GET_BYTES_FOR_ROW_HASH` that uses at most three user columns instead of the `instance identifier`, `chain identifier`, and `sequence number` to uniquely identify the row.
GET_BYTES_FOR_ROW_SIGNATURE Procedure	This procedure returns the bytes used to compute a user signature, a delegate signature, or a countersignature. For a user signature or a delegate signature, the procedure returns in `row_data` the bytes in the hash in the row without any metadata.
GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_KEY_COLUMNS Procedure	It is an extension of `GET_BYTES_FOR_ROW_SIGNATURE` that uses at most three user column values instead of the `instance identifier`, `chain identifier`, and `sequence number` to uniquely identify the row.
GET_SIGNED_BLOCKCHAIN_DIGEST Function	This function generates and returns the signed digest for a specified blockchain table using the table owner's private key stored in the database wallet. The `signed_bytes`, `signed_row_indexes`, and `schema_certificate_guid` are also returned.
GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function	This function generates and returns the signed digest for user-specified rows in a blockchain table using the table owner's private key stored in the database wallet. The particular rows in the digest are specified by the `row_selector` parameter. The `signed_bytes`, `signed_row_indexes`, and `schema_certificate_guid` are also returned.
SIGN_ROW Procedure	This procedure can be used by the current user to provide a signature on the row content of a previously inserted row.
SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure	This procedure allows an end user or a delegate to sign a row using a set of at most three user columns that uniquely identify a single row.
SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE Procedure	This procedure uses at most three user column names and values to uniquely identify a single row to sign and countersign.

Table 42-1 (Cont.) DBMS_BLOCKCHAIN_TABLE Package Subprograms

Subprogram	Description
SIGN_ROW_WITH_COUNTERSIGNATURE Procedure	It is an extension of `SIGN_ROW` that enables the user to request a countersignature from the database. The countersignature will be produced by signing the row data content using the table owner's private key stored in the database wallet.
VERIFY_ROWS Procedure	This procedure verifies all rows on all applicable system chains for the integrity of `HASH` column value for rows created in the range of `low_timestamp` to `high_timestamp`. Row signatures can be verified as an option.
VERIFY_TABLE_BLOCKCHAIN Procedure	This procedure verifies all rows whose creation-times fall between the minimum value for the row-creation time from `bytes_previous` and the maximum value for row-creation time from `bytes_latest` and returns the number of successfully verified rows.
VERIFY_USER_BLOCKCHAIN_ROWS Procedure	This procedure verifies rows of one or more user chains when the user chains feature is enabled on the blockchain table.

42.3.1 ADD_INTERVAL_PARTITIONING Procedure

This procedure adds interval partitioning to an existing, non-partitioned, V1 or V2 blockchain table.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.ADD_INTERVAL_PARTITIONING(
    schema_name               IN    VARCHAR2,
    table_name                IN    VARCHAR2,
    interval_number           IN    NUMBER,
    interval_frequency        IN    VARCHAR2,
    first_high_timestamp      IN    TIMESTAMP);
```

Parameters

Table 42-2 ADD_INTERVAL_PARTITIONING Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`interval_number`	Sets how often the database creates partitions for the blockchain table.
`interval_frequency`	Sets the frequency for the value that was set in the `interval_number` setting. Supported values are `YEAR`, `MONTH`, `DAY`, `HOUR`, and `MINUTE`.
`first_high_timestamp`	A timestamp that determines the upper boundary of the first partition in the blockchain table.

Usage Notes

- For an existing, non-partitioned, V1 or V2 immutable table, a procedure with the same name and the same parameters is provided in the `DBMS_IMMUTABLE_TABLE` package.

- Composite partitioning (that is, sub-partitioning) is not supported with the above interval partitioning.

42.3.2 COUNTERSIGN_ROW Procedure

This procedure procures a countersignature on a specified row in a blockchain table. The countersignature will be produced by signing the row data content using the table owner's private key stored in the database wallet. A row in a blockchain table can be countersigned only if the row belongs to the current epoch for the blockchain table.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.COUNTERSIGN_ROW (
   schema_name                      IN VARCHAR2,
   table_name                       IN VARCHAR2,
   instance_id                      IN NUMBER,
   chain_id                         IN NUMBER,
   sequence_id                      IN NUMBER,
   countersignature_algo            IN NUMBER DEFAULT
DBMS_BLOCKCHAIN_TABLE.SIGN_ALGO_DEFAULT,
   countersignature_signed_bytes    IN OUT BLOB,
   countersignature                 OUT RAW,
   countersignature_certificate_guid OUT RAW,
   countersignature_content_version  IN VARCHAR2 DEFAULT 'V2_DIGEST',
   pdb_guid                         IN RAW DEFAULT NULL);
```

Parameters

Table 42-3 COUNTERSIGN_ROW Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
instance_id	The instance that inserted the row. Valid values are 1, 2, and so on.
chain_id	The chain containing the row. By default, there are 32 chains in each instance, and they are numbered from 0 to 31.
sequence_id	The position of the row on the specified chain.
countersignature_algo	The digital signature algorithm for the countersignature. The parameter must be one of the following package constants: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
countersignature_signed_bytes	The bytes that are input to the algorithm that generates the countersignature. The caller must pass an empty BLOB for this parameter.
countersignature	The digital signature on the bytes returned in countersignature_signed_bytes.
countersignature_certificate_guid	A unique identifier for the certificate of the blockchain table owner stored in the database that may be used to verify the countersignature.
countersignature_content_version	The version of the data contents and layout that are used as input to the countersignature algorithm. Only 'V2_DIGEST' is supported in this release.
pdb_guid	The identifier of the local pluggable database. This parameter is used by Oracle GoldenGate replication and must be NULL.

42.3.3 COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure

It is an extension of `COUNTERSIGN_ROW` that uses at most three user columns to identify exactly one row in a blockchain table and procure a countersignature on that row. A row in a blockchain table can be countersigned only if the row belongs to the current epoch for the blockchain table.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS(
   schema_name                       IN      VARCHAR2,
   table_name                        IN      VARCHAR2,
   keycol1_name                      IN      VARCHAR2,
   keycol1_value                     IN      VARCHAR2,
   keycol2_name                      IN      VARCHAR2 DEFAULT NULL,
   keycol2_value                     IN      VARCHAR2 DEFAULT NULL,
   keycol3_name                      IN      VARCHAR2 DEFAULT NULL,
   keycol3_value                     IN      VARCHAR2 DEFAULT NULL,
   countersignature_algo             IN      NUMBER DEFAULT
DBMS_BLOCKCHAIN_TABLE.SIGN_ALGO_DEFAULT,
   countersignature_signed_bytes     IN OUT  BLOB,
   countersignature                  OUT     RAW,
   countersignature_certificate_guid OUT     RAW,
   countersignature_content_version  IN      VARCHAR2 DEFAULT 'V2_DIGEST',
   pdb_guid                          IN      RAW DEFAULT NULL);
```

Parameters

Table 42-4 COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`keycol1_name`	The name of the key column.
`keycol1_value`	The value of the key column.
`keycol2_name`	The name of the second column in a composite key.
`keycol2_value`	The value of the second column in a composite key.
`keycol3_name`	The name of the third column in a composite key.
`keycol3_value`	The value of the third column in a composite key.
`countersignature_algo`	The digital signature algorithm for the countersignature. The parameter must be one of the following package constants: • `SIGN_ALGO_RSA_SHA2_256` • `SIGN_ALGO_RSA_SHA2_384` • `SIGN_ALGO_RSA_SHA2_512`
`countersignature_signed_bytes`	The bytes that are input to the algorithm that generates the countersignature. The caller must pass an empty BLOB for this parameter.
`countersignature`	The digital signature on the bytes returned in `countersignature_signed_bytes`.

Table 42-4 (Cont.) COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure Parameters

Parameter	Description
countersignature_certificate_guid	A unique identifier for the certificate of the blockchain table owner stored in the database that may be used to verify the countersignature.
countersignature_content_version	The version of the data contents and layout that are used as input to the countersignature algorithm. Only 'V2_DIGEST' is supported in this release.
pdb_guid	The identifier of the local pluggable database. This parameter is used by Oracle GoldenGate replication and must be NULL.

42.3.4 DELETE_EXPIRED_ROWS Procedure

This procedure deletes rows outside the retention window created before_timestamp if the timestamp is specified; otherwise, this procedure deletes all rows outside the retention window. The number of rows deleted is returned in number_of_rows_deleted parameter. This procedure commits before deleting any expired rows and commits after deleting any expired rows.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.DELETE_EXPIRED_ROWS(
   schema_name             IN VARCHAR2,
   table_name                 IN VARCHAR2,
   before_timestamp         IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   number_of_rows_deleted   OUT NUMBER);
```

Parameters

Table 42-5 DELETE_EXPIRED_ROWS Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
before_timestamp	The end time for the range of rows deleted by the procedure, subject to the row retention time currently associated with the blockchain table. This is an optional parameter. The default value is NULL.
number_of_rows_deleted	The count of the number of rows deleted.

42.3.5 GET_BLOCKCHAIN_DIGEST Function

This function generates and returns a cryptographic hash of the digest for a specified blockchain table.

> **Note:**
>
> Rows inserted into a blockchain table that have not been committed are not guaranteed to be durable. Therefore, a blockchain digest does not include inserted rows that have not been committed.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_BLOCKCHAIN_DIGEST(
  schema_name             IN      VARCHAR2,
  table_name              IN      VARCHAR2,
  digest_bytes            IN OUT  BLOB,
  digest_rows_indexes     OUT     SYS.ORABCTAB_ROW_ARRAY_T,
  hash_algo               IN      NUMBER   DEFAULT
DBMS_BLOCKCHAIN_TABLE.HASH_ALGO_DEFAULT)
  RETURN RAW;
```

Parameters

Table 42-6 GET_BLOCKCHAIN_DIGEST Function Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
digest_bytes	The BLOB value that contains a header followed by an array of row-info. This sequence of bytes is input to the cryptographic hash function. The caller must pass an empty BLOB for this parameter.
digest_rows_indexes	This parameter specifies the rows in the blockchain table that were chosen for the digest.
hash_algo	The cryptographic hash algorithm to use. The parameter must be one of the following package constants: • HASH_ALGO_SHA2_256 • HASH_ALGO_SHA2_384 • HASH_ALGO_SHA2_512

Usage Notes

- A blockchain table digest created by the GET_BLOCKCHAIN_DIGEST function has table information specific to a pluggable database. Such a digest can be used only in the pluggable database in which the digest was created and only for the table that was used to create the digest. For DBMS_BLOCKCHAIN_TABLE.VERIFY_TABLE_BLOCKCHAIN, these requirements mean that both blockchain table digests must have been generated in the current pluggable database for the same blockchain table.

  For example, suppose you create a digest for a blockchain table in pluggable database A, use Data Pump to export the blockchain table, and use Data Pump to import the blockchain table into pluggable database B. The blockchain table digest created in

pluggable database A cannot be used in pluggable database B. You need to create a new blockchain table digest in pluggable database B.

> **See Also:**
>
> *Oracle Database Administrator's Guide*

42.3.6 GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function

This function generates and returns a cryptographic hash of the digest for user-specified rows in a blockchain table.

> **Note:**
>
> Rows inserted into a blockchain table that have not been committed are not guaranteed to be durable. Therefore, a blockchain digest does not include inserted rows that have not been committed.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS(
  schema_name             IN      VARCHAR2,
  table_name              IN      VARCHAR2,
  row_selector            IN      VARCHAR2,
  digest_bytes            IN OUT  BLOB,
  row_data_bytes          IN OUT  BLOB,
  digest_rows_indexes     OUT     SYS.ORABCTAB_ROW_ARRAY_T,
  hash_algo               IN      NUMBER   DEFAULT
DBMS_BLOCKCHAIN_TABLE.HASH_ALGO_DEFAULT)
  RETURN RAW;
```

Parameters

Table 42-7 GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`row_selector`	The condition, a where clause without the `WHERE` keyword.
`digest_bytes`	The BLOB value that contains a header followed by an array of row-info. This sequence of bytes is input to the cryptographic hash function. The caller must pass an empty BLOB for this parameter.
`row_data_bytes`	This parameter specifies the content of the rows in the blockchain table that were selected for the digest. The caller must pass an empty BLOB for this parameter.
`digest_rows_indexes`	This parameter specifies the rows in the blockchain table that were selected for the digest.

Table 42-7 (Cont.) GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function Parameters

Parameter	Description
hash_algo	The cryptographic hash algorithm to use. The parameter must be one of the following package constants: • HASH_ALGO_SHA2_256 • HASH_ALGO_SHA2_384 • HASH_ALGO_SHA2_512

Usage Notes

- A blockchain table digest created by the GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS function has table information specific to a pluggable database. Such a digest can be used only in the pluggable database in which the digest was created and only for the table that was used to create the digest. For DBMS_BLOCKCHAIN_TABLE.VERIFY_TABLE_BLOCKCHAIN, these requirements mean that both blockchain table digests must have been generated in the current pluggable database for the same blockchain table.

  For example, suppose you create a digest for a blockchain table in pluggable database A, use Data Pump to export the blockchain table, and use Data Pump to import the blockchain table into pluggable database B. The blockchain table digest created in pluggable database A cannot be used in pluggable database B. You need to create a new blockchain table digest in pluggable database B.

See Also:

Oracle Database Administrator's Guide

42.3.7 GET_BYTES_FOR_ROW_HASH Procedure

This procedure returns in row_data the bytes for the particular row identified (a series of meta-data-value, column-data-value pairs in column position order) followed by the hash for the previous row in the chain in the data format specified.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_BYTES_FOR_ROW_HASH(
   schema_name          IN VARCHAR2,
   table_name              IN VARCHAR2,
   instance_id          IN NUMBER,
   chain_id              IN NUMBER,
   sequence_id              IN NUMBER,
   data_format              IN NUMBER,
   row_data              IN OUT BLOB,
   chain_name            IN VARCHAR2 DEFAULT NULL,
   pdb_guid              IN RAW DEFAULT NULL);
```

Parameters

Table 42-8 GET_BYTES_FOR_ROW_HASH Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
instance_id	The instance that inserted the row. Valid values are 1, 2, and so on.
chain_id	The chain containing the row. By default, there are 32 chains in each instance, and they are numbered from 0 to 31.
sequence_id	The position of the row on the specified chain.
data_format	The version of the data layout for the hash in the specified row. Must be 1 in this release.
row_data	The bytes for the specified row in the specified data format that can be input to the cryptographic hash function to verify the value of the hash in the row. Any bytes in the BLOB are overwritten.
chain_name	The name of the user chain when the bytes for the cryptographic hash on the user chain are desired. Specify NULL when the bytes for the cryptographic hash on the system chain are desired.
pdb_guid	For a V2 blockchain table, the identifier of the pluggable database that inserted the row. Must be NULL for a V1 blockchain table.

Usage Notes

The metadata bytes for a column are 20 bytes that encode the blockchain algorithm version used to hash the row, the column position, the column data type, whether the column value is NULL, and the actual length of the column value in bytes.

For non-character columns, the column data bytes are the actual bytes representing the column value on the disk. For character and character LOB columns, the values are normalized to specific character sets. For CHAR and NCHAR columns, the number of blanks is normalized.

A few metadata bytes are reserved for future use.

> **See Also:**
>
> For more information on normalizations, see *Oracle Database Administrator's Guide*

42.3.8 GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_COLUMNS Procedure

It is an extension of `GET_BYTES_FOR_ROW_HASH` that uses at most three user columns instead of an `instance identifier`, `chain identifier`, and `sequence number` to uniquely identify the row.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_COLUMNS(
  schema_name    IN     VARCHAR2,
  table_name     IN     VARCHAR2,
  data_format    IN     NUMBER,
  row_data       IN OUT BLOB,
  chain_name     IN     VARCHAR2 DEFAULT NULL,
  keycol1_name   IN     VARCHAR2,
  keycol1_value  IN     VARCHAR2,
  keycol2_name   IN     VARCHAR2 DEFAULT NULL,
  keycol2_value  IN     VARCHAR2 DEFAULT NULL,
  keycol3_name   IN     VARCHAR2 DEFAULT NULL,
  keycol3_value  IN     VARCHAR2 DEFAULT NULL,
  pdb_guid       IN     RAW DEFAULT NULL);
```

Parameters

Table 42-9 GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_COLUMNS Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`data_format`	The version of the data layout for the hash in the specified row. Must be 1 in this release.
`row_data`	The bytes for the specified row in the specified data format that can be input to the cryptographic hash function to verify the value of the hash in the row. Any bytes in the BLOB are overwritten.
`chain_name`	The name of the user chain when the bytes for the cryptographic hash on the user chain are desired. Specify `NULL` when the bytes for the cryptographic hash on the system chain are desired.
`keycol1_name`	The name of the key column.
`keycol1_value`	The value of the key column.
`keycol2_name`	The name of the second column in a composite key.
`keycol2_value`	The value of the second column in a composite key.
`keycol3_name`	The name of the third column in a composite key.
`keycol3_value`	The value of the third column in a composite key.
`pdb_guid`	For a V2 blockchain table, the identifier of the pluggable database that inserted the row. Must be `NULL` for a V1 blockchain table.

42.3.9 GET_BYTES_FOR_ROW_SIGNATURE Procedure

This procedure returns the bytes used to compute a user signature, a delegate signature, or a countersignature. For a user signature or a delegate signature, the procedure returns in `row_data` the bytes in the hash in the row without any metadata. No other columns are involved, either in the row or in the previous row. For a countersignature, the routine returns in `row_data` the bytes that are input to the digital signature algorithm that computes a countersignature on the row.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_BYTES_FOR_ROW_SIGNATURE(
   schema_name          IN VARCHAR2,
   table_name              IN VARCHAR2,
   instance_id          IN NUMBER,
   chain_id          IN NUMBER,
   sequence_id              IN NUMBER,
   data_format              IN NUMBER,
   row_data              IN OUT BLOB,
   pdb_guid              IN RAW DEFAULT NULL,
   signature_type        IN VARCHAR2 DEFAULT 'USER');
```

Parameters

Table 42-10 GET_BYTES_FOR_ROW_SIGNATURE Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`instance_id`	The instance on which the row was inserted. Valid values are `1`, `2`, and so on.
`chain_id`	The chain on which the row was inserted. By default, there are 32 chains in each instance, and they are numbered from 0 to 31.
`sequence_id`	The position of the row on the chain.
`data_format`	The format of the data in `row_data`. Must be `1` in this release.
`row_data`	A sequence of bytes that must be signed. The caller must pass an empty BLOB for this parameter.
`pdb_guid`	For a V2 blockchain table, the identifier of the pluggable database that inserted the row. Must be `NULL` for a V1 blockchain table.
`signature_type`	The valid values for `signature_type` are `USER`, `DELEGATE`, and `COUNTERSIGNATURE`. `DELEGATE` and `USER` may be used interchangeably.

42.3.10 GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_KEY_COLUMNS Procedure

It is an extension of `GET_BYTES_FOR_ROW_SIGNATURE` that uses at most three user column values instead of an `instance identifier`, `chain identifier`, and `sequence number` to uniquely identify the row.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_KEY_COLUMNS(
  schema_name    IN     VARCHAR2,
  table_name     IN     VARCHAR2,
  data_format    IN     NUMBER,
  row_data       IN OUT BLOB,
  keycol1_name   IN     VARCHAR2,
  keycol1_value  IN     VARCHAR2,
  keycol2_name   IN     VARCHAR2 DEFAULT NULL,
  keycol2_value  IN     VARCHAR2 DEFAULT NULL,
  keycol3_name   IN     VARCHAR2 DEFAULT NULL,
  keycol3_value  IN     VARCHAR2 DEFAULT NULL,
  pdb_guid       IN     RAW DEFAULT NULL,
  signature_type IN     VARCHAR2 DEFAULT 'USER');
```

Parameters

Table 42-11 GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_KEY_COLUMNS Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`data_format`	The format of the data in `row_data`. Must be `1` in this release.
`row_data`	A sequence of bytes that must be signed. The caller must pass an empty BLOB for this parameter.
`keycol1_name`	The name of the key column.
`keycol1_value`	The value of the key column.
`keycol2_name`	The name of the second column in a composite key.
`keycol2_value`	The value of the second column in a composite key.
`keycol3_name`	The name of the third column in a composite key.
`keycol3_value`	The value of the third column in a composite key.
`pdb_guid`	For a V2 blockchain table, the identifier of the pluggable database that inserted the row. Must be `NULL` for a V1 blockchain table.
`signature_type`	The valid values for `signature_type` are `USER`, `DELEGATE`, and `COUNTERSIGNATURE`. `DELEGATE` and `USER` may be used interchangeably.

42.3.11 GET_SIGNED_BLOCKCHAIN_DIGEST Function

This function generates and returns the signed digest for a specified blockchain table using the table owner's private key stored in the database wallet. The `signed_bytes`, `signed_row_indexes`, and `schema_certificate_guid` are also returned.

> **Note:**
>
> Rows inserted into a blockchain table that have not been committed are not guaranteed to be durable. Therefore, a signed blockchain digest does not include inserted rows that have not been committed.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_SIGNED_BLOCKCHAIN_DIGEST(
   schema_name               IN      VARCHAR2,
   table_name                   IN      VARCHAR2,
   signed_bytes              IN OUT  BLOB,
   signed_rows_indexes       OUT     SYS.ORABCTAB_ROW_ARRAY_T,
   schema_certificate_guid   OUT     RAW,
   signature_algo            IN      NUMBER DEFAULT
DBMS_BLOCKCHAIN_TABLE.SIGN_ALGO_DEFAULT)
   RETURN RAW;
```

Parameters

Table 42-12 GET_SIGNED_BLOCKCHAIN_DIGEST Function Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`signed_bytes`	The BLOB value that contains a header followed by an array of row-info. This sequence of bytes is the digest that is digitally signed. The caller must pass an empty BLOB for this parameter.
`signed_rows_indexes`	This parameter specifies the rows in the blockchain table that were digitally signed.
`schema_certificate_guid`	A unique identifier for the certificate of the blockchain table owner stored in the database that may be used to verify the digital signature.
`signature_algo`	The digital signature algorithm to use. The parameter must be one of the following package constants: • `SIGN_ALGO_RSA_SHA2_256` • `SIGN_ALGO_RSA_SHA2_384` • `SIGN_ALGO_RSA_SHA2_512`

Usage Notes

- The database computes the signature on `signed_bytes` using the PKI private key of blockchain table owner.
- The certificate of blockchain table owner must be added to the database using `DBMS_USER_CERTS.ADD_CERTIFICATE`.

- The PKI private key and certificate of blockchain table owner must exist in a wallet located under the `<WALLET_ROOT>/bctable/` directory for a non-container database.
- The PKI private key and certificate of blockchain table owner must exist in a wallet located under the `<WALLET_ROOT>/pdb_guid/bctable/` directory for a container database.
- A blockchain table digest created by the `GET_SIGNED_BLOCKCHAIN_DIGEST` function has table information specific to a pluggable database. Such a digest can be used only in the pluggable database in which the digest was created and only for the table that was used to create the digest. For `DBMS_BLOCKCHAIN_TABLE.VERIFY_TABLE_BLOCKCHAIN`, these requirements mean that both blockchain table digests must have been generated in the current pluggable database for the same blockchain table.

  For example, suppose you create a digest for a blockchain table in pluggable database A, use Data Pump to export the blockchain table, and use Data Pump to import the blockchain table into pluggable database B. The blockchain table digest created in pluggable database A cannot be used in pluggable database B. You need to create a new blockchain table digest in pluggable database B.

Note:

The `bctable` subdirectory is the name of a database component that uses wallets. It is not the name of a blockchain table.

See Also:

Oracle Database Administrator's Guide

42.3.12 GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function

This function generates and returns the signed digest for user-specified rows in a blockchain table using the table owner's private key stored in the database wallet. The particular rows in the digest are specified by the `row_selector` parameter. The `signed_bytes`, `signed_row_indexes`, and `schema_certificate_guid` are also returned.

Note:

Rows inserted into a blockchain table that have not been committed are not guaranteed to be durable. Therefore, a signed blockchain digest does not include inserted rows that have not been committed.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS(
   schema_name               IN     VARCHAR2,
   table_name                   IN     VARCHAR2,
   row_selector              IN     VARCHAR2,
   signed_bytes              IN OUT BLOB,
   row_data_bytes            IN OUT BLOB,
```

```
   signed_rows_indexes        OUT     SYS.ORABCTAB_ROW_ARRAY_T,
   schema_certificate_guid    OUT     RAW,
   signature_algo             IN      NUMBER DEFAULT
DBMS_BLOCKCHAIN_TABLE.SIGN_ALGO_DEFAULT)
   RETURN RAW;
```

Parameters

Table 42-13 GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
row_selector	The condition, a where clause without the WHERE keyword.
signed_bytes	The BLOB value that contains a header followed by an array of row-info. This sequence of bytes is the digest that is digitally signed. The caller must pass an empty BLOB for this parameter.
row_data_bytes	This parameter specifies the content of the rows in the blockchain table that were selected for the digest. The caller must pass an empty BLOB for this parameter.
signed_rows_indexes	This parameter specifies the rows in the blockchain table that were digitally signed.
schema_certificate_guid	A unique identifier for the certificate of the blockchain table owner stored in the database that may be used to verify the digital signature.
signature_algo	The algorithm used to create the digital signature. The algorithm must be one of the following constants defined in the DBMS_BLOCKCHAIN_TABLE package: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512

Usage Notes

- The database computes the signature on signed_bytes using the PKI private key of blockchain table owner.
- The certificate of blockchain table owner must be added to the database using DBMS_USER_CERTS.ADD_CERTIFICATE.
- The PKI private key and certificate of blockchain table owner must exist in a wallet located under the <WALLET_ROOT>/bctable/ directory for a non-container database.
- The PKI private key and certificate of blockchain table owner must exist in a wallet located under the <WALLET_ROOT>/pdb_guid/bctable/ directory for a container database.
- A blockchain table digest created by the GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS function has table information specific to a pluggable database. Such a digest can be used only in the pluggable database in which the digest was created and only for the table that was used to create the digest. For DBMS_BLOCKCHAIN_TABLE.VERIFY_TABLE_BLOCKCHAIN, these requirements mean that both blockchain table digests must have been generated in the current pluggable database for the same blockchain table.

  For example, suppose you create a digest for a blockchain table in pluggable database A, use Data Pump to export the blockchain table, and use Data Pump to import the blockchain table into pluggable database B. The blockchain table digest created in

pluggable database A cannot be used in pluggable database B. You need to create a new blockchain table digest in pluggable database B.

> **Note:**
>
> The `bctable` subdirectory is the name of a database component that uses wallets. It is not the name of a blockchain table.

> **See Also:**
>
> *Oracle Database Administrator's Guide*

42.3.13 SIGN_ROW Procedure

This procedure can be used by the current user to provide a signature on row content of a previously inserted row. The transaction that inserted the row into the blockchain table must have committed before the `SIGN_ROW` procedure is called.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.SIGN_ROW(
   schema_name               IN VARCHAR2,
   table_name                IN VARCHAR2,
   instance_id               IN NUMBER,
   chain_id               IN NUMBER,
   sequence_id               IN NUMBER,
   hash                      IN RAW DEFAULT NULL,
   signature                 IN RAW,
   certificate_guid          IN RAW,
   signature_algo            IN NUMBER,
   delegate                  IN BOOLEAN DEFAULT FALSE,
   pdb_guid                  IN RAW DEFAULT NULL);
```

Parameters

Table 42-14 SIGN_ROW Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`instance_id`	The instance on which the row was inserted.
`chain_id`	The chain containing the row to be signed. There are 32 chains in each instance, and they are numbered from 0 to 31.
`sequence_id`	The position of the row on the chain. Valid values are `1`, `2`, and so on.
`hash`	If `non-NULL`, the expected value of the hash in the row to be signed. If `NULL`, the hash in the row to be signed is not checked.
`signature`	The user's digital signature on the hash value stored in the row.
`certificate_guid`	A unique identifier for the certificate stored in the database that may be used to verify the digital signature.

Table 42-14 (Cont.) SIGN_ROW Procedure Parameters

Parameter	Description
signature_algo	The algorithm used to create the digital signature. The algorithm must be one of the following constants defined in the DBMS_BLOCKCHAIN_TABLE package: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
delegate	If TRUE, then the row is being signed by a delegate. If FALSE, then the row is being signed by the user that inserted the row.
pdb_guid	The identifier of the local pluggable database. This parameter is used by Oracle GoldenGate replication and must be NULL.

> **Note:**
>
> For information on hidden columns in blockchain tables, see Hidden Columns in Blockchain Tables

Usage Notes

- The database will verify that:
  - if delegate is FALSE, the current user's obj# matches the value in the hidden column ORABCTAB_USER_NUMBER$ (ensures that the user owns the row)
  - if delegate is TRUE and the hidden column ORABCTAB_DELEGATE_USER_NUMBER$ is not NULL, the current user's obj# matches the value in this column
  - if delegate is TRUE, the current user has SIGN privilege on the blockchain table
  - the hash (if provided) matches the hash column content for the row
  - the signature column value for the specific row identified by 'instance_id', 'chain_id', and 'sequence_id' is NULL
  - if the verification succeeds, the signature value is stored for the row.
  - The SIGN_ROW procedure depends on information specific to a pluggable database and is applicable only to rows that were inserted in the current pluggable database by users, applications, or utilities other than Data Pump.

    For example, suppose you insert a row into a blockchain table in pluggable database A, commit the INSERT transaction, use Data Pump to export the blockchain table, and use Data Pump to import the blockchain table into pluggable database B. If you try to sign this row in pluggable database B, DBMS_BLOCKCHAIN_TABLE.SIGN_ROW will raise an exception. Hence you should sign all rows in a blockchain table that need to be signed before using Data Pump to create a copy of the blockchain table.

42.3.14 SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure

This procedure allows an end user or a delegate to sign a row using a set of at most three user columns that uniquely identify a single row.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS(
    schema_name      IN      VARCHAR2,
    table_name       IN      VARCHAR2,
    hash             IN      RAW DEFAULT NULL,
    signature        IN      RAW,
    certificate_guid IN      RAW,
    signature_algo   IN      NUMBER,
    delegate         IN      BOOLEAN DEFAULT FALSE,
    keycol1_name     IN      VARCHAR2,
    keycol1_value    IN      VARCHAR2,
    keycol2_name     IN      VARCHAR2 DEFAULT NULL,
    keycol2_value    IN      VARCHAR2 DEFAULT NULL,
    keycol3_name     IN      VARCHAR2 DEFAULT NULL,
    keycol3_value    IN      VARCHAR2 DEFAULT NULL,
    pdb_guid         IN      RAW DEFAULT NULL);
```

Parameters

Table 42-15 SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
hash	If non-NULL, the expected value of the hash in the row to be signed. If NULL, the hash in the row to be signed is not checked.
signature	The user's digital signature on the hash value stored in the row.
certificate_guid	A unique identifier for the certificate stored in the database that may be used to verify the digital signature.
signature_algo	The algorithm used to create the digital signature. The algorithm must be one of the following constants defined in the DBMS_BLOCKCHAIN_TABLE package: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
delegate	If TRUE, then the row is being signed by a delegate. If FALSE, then the row is being signed by the user that inserted the row.
keycol1_name	The name of the key column.
keycol1_value	The value of the key column.
keycol2_name	The name of the second column in a composite key.
keycol2_value	The value of the second column in a composite key.
keycol3_name	The name of the third column in a composite key.
keycol3_value	The value of the third column in a composite key.
pdb_guid	The identifier of the local pluggable database. This parameter is used by Oracle GoldenGate replication and must be NULL.

Usage Notes

Only `non-NULL` values can be used for `keycol1_value`, `keycol2_value`, and `keycol3_value`. If a `NULL` value is specified for `keycol2_value` or `keycol3_value`, then the corresponding parameter `keycol2_name` or `keycol3_name` must be `NULL`. Any other combination raises an exception.

42.3.15 SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE Procedure

This procedure uses at most three user column names and values to uniquely identify a single row to sign and countersign. A row in a blockchain table can be countersigned only if the row belongs to the current epoch for the blockchain table.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE (
   schema_name                       IN      VARCHAR2,
   table_name                        IN      VARCHAR2,
   hash                              IN      RAW DEFAULT NULL,
   signature                         IN      RAW,
   certificate_guid                  IN      RAW,
   signature_algo                    IN      NUMBER,
   delegate                          IN      BOOLEAN DEFAULT FALSE,
   keycol1_name                      IN      VARCHAR2,
   keycol1_value                     IN      VARCHAR2,
   keycol2_name                      IN      VARCHAR2 DEFAULT NULL,
   keycol2_value                     IN      VARCHAR2 DEFAULT NULL,
   keycol3_name                      IN      VARCHAR2 DEFAULT NULL,
   keycol3_value                     IN      VARCHAR2 DEFAULT NULL,
   countersignature_algo             IN      NUMBER DEFAULT
DBMS_BLOCKCHAIN_TABLE.SIGN_ALGO_DEFAULT,
   countersignature_signed_bytes     IN      OUT BLOB,
   countersignature                  OUT     RAW,
   countersignature_certificate_guid OUT     RAW,
   countersignature_content_version  IN      VARCHAR2 DEFAULT 'V2_DIGEST',
   pdb_guid                          IN      RAW DEFAULT NULL);
```

Parameters

Table 42-16 SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the blockchain table.
`hash`	If `non-NULL`, the expected value of the hash in the row to be signed. If `NULL`, the hash in the row to be signed is not checked
`signature`	The user's digital signature on the hash value stored in the row.
`certificate_guid`	A unique identifier for the certificate stored in the database that may be used to verify the digital signature.

Table 42-16 (Cont.) SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE Procedure Parameters

Parameter	Description
signature_algo	The algorithm used to create the digital signature. The algorithm must be one of the following constants defined in the DBMS_BLOCKCHAIN_TABLE package: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
delegate	If TRUE, then the row is being signed by a delegate. If FALSE, then the row is being signed by the user that inserted the row.
keycol1_name	The name of the key column.
keycol1_value	The value of the key column.
keycol2_name	The name of the second column in a composite key.
keycol2_value	The value of the second column in a composite key.
keycol3_name	The name of the third column in a composite key.
keycol3_value	The value of the third column in a composite key.
countersignature_algo	The digital signature algorithm for the countersignature. The parameter must be one of the following package constants: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
countersignature_signed_bytes	The bytes that are input to the algorithm that generates the countersignature. The caller must pass an empty BLOB for this parameter.
countersignature	The digital signature on the bytes returned in countersignature_signed_bytes.
countersignature_certificate_guid	A unique identifier for the certificate of the blockchain table owner stored in the database that may be used to verify the countersignature.
countersignature_content_version	The version of the data contents and layout that are used as input to the countersignature algorithm. Only 'V2_DIGEST' is supported in this release.
pdb_guid	The identifier of the local pluggable database. This parameter is used by Oracle GoldenGate replication and must be NULL.

42.3.16 SIGN_ROW_WITH_COUNTERSIGNATURE Procedure

It is an extension of SIGN_ROW that enables the user to request a countersignature from the database. The countersignature will be produced by signing the row data content using the table owner's private key stored in the database wallet. A row in a blockchain table can be countersigned only if the row belongs to the current epoch for the blockchain table.

If delegate is TRUE, then it is the delegate user that is requesting the countersignature.

The countersignature is stored in the row along with all the meta-information needed to recompute the signed bytes used in its computation.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.SIGN_ROW_WITH_COUNTERSIGNATURE(
   schema_name                          IN VARCHAR2,
   table_name                           IN VARCHAR2,
```

```
   instance_id                       IN NUMBER,
   chain_id                          IN NUMBER,
   sequence_id                       IN NUMBER,
   hash                              IN RAW DEFAULT NULL,
   signature                         IN RAW,
   certificate_guid                  IN RAW,
   signature_algo                    IN NUMBER,
   delegate                          IN BOOLEAN DEFAULT FALSE,
   countersignature_algo             IN NUMBER DEFAULT
DBMS_BLOCKCHAIN_TABLE.SIGN_ALGO_DEFAULT,
   countersignature_signed_bytes     IN OUT BLOB,
   countersignature                  OUT RAW,
   countersignature_certificate_guid OUT RAW,
   countersignature_content_version  IN VARCHAR2 DEFAULT 'V2_DIGEST',
   pdb_guid                          IN RAW DEFAULT NULL);
```

Parameters

Table 42-17 SIGN_ROW_WITH_COUNTERSIGNATURE Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
instance_id	The instance on which the row was inserted.
chain_id	The chain containing the row to be signed. By default, there are 32 chains in each instance, and they are numbered from 0 to 31.
sequence_id	The position of the row on the chain. Valid values are 1, 2, and so on.
hash	If non-NULL, the expected value of the hash in the row to be signed. If NULL, the hash in the row to be signed is not checked.
signature	The user's digital signature on the hash value stored in the row.
certificate_guid	A unique identifier for the certificate stored in the database that may be used to verify the digital signature.
signature_algo	The algorithm used to create the digital signature. The algorithm must be one of the following constants defined in the DBMS_BLOCKCHAIN_TABLE package: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
delegate	If TRUE, then the row is being signed by a delegate. If FALSE, then the row is being signed by the user that inserted the row.
countersignature_algo	The digital signature algorithm for the countersignature. The parameter must be one of the following package constants: • SIGN_ALGO_RSA_SHA2_256 • SIGN_ALGO_RSA_SHA2_384 • SIGN_ALGO_RSA_SHA2_512
countersignature_signed_bytes	The bytes that are input to the algorithm that generates the countersignature. The caller must pass an empty BLOB for this parameter.
countersignature	The digital signature on the bytes returned in countersignature_signed_bytes.
countersignature_certificate_guid	A unique identifier for the certificate of the blockchain table owner stored in the database that may be used to verify the countersignature.

Table 42-17 (Cont.) SIGN_ROW_WITH_COUNTERSIGNATURE Procedure Parameters

Parameter	Description
countersignature_content_version	The version of the data contents and layout that are used as input to the countersignature algorithm. Only 'V2_DIGEST' is supported in this release.
pdb_guid	The identifier of the local pluggable database. This parameter is used by Oracle GoldenGate replication and must be NULL.

Usage Notes

- SYS can always countersign any row of a blockchain table.
- An end-user or a delegate that signs and procures a countersignature using SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_COUNTERSIGNATURE or SIGN_ROW_WITH_COUNTERSIGNATURE does not need any additional privileges to procure a countersignature.

42.3.17 VERIFY_ROWS Procedure

Verifies all rows on all applicable system chains for integrity of HASH column value and optionally the SIGNATURE column value for rows created in the range of low_timestamp to high_timestamp. An appropriate exception is thrown if the integrity of chains is compromised.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.VERIFY_ROWS(
   schema_name                  IN VARCHAR2,
   table_name                   IN VARCHAR2,
   low_timestamp                IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   high_timestamp               IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   instance_id                  IN NUMBER DEFAULT NULL,
   chain_id                     IN NUMBER DEFAULT NULL,
   number_of_rows_verified      OUT NUMBER,
   verify_signature             IN BOOLEAN DEFAULT TRUE,
   verify_delegate_signature    IN BOOLEAN DEFAULT TRUE,
   verify_countersignature      IN BOOLEAN DEFAULT TRUE,
   pdb_guid                     IN RAW DEFAULT NULL);
```

Parameters

Table 42-18 VERIFY_ROWS Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
low_timestamp	If specified, the low end of the time range for verifying rows. The default value is NULL.
high_timestamp	If specified, the high end of the time range for verifying rows. The default value is NULL.
instance_id	If specified, restricts row verification to rows inserted on the specified instance.
chain_id	If specified, restricts row verification to rows on the specified chain. By default, there are 32 chains in each instance, and they are numbered from 0 to 31.

Table 42-18 (Cont.) VERIFY_ROWS Procedure Parameters

Parameter	Description
number_of_rows_verified	The number of rows verified.
verify_signature	If verify_signature is TRUE, both the hash on each row and any user signature on the row are verified. If verify_signature is FALSE, user signatures are not verified.
verify_delegate_signature	If verify_delegate_signature is TRUE, both the hash on each row and any delegate signature on the row are verified. If verify_delegate_signature is FALSE, delegate signatures are not verified.
verify_countersignature	If verify_countersignature is TRUE, both the hash on each row and any countersignature on the row are verified. If verify_countersignature is FALSE, countersignatures are not verified.
pdb_guid	For a V2 blockchain table, if not NULL, restricts attention to system chains inserted by the specified pluggable database. Must be NULL for a V1 blockchain table.

Usage Notes

- schema_name and table_name are the required input parameters.
- All other input parameters are optional, with the following exception:
  - If chain_id is specified, instance_id must be specified.
- Valid values for instance_id are 1, 2, ... etc.
- If neither instance_id, nor chain_id is specified, then it implies all chains. If only instance_id is specified, then it implies all chains on that instance. If both are specified, then it implies the specific chain provided by the combination.
- If both low_timestamp and high_timestamp are specified, then high_timestamp must be later than low_timestamp.

  If low_timestamp is not specified, then the range is the oldest row in the blockchain table to high_timestamp.

  If high_timestamp is not specified, then the range is low_timestamp to the timestamp of the last row inserted in the table.

42.3.18 VERIFY_TABLE_BLOCKCHAIN Procedure

This procedure verifies signatures and system chains for all rows whose creation-times fall between the minimum value for the row-creation time from bytes_previous and the maximum value for row-creation time from bytes_latest. The OUT parameter number_of_rows_verified returns the number of successfully verified rows.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.VERIFY_TABLE_BLOCKCHAIN(
  bytes_latest              IN    BLOB      DEFAULT NULL,
  bytes_previous            IN    BLOB      DEFAULT NULL,
  number_of_rows_verified   OUT   NUMBER,
  verify_signature          IN    BOOLEAN   DEFAULT TRUE,
  verify_delegate_signature IN    BOOLEAN   DEFAULT TRUE,
  verify_countersignature   IN    BOOLEAN   DEFAULT TRUE,
```

```
  signed_bytes_latest       IN    BLOB      DEFAULT NULL,
  signed_bytes_previous     IN    BLOB      DEFAULT NULL);
```

Parameters

Table 42-19 VERIFY_TABLE_BLOCKCHAIN Procedure Parameters

Parameter	Description
bytes_latest	A digest populated by a call to GET_SIGNED_BLOCKCHAIN_DIGEST, GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS, GET_BLOCKCHAIN_DIGEST, or GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS.
bytes_previous	A digest populated by a call to GET_SIGNED_BLOCKCHAIN_DIGEST, GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS, GET_BLOCKCHAIN_DIGEST, or GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS before the bytes_latest BLOB was populated.
number_of_rows_verified	The count of the rows in the blockchain table that were verified.
verify_signature	If verify_signature is TRUE, both the hash on each row and any user signature on the row are verified. If verify_signature is FALSE, user signatures are not verified.
verify_delegate_signature	If verify_delegate_signature is TRUE, both the hash on each row and any delegate signature on the row are verified. If verify_delegate_signature is FALSE, delegate signatures are not verified.
verify_countersignature	If verify_countersignature is TRUE, both the hash on each row and any countersignature on the row are verified. If verify_countersignature is FALSE, countersignatures are not verified.
signed_bytes_latest	signed_bytes_latest has been deprecated.
signed_bytes_previous	signed_bytes_previous has been deprecated.

Usage Notes

The BLOBs in bytes_latest and bytes_previous must be associated with the same blockchain table. For GET_SIGNED_BLOCKCHAIN_DIGEST and GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS, the relevant BLOB parameter is called signed_bytes. For GET_BLOCKCHAIN_DIGEST and GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS, the relevant BLOB parameter is called digest_bytes.

42.3.19 VERIFY_USER_BLOCKCHAIN_ROWS Procedure

This procedure verifies rows of one or more user chains when the user chains feature is enabled for a blockchain table.

Syntax

```
DBMS_BLOCKCHAIN_TABLE.verify_user_blockchain_rows (
   schema_name               IN    VARCHAR2,
   table_name                IN    VARCHAR2,
   row_version_name          IN    VARCHAR2,
   number_of_rows_verified   OUT   NUMBER,
   keycol1_value             IN    VARCHAR2 DEFAULT NULL,
   keycol2_value             IN    VARCHAR2 DEFAULT NULL,
```

```
keycol3_value             IN    VARCHAR2 DEFAULT NULL,
low_timestamp             IN    TIMESTAMP WITH TIME ZONE DEFAULT NULL,
high_timestamp            IN    TIMESTAMP WITH TIME ZONE DEFAULT NULL,
verify_signature          IN    BOOLEAN DEFAULT TRUE,
verify_delegate_signature IN    BOOLEAN DEFAULT TRUE,
verify_countersignature   IN    BOOLEAN DEFAULT TRUE,
pdb_guid                  IN    RAW DEFAULT NULL);
```

Parameters

Table 42-20 VERIFY_USER_BLOCKCHAIN_ROWS Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the blockchain table.
row_version_name	The name of the row version given when the blockchain table was created.
number_of_rows_verified	The number of rows verified.
keycol1_value	The value of the key column.
keycol2_value	The value of the second column in a composite key.
keycol3_value	The value of the third column in a composite key.
low_timestamp	If specified, the low end of the time range for verifying rows. The default value is NULL.
high_timestamp	If specified, the high end of the time range for verifying rows. The default value is NULL
verify_signature	If verify_signature is TRUE, both the hash on each row and any user signature on the row are verified. If verify_signature is FALSE, user signatures are not verified.
verify_delegate_signature	If verify_delegate_signature is TRUE, both the hash on each row and any delegate signature on the row are verified. If verify_delegate_signature is FALSE, delegate signatures are not verified.
verify_countersignature	If verify_countersignature is TRUE, both the hash on each row and any countersignature on the row are verified. If verify_countersignature is FALSE, countersignatures are not verified.
pdb_guid	For a V2 blockchain table, if not NULL, restricts attention to user chains inserted by the specified pluggable database. Must be NULL for a V1 blockchain table.

Usage Notes

The parameters keycol1_value, keycol2_value, and keycol3_value may be used to limit the user chains that are verified. If all three of these parameters are NULL, all user chains are verified. If one or more of these parameters are not NULL, then only user chains matching the non-NULL parameters are verified.

43
DBMS_BLOCKER_RESOLVER

The DBMS_BLOCKER_RESOLVER package provides a method of changing some Blocker Resolver configuration parameters.

This chapter contains the following topics:

- DBMS_BLOCKER_RESOLVER Overview
- DBMS_BLOCKER_RESOLVER Security Model
- DBMS_BLOCKER_RESOLVER Constants
- DBMS_BLOCKER_RESOLVER Exceptions
- Summary of DBMS_BLOCKER_RESOLVER Subprograms

43.1 DBMS_BLOCKER_RESOLVER Overview

The `DBMS_BLOCKER_RESOLVER` package provides a method of changing some Blocker Resolver configuration parameters.

> **Note:**
>
> The `DBMS_HANG_MANAGER` package is deprecated in Oracle Database 23ai. Use `DBMS_BLOCKER_RESOLVER` instead.
> The `DBMS_HANG_MANAGER` package provides a method of changing some configuration parameters and constraints to address session issues. This package is being replaced with `DBMS_BLOCKER_RESOLVER`. `DBMS_HANG_MANAGER` can be removed in a future release.

Blocker Resolver is always enabled on all databases, Automatic Storage Management (ASM) disk groups, and single instance systems (when not using Oracle RAC). Delay resolution is enabled on all instances by default except for single instance systems (when not using Oracle RAC) where it is disabled and cannot be enabled at this time.

> **Note:**
>
> The `DBMS_BLOCKER_RESOLVER` package is intended for use by database administrators (DBA) only. It is not intended for use by application developers.

43.2 DBMS_BLOCKER_RESOLVER Security Model

`DBMS_BLOCKER_RESOLVER` can only be run as the `SYS` user.

43.3 DBMS_BLOCKER_RESOLVER Constants

The table in this topic lists DBMS_BLOCKER_RESOLVER package constants.

Table 43-1 DBMS_BLOCKER_RESOLVER Constants

Name	Type	Value	Description
RESOLUTION_SCOPE	VARCHAR2	resolution scope	Resolution scope parameter name
SENSITIVITY	VARCHAR2	sensitivity	Sensitivity parameter name
BASE_FILE_SIZE_LIMIT	VARCHAR2	base file size limit	Base file size limit parameter name. The default value of this parameter is 100 MB.
BASE_FILE_SET_COUNT	VARCHAR2	base file set count	Base file set count parameter name. The default value is 5 trace files in the circular trace file set.
LWS_FILE_SIZE_LIMIT	VARCHAR2	long waiting session file size limit	Long waiting session file size limit parameter name. The default value of this parameter is 100 MB.
LWS_FILE_SET_COUNT	VARCHAR2	long waiting session set count	Long waiting session set count parameter name. The default value is 5 trace files in the circular trace file set.
RESOLUTION_SCOPE_PROCESS	VARCHAR2	PROCESS	Enables only session and/or process termination
RESOLUTION_SCOPE_INSTANCE	VARCHAR2	INSTANCE	Enables session, process, and instance termination. The latter is used in specific scenarios. The default value is INSTANCE.
SENSITIVITY_NORMAL	VARCHAR2	NORMAL	Sets various Blocker Resolver intervals to their default values. The default value is NORMAL.
SENSITIVITY_HIGH	VARCHAR2	HIGH	Sets the various Blocker Resolver intervals to half of their default values

43.4 DBMS_BLOCKER_RESOLVER Exceptions

The table in this topic lists the exceptions raised by the `DBMS_BLOCKER_RESOLVER` package.

Table 43-2 DBMS_BLOCKER_RESOLVER Package Exceptions

Exception	Error Code	Description
`ERRNUM_INPUT_ERROR`	`-32706`	Invalid user input
`ERRNUM_INTERNAL_ERROR`	`-32707`	Database experienced an error setting a parameter
`ERRNUM_UNSUPPORTED_ERROR`	`-32708`	Operation is not supported on this instance type

43.5 Summary of DBMS_BLOCKER_RESOLVER Subprograms

The `DBMS_BLOCKER_RESOLVER` package includes the `SET` procedure subprogram.

Table 43-3 DBMS_BLOCKER_RESOLVER Package Subprograms

Subprogram	Description
SET Procedure	Sets the specified parameter to the specified value

43.5.1 SET Procedure

This procedure sets the specified parameters to specified values.

Syntax

```
DBMS_BLOCKER_RESOLVER.SET (
   pname   IN  VARCHAR2,
   pvalue  IN  VARCHAR2);
```

Parameters

Table 43-4 SET Procedure Parameters

Parameter	Description
`pname`	Contains the parameter name that needs to be modified.
`pvalue`	Contains the new value for the parameter name specified in `pname`.

44
DBMS_CAPTURE_ADM

The `DBMS_CAPTURE_ADM` package, one of a set of Oracle Replication packages, provides subprograms for starting, stopping, and configuring a capture process. The source of the captured changes is the redo logs, and the repository for the captured changes is a queue.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_CAPTURE_ADM Subprograms

44.1 DBMS_CAPTURE_ADM Overview

This package provides interfaces to start, stop, and configure a capture process or a synchronous capture. This package includes subprograms for preparing database objects for instantiation.

Capture processes can be used in an XStream configuration in a multitenant container database (CDB). A CDB is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs).

> **See Also:**
>
> *Oracle Database Concepts* for more information about CDBs and PDBs

44.2 DBMS_CAPTURE_ADM Security Model

The DBMS_CAPTURE_ADM security can be controlled in one of two ways.

- Granting `EXECUTE` on this package to selected users or roles.
- Granting `EXECUTE_CATALOG_ROLE` to selected users or roles.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

When the `DBMS_CAPTURE_ADM` package is used to manage an Oracle Replication configuration, it requires that the user is granted the privileges of an Oracle Replication administrator.

When the `DBMS_CAPTURE_ADM` package is used to manage an XStream configuration, it requires that the user is granted the privileges of an XStream administrator.

Note:

- The user must be granted additional privileges to perform some administrative tasks using the subprograms in this package, such as setting a capture user. If additional privileges are required for a subprogram, then the privileges are documented in the section that describes the subprogram.
- Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

See Also:

Oracle Database XStream Guide for information about configuring an XStream administrator

44.3 Summary of DBMS_CAPTURE_ADM Subprograms

This table lists the `DBMS_CAPTURE_ADM` subprograms and briefly describes them.

Table 44-1 DBMS_CAPTURE_ADM Package Subprograms

Subprogram	Description
ABORT_GLOBAL_INSTANTIATION Procedure	Reverses the effects of running the `PREPARE_GLOBAL_INSTANTIATION`, `PREPARE_SCHEMA_INSTANTIATION`, and `PREPARE_TABLE_INSTANTIATION` procedures
ABORT_SCHEMA_INSTANTIATION Procedure	Reverses the effects of running the `PREPARE_SCHEMA_INSTANTIATION` and `PREPARE_TABLE_INSTANTIATION` procedures
ABORT_SYNC_INSTANTIATION Procedure	Reverses the effects of running the `PREPARE_SYNC_INSTANTIATION` procedure
ABORT_TABLE_INSTANTIATION Procedure	Reverses the effects of running the `PREPARE_TABLE_INSTANTIATION` procedure
ALTER_CAPTURE Procedure	Alters a capture process
ALTER_SYNC_CAPTURE Procedure	Alters a synchronous capture

Table 44-1 (Cont.) DBMS_CAPTURE_ADM Package Subprograms

Subprogram	Description
BUILD Procedure	Extracts the data dictionary of the current database to the redo logs and automatically specifies database supplemental logging for all primary key and unique key columns
CREATE_CAPTURE Procedure	Creates a capture process
CREATE_SYNC_CAPTURE Procedure	Creates a synchronous capture
DROP_CAPTURE Procedure	Drops a capture process
INCLUDE_EXTRA_ATTRIBUTE Procedure	Includes or excludes an extra attribute in logical change records (LCRs) captured by the specified capture process or synchronous capture
PREPARE_GLOBAL_INSTANTIATION Procedure	Performs the synchronization necessary for instantiating all the tables in the database at another database and can enable supplemental logging for key columns or all columns in these tables
PREPARE_SCHEMA_INSTANTIATION Procedure	Performs the synchronization necessary for instantiating all tables in the schema at another database and can enable supplemental logging for key columns or all columns in these tables
PREPARE_SYNC_INSTANTIATION Function	Performs the synchronization necessary for instantiating one or more tables at another database and returns the prepare SCN
PREPARE_TABLE_INSTANTIATION Procedure	Performs the synchronization necessary for instantiating the table at another database and can enable supplemental logging for key columns or all columns in the table
SET_PARAMETER Procedure	Sets a capture process parameter to the specified value
START_CAPTURE Procedure	Starts the capture process, which mines redo logs and enqueues the mined redo information into the associated queue
STOP_CAPTURE Procedure	Stops the capture process from mining redo logs

> **Note:**
>
> All subprograms commit unless specified otherwise.

44.3.1 ABORT_GLOBAL_INSTANTIATION Procedure

This procedure reverses the effects of running the `PREPARE_GLOBAL_INSTANTIATION`, `PREPARE_SCHEMA_INSTANTIATION`, and `PREPARE_TABLE_INSTANTIATION` procedures.

Specifically, this procedure performs the following actions:

- Removes data dictionary information related to the database, schema, and table instantiations
- Removes any supplemental logging enabled by the `PREPARE_GLOBAL_INSTANTIATION`, `PREPARE_SCHEMA_INSTANTIATION`, and `PREPARE_TABLE_INSTANTIATION` procedures

Syntax

```
DBMS_CAPTURE_ADM.ABORT_GLOBAL_INSTANTIATION(
   container    IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameter

Table 44-2 ABORT_GLOBAL_INSTANTIATION Procedure Parameter

Parameter	Description
container	Either CURRENT , ALL, or *pdb_name*. If CURRENT is specified, then this procedure removes supplemental logging for the current container. If ALL is specified, then this procedure removes supplemental logging for all of the containers in the current CDB. If *pdb_name* is specified, then this procedure removes supplemental logging for the specified PDB. ALL and *pdb_name* are valid only when you invoke the procedure from the root.

44.3.2 ABORT_SCHEMA_INSTANTIATION Procedure

This procedure reverses the effects of running the PREPARE_SCHEMA_INSTANTIATION procedure. It also reverses the effects of running the PREPARE_TABLE_INSTANTIATION procedure on tables in the specified schema.

Specifically, this procedure performs the following actions:

- Removes data dictionary information related to schema instantiations and table instantiations of tables in the schema
- Removes any supplemental logging enabled by the PREPARE_SCHEMA_INSTANTIATION procedure
- Removes any supplemental logging enabled by the PREPARE_TABLE_INSTANTIATION procedure for tables in the specified schema

Syntax

```
DBMS_CAPTURE_ADM.ABORT_SCHEMA_INSTANTIATION(
   schema_name  IN  VARCHAR2,
   container    IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameter

Table 44-3 ABORT_SCHEMA_INSTANTIATION Procedure Parameter

Parameter	Description
schema_name	The name of the schema for which to abort the effects of preparing instantiation

Table 44-3 (Cont.) ABORT_SCHEMA_INSTANTIATION Procedure Parameter

Parameter	Description
container	Either CURRENT , ALL, or *pdb_name*. If CURRENT is specified, then this procedure removes supplemental logging for the current container. If ALL is specified, then this procedure removes supplemental logging for all of the containers in the current CDB. If *pdb_name* is specified, then this procedure removes supplemental logging for the specified PDB. ALL and *pdb_name* are valid only when you invoke the procedure from the root.

44.3.3 ABORT_SYNC_INSTANTIATION Procedure

This procedure reverses the effects of running the PREPARE_SYNC_INSTANTIATION procedure. Specifically, this procedure removes data dictionary information related to the table instantiation.

This procedure is overloaded. The table_names parameter is VARCHAR2 datatype in one version and DBMS_UTILITY.UNCL_ARRAY datatype in the other version.

Syntax

```
DBMS_CAPTURE_ADM.ABORT_SYNC_INSTANTIATION(
   table_names IN  VARCHAR2);

DBMS_CAPTURE_ADM.ABORT_SYNC_INSTANTIATION(
   table_names IN  DBMS_UTILITY.UNCL_ARRAY);
```

Parameters

Table 44-4 ABORT_SYNC_INSTANTIATION Procedure Parameter

Parameter	Description
table_names	When the table_names parameter is VARCHAR2 datatype, a comma-delimited list of the tables for which to abort the effects of preparing instantiation. There must be no spaces between entries. When the table_names parameter is DBMS_UTILITY.UNCL_ARRAY datatype, specify a PL/SQL associative array of this type that contains the names of the tables for which to abort the effects of preparing instantiation. The first table name is at position 1, the second at position 2, and so on. The table does not need to be NULL terminated. In either version of the procedure, specify the name of each table in the form [*schema_name.*]*table_name*. For example, hr.employees. If the schema is not specified, then the current user is the default.

44.3.4 ABORT_TABLE_INSTANTIATION Procedure

This procedure reverses the effects of running the PREPARE_TABLE_INSTANTIATION procedure.

Specifically, this procedure performs the following actions:

- Removes data dictionary information related to the table instantiation

- Removes any supplemental logging enabled by the `PREPARE_TABLE_INSTANTIATION` procedure

Syntax

```
DBMS_CAPTURE_ADM.ABORT_TABLE_INSTANTIATION(
   table_name  IN  VARCHAR2);
   container   IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameter

Table 44-5 ABORT_TABLE_INSTANTIATION Procedure Parameter

Parameter	Description
`table_name`	The name of the table for which to abort the effects of preparing instantiation, specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default.
`container`	Either `CURRENT` , `ALL`, or *`pdb_name`*. If `CURRENT` is specified, then this procedure removes supplemental logging for the current container. If `ALL` is specified, then this procedure removes supplemental logging for all of the containers in the current CDB. If *`pdb_name`* is specified, then this procedure removes supplemental logging for the specified PDB. `ALL` and *`pdb_name`* are valid only when you invoke the procedure from the root.

44.3.5 ALTER_CAPTURE Procedure

This procedure alters a capture process.

Syntax

```
DBMS_CAPTURE_ADM.ALTER_CAPTURE(
   capture_name               IN  VARCHAR2,
   rule_set_name              IN  VARCHAR2   DEFAULT NULL,
   remove_rule_set            IN  BOOLEAN    DEFAULT FALSE,
   start_scn                  IN  NUMBER     DEFAULT NULL,
   use_database_link          IN  BOOLEAN    DEFAULT NULL,
   first_scn                  IN  NUMBER     DEFAULT NULL,
   negative_rule_set_name     IN  VARCHAR2   DEFAULT NULL,
   remove_negative_rule_set   IN  BOOLEAN    DEFAULT FALSE,
   capture_user               IN  VARCHAR2   DEFAULT NULL,
   checkpoint_retention_time  IN  NUMBER     DEFAULT NULL,
   start_time                 IN  TIMESTAMP  DEFAULT NULL,
   oldest_scn                 IN  NUMBER     DEFAULT NULL);
```

Parameters

Table 44-6 ALTER_CAPTURE Procedure Parameters

Parameter	Description
`capture_name`	The name of the capture process being altered. You must specify an existing capture process name. Do not specify an owner.

Table 44-6 (Cont.) ALTER_CAPTURE Procedure Parameters

Parameter	Description
rule_set_name	The name of the positive rule set for the capture process. The positive rule set contains the rules that instruct the capture process to capture changes. To change the positive rule set for the capture process, specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a positive rule set in the hr schema named job_capture_rules, enter hr.job_capture_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify NULL and the remove_rule_set parameter is set to FALSE, then the procedure retains any existing positive rule set. If you specify NULL and the remove_rule_set parameter is set to TRUE, then the procedure removes any existing positive rule set.
remove_rule_set	If TRUE, then the procedure removes the positive rule set for the specified capture process. If you remove a positive rule set for a capture process, and the capture process does not have a negative rule set, then the capture process captures all supported changes to all objects in the database, excluding database objects in the SYS and SYSTEM schemas. If you remove a positive rule set for a capture process, and the capture process has a negative rule set, then the capture process captures all supported changes that are not discarded by the negative rule set. If FALSE, then the procedure retains the positive rule set for the specified capture process. If the rule_set_name parameter is non-NULL, then ensure that this parameter is set to FALSE.
start_scn	A valid SCN for the database from which the capture process starts capturing changes. The SCN value must be greater than or equal to the first SCN for the capture process. Also, the capture process must be stopped before resetting its start SCN. An error is returned if an invalid SCN is specified or if the capture process is enabled.

Table 44-6 (Cont.) ALTER_CAPTURE Procedure Parameters

Parameter	Description
use_database_link	If TRUE, then the capture process at a downstream database uses a database link to the source database for administrative purposes relating to the capture process. If you want a capture process that is not using a database link currently to begin using a database link, then specify TRUE. In this case, a database link with the same name as the global name of the source database must exist at the downstream database. If FALSE, then either the capture process is running on the source database, or the capture process at a downstream database does not use a database link to the source database. If you want a capture process that is using a database link currently to stop using a database link, then specify FALSE. In this case, you must prepare source database objects for instantiation manually when you add or change capture process rules that pertain to these objects. If NULL, then the current value of this parameter for the capture process is not changed.
first_scn	The lowest SCN in the redo log from which a capture process can capture changes. If you specify a new first SCN for the capture process, then the specified first SCN must meet the following requirements: • It must be greater than the current first SCN for the capture process. • It must be less than or equal to the current applied SCN for the capture process. However, this requirement does not apply if the current applied SCN for the capture process is zero. • It must be less than or equal to the required checkpoint SCN for the capture process. An error is returned if the specified SCN does not meet the first three requirements. See "Usage Notes" for information about determining an SCN value that meets all of these conditions. When the first SCN is modified, the capture process purges information from its LogMiner data dictionary that is required to restart it at an earlier SCN. See BUILD Procedure for more information about a LogMiner data dictionary. If the specified first SCN is higher than the current start SCN for the capture process, then the start SCN is set automatically to the new value of the first SCN.

Table 44-6 (Cont.) ALTER_CAPTURE Procedure Parameters

Parameter	Description
negative_rule_set_name	The name of the negative rule set for the capture process. The negative rule set contains the rules that instruct the capture process to discard changes. To change the negative rule set for the capture process, specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a negative rule set in the hr schema named neg_capture_rules, enter hr.neg_capture_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify NULL and the remove_negative_rule_set parameter is set to FALSE, then the procedure retains any existing negative rule set. If you specify NULL and the remove_negative_rule_set parameter is set to TRUE, then the procedure removes any existing negative rule set. If you specify both a positive and a negative rule set for a capture process, then the negative rule set is always evaluated first.
remove_negative_rule_set	If TRUE, then the procedure removes the negative rule set for the specified capture process. If you remove a negative rule set for a capture process, and the capture process does not have a positive rule set, then the capture process captures all supported changes to all objects in the database, excluding database objects in the SYS and SYSTEM schemas. If you remove a negative rule set for a capture process, and a positive rule set exists for the capture process, then the capture process captures all changes that are not discarded by the positive rule set. If FALSE, then the procedure retains the negative rule set for the specified capture process. If the negative_rule_set_name parameter is non-NULL, then ensure that this parameter is set to FALSE.

Table 44-6 (Cont.) ALTER_CAPTURE Procedure Parameters

Parameter	Description
capture_user	The user in whose security domain a capture process captures changes that satisfy its rule sets and runs custom rule-based transformations configured for capture process rules. If NULL, then the capture user is not changed. To change the capture user, the user who invokes the ALTER_CAPTURE procedure must be granted the DBA role. Only the SYS user can set the capture_user to SYS. If you change the capture user, then this procedure grants the new capture user enqueue privilege on the queue used by the capture process and configures the user as a secure queue user of the queue. In addition, ensure that the capture user has the following privileges: • EXECUTE privilege on the rule sets used by the capture process • EXECUTE privilege on all rule-based transformation functions used in the rule set These privileges can be granted directly to the capture user, or they can be granted through roles. In addition, the capture user must be granted EXECUTE privilege on all packages, including Oracle-supplied packages, that are invoked in rule-based transformations run by the capture process. These privileges must be granted directly to the capture user. They cannot be granted through roles. The capture process is stopped and restarted automatically when you change the value of this parameter. **Note:** If the capture user for a capture process is dropped using DROP USER . . . CASCADE, then the capture process is also dropped automatically.
checkpoint_retention_time	Either the number of days that a capture process retains checkpoints before purging them automatically, or DBMS_CAPTURE_ADM.INFINITE if checkpoints should not be purged automatically. If NULL, then the checkpoint retention time is not changed. If a number is specified, then a capture process purges a checkpoint the specified number of days after the checkpoint was taken. Partial days can be specified using decimal values. For example, .25 specifies 6 hours. When a checkpoint is purged, LogMiner data dictionary information for the archived redo log file that corresponds to the checkpoint is purged, and the first_scn of the capture process is reset to the SCN value corresponding to the first change in the next archived redo log file.
start_time	A valid time from which the capture process starts capturing changes. The capture process must be stopped before resetting its start time. An error is returned if an invalid time is specified or if the capture process is enabled. The start_scn and start_time parameters are mutually exclusive.
oldest_scn	The oldest SCN of the transactions currently being processed.

Usage Notes

If you want to alter the first SCN for a capture process, then the value specified must meet the conditions in the description for the `first_scn` parameter.

Examples

The following query determines the current first SCN, applied SCN, and required checkpoint SCN for each capture process in a database:

```
SELECT CAPTURE_NAME, FIRST_SCN, APPLIED_SCN, REQUIRED_CHECKPOINT_SCN
   FROM DBA_CAPTURE;
```

44.3.6 ALTER_SYNC_CAPTURE Procedure

This procedure alters a synchronous capture.

Syntax

```
DBMS_CAPTURE_ADM.ALTER_SYNC_CAPTURE(
   capture_name    IN  VARCHAR2,
   rule_set_name   IN  VARCHAR2  DEFAULT NULL,
   capture_user    IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 44-7 ALTER_SYNC_CAPTURE Procedure Parameters

Parameter	Description
`capture_name`	The name of the synchronous capture being altered. You must specify an existing synchronous capture name. Do not specify an owner.
`rule_set_name`	The name of the positive rule set for the synchronous capture. The positive rule set contains the rules that instruct the synchronous capture to capture changes. To change the rule set for the synchronous capture, specify an existing rule set in the form `[`*`schema_name`*`.]`*`rule_set_name`*. For example, to specify a positive rule set in the `strmadmin` schema named `sync_cap_rules`, enter `strmadmin.sync_cap_rules`. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. If `NULL`, then the rule set is not changed.

Table 44-7 (Cont.) ALTER_SYNC_CAPTURE Procedure Parameters

Parameter	Description
capture_user	The user in whose security domain a synchronous capture captures changes that satisfy its rule set and runs custom rule-based transformations configured for synchronous capture rules. If NULL, then the capture user is not changed. To change the capture user, the user who invokes the ALTER_SYNC_CAPTURE procedure must be granted the DBA role. Only the SYS user can set the capture_user to SYS. If you change the capture user, then this procedure grants the new capture user enqueue privilege on the queue used by the synchronous capture and configures the user as a secure queue user of the queue. In addition, ensure that capture user has the following privileges: • EXECUTE privilege on the rule sets used by the synchronous capture • EXECUTE privilege on all rule-based transformation functions used in the rule set These privileges can be granted directly to the capture user, or they can be granted through roles. In addition, the capture user must be granted EXECUTE privilege on all packages, including Oracle-supplied packages, that are invoked in rule-based transformations run by the synchronous capture. These privileges must be granted directly to the capture user. They cannot be granted through roles.

Usage Notes

If the capture user for a synchronous capture is dropped using DROP USER . . . CASCADE, then the synchronous capture is also dropped automatically.

44.3.7 BUILD Procedure

This procedure extracts the data dictionary of the current database to the redo log and automatically specifies database supplemental logging by running the SQL statement ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This procedure is overloaded. One version of this procedure contains the OUT parameter first_scn, and the other does not.

Syntax

```
DBMS_CAPTURE_ADM.BUILD(
   first_scn OUT NUMBER);

DBMS_CAPTURE_ADM.BUILD;
```

Parameters

Table 44-8 BUILD Procedure Parameter

Parameter	Description
first_scn	Contains the lowest SCN value corresponding to the data dictionary extracted to the redo log that can be specified as a first SCN for a capture process

Usage Notes

The following usage notes apply to this procedure:

- You can run this procedure multiple times at a source database.
- If you plan to capture changes originating at a source database with a capture process, then this procedure must be executed at the source database at least once. When the capture process is started, either at a local source database or at a downstream database, the capture process uses the extracted information in the redo log to create a LogMiner data dictionary.
- A LogMiner data dictionary is a separate data dictionary used by a capture process to determine the details of a change that it is capturing. The LogMiner data dictionary is necessary because the primary data dictionary of the source database might not be synchronized with the redo data being scanned by a capture process.
- After executing this procedure, you can query the `FIRST_CHANGE#` column of the `V$ARCHIVED_LOG` dynamic performance view where the `DICTIONARY_BEGIN` column is `YES` to determine the lowest SCN value for the database that can be specified as a first SCN for a capture process. The first SCN for a capture process is the lowest SCN in the redo log from which the capture process can capture changes.You can specify the first SCN for a capture process when you run the `CREATE_CAPTURE` or `ALTER_CAPTURE` procedure in the `DBMS_CAPTURE_ADM` package.
- In a CDB, the `BUILD` procedure must be executed from the root.

44.3.8 CREATE_CAPTURE Procedure

This procedure creates a capture process.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

> **See Also:**
>
> DBMS_RULE_ADM for more information about rules and rule sets

Syntax

```
DBMS_CAPTURE_ADM.CREATE_CAPTURE(
   queue_name                 IN  VARCHAR2,
   capture_name               IN  VARCHAR2,
   rule_set_name              IN  VARCHAR2   DEFAULT NULL,
   start_scn                  IN  NUMBER     DEFAULT NULL,
   source_database            IN  VARCHAR2   DEFAULT NULL,
   use_database_link          IN  BOOLEAN    DEFAULT FALSE,
   first_scn                  IN  NUMBER     DEFAULT NULL,
   logfile_assignment         IN  VARCHAR2   DEFAULT 'implicit',
   negative_rule_set_name     IN  VARCHAR2   DEFAULT NULL,
   capture_user               IN  VARCHAR2   DEFAULT NULL,
   checkpoint_retention_time  IN  NUMBER     DEFAULT 60,
   start_time                 IN  TIMESTAMP  DEFAULT NULL,
   source_root_name           IN  VARCHAR2   DEFAULT NULL,
   capture_class              IN  VARCHAR2   DEFAULT 'Streams');
```

Parameters

Table 44-9 CREATE_CAPTURE Procedure Parameters

Parameter	Description
queue_name	The name of the queue into which the capture process enqueues changes. You must specify an existing queue in the form [*schema_name.*]*queue_name*. For example, to specify a queue in the hr schema named streams_queue, enter hr.streams_queue. If the schema is not specified, then the current user is the default. **Note:** The queue_name setting cannot be altered after the capture process is created.
capture_name	The name of the capture process being created. A NULL specification is not allowed. Do not specify an owner. **Note:** The capture_name setting cannot be altered after the capture process is created.
rule_set_name	The name of the positive rule set for the capture process. The positive rule set contains the rules that instruct the capture process to capture changes. If you want to use a positive rule set for the capture process, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a positive rule set in the hr schema named job_capture_rules, enter hr.job_capture_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify NULL, and no negative rule set is specified, then the capture process captures all supported changes to all objects in the database, excluding database objects in the SYS and SYSTEM schemas. If you specify NULL, and a negative rule set exists for the capture process, then the capture process captures all changes that are not discarded by the negative rule set.

Table 44-9 (Cont.) CREATE_CAPTURE Procedure Parameters

Parameter	Description
start_scn	A valid SCN for the database from which the capture process starts capturing changes. An error is returned if an invalid SCN is specified. The start_scn and start_time parameters are mutually exclusive. **See Also:** "Usage Notes" for more information setting the start_scn parameter
source_database	The global name of the source database. The source database is where the changes to be captured originated. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. If NULL, or if the specified name is the same as the global name of the current database, then local capture is assumed and only the default values for use_database_link and first_scn can be specified.
use_database_link	If TRUE, then the capture process at a downstream database uses a database link to the source database for administrative purposes relating to the capture process. A database link with the same name as the global name of the source database must exist at the downstream database. The capture process uses the database link to prepare database objects for instantiation at the source database and run the DBMS_CAPTURE_ADM.BUILD procedure at the source database, if necessary. During the creation of a downstream capture process, if the first_scn parameter is set to NULL, then the use_database_link parameter must be set to TRUE. Otherwise, an error is returned. If FALSE, then either the capture process is running on the source database, or the capture process at a downstream database does not use a database link to the source database. In this case, you must perform the following administrative tasks manually: • Run the DBMS_CAPTURE_ADM.BUILD procedure at the source database to extract the data dictionary at the source database to the redo log when a capture process is created. • Obtain the first SCN for the downstream capture process if the first SCN is not specified during capture process creation. The first SCN is needed to create and maintain a capture process. • Prepare source database objects for instantiation.

Table 44-9 (Cont.) CREATE_CAPTURE Procedure Parameters

Parameter	Description
first_scn	The lowest SCN in the redo log from which a capture process can capture changes. A non-NULL value for this parameter is valid only if the DBMS_CAPTURE_ADM.BUILD procedure has been run at least once at the source database. You can query the FIRST_CHANGE# column of the V$ARCHIVED_LOG dynamic performance view where the DICTIONARY_BEGIN column is YES to determine whether the DBMS_CAPTURE_ADM.BUILD procedure has been run on a source database. Any of the values returned by such a query can be used as a first_scn value if the redo log containing that SCN value is still available. **See Also:** "Usage Notes" for more information setting the first_scn parameter
logfile_assignment	If implicit, which is the default, then the capture process at a downstream database scans all redo log files added by redo transport services or manually from the source database to the downstream database. If explicit, then a redo log file is scanned by a capture process at a downstream database only if the capture process name is specified in the FOR *logminer_session_name* clause. If explicit, then the redo log file must be added manually to the downstream database, and redo transport services cannot be used to add redo log files to the capture process being created. If you specify explicit for this parameter for a local capture process, then the local capture process cannot use the online redo log to find changes. In this case, the capture process must use the archived redo log. **See Also:** "Usage Notes" for information about adding redo log files manually

Table 44-9 (Cont.) CREATE_CAPTURE Procedure Parameters

Parameter	Description
negative_rule_set_name	The name of the negative rule set for the capture process. The negative rule set contains the rules that instruct the capture process to discard changes. If you want to use a negative rule set for the capture process, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a negative rule set in the hr schema named neg_capture_rules, enter hr.neg_capture_rules. If the schema is not specified, then the current user is the default. If you specify NULL, and no positive rule set is specified, then the capture process captures all supported changes to all objects in the database, excluding database objects in the SYS and SYSTEM schemas. If you specify NULL, and a positive rule set exists for the capture process, then the capture process captures all changes that are not discarded by the positive rule set. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify both a positive and a negative rule set for a capture process, then the negative rule set is always evaluated first.
capture_user	The user in whose security domain a capture process captures changes that satisfy its rule sets and runs custom rule-based transformations configured for capture process rules. If NULL, then the user who runs the CREATE_CAPTURE procedure is used. **Note:** If the capture user for a capture process is dropped using DROP USER . . . CASCADE, then the capture process is also dropped automatically. **See Also:** "Usage Notes" for more information about this parameter.
checkpoint_retention_time	Either specify the number of days that a capture process retains checkpoints before purging them automatically, or specify DBMS_CAPTURE_ADM.INFINITE if checkpoints should not be purged automatically. If a number is specified, then a capture process purges a checkpoint the specified number of days after the checkpoint was taken. Partial days can be specified using decimal values. For example, .25 specifies 6 hours. When a checkpoint is purged, LogMiner data dictionary information for the archived redo log file that corresponds to the checkpoint is purged, and the first_scn of the capture process is reset to the SCN value corresponding to the first change in the next archived redo log file.

Table 44-9 (Cont.) CREATE_CAPTURE Procedure Parameters

Parameter	Description
start_time	A valid time from which the capture process starts capturing changes. An error is returned if an invalid time is specified. The start_scn and start_time parameters are mutually exclusive. **See Also:** "Usage Notes" for more information setting the start_time parameter
source_root_name	The global name of the root in the source CDB. If you specify NULL, or if the specified name is the same as the global name of the current root, then local capture is assumed. If not NULL, then remote capture is assumed and a condition is added to the generated rules to filter the LCRs based on the root in the source CDB. **Note:** In a downstream capture configuration, if the capture database is CDB and the source database is a non-CDB, then specify the same value for source_root_name and source_database. The source_root_name parameter does not need to be specified for local capture for either a CDB or a non-CDB.
capture_class	The valid values are 'Streams', 'XStream' or 'GoldenGate'. If NULL is specified, then 'Streams' is assumed. **Note:** The capture_class parameter cannot be set to 'Streams' or NULL when the capture database is a CDB.

Usage Notes

Consider the following usage notes when you run this procedure:

- DBA Role Requirement
- Capture User Requirements
- First SCN and Start SCN Settings
- Explicit Log File Assignment

DBA Role Requirement

If the user who invokes this procedure is different from the user specified in the capture_user parameter, then the invoking user must be granted the DBA role. If the user who invokes this procedure is the same as the user specified in the capture_user parameter, then the DBA role is not required for the invoking user. Only the SYS user can set the capture_user to SYS.

Capture User Requirements

The capture_user parameter specifies the user who captures changes that satisfy the capture process rule sets. This user must have the necessary privileges to capture changes. This procedure grants the capture user enqueue privilege on the queue used by the capture process and configures the user as a secure queue user of the queue.

In addition, ensure that the capture user has the following privileges:

- `EXECUTE` privilege on the rule sets used by the capture process
- `EXECUTE` privilege on all rule-based transformation functions used in the positive rule set

These privileges can be granted directly to the capture user, or they can be granted through roles.

In addition, the capture user must be granted `EXECUTE` privilege on all packages, including Oracle-supplied packages, that are invoked in rule-based transformations run by the capture process. These privileges must be granted directly to the capture user. They cannot be granted through roles.

> **Note:**
>
> - A capture user does not require privileges on a database object to capture changes to the database object. The capture process can pass these changes to a rule-based transformation function. Therefore, ensure that you consider security implications when you configure a capture process.
> - Creation of the first capture process in a database might take some time because the data dictionary is duplicated during this creation.

First SCN and Start SCN Settings

When you create a capture process using this procedure, you can specify the first SCN and start SCN for the capture process. A capture process scans the redo data from the first SCN or an existing capture process checkpoint forward, even if the start SCN is higher than the first SCN or the checkpoint SCN. In this case, the capture process does not capture any changes in the redo data before the start SCN. Oracle recommends that, at capture process creation time, the difference between the first SCN and start SCN be as small as possible to keep the amount of redo scanned by the capture process to a minimum.

> **Note:**
>
> When you specify the `start_time` parameter instead of the `start_scn` parameter, the `start_time` corresponds with a specific SCN. In this case, the information in this section also applies to the SCN that corresponds with the specified `start_time`.

In some cases, the behavior of the capture process is different depending on the settings of these SCN values and on whether the capture process is local or downstream.

The following table describes capture process behavior for SCN value settings:

first_scn Setting	start_scn Setting	Capture Process Type	Description
Non-NULL	NULL	Local or Downstream	The new capture process is created at the local database with a new LogMiner session starting from the value specified for the first_scn parameter. The start SCN is set to the specified first SCN value automatically, and the new capture process does not capture changes that were made before this SCN. The BUILD procedure in the DBMS_CAPTURE_ADM package is not run automatically. This procedure must have been run at least once before on the source database, and the specified first SCN must correspond to the SCN value of a previous build that is still available in the redo log. When the new capture process is started for the first time, it creates a new LogMiner data dictionary using the data dictionary information in the redo log. If the BUILD procedure has not been run at least once on the source database, then an error is raised when the capture process is started. Capture process behavior is the same for a local capture process and a downstream capture process created with these SCN settings, except that a local capture process is created at the source database and a downstream capture process is created at the downstream database.
Non-NULL	Non-NULL	Local or Downstream	If the specified value for the start_scn parameter is greater than or equal to the specified value for the first_scn parameter, then the new capture process is created at the local database with a new LogMiner session starting from the specified first SCN. In this case, the new capture process does not capture changes that were made before the specified start SCN. If the specified value for the start_scn parameter is less than the specified value for the first_scn parameter, then an error is raised. The BUILD procedure in the DBMS_CAPTURE_ADM package is not run automatically. This procedure must have been called at least once before on the source database, and the specified first_scn must correspond to the SCN value of a previous build that is still available in the redo log. When the new capture process is started for the first time, it creates a new LogMiner data dictionary using the data dictionary information in the redo log. If the BUILD procedure has not been run at least once on the source database, then an error is raised. Capture process behavior is the same for a local capture process and a downstream capture process created with these SCN settings, except that a local capture process is created at the source database and a downstream capture process is created at the downstream database.

first_scn Setting	start_scn Setting	Capture Process Type	Description
NULL	Non-NULL	Local	The new capture process creates a new LogMiner data dictionary if either one of the following conditions is true: • There is no existing capture process for the local source database, and the specified value for the start_scn parameter is greater than or equal to the current SCN for the database. • There are existing capture processes, but none of the capture processes have taken a checkpoint yet, and the specified value for the start_scn parameter is greater than or equal to the current SCN for the database. In either of these cases, the BUILD procedure in the DBMS_CAPTURE_ADM package is run during capture process creation. The new capture process uses the resulting build of the source data dictionary in the redo log to create a LogMiner data dictionary the first time it is started, and the first SCN corresponds to the SCN of the data dictionary build. If there are any in-flight transactions, then the BUILD procedure waits until these transactions commit before completing. An in-flight transaction is one that is active during capture process creation or a data dictionary build. However, if there is at least one existing local capture process for the local source database that has taken a checkpoint, then the new capture process shares an existing LogMiner data dictionary with one or more of the existing capture processes. In this case, a capture process with a first SCN that is lower than or equal to the specified start SCN must have been started successfully at least once. Also, if there are any in-flight transactions, then the capture process is created after these transactions commit. If there is no existing capture process for the local source database (or if no existing capture processes have taken a checkpoint yet), and the specified start SCN is less than the current SCN for the database, then an error is raised.

first_scn Setting	start_scn Setting	Capture Process Type	Description
NULL	Non-NULL	Downstream	When the CREATE_CAPTURE procedure creates a downstream capture process, the use_database_link parameter must be set to TRUE when the first_scn parameter is set to NULL. Otherwise, an error is raised. The database link is used to obtain the current SCN of the source database. The new capture process creates a new LogMiner data dictionary if either one of the following conditions is true: • There is no existing capture process that captures changes to the source database at the downstream database, and the specified value for the start_scn parameter is greater than or equal to the current SCN for the source database. • There are existing capture processes that capture changes to the source database at the downstream database, but none of the capture processes have taken a checkpoint yet, and the specified value for the start_scn parameter is greater than or equal to the current SCN for the source database. In either of these cases, the BUILD procedure in the DBMS_CAPTURE_ADM package is run during capture process creation. The first time you start the new capture process, it uses the resulting build of the source data dictionary in the redo log files copied to the downstream database to create a LogMiner data dictionary. Here, the first SCN for the new capture process corresponds to the SCN of the data dictionary build. If there are any in-flight transactions, then the BUILD procedure waits until these transactions commit before completing. However, if at least one existing capture process has taken a checkpoint and captures changes to the source database at the downstream database, then the new capture process shares an existing LogMiner data dictionary with one or more of these existing capture processes. In this case, one of these existing capture processes with a first SCN that is lower than or equal to the specified start SCN must have been started successfully at least once. Also, if there are any in-flight transactions, then the capture process is created after these transactions commit. If there is no existing capture process that captures changes to the source database at the downstream database (or no existing capture process has taken a checkpoint), and the specified start_scn parameter value is less than the current SCN for the source database, then an error is raised.
NULL	NULL	Local or Downstream	The behavior is the same as setting the first_scn parameter to NULL and setting the start_scn parameter to the current SCN of the source database.

> **See Also:**
>
> BUILD Procedure for more information about the BUILD procedure and the LogMiner data dictionary

Explicit Log File Assignment

If you specify explicit for the logfile_assignment parameter, then you add a redo log file manually to a downstream database using the following statement:

```
ALTER DATABASE REGISTER LOGICAL LOGFILE
   file_name FOR capture_process;
```

Here, *file_name* is the name of the redo log file being added and *capture_process* is the name of the capture process that will use the redo log file at the downstream database. The *capture_process* is equivalent to the *logminer_session_name* and must be specified. The redo log file must be present at the site running the downstream database. You must transfer this file manually to the site running the downstream database using the DBMS_FILE_TRANSFER package, FTP, or some other transfer method.

> **See Also:**
>
> *Oracle Database SQL Language Reference* for more information about the ALTER DATABASE statement and *Oracle Data Guard Concepts and Administration* for more information registering redo log files

44.3.9 CREATE_SYNC_CAPTURE Procedure

This procedure creates a synchronous capture.

Syntax

```
DBMS_CAPTURE_ADM.CREATE_SYNC_CAPTURE(
   queue_name     IN  VARCHAR2,
   capture_name   IN  VARCHAR2,
   rule_set_name  IN  VARCHAR2,
   capture_user   IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 44-10 CREATE_SYNC_CAPTURE Procedure Parameters

Parameter	Description
queue_name	The name of the queue into which the synchronous capture enqueues changes. You must specify an existing queue in the form [*schema_name.*]*queue_name*. For example, to specify a queue in the strmadmin schema named streams_queue, enter strmadmin.streams_queue. If the schema is not specified, then the current user is the default. **Note:** The queue_name setting cannot be altered after the synchronous capture is created.
capture_name	The name of the synchronous capture being created. A NULL specification is not allowed. Do not specify an owner. **Note:** The capture_name setting cannot be altered after the synchronous capture is created.
rule_set_name	The name of the positive rule set for the synchronous capture. The positive rule set contains the rules that instruct the synchronous capture to capture changes. Specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a positive rule set in the strmadmin schema named sync_cap_rules, enter strmadmin.sync_cap_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. If NULL, then an error is returned.

Table 44-10 (Cont.) CREATE_SYNC_CAPTURE Procedure Parameters

Parameter	Description
capture_user	The user in whose security domain the synchronous capture captures changes that satisfy its rule set and runs custom rule-based transformations configured for synchronous capture rules. If NULL, then the user who runs the CREATE_SYNC_CAPTURE procedure is used. Only a user who is granted the DBA role can set a capture user. Only the SYS user can set the capture_user to SYS. **Note:** If the capture user for a synchronous capture is dropped using DROP USER . . . CASCADE, then the synchronous capture is also dropped automatically. **See Also:** "Usage Notes" for more information about this parameter.

Usage Notes

When the CREATE_SYNC_CAPTURE procedure creates a synchronous capture, the procedure must obtain an exclusive lock on each table for which it will capture changes. The rules in the specified rule set for the synchronous capture determine these tables. If there are outstanding transactions on a table for which the synchronous capture will capture changes, then the procedure waits until it can obtain a lock.

The capture_user parameter specifies the user who captures changes that satisfy the synchronous capture rule set. This user must have the necessary privileges to capture changes.

In addition, ensure that the capture user has the following privileges:

- ENQUEUE privilege on the queue specified in the queue_name parameter
- EXECUTE privilege on the rule set used by the synchronous capture
- EXECUTE privilege on all rule-based transformation functions used in the rule set

These privileges can be granted directly to the capture user, or they can be granted through roles.

In addition, the capture user must be granted EXECUTE privilege on all packages, including Oracle-supplied packages, that are invoked in rule-based transformations run by the synchronous capture. These privileges must be granted directly to the capture user. These privileges cannot be granted through roles.

> **Note:**
>
> A capture user does not require privileges on a database object to capture changes to the database object. The synchronous capture can pass these changes to a rule-based transformation function. Therefore, ensure that you consider security implications when you configure a synchronous capture.

44.3.10 DROP_CAPTURE Procedure

This procedure drops a capture process.

Syntax

```
DBMS_CAPTURE_ADM.DROP_CAPTURE(
   capture_name           IN  VARCHAR2,
   drop_unused_rule_sets  IN  BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 44-11 DROP_CAPTURE Procedure Parameters

Parameter	Description
capture_name	The name of the capture process being dropped. Specify an existing capture process name. Do not specify an owner.
drop_unused_rule_sets	If TRUE, then the procedure drops any rule sets, positive and negative, used by the specified capture process if these rule sets are not used by any other Oracle Replication client. Oracle Replication clients include capture processes, propagations, apply processes, and messaging clients. If this procedure drops a rule set, then this procedure also drops any rules in the rule set that are not in another rule set. If FALSE, then the procedure does not drop the rule sets used by the specified capture process, and the rule sets retain their rules.

Usage Notes

The following usage notes apply to this procedure:

- The Capture Process Must Be Stopped Before It Is Dropped
- The DROP_CAPTURE Procedure and Rules-related Information

The Capture Process Must Be Stopped Before It Is Dropped

A capture process must be stopped before it can be dropped.

> **See Also:**
>
> STOP_CAPTURE Procedure

The DROP_CAPTURE Procedure and Rules-related Information

When you use this procedure to drop a capture process, rules-related information for the capture process is removed from the data dictionary views for Oracle Replication rules. Information about such a rule is removed even if the rule is not in either rule set for the capture process.

The following are the data dictionary views for Oracle Replication rules:

- ALL_STREAMS_GLOBAL_RULES
- DBA_STREAMS_GLOBAL_RULES

- ALL_STREAMS_SCHEMA_RULES
- DBA_STREAMS_SCHEMA_RULES
- ALL_STREAMS_TABLE_RULES
- DBA_STREAMS_TABLE_RULES
- DBA_STREAMS_RULES

44.3.11 INCLUDE_EXTRA_ATTRIBUTE Procedure

This procedure includes or excludes an extra attribute in logical change records (LCRs) captured by the specified capture process or synchronous capture.

Syntax

```
DBMS_CAPTURE_ADM.INCLUDE_EXTRA_ATTRIBUTE(
   capture_name    IN  VARCHAR2,
   attribute_name  IN  VARCHAR2,
   include         IN  BOOLEAN   DEFAULT TRUE);
```

Parameters

Table 44-12 INCLUDE_EXTRA_ATTRIBUTE Procedure Parameters

Parameter	Description
capture_name	The name of the capture process or synchronous capture. Specify an existing capture process name or synchronous capture name. Do not specify an owner.
attribute_name	The name of the attribute to be included in or excluded from LCRs captured by the capture process or synchronous capture. The following names are valid settings: • row_id The rowid of the row changed in a row LCR. This attribute is not included in DDL LCRs, or in row LCRs for index-organized tables. The type is VARCHAR2. • serial# The serial number of the session that performed the change captured in the LCR. The type is NUMBER. • session# The identifier of the session that performed the change captured in the LCR. The type is NUMBER. • thread# The thread number of the instance in which the change captured in the LCR was performed. Typically, the thread number is relevant only in an Oracle Real Application Clusters (Oracle RAC) environment. The type is NUMBER. • tx_name The name of the transaction that includes the LCR. The type is VARCHAR2. • username The name of the user who performed the change captured in the LCR. The type is VARCHAR2.
include	If TRUE, then the specified attribute is included in LCRs captured by the capture process or synchronous capture. If FALSE, then the specified attribute is excluded from LCRs captured by the capture process or synchronous capture.

Usage Notes

Some information is not captured by a capture process or synchronous capture unless you use this procedure to specify that the information should be captured. If you want to exclude an extra attribute that is being captured by a capture process or synchronous capture, then specify the attribute and specify `FALSE` for the `include` parameter.

44.3.12 PREPARE_GLOBAL_INSTANTIATION Procedure

This procedure performs the synchronization necessary for instantiating all the tables in the database at another database and can enable supplemental logging for key columns or all columns in these tables.

This procedure prepares the tables in the database for instantiation when a capture process will be used to capture changes to the tables in the database.

This procedure records the lowest SCN of each object in the database for instantiation. SCNs after the lowest SCN for an object can be used for instantiating the object. Running this procedure prepares all current and future objects in the database for instantiation.

Syntax

```
DBMS_CAPTURE_ADM.PREPARE_GLOBAL_INSTANTIATION
   supplemental_logging  IN  VARCHAR2  DEFAULT 'KEYS',
   container             IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameter

Table 44-13 PREPARE_GLOBAL_INSTANTIATION Procedure Parameter

Parameter	Description
`supplemental_logging`	Either `NONE`, `KEYS`, or `ALL`. If `NONE` is specified, then this procedure does not enable supplemental logging for any columns in the tables in the database. This procedure does not remove existing supplemental logging specifications for these tables. If `KEYS` is specified, then this procedure enables supplemental logging for primary key, unique key, bitmap index, and foreign key columns in the tables in the database and for any table added to the database in the future. Primary key columns are logged unconditionally. Unique key, bitmap index, and foreign key columns are logged conditionally. Specifying `KEYS` does not enable supplemental logging of bitmap join index columns. If `ALL` is specified, then this procedure enables supplemental logging for all columns in the tables in the database and for any table added to the database in the future. The columns are logged unconditionally. Supplemental logging is not enabled for columns of the following types: LOB, `LONG`, `LONG RAW`, user-defined types, and Oracle-supplied types.

Table 44-13 (Cont.) PREPARE_GLOBAL_INSTANTIATION Procedure Parameter

Parameter	Description
container	Either CURRENT , ALL, or *pdb_name*. If CURRENT is specified, then this procedure adds supplemental logging for the current container. If ALL is specified, then this procedure adds supplemental logging for all of the containers in the current CDB. If *pdb_name* is specified, then this procedure adds supplemental logging for the specified PDB. ALL and *pdb_name* are valid only when you invoke the procedure from the root.

Usage Notes

Run this procedure at the source database.

If you use a capture process to capture all of the changes to a database, then use this procedure to prepare the tables in the database for instantiation after the capture process has been configured.

44.3.13 PREPARE_SCHEMA_INSTANTIATION Procedure

This procedure performs the synchronization necessary for instantiating all tables in the schema at another database and can enable supplemental logging for key columns or all columns in these tables.

This procedure prepares the tables in the schema for instantiation when a capture process will be used to capture changes to the tables in the schema.

This procedure records the lowest SCN of each object in the schema for instantiation. SCNs after the lowest SCN for an object can be used for instantiating the object. Running this procedure prepares all current and future objects in the schema for instantiation.

Syntax

```
DBMS_CAPTURE_ADM.PREPARE_SCHEMA_INSTANTIATION(
   schema_name           IN  VARCHAR2,
   supplemental_logging  IN  VARCHAR2  DEFAULT 'KEYS',
   container             IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameters

Table 44-14 PREPARE_SCHEMA_INSTANTIATION Procedure Parameters

Parameter	Description
schema_name	The name of the schema. For example, hr.

Table 44-14 (Cont.) PREPARE_SCHEMA_INSTANTIATION Procedure Parameters

Parameter	Description
supplemental_logging	Either NONE, KEYS, ALL, PARTIAL_JSON_OFF, or PARTIAL_JSON_ON. If NONE is specified, then this procedure does not enable supplemental logging for any columns in the tables in the schema. This procedure does not remove existing supplemental logging specifications for these tables. If KEYS is specified, then this procedure enables supplemental logging for primary key, unique key, bitmap index, and foreign key columns in the tables in the schema and for any table added to this schema in the future. Primary key columns are logged unconditionally. Unique key, bitmap index, and foreign key columns are logged conditionally. Specifying KEYS does not enable supplemental logging of bitmap join index columns. If ALL is specified, then this procedure enables supplemental logging for all columns in the tables in the schema and for any table added to this schema in the future. The columns are logged unconditionally. Supplemental logging is not enabled for columns of the following types: LOB, LONG, LONG RAW, user-defined types, and Oracle-supplied types. If PARTIAL_JSON_OFF is specified, partial JSON update is disabled at schema level. If PARTIAL_JSON_ON is specified, partial JSON update is enabled at schema level. If PARTIAL_JSON_ON or PARTIAL_JSON_OFF is unspecified, schema level partial JSON setting will be unchanged. If PARTIAL_JSON_OFF has never been specified, by default, partial JSON update is disabled at schema level.
container	Either CURRENT , ALL, or *pdb_name*. If CURRENT is specified, then this procedure adds supplemental logging for the current container. If ALL is specified, then this procedure adds supplemental logging for all of the containers in the current CDB. If *pdb_name* is specified, then this procedure adds supplemental logging for the specified PDB ALL and *pdb_name* are valid only when you invoke the procedure from the root.

Usage Notes

Run this procedure at the source database. If you use a capture process to capture all of the changes to a schema, then use this procedure to prepare the tables in the schema for instantiation after the capture process has been configured.

44.3.14 PREPARE_SYNC_INSTANTIATION Function

This function performs the synchronization necessary for instantiating one or more tables at another database. This function returns the prepare system change number (SCN) for the table or tables being prepared for instantiation.

This function prepares one or more tables for instantiation when a synchronous capture will be used to capture changes to the tables.

This function records the lowest SCN of each table for instantiation (prepare SCN). SCNs after the lowest SCN for an object can be used for instantiating the object.

This function is overloaded. The `table_names` parameter is `VARCHAR2` datatype in one version and `DBMS_UTILITY.UNCL_ARRAY` datatype in the other version.

Syntax

```
DBMS_CAPTURE_ADM.PREPARE_SYNC_INSTANTIATION(
   table_names  IN  VARCHAR2)
RETURN NUMBER;

DBMS_CAPTURE_ADM.PREPARE_SYNC_INSTANTIATION(
   table_names  IN  DBMS_UTILITY.UNCL_ARRAY)
RETURN NUMBER;
```

Parameters

Table 44-15 PREPARE_SYNC_INSTANTIATION Function Parameter

Parameter	Description
`table_names`	When the `table_names` parameter is `VARCHAR2` datatype, a comma-delimited list of the tables to prepare for instantiation. There must be no spaces between entries. When the `table_names` parameter is `DBMS_UTILITY.UNCL_ARRAY` datatype, specify a PL/SQL associative array of this type that contains the names of the tables to prepare for instantiation. The first table name is at position 1, the second at position 2, and so on. The table does not need to be `NULL` terminated. In either version of the function, specify the name of each table in the form `[`*`schema_name`*`.]`*`table_name`*. For example, `hr.employees`. If the schema is not specified, then the current user is the default.

44.3.15 PREPARE_TABLE_INSTANTIATION Procedure

This procedure performs the synchronization necessary for instantiating the table at another database and can enable supplemental logging for key columns or all columns in the table

This procedure prepares the table for instantiation when a capture process will be used to capture changes to the table.

This procedure records the lowest SCN of the table for instantiation. SCNs after the lowest SCN for an object can be used for instantiating the object.

Syntax

```
DBMS_CAPTURE_ADM.PREPARE_TABLE_INSTANTIATION(
   table_name            IN  VARCHAR2,
   supplemental_logging  IN  VARCHAR2  DEFAULT 'KEYS',
   container             IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameters

Table 44-16 PREPARE_TABLE_INSTANTIATION Procedure Parameters

Parameter	Description
`table_name`	The name of the table specified as `[`*`schema_name`*`.]`*`object_name`*. For example, `hr.employees`. If the schema is not specified, then the current user is the default.

Table 44-16 (Cont.) PREPARE_TABLE_INSTANTIATION Procedure Parameters

Parameter	Description
supplemental_logging	Either NONE, KEYS, or ALL. If NONE is specified, then this procedure does not enable supplemental logging for any columns in the table. This procedure does not remove existing supplemental logging specifications for the table. If KEYS is specified, then this procedure enables supplemental logging for primary key, unique key, bitmap index, and foreign key columns in the table. The procedure places the key columns for the table in three separate log groups: the primary key columns in an unconditional log group, the unique key columns and bitmap index columns in a conditional log group, and the foreign key columns in a conditional log group. Specifying KEYS does not enable supplemental logging of bitmap join index columns. If ALL is specified, then this procedure enables supplemental logging for all columns in the table. The procedure places all of the columns for the table in an unconditional log group. Supplemental logging is not enabled for columns of the following types: LOB, LONG, LONG RAW, user-defined types, and Oracle-supplied types.
container	Either CURRENT , ALL, or *pdb_name*. If CURRENT is specified, then this procedure adds supplemental logging for the current container. If ALL is specified, then this procedure adds supplemental logging for all of the containers in the current CDB. If *pdb_name* is specified, then this procedure adds supplemental logging for the specified PDB. ALL and *pdb_name* are valid only when you invoke the procedure from the root.

Usage Notes

Run this procedure at the source database. If you use a capture process to capture all of the changes to a table, then use this procedure to prepare the table for instantiation after the capture process has been configured.

44.3.16 SET_PARAMETER Procedure

This procedure sets a capture process parameter to the specified value.

Syntax

```
DBMS_CAPTURE_ADM.SET_PARAMETER(
   capture_name  IN  VARCHAR2,
   parameter     IN  VARCHAR2,
   value         IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 44-17 SET_PARAMETER Procedure Parameters

Parameter	Description
capture_name	The name of the capture process. Do not specify an owner.

Table 44-17 (Cont.) SET_PARAMETER Procedure Parameters

Parameter	Description
parameter	The name of the parameter you are setting.
value	The value to which the parameter is set. If NULL, then the parameter is set to its default value.

Capture Process Parameters

The following table lists the parameters for the capture process.

Table 44-18 Capture Process Parameters

Parameter Name	Possible Values	Default	Description
annotateddlsize	0 or a positive integer up to 4294967296 (4 GB)	0	The maximum length of the DDL text, in bytes, up to which the DDL can be chosen for annotation, if applicable for the DDL. Any DDL text whose size is greater than the specified size is not annotated. The default, 0, indicates that no DDL text processed by this capture process is annotated. Annotating DDL can simplify filtering and transformation of the DDL within Oracle GoldenGate. **Note:** This parameter is intended for Oracle GoldenGate. Do not use this parameter in an Oracle Replication environment or in an XStream environment.
capture_idkey_objects	Y or N	N	If Y, then the capture process captures ID key logical change records (LCRs). If N, then the capture process does not capture ID key LCRs. Capture processes do not fully support capturing changes to some data types from the redo log. ID key LCRs enable an XStream configuration to capture these changes and process them with an XStream client application. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter and *Oracle Database XStream Guide* for more information about ID key LCRs

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
capture_sequence_nextval	Y or N	N	If Y, then the capture process captures sequence LCRs for all of the sequences in the database, except for sequences in Oracle-supplied, administrative schemas such as SYS and SYSTEM. The capture process's rule sets can filter sequence LCRs in the same way that they filter row LCRs and DDL LCRs. If N, then the capture process does not capture sequence LCRs. An apply process or XStream inbound server can use sequence LCRs to ensure that the sequence values at a destination database use the appropriate values. For increasing sequences, the sequence values at the destination are equal to or greater than the sequence values at the source database. For decreasing sequences, the sequence values at the destination are less than or equal to the sequence values at the source database. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter and "SET_PARAMETER Procedure" for information about the apply_sequence_nextval apply process parameter
disable_on_limit	Y or N	N	If Y, then the capture process is disabled because it reached a value specified by the time_limit parameter or message_limit parameter. If N, then the capture process is restarted immediately after stopping because it reached a limit. When a capture process is restarted, it starts to capture changes at the point where it last stopped. A restarted capture process gets a new session identifier, and the processes associated with the capture process also get new session identifiers. However, the capture process number (CP*nn*) remains the same.

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
downstream_real_time_mine	Y or N	Y for local capture processes N for downstream capture processes	If Y, then the capture process is a real-time downstream capture process. After setting this parameter to Y, switch the redo log file at the source database using the SQL statement ALTER SYSTEM ARCHIVE LOG CURRENT to begin real-time downstream capture. If this parameter is set to Y, then redo data from the source database must be sent to the standby redo log at the downstream database. If N, then the capture process is an archived-log downstream capture process. An error is raised if an attempt is made to set this parameter for a local capture process.
excludetag	Comma-delimited list of Oracle Replication tags	NULL	Controls whether the capture process captures DML changes that are tagged with one of the specified Oracle Replication tags. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
excludetrans	Comma-delimited list of transaction names	NULL	Controls whether the capture process captures DML changes in the specified transaction names. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
excludeuser	Comma-delimited list of user names	NULL	Controls whether the capture process captures DML changes made by the specified users. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. Specify an exact pattern match for each user name. The pattern match is case sensitive. For example, specify HR for the hr user. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
excludeuserid	Comma-delimited list of user ID values	NULL	Controls whether the capture process captures data manipulation language (DML) changes made by the specified users. Whether the capture process captures these changes depends on the settings for the getapplops and getreplicates parameters. To view the user ID for a user, query the USER_ID column in the ALL_USERS data dictionary view. If NULL, then the capture process ignores this parameter. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
getapplops	Y or N	Y	If Y, then the capture process captures DML changes if the original user is not specified in the excludeuserid or excludeuser parameters and the transaction name is not specified in the excludetrans parameter. If N, then the capture process ignores DML changes if the original user is not specified in the excludeuserid or excludeuser parameters and the transaction name is not specified in the excludetrans parameter. In either case, the capture process captures a DML change only if it satisfies the capture process's rule sets. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
getreplicates	Y or N	N	If Y, then the capture process captures DML changes if the original user is specified in the excludeuserid or excludeuser parameters and the transaction name is specified in the excludetrans parameter. If N, then the capture process ignores DML changes if the original user is specified in the excludeuserid or excludeuser parameters and the transaction name is specified in the excludetrans parameter. In either case, the capture process captures a DML change only if it satisfies the capture process's rule sets. **Note:** This parameter is intended for XStream. Do not use this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
ignore_transaction	A valid transaction ID or NULL	NULL	Instructs the capture process to ignore the specified transaction from the source database, effective immediately. The capture process eliminates all subsequent LCRs for the transaction. If the specified transaction is committed successfully at the source database, the destination database will receive a ROLLBACK statement instead, and any LCRs from the transaction that were enqueued before the ignore transaction request are rolled backed at the destination database. If NULL, then the capture process ignores this parameter. Use caution when setting this parameter because ignoring a transaction might lead to data divergence between the source database and destination database. To ignore multiple transactions, specify each transaction in a separate call to the SET_PARAMETER procedure. The DBA_CAPTURE_PARAMETERS view displays a comma-delimited list of all transactions to be ignored. To clear the list of ignored transactions, run the SET_PARAMETER procedure and specify NULL for the ignore_transaction parameter.
ignore_unsupported_table	A fully qualified table name, *, or -	*	Controls the behavior of the capture process when it tries to capture changes to a specified table or to an unsupported table. A capture process tries to capture changes to an unsupported table when its rule sets instruct it to do so. If you do not want the capture process to try to capture changes to unsupported tables, then ensure that the capture process's rule sets exclude unsupported tables. When a table name is specified, the capture process does not capture changes to the specified table. The table name must be entered in the form *table_owner. table_name*. For example. hr.employees. To specify multiple tables, specify each table in a separate call to the SET_PARAMETER procedure. When * is specified and the capture process tries to capture a change to an unsupported table, the capture process ignores the change and continues to run. The change to the unsupported table is not captured, and the capture process records the unsupported table in the alert log. When - is specified and the capture process tries to capture a change to an unsupported table, the capture process aborts.

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
include_objects	A list of tables or schema names separated by commas	none	Directs capture to include changes from the specified tables or schemas. An LCR that is selected by include_objects is passed to the outbound server regardless of any further filtering that is specified. **Note:** This parameter is intended for XStream. Do not use or attempt to set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
inline_lob_optimization	Y or N	N	If Y, then LOBs that can be processed inline (such as small LOBs) are included in the LCR directly, rather than sending LOB chunk LCRs. If N, then each LOB column is sent as NULL followed by LOB chunk LCRs to update the LOB column. **Note:** This parameter is intended for XStream. Do not use or attempt to set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
maximum_scn	A valid SCN or INFINITE	INFINITE	The capture process is disabled before capturing a change record with an SCN greater than or equal to the value specified. If INFINITE, then the capture process runs regardless of the SCN value.

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
max_sga_size	A positive integer	INFINITE	Controls the amount of system global area (SGA) memory allocated specifically to the capture process, in megabytes. The capture process attempts to allocate memory up to this limit. A capture process uses Oracle LogMiner to scan for changes in the redo log. The memory is allocated for the duration of the capture process session and is released when the capture process becomes disabled. **Note:** The sum of system global area (SGA) memory allocated for all components on a database must be less than the value set for the STREAMS_POOL_SIZE initialization parameter. If NULL, then the capture component uses the original default value. A NULL value has the same effect as resetting the parameter to its default value. **Note:** This parameter is intended for XStream. Do not use or attempt to set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
merge_threshold	A negative integer, 0, a positive integer, or INFINITE	60	The amount of time, in seconds, between the message creation time of the original capture process and the message creation time of the cloned capture process. Specifically, if the difference, in seconds, between the CAPTURE_MESSAGE_CREATE_TIME of the cloned capture process and the original capture process is less than or equal to the value specified for this parameter, then automatic merge begins by running the MERGE_STREAMS procedure. This parameter is relevant only when changes captured by the capture process are applied by two or more apply processes and the split_threshold parameter is set to a value other than INFINITE. If a negative value is specified, then automatic merge is disabled. If 0 (zero) is specified, then there must be no lag between the original capture process and the cloned capture process to begin the merge. If INFINITE is specified, then automatic merging starts immediately.

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
message_limit	A positive integer or INFINITE	INFINITE	The capture process stops after capturing the specified number of messages. If INFINITE, then the capture process continues to run regardless of the number of messages captured.
message_tracking_frequency	0 or a positive integer	2000000	The frequency at which messages captured by the capture process are tracked automatically. For example, if this parameter is set to the default value of 2000000, then every two-millionth message is tracked automatically. The tracking label used for automatic message tracking is *capture_process_name*:AUTOTRACK, where *capture_process_name* is the name of the capture process. Only the first 20 bytes of the capture process name are used; the rest is truncated if it exceeds 20 bytes. If 0 (zero), then no messages are tracked automatically.
parallelism	A positive integer	1	The number of preparer servers that can concurrently mine the redo log for the capture process. A capture process consists of one reader server, one or more preparer servers, and one builder server. The preparer servers concurrently format changes found in the redo log into logical change records (LCRs). Each reader server, preparer server, and builder server is a process, and the number of preparer servers equals the number specified for the parallelism capture process parameter. So, if parallelism is set to 5, then a capture process uses a total of seven processes: one reader server, five preparer servers, and one builder server. Setting the parallelism parameter to a number higher than the number of available parallel execution servers might disable the capture process. Ensure that the PROCESSES initialization parameter is set appropriately when you set the parallelism capture process parameter. **Note:** When you change the value of this parameter, the capture process is stopped and restarted automatically.
skip_autofiltered_table_ddl	Y or N	Y	If Y, then the capture process does not capture data definition language (DDL) changes to tables that are automatically filtered by the capture process. If N, then the capture process can capture DDL changes to tables that are automatically filtered by the capture process if the DDL changes satisfy the capture process rule sets.

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
split_threshold	0, a positive integer, or INFINITE	1800	The amount of time, in seconds, that a stream is broken before the stream is automatically split from other streams that flow from the capture process. When a stream is split, the capture process, queue, and propagation are cloned. In this case, a stream is a flow of logical change records (LCRs) that flows from a capture process to an apply. A stream is broken when LCRs captured by the capture process cannot reach the apply process. For example, a stream is broken when the relevant propagation or apply process is disabled. This parameter is relevant only when changes captured by the capture process are applied by two or more apply processes. If 0 (zero), then automatic splitting starts immediately. If INFINITE, then automatic splitting is disabled. The stream is not split regardless of the amount of time that it is broken. This parameters is designed to be used with the merge_threshold parameter.
startup_seconds	0, a positive integer, or INFINITE	0	The maximum number of seconds to wait for another instantiation of the same capture process to finish. If the other instantiation of the same capture process does not finish within this time, then the capture process does not start. This parameter is useful only if you are starting the capture process manually. If INFINITE, then the capture process does not start until another instantiation of the same capture process finishes.
time_limit	A positive integer or INFINITE	INFINITE	The capture process stops as soon as possible after the specified number of seconds since it started. If INFINITE, then the capture process continues to run until it is stopped explicitly.
trace_level	0 or a positive integer	0	Set this parameter only under the guidance of Oracle Support Services.

Table 44-18 (Cont.) Capture Process Parameters

Parameter Name	Possible Values	Default	Description
use_rac_service	Y or N	Y if Replication N if XStream	If Y, then the capture process is run in the owning instance of its queue. If N, then the client specifies where the capture process is to run. **Note:** This parameter is intended for XStream. Do not use or attempt to set this parameter in an Oracle Replication environment unless XStream optimizations are enabled by the DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS procedure. See "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure" for information about enabling XStream optimizations. **See Also:** "Usage Notes" for more information about this parameter
write_alert_log	Y or N	Y	If Y, then the capture process writes a message to the alert log on exit. If N, then the capture process does not write a message to the alert log on exit. The message specifies the reason the capture process stopped.
xout_client_exists	Y or N	Y if the capture process sends LCRs to XStream outbound servers N if the capture process sends LCRs to Oracle Apply processes	Y indicates that the capture process sends LCRs to one or more XStream outbound servers. N indicates that the capture process sends LCRs to one or more Oracle Apply processes. A single capture process cannot send LCRs to both outbound servers and apply processes. In an XStream configuration where an outbound server runs on a different database than its capture process, set this parameter to Y to enable the capture process to send LCRs to the outbound server. **Note:** Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

Usage Notes

The following usage notes apply to the SET_PARAMETER procedure:

- Delays Are Possible Before New Parameter Settings Take Effect
- Parameters Interpreted as Positive Integers
- Parameters with a System Change Number (SCN) Setting
- Parameters that Require XStream Optimizations
- XStream or Oracle GoldenGate Integrated Capture Configurations

Delays Are Possible Before New Parameter Settings Take Effect

When you alter a parameter value, a short amount of time might pass before the new value for the parameter takes effect.

Parameters Interpreted as Positive Integers

For all parameters that are interpreted as positive integers, the maximum possible value is `4,294,967,295`. Where applicable, specify `INFINITE` for larger values.

Parameters with a System Change Number (SCN) Setting

For parameters that require an SCN setting, any valid SCN value can be specified.

Parameters that Require XStream Optimizations

A capture process uses the following parameters only when the capture process is sending logical change records (LCRs) to an XStream outbound server or when XStream optimizations are enabled for Oracle Replication components:

- `capture_idkey_objects`
- `capture_sequence_nextval`
- `excludetag`
- `excludetrans`
- `excludeuser`
- `excludeuserid`
- `getapplops`
- `getreplicates`
- `include_objects`
- `inline_lob_optimization`
- `max_sga_size`
- `use_rac_services`

The `DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS` procedure enables XStream optimizations for Oracle Replication. When XStream optimizations are not enabled by the `DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS` procedure, a capture process raises an error if one of these parameters is set to any value other than its default value.

When XStream optimizations are enabled for Oracle Replication and the `capture_idkey_objects` parameter is set to `Y`, a capture process can capture ID key LCRs. ID key LCRs do not contain all of the columns for a row change. Instead, they contain the rowid of the changed row, a group of key columns to identify the row in the table, and the data for the scalar columns of the table that are supported by capture processes. An apply process can apply these changes using the information available the ID key LCRs.

To determine the database objects for which a capture process will capture ID key LCRs, run the following query on the source database:

```
SELECT OWNER, OBJECT_NAME
  FROM DBA_XSTREAM_OUT_SUPPORT_MODE
  WHERE SUPPORT_MODE='ID KEY';
```

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about ID key LCRs

XStream or Oracle GoldenGate Integrated Capture Configurations

In an XStream or Oracle GoldenGate integrated capture configuration, the following parameters control which changes are captured by a capture process:

- `capture_idkey_objects`
- `capture_sequence_nextval`
- `excludetag`
- `excludetrans`
- `excludeuser`
- `excludeuserid`
- `getapplops`
- `getreplicates`
- `include_objects`
- `inline_lob_optimization`
- `max_sga_size`
- `use_rac_services`

You can set these parameters to avoid change cycling. Change cycling sends a change back to the database where it originated. Typically, change cycling should be avoided in a replication environment so that the same change is not made to a database more than once.

In an XStream or Oracle GoldenGate integrated capture configuration that performs bi-directional replication, a GoldenGate Replicat process runs on the source database for a capture process. Therefore, the changes made by the GoldenGate Replicat are recorded in the redo log.

If an integrated configuration performs bi-directional replication, then, to avoid change cycling, the capture process should not capture the changes made by the Oracle GoldenGate Replicat process. To accomplish this goal, use the default settings for the `getapplops` and `getreplicates` parameters and exclude changes made by the user running the Replicat process. To exclude these changes, specify this user in the `excludeuserid` or `excludeuser` parameter. Typically, the user running the Oracle GoldenGate Replicat process is the XStream administrator.

In some configurations, the goal might be to capture or exclude changes made by applications or by the Replicat process. For example, an intermediate database in a replication environment might capture all of the changes made to the database, including both application changes and Replicat process changes, and send these changes to a different destination database.

Table 44-19 describes the capture process behavior when at least one of the exclude parameters is non-`NULL`.

Table 44-19 Behavior When at Least One exclude Parameter Is Non-NULL

getapplops Setting	getreplicates Setting	Description
Y	Y	The capture process captures all DML changes.
Y	N	The capture process captures the DML changes made by the users that are not in the `excludeuserid` or `excludeuser` parameters. The capture process captures the DML changes that are not in the transactions in the `excludetrans` parameter. The capture process captures only the DML changes that do not have a tag that is in the `excludetags` parameter.
N	Y	The capture process captures only the DML changes made by the users that are in the `excludeuserid` or `excludeuser` parameters. The capture process captures only the DML changes that are in the transactions in the `excludetrans` parameter. The capture process captures only the DML changes that have a tag that is in the `excludetags` parameter.
N	N	The capture process does not capture any DML changes.

Table 44-19 describes the capture process behavior when all of the exclude parameters are set to `NULL`.

Table 44-20 Behavior When All exclude Parameters Are Set to NULL

getapplops Setting	getreplicates Setting	Description
Y	Y	The capture process captures all DML changes.
Y	N	The capture process captures all DML changes.
N	Y	The capture process does not capture any DML changes.
N	N	The capture process does not capture any DML changes.

See the documentation for the Oracle GoldenGate product for more information:

http://docs.oracle.com/cd/E15881_01/index.htm

> **Note:**
>
> A capture process evaluates a change using these parameters before it evaluates a change using its rule sets. Therefore, a capture process can discard a change before the change is evaluated against the capture process's rule sets. Also, regardless of the settings for these parameters, a capture process captures a change only if the change satisfies the capture process's rule sets.

> **See Also:**
>
> *Oracle Database XStream Guide*

44.3.17 START_CAPTURE Procedure

This procedure starts the capture process, which mines redo logs and enqueues the mined redo information into the associated queue.

The start status is persistently recorded. Hence, if the status is `ENABLED`, then the capture process is started upon database instance startup.

The capture process is a background Oracle process and is prefixed by `c`.

The enqueue and dequeue state of `DBMS_AQADM.START_QUEUE` and `DBMS_AQADM.STOP_QUEUE` have no effect on the start status of a capture process.

Syntax

```
DBMS_CAPTURE_ADM.START_CAPTURE(
   capture_name  IN  VARCHAR2);
```

Parameters

Table 44-21 START_CAPTURE Procedure Parameter

Parameter	Description
`capture_name`	The name of the capture process. Do not specify an owner. The capture process uses LogMiner to capture changes in the redo information. A `NULL` setting is not allowed.

Usage Notes

The capture process status is persistently recorded. Hence, if the status is `ENABLED`, then the capture process is started upon database instance startup. A capture process (*`cnnn`*) is an Oracle background process.

44.3.18 STOP_CAPTURE Procedure

This procedure stops the capture process from mining redo logs.

Syntax

```
DBMS_CAPTURE_ADM.STOP_CAPTURE(
   capture_name  IN  VARCHAR2,
   force         IN  BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 44-22 STOP_CAPTURE Procedure Parameters

Parameter	Description
capture_name	The name of the capture process. A NULL setting is not allowed. Do not specify an owner.
force	If TRUE, then the procedure stops the capture process as soon as possible. If the capture process cannot stop normally, then it aborts. If FALSE, then the procedure stops the capture process as soon as possible. If the capture process cannot stop normally, then an ORA-26672 error is returned, and the capture process might continue to run.

Usage Notes

The following usage notes apply to this procedure:

- The capture process status is persistently recorded. Hence, if the status is DISABLED or ABORTED, then the capture process is not started upon database instance startup.
- A capture process is an Oracle background process with a name in the form CP*nn*, where *nn* can include letters and numbers.
- The enqueue and dequeue state of DBMS_AQADM.START_QUEUE and DBMS_AQADM.STOP_QUEUE have no effect on the stop status of a capture process.

45
DBMS_CACHEUTIL

The DBMS_CACHEUTIL package provides an interface to control object-level cache fusion locking policies in an Oracle Real Application Clusters (Oracle RAC) environment and to manage assignments to the KEEP buffer pool for Oracle True Cache.

45.1 DBMS_CACHEUTIL Overview

DBMS_CACHEUTIL provides a collection of procedures related to Oracle Real Application Clusters (Oracle RAC) and Oracle True Cache.

With DBMS_CACHEUTIL, you can do the following:

- Control object-level cache fusion locking policies, such as object affinity, read-mostly, and object down-convert in an Oracle RAC environment.

  For example, you can fine-tune cache fusion locking policies for a specific configuration or workload at the object level. This might reduce locking overhead in an Oracle RAC environment and improve performance.

  - Affinity specifies an implicit, exclusive lock for an object to an instance, so that read accesses and updates to the object on this instance become more optimized. However, updates to the object on other instances become less optimized.
  - Read-mostly specifies an implicit, shared lock for an object in the cluster, so that read access to the object becomes more optimized. However, updates to the object on any instances become less optimized.
  - A down-convert converts an object's exclusive locks on an instance to shared mode for all instances in an Oracle RAC environment.

  > **Note:**
  >
  > Object affinity and read-mostly locking optimization works only for the primary RAC database, and it doesn't work for Oracle Active Data Guard.

- Manage object assignments to the KEEP buffer cache for an Oracle True Cache instance. True Cache is an in-memory, consistent, and automatically managed SQL and key-value (object or JSON) cache for Oracle Database.

You can apply the procedures in this package to tables, indexes, lobs, partitions (for partitioned objects), and subpartitions (for composite partitioned objects).

For non-partitioned objects, the input (schema and obj corresponds to OWNER and OBJECT_NAME in DBA_OBJECTS and should uniquely define a segment. That means it should have, for example, a valid DATA_OBJECT_ID in DBA_OBJECTS.

For partitioned or composite partitioned objects, the input (schema, obj, and partition) corresponds to OWNER, OBJECT_NAME, and SUBOBJECT_NAME in DBA_OBJECTS and should uniquely define a segment. That means it should have a valid DATA_OBJECT_ID in DBA_OBJECTS.

Except for depended indexes, this procedure does not apply for the dependent object segments, such as lob and partition or subpartition segments. It should be explicitly called on the segment if necessary.

Related Topics

- Real Application Clusters Administration and Deployment Guide

Related Topics

- Oracle True Cache User's Guide

45.2 DBMS_CACHEUTIL Security Model

The security model describes the privileges needed for using the `DBMS_CACHEUTIL` package.

Use the following for this package:

```
CREATE OR REPLACE PUBLIC SYNONYM dbms_cacheutil FOR sys.dbms_cacheutil


GRANT EXECUTE ON dbms_cacheutil TO dba
```

45.3 Summary of DBMS_CACHEUTIL Subprograms

This table lists the package subprograms in alphabetical order.

Table 45-1 DBMS_CACHEUTIL Package Subprograms

Subprogram	Description
DISSOLVE_AFFINITY Procedure	This procedure dissolves the object affinity from a node in an Oracle RAC environment.
DISSOLVE_READMOSTLY Procedure	This procedure dissolves an object's read-mostly setting in an Oracle RAC environment.
GRAB_AFFINITY Procedure	This procedure sets the object affinity to a node in an Oracle RAC environment.
GRAB_READMOSTLY Procedure	This procedure grabs an object's read-mostly setting in an Oracle RAC environment.
LIST_READMOSTLY Procedure	This procedure lists objects that have the read-mostly property set in an Oracle RAC environment.
OBJECT_DOWNCONVERT Procedure	This procedure down-converts an object's exclusive locks on an instance to shared mode for all instances in an Oracle RAC environment.
TRUE_CACHE_KEEP Procedure	When you call this procedure on a True Cache instance, it assigns the object to the `KEEP` buffer pool on that instance.
TRUE_CACHE_UNKEEP Procedure	When an object on a True Cache instance no longer needs to be in the `KEEP` buffer pool, use this procedure to remove the object's `KEEP` assignment.

45.3.1 DISSOLVE_AFFINITY Procedure

`DISSOLVE_AFFINITY` dissolves the object affinity from a node in an Oracle RAC environment.

Syntax

```
DBMS_CACHEUTIL.DISSOLVE_AFFINITY(
   schema              IN VARCHAR2,
   obj                 IN VARCHAR2,
   partition           IN VARCHAR2 := NULL,
   dissolve_index      IN BOOLEAN := TRUE,
   active_drm          IN BOOLEAN := FALSE);
```

Parameters

Table 45-2 DISSOLVE_AFFINITY Procedure Parameters

Parameter	Description
`schema`	The name of the schema for the object.
`obj`	The name of the object.
`partition`	• If the object is not partitioned, then this is `NULL`. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.
`dissolve_index`	To dissolve affinity for dependent indexes, use `TRUE`.
`active_drm`	For internal use only. Do not use.

45.3.2 DISSOLVE_READMOSTLY Procedure

`DISSOLVE_READMOSTLY` dissolves an object's read-mostly setting in an Oracle RAC environment.

Syntax

```
DBMS_CACHEUTIL.DISSOLVE_READMOSTLY(
   schema              IN VARCHAR2,
   obj                 IN VARCHAR2,
   partition           IN VARCHAR2 := NULL,
   dissolve_index      IN BOOLEAN := TRUE);
```

Parameters

Table 45-3 DISSOLVE_READMOSTLY Procedure Parameters

Parameter	Description
`schema`	The name of the schema for the object.
`obj`	The name of the object.

Table 45-3 (Cont.) DISSOLVE_READMOSTLY Procedure Parameters

Parameter	Description
partition	• If the object is not partitioned, then this is NULL. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.
dissolve_index	To dissolve read-mostly for dependent indexes, use TRUE.

45.3.3 GRAB_AFFINITY Procedure

GRAB_AFFINITY sets the object affinity to a node in an Oracle RAC environment.

This is useful when you know that, for a particular workload, an object will be heavily accessed by one particular Oracle RAC node. Grabbing an object affinity to that node reduces locking overhead and could improve performance.

> **Note:**
>
> The affinity that this procedure grabs can still change by the automatic object affinity policy if the policy decides that it's not optimal to keep the affinity. The CURRENT_MASTER column in the V$GCSPFMASTER_INFO view indicates the current affinity of the particular object segment.

> **Note:**
>
> Object affinity locking optimization works only for the primary RAC database, and it doesn't work for Oracle Active Data Guard.

Syntax

```
DBMS_CACHEUTIL.GRAB_AFFINITY(
   schema              IN VARCHAR2,
   obj                 IN VARCHAR2,
   partition           IN VARCHAR2 := NULL,
   grab_index          IN BOOLEAN := TRUE,
   active_drm          IN BOOLEAN := FALSE);
```

Parameters

Table 45-4 GRAB_AFFINITY Procedure Parameters

Parameter	Description
schema	The name of the schema for the object.
obj	The name of the object.

Table 45-4 (Cont.) GRAB_AFFINITY Procedure Parameters

Parameter	Description
partition	• If the object is not partitioned, then this is NULL. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.
grab_index	To grab affinity for dependent indexes, use TRUE.
active_drm	For internal use only. Do not use.

45.3.4 GRAB_READMOSTLY Procedure

GRAB_READMOSTLY sets an object to read-mostly in an Oracle RAC environment.

This is useful when you know that, for a particular workload, accesses to an object will mostly be READ instead of INSERT or UPDATE. Grabbing read-mostly for the object would reduce locking overhead and could improve performance.

> **Note:**
>
> The read-mostly that this procedure grabs can still change by the automatic read-mostly policy if the policy decides that it's not optimal to keep the read-mostly set. The read-mostly property is persistent across instance lifetimes. You can use the LIST_READMOSTLY procedure included in this package to list the current set of object segments with read-mostly set.

> **Note:**
>
> Read-mostly locking optimization works only for the primary RAC database, and it doesn't work for Oracle Active Data Guard.

Syntax

```
DBMS_CACHEUTIL.GRAB_READMOSTLY(
   schema              IN VARCHAR2,
   obj                 IN VARCHAR2,
   partition           IN VARCHAR2 := NULL,
   grab_index          IN BOOLEAN := TRUE);
```

Parameters

Table 45-5 GRAB_READMOSTLY Procedure Parameters

Parameter	Description
schema	The name of the schema for the object.
obj	The name of the object.

Table 45-5 (Cont.) GRAB_READMOSTLY Procedure Parameters

Parameter	Description
partition	• If the object is not partitioned, then this is NULL. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.
grab_index	To grab read-mostly for dependent indexes, use TRUE.

45.3.5 LIST_READMOSTLY Procedure

This procedure lists objects that have the read-mostly property set in an Oracle RAC environment.

Syntax

```
DBMS_CACHEUTIL.LIST_READMOSTLY;
```

45.3.6 OBJECT_DOWNCONVERT Procedure

OBJECT_DOWNCONVERT down-converts an object's exclusive locks on an instance to shared mode for all instances in an Oracle RAC environment.

This is useful when a particular object is loaded through the buffer cache by an Oracle RAC node and this node holds many exclusive locks for this object, which has caused extra locking overhead. Down-converting locks for this object to shared locks on the RAC node could reduce the locking overhead and improve performance.

Syntax

```
DBMS_CACHEUTIL.OBJECT_DOWNCONVERT(
   schema                IN VARCHAR2,
   obj                   IN VARCHAR2,
   partition             IN VARCHAR2 := NULL,
   downconvert_index     IN BOOLEAN := TRUE);
```

Parameters

Table 45-6 OBJECT_DOWNCONVERT Procedure Parameters

Parameter	Description
schema	The name of the schema for the object.
obj	The name of the object.
partition	• If the object is not partitioned, then this is NULL. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.
downconvert_index	To down-convert dependent indexes, use TRUE.

45.3.7 TRUE_CACHE_KEEP Procedure

When you call this procedure on True Cache, it assigns the object to the `KEEP` buffer pool on that cache.

To use this procedure, the `DB_KEEP_CACHE_SIZE` initialization parameter must be configured on True Cache.

Syntax

```
DBMS_CACHEUTIL.TRUE_CACHE_KEEP(
   schema                 IN VARCHAR2,
   obj                    IN VARCHAR2,
   partition              IN VARCHAR2 := NULL);
```

Parameters

Table 45-7 TRUE_CACHE_KEEP Procedure Parameters

Parameter	Description
`schema`	The name of the schema for the object.
`obj`	The name of the object.
`partition`	• If the object is not partitioned, then this is `NULL`. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.

Related Topics

- Oracle True Cache User's Guide

45.3.8 TRUE_CACHE_UNKEEP Procedure

When an object on True Cache no longer needs to be in the `KEEP` buffer pool, use this procedure to remove the object's `KEEP` assignment.

Syntax

```
DBMS_CACHEUTIL.TRUE_CACHE_UNKEEP(
   schema                 IN VARCHAR2,
   obj                    IN VARCHAR2,
   partition              IN VARCHAR2 := NULL);
```

Parameters

Table 45-8 TRUE_CACHE_UNKEEP Procedure Parameters

Parameter	Description
`schema`	The name of the schema for the object.
`obj`	The name of the object.

Table 45-8 (Cont.) TRUE_CACHE_UNKEEP Procedure Parameters

Parameter	Description
`partition`	• If the object is not partitioned, then this is `NULL`. • If the object is partitioned, use the name of the partition segment. • If it's a composited partitioned object, use the name of the subpartition segment.

Related Topics

- Oracle True Cache User's Guide

46
DBMS_CLOUD

The `DBMS_CLOUD` package provides comprehensive support for working with data in Object Storage.

This chapter contains the following topics:

- DBMS_CLOUD Overview
- Installing the DBMS_CLOUD Package
- DBMS_CLOUD Security Model
- DBMS_CLOUD Exceptions
- Summary of DBMS_CLOUD Subprograms
- DBMS_CLOUD Package File URI Formats
- DBMS_CLOUD Package Format Options
- DBMS_CLOUD Avro and Parquet Support

46.1 DBMS_CLOUD Overview

Using `DBMS_CLOUD` to work with data in Object Storage you can do the following:

- Manage your credentials for accessing Object Store.
- Copy data as-is: External data is loaded into the database in a single one-step operation, without the need to create any objects in the database. The load operation is monitored and tracked inside the database.
- Manage the information about your load operations.
- Validate the content of external data prior to consumption. Prior to consumption of external data you verify its content and identify rows that do not conform to the expected external data format.
- Create an external table. Create the necessary object in the database for flexible and continuous consumption.
- Work with objects in object stores and directories.

46.2 Installing the DBMS_CLOUD Package

The `DBMS_CLOUD` package provides a comprehensive support for working with data in Object Storage starting with Oracle Database 19c and later releases. However, this package is not pre-installed with Oracle Database. You need to manually install the `DBMS_CLOUD` package and also configure users or roles to use this package.

For information on installing the `DBMS_CLOUD` package and configuring users/roles, see the MOS-NOTE with Doc ID 2748362.1.

46.3 DBMS_CLOUD Security Model

Security on this package can be controlled by granting `EXECUTE` on this package to selected users or roles.

46.4 DBMS_CLOUD Exceptions

The following table describes exceptions for `DBMS_CLOUD`.

Exception	Code	Description
`reject_limit`	20003	The reject limit of an external table was reached.
`credential_not_exist`	20004	A credential object does not exist.
`table_not_exist`	20005	A table does not exist.
`unsupported_obj_store`	20006	An unsupported object store URI was provided.
`iden_too_long`	20008	An identifier is too long.
`invalid_format`	20009	A format argument is not valid.
`missing_credential`	20010	Mandatory credential object information was not specified.
`invalid_object_uri`	20011	An invalid object URI was provided.
`invalid_partitioning_clause`	20012	An partitioning clause is missing or was not provided.
`unsupported_feature`	20013	An unsupported feature was used that is not existent in the current database version.
`part_not_exist`	20014	A partition or subpartition does not exist, or a table is not a partitioned external table or hybrid partitioned table.
`invalid_table_name`	20016	An invalid table name was used.
`invalid_schema_name`	20017	An invalid schema name was used.
`invalid_dir_name`	20018	An invalid directory name was used.
`invalid_file_name`	20019	An invalid file name was used.
`invalid_cred_attribute`	20020	Invalid credential attributes were specified.
`table_exist`	20021	A table already exists.
`credential_exist`	20022	A credential object already exists.
`invalid_req_method`	20023	A request method is either too long or invalid.
`invalid_req_header`	20024	An invalid request header was specified.
`file_not_exist`	20025	A file does not exist.
`invalid_response`	20026	An HTTP response was not valid.
`invalid_operation`	20027	An invalid task class or ID was specified.
`invalid_user_name`	20028	An invalid username was specified.

46.5 Summary of DBMS_CLOUD Subprograms

This section covers the `DBMS_CLOUD` subprograms provided with Oracle Database.

Subprogram	Description
COPY_DATA Procedure	This procedure loads data into existing Oracle Database tables from files in the Cloud.

Subprogram	Description
COPY_DATA Procedure for Avro or Parquet Files	This procedure with the `format` parameter `type` set to the value `avro` or `parquet` loads data into existing Oracle Database tables from Avro or Parquet files in the Cloud. Similar to text files, the data is copied from the source Avro or Parquet file into the preexisting internal table.
CREATE_CREDENTIAL Procedure	This procedure stores cloud service credentials in Oracle Database.
CREATE_EXTERNAL_TABLE Procedure	This procedure creates an external table on files in the Cloud. This allows you to run queries on external data from Oracle Database.
CREATE_EXTERNAL_TABLE Procedure for Avro or Parquet Files	This procedure with the `format` parameter `type` set to the value `avro` or `parquet` creates an external table with either Avro or Parquet format files in the Cloud. This allows you to run queries on external data from Oracle Database.
CREATE_EXTERNAL_PART_TABLE Procedure	This procedure creates an external partitioned table on files in the Cloud. This allows you to run queries on external data from Oracle Database.
CREATE_HYBRID_PART_TABLE Procedure	This procedure creates a hybrid partitioned table. This allows you to run queries on hybrid partitioned data from Oracle Database.
DELETE_ALL_OPERATIONS Procedure	This procedure clears either all data load operations logged in the `user_load_operations` table in your schema or clears all the data load operations of the specified type, as indicated with the `type` parameter.
DELETE_FILE Procedure	This procedure removes the specified file from the specified directory on Oracle Database
DELETE_OBJECT Procedure	This procedure deletes the specified object on object store.
DROP_CREDENTIAL Procedure	This procedure removes an existing credential from Oracle Database.
GET_OBJECT Procedure and Function	This procedure is overloaded. The procedure form reads an object from Cloud Object Storage and copies it to Oracle Database. The function form reads an object from Cloud Object Storage and returns a `BLOB` to Oracle Database.
LIST_FILES Function	This function lists the files in the specified directory. The results include the file names and additional metadata about the files such as file size in bytes, creation timestamp, and the last modification timestamp.
LIST_OBJECTS Function	This function lists objects in the specified location on object store. The results include the object names and additional metadata about the objects such as size, checksum, creation timestamp, and the last modification timestamp.
PUT_OBJECT Procedure	This procedure is overloaded. In one form the procedure copies a file from Oracle Database to the Cloud Object Storage. In another form the procedure copies a `BLOB` from Oracle Database to the Cloud Object Storage.
UPDATE_CREDENTIAL Procedure	This procedure updates cloud service credential attributes in Oracle Database.
VALIDATE_EXTERNAL_TABLE Procedure	This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Oracle Database.

Subprogram	Description
VALIDATE_EXTERNAL_PART_TABLE Procedure	This procedure validates the source files for an external partitioned table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Oracle Database.
VALIDATE_HYBRID_PART_TABLE Procedure	This procedure validates the source files for a hybrid partitioned table, generates log information, and stores the rows that do not match the format options specified for the hybrid table in a *badfile* table on Oracle Database.

46.5.1 COPY_DATA Procedure

This procedure loads data into existing Oracle Database tables from files in the Cloud, or from files in a directory. The overloaded form enables you to use the `operation_id` parameter.

Syntax

```
DBMS_CLOUD.COPY_DATA (
    table_name         IN VARCHAR2,
    credential_name    IN VARCHAR2 DEFAULT NULL,
    file_uri_list      IN CLOB,
    schema_name        IN VARCHAR2,
    field_list         IN CLOB,
    format             IN CLOB);

DBMS_CLOUD.COPY_DATA (
    table_name         IN VARCHAR2,
    credential_name    IN VARCHAR2 DEFAULT NULL,
    file_uri_list      IN CLOB DEFAULT NULL,
    schema_name        IN VARCHAR2 DEFAULT NULL,
    field_list         IN CLOB DEFAULT NULL,
    format             IN CLOB DEFAULT NULL
    operation_id       OUT NOCOPY NUMBER);
```

Parameters

Parameter	Description
`table_name`	The name of the target table on the database. The target table needs to be created before you run `COPY_DATA`.
`credential_name`	The name of the credential to access the Cloud Object Storage. This parameter is not used when you specify a directory with `file_uri_list`.

Parameter	Description
file_uri_list	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards in the file names in your URIs. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is: '*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. You can use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. For example: '*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example: '*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example: '"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example: '*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote (') .
schema_name	The name of the schema where the target table resides. The default value is NULL meaning the target table is in the same schema as the user running the procedure.
field_list	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the target table definition. This argument's syntax is the same as the field_list clause in regular Oracle external tables. For more information about field_list see *Oracle® Database Utilities*. For an example using field_list, see CREATE_EXTERNAL_TABLE Procedure.
format	The options describing the format of the source, log, and bad files. For the list of the options and how to specify the values see DBMS_CLOUD Package Format Options. For Avro or Parquet file format options, see DBMS_CLOUD Package Format Options for Avro or Parquet.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.

Usage Note

The default record delimiter is detected newline. With detected newline, DBMS_CLOUD tries to automatically find the correct newline character to use as the record delimiter. DBMS_CLOUD first searches for the Windows newline character \r\n. If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, DBMS_CLOUD searches for the UNIX/Linux newline character \n, and if it finds one it uses \n as the record delimiter for all files in the procedure. If the source files use a combination of different record delimiters, you may encounter an error such as, "KUP-04020:

`found record longer than buffer size supported`". In this case, you need to either modify the source files to use the same record delimiter or only specify the source files that use the same record delimiter.

See DBMS_CLOUD Package Format Options for information on the `recorddelmiter` format option.

46.5.2 COPY_DATA Procedure for Avro or Parquet Files

This procedure with the `format` parameter `type` set to the value `avro` or `parquet` loads data into existing Oracle Database tables from Avro or Parquet files in the Cloud or from files in a directory.

Similar to text files, the data is copied from the source Avro or Parquet file into the preexisting internal table.

Syntax

```
DBMS_CLOUD.COPY_DATA (
    table_name         IN VARCHAR2,
    credential_name    IN VARCHAR2 DEFAULT NULL,
    file_uri_list      IN CLOB,
    schema_name        IN VARCHAR2 DEFAULT,
    field_list         IN CLOB DEFAULT,
    format             IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the target table on the database. The target table needs to be created before you run `COPY_DATA`.
`credential_name`	The name of the credential to access the Cloud Object Storage. This parameter is not used when you specify a directory with `file_uri_list`.

Parameter	Description
file_uri_list	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards in the file names in your URIs. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is: '*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. You can use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. For example: '*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example: '*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example: '"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example: '*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote (') .
schema_name	The name of the schema where the target table resides. The default value is NULL meaning the target table is in the same schema as the user running the procedure.
field_list	Ignored for Avro or Parquet files. The fields in the source match the external table columns by name. Source data types are converted to the external table column data type. For Parquet files, see DBMS_CLOUD Package Parquet to Oracle Data Type Mapping for details on mapping. For Avro files, see DBMS_CLOUD Package Avro to Oracle Data Type Mapping for details on mapping.
format	The options describing the format of the source files. For Avro or Parquet files, only two options are supported: see DBMS_CLOUD Package Format Options for Avro or Parquet.

Usage Notes

- As with other data files, Avro and Parquet data loads generate logs that are viewable in the tables dba_load_operations and user_load_operations. Each load operation adds a record to dba[*user*]_load_operations that indicates the table containing the logs.

  The log table provides summary information about the load.

- For Avro or Parquet, when the format parameter type is set to the value avro or parquet, the BADFILE_TABLE table is always empty.

  – For Parquet files, PRIMARY KEY constraint errors throw an ORA error.

- If data for a column encounters a conversion error, for example, the target column is not large enough to hold the converted value, the value for the column is set to `NULL`. This does not produce a rejected record.

46.5.3 CREATE_CREDENTIAL Procedure

This procedure stores cloud service credentials in Oracle Database.

Use stored cloud service credentials to access the cloud service for data loading, for querying external data residing in the cloud, or for other cases when you use `DBMS_CLOUD` procedures with a `credential_name` parameter. This procedure is overloaded:

- Use the Oracle Cloud Infrastructure-related parameters, including: `user_ocid`, `tenancy_ocid`, `private_key`, and `fingerprint` only when you are using Oracle Cloud Infrastructure Signing Keys authentication.
- Use the AWS ARN-related parameter `params`, only when you are using Amazon Resource Names (ARNs) credentials.

Syntax

```
DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name   IN VARCHAR2,
    username          IN VARCHAR2,
    password          IN VARCHAR2 DEFAULT NULL);


DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name IN VARCHAR2,
    user_ocid       IN VARCHAR2,
    tenancy_ocid    IN VARCHAR2,
    private_key     IN VARCHAR2,
    fingerprint     IN VARCHAR2);


DBMS_CLOUD.CREATE_CREDENTIAL (
    credential_name  IN VARCHAR2,
    params           IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to be stored. The `credential_name` parameter must conform to Oracle object naming conventions, which do not allow spaces or hyphens.
`username`	The `username` and `password` arguments together specify your cloud service credentials. See the usage notes for what to specify for the `username` and `password` for different cloud services.
`password`	The `username` and `password` arguments together specify your cloud service credentials.
`user_ocid`	Specifies the user's OCID. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the User's OCID.
`tenancy_ocid`	Specifies the tenancy's OCID. See Where to Get the Tenancy's OCID and User's OCID for details on obtaining the Tenancy's OCID.

Parameter	Description
private_key	Specifies the generated private key. Private keys generated with a passphrase are not supported. You need to generate the private key without a passphrase. See How to Generate an API Signing Key for details on generating a key pair in PEM format.
fingerprint	Specifies a fingerprint. After a generated public key is uploaded to the user's account the fingerprint is displayed in the console. Use the displayed fingerprint for this argument. See How to Get the Key's Fingerprint and How to Generate an API Signing Key for more details.
params	Specifies credential parameters using Amazon Resource Names (ARNs) credentials.

Usage Notes

- This operation stores the credentials in the database in an encrypted format.
- You can see the credentials in your schema by querying the `user_credentials` table.
- The `ADMIN` user can see all the credentials by querying the `dba_credentials` table.
- You only need to create credentials once unless your cloud service credentials change. Once you store the credentials you can then use the same credential name for `DBMS_CLOUD` procedures that require a `credential_name` parameter.
- This procedure is overloaded. If you provide one of the key based authentication attributes, `user_ocid`, `tenancy_ocid`, `private_key`, or `fingerprint`, the call is assumed to be an Oracle Cloud Infrastructure Signing Key based credential.
- You can list credentials from the view `ALL_CREDENTIALS`. For example, run the following command to list credentials:

```
SELECT credential_name, username, comments FROM all_credentials;
```

Oracle Cloud Infrastructure Credentials (Auth Tokens)

For Oracle Cloud Infrastructure the `username` is your Oracle Cloud Infrastructure user name. The `password` is your Oracle Cloud Infrastructure auth token. See Working with Auth Tokens.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'DEF_CRED_NAME',
    username => 'adwc_user@example.com',
    password => 'password' );
END;
/
```

Oracle Cloud Infrastructure Signing Key Based Credentials

Use the Oracle Cloud Infrastructure signing key related parameters, including: `user_ocid`, `tenancy_ocid`, `private_key`, and `fingerprint` with Oracle Cloud Infrastructure Signing Keys authentication.

For example:

```
BEGIN
   DBMS_CLOUD.CREATE_CREDENTIAL (
       credential_name => 'OCI_KEY_CRED',
       user_ocid       =>
'ocid1.user.oc1..aaaaaaaauq54mi7zdyfhw33ozkwuontjceel7fok5nq3bf2vwetkpqsoa',
       tenancy_ocid    =>
'ocid1.tenancy.oc1..aabbbbbbaafcue47pqmrf4vigneebgbcmmoy5r7xvoypicjqqge32ewnrc
yx2a',
       private_key     =>
'MIIEogIBAAKCAQEAtUnxbmrekwgVac6FdWeRzoXvIpA9+0r1.....wtnNpESQQQ0QLGPD8NM//
JEBg=',
       fingerprint     => 'f2:db:f9:18:a4:aa:fc:94:f4:f6:6c:39:96:16:aa:27');
END;
/
```

Private keys generated with a passphrase are not supported. You need to generate the private key without a passphrase. See How to Generate an API Signing Key for more information.

Oracle Cloud Infrastructure Object Storage Classic Credentials

If your source files reside in Oracle Cloud Infrastructure Object Storage Classic, the `username` is your Oracle Cloud Infrastructure Classic user name and the `password` is your Oracle Cloud Infrastructure Classic password.

Amazon Web Services (AWS) Credentials

If your source files reside in Amazon S3 or you are calling an AWS API, the `username` is your AWS access key ID and the `password` is your AWS secret access key. See AWS Identity and Access Management.

Microsoft Azure Credentials

If your source files reside in Azure Blob Storage or you are calling an Azure API, the `username` is your Azure storage account name and the `password` is an Azure storage account access key. See About Azure storage accounts.

AWS Amazon Resource Names (ARN) Credentials

If your source files reside in Amazon S3 or you are calling an AWS API, use `params` to specify the parameters for the Amazon Resource Names (ARN).

Parameter	Value
`aws_role_arn`	Specifies the Amazon Resource Name (ARN) that identifies the AWS role. If this parameter is not supplied when creating the credential, `ORA-20041` is raised.
`external_id_type`	Optionally set the `external_id_type` to use the Oracle Database compartment OCID, database OCID, or tenancy OCID by supplying one of: `compartment_ocid`, `database_ocid`, or `tenant_ocid`. If this parameter is not given when creating the credential, the default value is `database_ocid`.

For example:

```
BEGIN DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name      => 'MY_CRED',
        params               => JSON_OBJECT(
               'aws_role_arn'      value 'arn:aws:iam::123456:role/
AWS_ROLE_ARN',
               'external_id_type'  value 'database_ocid'));
END;
/
```

GitHub Personal Access Token

If your source files reside in a GitHub repository or you are calling a GitHub API, the `username` is your GitHub email and the `password` is your GitHub personal access token. See Creating a personal access token for more information.

For example:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'MY_GITHUB_CRED',
    username => 'user@example.com',
    password => 'your_personal_access_token' );
END;
/
```

46.5.4 CREATE_EXTERNAL_TABLE Procedure

This procedure creates an external table on files in the Cloud or from files in a directory. This allows you to run queries on external data from Oracle Database.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       IN VARCHAR2,
    credential_name  IN VARCHAR2 DEFAULT NULL,
    file_uri_list    IN CLOB,
    column_list      IN CLOB,
    field_list       IN CLOB DEFAULT,
    format           IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage. This parameter is not used when you specify a directory with `file_uri_list`.

Parameter	Description
file_uri_list	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards in the file names in your URIs. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is:'*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. You can use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. For example:'*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example:'*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example:'"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example:'*MY_DIR*:''filename.ext'. This specifies the filename starts with a quote ('). **Usage Notes** • The procedure DBMS_CLOUD.CREATE_EXTERNAL_TABLE supports external partitioned files in the supported cloud object storage services, including: – Oracle Cloud Infrastructure Object Storage – Azure Blob Storage – Amazon S3 – GitHub Repository The credential is a table level property; therefore, the external files must be on the same object store. See DBMS_CLOUD Package File URI Formats for more information.
column_list	Comma-delimited list of column names and data types for the external table.
field_list	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the *column_list* parameter. This argument's syntax is the same as the field_list clause in regular Oracle external tables. For more information about field_list see *Oracle® Database Utilities*.
format	The options describing the format of the source files. For the list of the options and how to specify the values see DBMS_CLOUD Package Format Options. For Avro or Parquet format files, see CREATE_EXTERNAL_TABLE Procedure for Avro or Parquet Files.

- The default record delimiter is `detected newline`. With `detected newline`, `DBMS_CLOUD` tries to automatically find the correct newline character to use as the record delimiter. `DBMS_CLOUD` first searches for the Windows newline character `\r\n`. If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, `DBMS_CLOUD` searches for the UNIX/Linux newline character `\n`, and if it finds one it uses `\n` as the record delimiter for all files in the procedure. If the source files use a combination of different record delimiters, you may encounter an error such as, "`KUP-04020: found record longer than buffer size supported`". In this case, you need to either modify the source files to use the same record delimiter or only specify the source files that use the same record delimiter.

  See DBMS_CLOUD Package Format Options for information on the `recorddelmiter` format option.

Example

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_TABLE(
      table_name =>'WEATHER_REPORT_DOUBLE_DATE',
      credential_name =>'OBJ_STORE_CRED',
      file_uri_list =>'&base_URL/
Charlotte_NC_Weather_History_Double_Dates.csv',
      format => json_object('type' value 'csv', 'skipheaders' value '1'),
      field_list => 'REPORT_DATE DATE''mm/dd/yy'',
                     REPORT_DATE_COPY DATE ''yyyy-mm-dd'',
                     ACTUAL_MEAN_TEMP,
                     ACTUAL_MIN_TEMP,
                     ACTUAL_MAX_TEMP,
                     AVERAGE_MIN_TEMP,
                     AVERAGE_MAX_TEMP,
                     AVERAGE_PRECIPITATION',
      column_list => 'REPORT_DATE DATE,
                     REPORT_DATE_COPY DATE,
                     ACTUAL_MEAN_TEMP NUMBER,
                     ACTUAL_MIN_TEMP NUMBER,
                     ACTUAL_MAX_TEMP NUMBER,
                     AVERAGE_MIN_TEMP NUMBER,
                     AVERAGE_MAX_TEMP NUMBER,
                     AVERAGE_PRECIPITATION NUMBER');
   END;
/

SELECT * FROM WEATHER_REPORT_DOUBLE_DATE where
   actual_mean_temp > 69 and actual_mean_temp < 74
```

46.5.5 CREATE_EXTERNAL_TABLE Procedure for Avro or Parquet Files

This procedure with the `format` parameter `type` set to the value `avro` or `parquet` creates an external table with either Avro or Parquet format files in the Cloud or from files in a directory.

This allows you to run queries on external data from Oracle Database.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_TABLE (
    table_name       IN VARCHAR2,
    credential_name  IN VARCHAR2 DEFAULT NULL,
    file_uri_list    IN CLOB,
    column_list      IN CLOB,
    field_list       IN CLOB DEFAULT,
    format           IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage. This parameter is not used when you specify a directory with `file_uri_list`.
`file_uri_list`	This parameter specifies either a comma-delimited list of source file URIs or one or more directories and source files. **Cloud source file URIs** You can use wildcards in the file names in your URIs. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. The format of the URIs depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats. **Directory** You can specify one directory and one or more file names or use a comma separated list of directories and file names. The format to specify a directory is:'*MY_DIR*:*filename.ext*'. By default the directory name *MY_DIR* is a database object and is case-insensitive. The file name is case sensitive. You can use wildcards to specify file names in a directory. The character "*" can be used as the wildcard for multiple characters, the character "?" can be used as the wildcard for a single character. For example:'*MY_DIR*:*" or '*MY_DIR*:test?' To specify multiple directories, use a comma separated list of directories: For example:'*MY_DIR1*:*, *MY_DIR2*:test?' Use double quotes to specify a case-sensitive directory name. For example:'"*my_dir1*":*, "*my_dir2*":Test?' To include a quote character, use two quotes. For example:'*MY_DIR*:''filename.ext'. This specifies the `filename` starts with a quote (`'`).

Parameter	Description
column_list	(Optional) This field, when specified, overrides the format->schema parameter which specifies that the schema, columns, and data types, are derived automatically. See the format parameter for details. When the column_list is specified for Avro or Parquet source, the column names must match those columns found in the file. Oracle data types must map appropriately to the Avro or Parquet data types. For Parquet files, see DBMS_CLOUD Package Parquet to Oracle Data Type Mapping for details. For Avro files, see DBMS_CLOUD Package Avro to Oracle Data Type Mapping for details.
field_list	Ignored for Avro or Parquet files. The fields in the source match the external table columns by name. Source data types are converted to the external table column data type. For Parquet files, see DBMS_CLOUD Package Parquet to Oracle Data Type Mapping for details. For Avro files, see DBMS_CLOUD Package Avro to Oracle Data Type Mapping for details.
format	For Avro or Parquet, there are only two supported parameters. See DBMS_CLOUD Package Format Options for Avro or Parquet for details.

Examples Avro

```
format => '{"type":"avro", "schema": "all"}'
```

```
format => json_object('type' value 'avro', 'schema' value 'first')
```

Examples Parquet

```
format => '{"type":"parquet", "schema": "all"}'
```

```
format => json_object('type' value 'parquet', 'schema' value 'first')
```

Avro or Parquet Column Name Mapping to Oracle Column Names

See DBMS_CLOUD Package Avro and Parquet to Oracle Column Name Mapping for information on column name mapping and column name conversion usage in Oracle SQL.

46.5.6 CREATE_EXTERNAL_PART_TABLE Procedure

This procedure creates an external partitioned table on files in the Cloud. This allows you to run queries on external data from Oracle Database.

Syntax

```
DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
    table_name           IN VARCHAR2,
    credential_name      IN VARCHAR2,
    partitioning_clause  IN CLOB,
    column_list          IN CLOB,
```

```
    field_list           IN CLOB DEFAULT,
    format               IN CLOB DEFAULT);


DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE (
    table_name           IN VARCHAR2,
    credential_name      IN VARCHAR2,
    file_uri_list        IN VARCHAR2,
    column_list          IN CLOB,
    field_list           IN CLOB DEFAULT,
    format               IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage.
`partitioning_clause`	Specifies the complete partitioning clause, including the location information for individual partitions. If you use the `partitioning_clause` parameter, the `file_uri_list` parameter is not allowed.
`file_uri_list`	There are two options for the `file_uri_list` parameter: • A common-delimited list of individual file URIs without wildcards. • A single file URI with wildcards, and the wildcards can only be after the last slash "/". If you use the parameter `file_uri_list`, the `partitioning_clause` parameter is not allowed. The format of the URIs depends on the Cloud Object Storage service. See DBMS_CLOUD Package File URI Formats for more information.
`column_list`	Comma-delimited list of column names and data types for the external table. This parameter has the following requirements, depending on the type of the data files specified with the `file_uri_list` parameter: • The `column_list` parameter is required with unstructured files. Using unstructured files, for example with CSV text files, the `column_list` parameter must specify all the column names and data types inside the data file as well as the partition columns derived from the object name. • The `column_list` parameter is optional with structured files. For example, with Avro, ORC, or Parquet data files, the `column_list` is not required. When the `column_list` is not included, the `format` parameter `partition_columns` option must include specifications for both column names (`name`) and data types (`type`).
`field_list`	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the *column_list* parameter. This argument's syntax is the same as the `field_list` clause in regular Oracle external tables. For more information about `field_list` see *Oracle® Database Utilities*.

Parameter	Description
format	The format option partition_columns specifies the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE column names and data types of partition columns when the partition columns are derived from the file path, depending on the type of data file, structured or unstructured: • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE includes the column_list parameter and the data files are unstructured, such as with CSV text files, partition_columns does not include the data type. For example, use a format such as the following for this type of partition_columns specification: '"partition_columns":["state","zipcode"]' The data type is not required because it is specified in the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE column_list parameter. • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE does not include the column_list parameter and the data files are structured, such as Avro, ORC, or Parquet files, the partition_columns option includes both the column name, name sub-clause, and the data type, type sub-clause. For example, the following shows a partition_columns specification: '"partition_columns":[{"name":"country", "type":"varchar2(10)"}, {"name":"year", "type":"number"}, {"name":"month", "type":"varchar2(10)"}]' If the data files are unstructured and the type sub-clause is specified with partition_columns, the type sub-clause is ignored. For object names that are not based on hive format, the order of the partition_columns specified columns must match the order as they appear in the object name in the file path specified in the file_uri_list parameter. To see all the format parameter options describing the format of the source files, see DBMS_CLOUD Package Format Options.

Usage Notes

- You cannot call this procedure with both partitioning_clause and file_uri_list parameters.
- Specifying the column_list parameter is optional with structured data files, including Avro, Parquet, or ORC data files. If column_list is not specified, the format parameter partition_columns option must include both name and type.
- The column_list parameter is required with unstructured data files, such as CSV text files.
- The procedure DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE supports external partitioned files in the supported cloud object storage services, including:
  - Oracle Cloud Infrastructure Object Storage

  - Azure Blob Storage
  - Amazon S3
  - GitHub Repository

  See DBMS_CLOUD Package File URI Formats for more information.

- When you call `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` with the `file_uri_list` parameter, the types for columns specified in the Cloud Object Store file name must be one of the following types:

  ```
  VARCHAR2(n)
  NUMBER(n)
  NUMBER(p,s)
  NUMBER
  DATE
  TIMESTAMP(9)
  ```

- The default record delimiter is `detected newline`. With `detected newline`, `DBMS_CLOUD` tries to automatically find the correct newline character to use as the record delimiter. `DBMS_CLOUD` first searches for the Windows newline character `\r\n`. If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, `DBMS_CLOUD` searches for the UNIX/Linux newline character `\n`, and if it finds one it uses `\n` as the record delimiter for all files in the procedure. If the source files use a combination of different record delimiters, you may encounter an error such as, "`KUP-04020: found record longer than buffer size supported`". In this case, you need to either modify the source files to use the same record delimiter or only specify the source files that use the same record delimiter.

  See DBMS_CLOUD Package Format Options for information on the `recorddelmiter` format option.

- The external partitioned tables you create with `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from.

  - `file$path`: Specifies the file path text up to the beginning of the object name.
  - `file$name`: Specifies the object name, including all the text that follows the bucket name.

Examples

Example using the `partitioning_clause` parameter:

```
BEGIN
   DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
      table_name =>'PET1',
      credential_name =>'OBJ_STORE_CRED',
      format => json_object('delimiter' value ',', 'recorddelimiter' value
'newline', 'characterset' value 'us7ascii'),
      column_list => 'col1 number, col2 number, col3 number',
      partitioning_clause => 'partition by range (col1)
                              (partition p1 values less than (1000) location
                                  ( ''&base_URL//file_11.txt'')
                               ,
                               partition p2 values less than (2000) location
                                  ( ''&base_URL/file_21.txt'')
```

```
                                        ,
                                        partition p3 values less than (3000)
               location
                                           ( ''&base_URL/file_31.txt'')
                                        )'
                    );
                 END;
               /
```

Example using the `file_uri_list` and `column_list` parameters with unstructured data files:

```
BEGIN
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
   table_name => 'MYSALES',
   credential_name => 'DEF_CRED_NAME',
   file_uri_list     => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/*.csv',
   column_list       => 'product varchar2(100), units number, country varchar2(100), year
number, month varchar2(2)',
   field_list        => 'product, units', --[Because country, year and month are not in the
file, they are not listed in the field list]
   format            => '{"type":"csv", "partition_columns":["country","year","month"]}');
END;
/
```

Example using the `file_uri_list` without the `column_list` parameter with structured data files:

```
BEGIN
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
  table_name => 'MYSALES',
  credential_name => 'DEF_CRED_NAME',
  DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE(
    table_name      => 'MYSALES',
    credential_name => 'DEF_CRED_NAME',
    file_uri_list   => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/*.parquet',
    format          =>
        json_object('type' value 'parquet', 'schema' value 'first',
                    'partition_columns' value
                          json_array(
                                json_object('name' value 'country', 'type' value
'varchar2(100)'),
                                json_object('name' value 'year', 'type' value 'number'),
                                json_object('name' value 'month', 'type' value 'varchar2(2)')
                          )
         )
    );
END;
/
```

46.5.7 CREATE_HYBRID_PART_TABLE Procedure

This procedure creates a hybrid partitioned table. This allows you to run queries on hybrid partitioned data from Oracle Database.

Syntax

```
DBMS_CLOUD.CREATE_HYBRID_PART_TABLE (
    table_name           IN VARCHAR2,
    credential_name      IN VARCHAR2,
    partitioning_clause  IN CLOB,
    column_list          IN CLOB,
    field_list           IN CLOB DEFAULT,
    format               IN CLOB DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`credential_name`	The name of the credential to access the Cloud Object Storage.
`partitioning_clause`	Specifies the complete partitioning clause, including the location information for individual partitions.
`column_list`	Comma-delimited list of column names and data types for the external table.
`field_list`	Identifies the fields in the source files and their data types. The default value is NULL meaning the fields and their data types are determined by the *column_list* parameter. This argument's syntax is the same as the `field_list` clause in regular Oracle external tables. For more information about `field_list` see *Oracle® Database Utilities*.
`format`	The options describing the format of the source files. For the list of the options and how to specify the values see DBMS_CLOUD Package Format Options.

Usage Note

- The procedure `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` supports external partitioned files in the supported cloud object storage services, including:
  - Oracle Cloud Infrastructure Object Storage
  - Azure Blob Storage
  - Amazon S3
  - GitHub Repository

  The credential is a table level property; therefore, the external files must be on the same object store.

  See DBMS_CLOUD Package File URI Formats for more information.
- The external partitioned tables you create with `DBMS_CLOUD.CREATE_HYBRID_PART_TABLE` include two invisible columns `file$path` and `file$name`. These columns help identify which file a record is coming from.
  - `file$path`: Specifies the file path text up to the beginning of the object name.

- `file$name`: Specifies the object name, including all the text that follows the bucket name.

Example

```
BEGIN
   DBMS_CLOUD.CREATE_HYBRID_PART_TABLE(
      table_name =>'HPT1',
      credential_name =>'OBJ_STORE_CRED',
      format => json_object('delimiter' value ',', 'recorddelimiter' value
'newline', 'characterset' value 'us7ascii'),
      column_list => 'col1 number, col2 number, col3 number',
      partitioning_clause => 'partition by range (col1)
                                (partition p1 values less than (1000)
external location
                                    ( ''&base_URL/file_11.txt'')
                                ,
                                partition p2 values less than (2000)
external location
                                    ( ''&base_URL/file_21.txt'')
                                ,
                                partition p3 values less than (3000)
                                )'
     );
   END;
/
```

46.5.8 DELETE_ALL_OPERATIONS Procedure

This procedure clears either all data load operations logged in the `user_load_operations` table in your schema or clears all the data load operations of the specified type, as indicated with the `type` parameter.

Syntax

```
DBMS_CLOUD.DELETE_ALL_OPERATIONS (
    type      IN VARCHAR DEFAULT NULL);
```

Parameters

Parameter	Description
`type`	Specifies the type of operation to delete. Type values can be found in the `TYPE` column in the `user_load_operations` table. If no `type` is specified all rows are deleted.

Usage Note

- `DBMS_CLOUD.DELETE_ALL_OPERATIONS` does not delete currently running operations (operations in a "Running" status).

46.5.9 DELETE_FILE Procedure

This procedure removes the specified file from the specified directory on Oracle Database.

Syntax

```
DBMS_CLOUD.DELETE_FILE (
      directory_name     IN VARCHAR2,
      file_name          IN VARCHAR2);
```

Parameters

Parameter	Description
directory_name	The name of the directory on the Oracle Database instance.
file_name	The name of the file to be removed.

> **Note:**
>
> To run DBMS_CLOUD.DELETE_FILE, you need to grant write privileges on the directory that contains the file to the user. For example, run the following command as ADMIN to grant write privileges to db_user:
>
> ```
> GRANT WRITE ON DIRECTORY data_pump_dir TO db_user;
> ```

Example

```
BEGIN
   DBMS_CLOUD.DELETE_FILE(
      directory_name =>  'DATA_PUMP_DIR',
      file_name => 'exp1.dmp' );
   END;
/
```

46.5.10 DELETE_OBJECT Procedure

This procedure deletes the specified object on object store.

Syntax

```
DBMS_CLOUD.DELETE_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       force                IN BOOLEAN DEFAULT FALSE);
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage.
object_uri	Object or file URI for the object to delete. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats.
force	Ignore and do not report errors if object does not exist. Valid values are: TRUE and FALSE. The default value is FALSE.

Example

```
BEGIN
   DBMS_CLOUD.DELETE_OBJECT(
       credential_name => 'DEF_CRED_NAME',
       object_uri => 'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/exp1.dmp' );
   END;
/
```

46.5.11 DROP_CREDENTIAL Procedure

This procedure removes an existing credential from Oracle Database.

Syntax

```
DBMS_CLOUD.DROP_CREDENTIAL (
   credential_name     IN VARCHAR2);
```

Parameters

Parameter	Description
credential_name	The name of the credential to be removed.

46.5.12 GET_OBJECT Procedure and Function

This procedure is overloaded. The procedure form reads an object from Cloud Object Storage and copies it to Oracle Database. The function form reads an object from Cloud Object Storage and returns a BLOB to Oracle Database.

Syntax

```
DBMS_CLOUD.GET_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       directory_name       IN VARCHAR2,
       file_name            IN VARCHAR2 DEFAULT  NULL,
       startoffset          IN NUMBER DEFAULT  0,
       endoffset            IN NUMBER DEFAULT  0,
       compression          IN VARCHAR2 DEFAULT  NULL);
```

```
DBMS_CLOUD.GET_OBJECT(
       credential_name      IN VARCHAR2 DEFAULT NULL,
       object_uri           IN VARCHAR2,
       startoffset          IN NUMBER DEFAULT  0,
       endoffset            IN NUMBER DEFAULT  0,
       compression          IN VARCHAR2 DEFAULT  NULL)
RETURN BLOB;
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage.
object_uri	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats.
directory_name	The name of the directory on the database. [1]
file_name	Specifies the name of the file to create. If file name is not specified, the file name is taken from after the last slash in the object_uri parameter. For special cases, for example when the file name contains slashes, use the file_name parameter.
startoffset	The offset, in bytes, from where the procedure starts reading.
endoffset	The offset, in bytes, until where the procedure stops reading.
compression	Specifies the compression used to store the object. When compression is set to 'AUTO' the file is uncompressed (the value 'AUTO' implies the object specified with object_uri is compressed with Gzip).

> **Note:**
>
> To run DBMS_CLOUD.GET_OBJECT, you need to grant WRITE privileges on the directory to the user. For example, run the following command as ADMIN to grant write privileges to db_user:
>
> ```
> GRANT WRITE ON DIRECTORY data_pump_dir TO db_user;
> ```

Return Values

The function form reads from Object Store and DBMS_CLOUD.GET_OBJECT returns a BLOB.

Examples

```
BEGIN
   DBMS_CLOUD.GET_OBJECT(
     credential_name => 'OBJ_STORE_CRED',
     object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/file.txt',
     directory_name => 'DATA_PUMP_DIR');
```

```
END;
/
```

To read character data from a file in Object Store:

```
SELECT to_clob(
     DBMS_CLOUD.GET_OBJECT(
       credential_name => 'OBJ_STORE_CRED',
       object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/file.txt'))
FROM DUAL;
```

To add an image stored on Object Store in a `BLOB` in the database:

```
DECLARE
   l_blob BLOB := NULL;
BEGIN
   l_blob := DBMS_CLOUD.GET_OBJECT(
     credential_name => 'OBJ_STORE_CRED',
     object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/MyImage.gif' );
END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

46.5.13 LIST_FILES Function

This function lists the files in the specified directory. The results include the file names and additional metadata about the files such as file size in bytes, creation timestamp, and the last modification timestamp.

Syntax

```
DBMS_CLOUD.LIST_FILES (
    directory_name      IN VARCHAR2)
       RETURN TABLE;
```

Parameters

Parameter	Description
`directory_name`	The name of the directory on the database.

Usage Notes

- `DBMS_CLOUD.LIST_FILES` is only supported for directory objects mapping to Oracle File System (OFS) or Database File System (DBFS) file systems.

- To run DBMS_CLOUD.LIST_FILES, you need to grant read privileges on the directory to the user. For example, run the following command as ADMIN to grant read privileges to db_user:

```
GRANT READ ON DIRECTORY data_pump_dir TO db_user;
```

- This is a pipelined table function with return type as DBMS_CLOUD_TYPES.list_object_ret_t.
- DBMS_CLOUD.LIST_FILES does not obtain the checksum value and returns NULL for this field.

Example

This is a pipelined function that returns a row for each file. For example, use the following query to use this function:

```
SELECT * FROM DBMS_CLOUD.LIST_FILES('DATA_PUMP_DIR');

OBJECT_NAME       BYTES   CHECKSUM      CREATED              LAST_MODIFIED
------------ ---------- ----------    ---------------------
---------------------
cwallet.sso        2965               2018-12-12T18:10:47Z
2019-11-23T06:36:54Z
```

46.5.14 LIST_OBJECTS Function

This function lists objects in the specified location on object store. The results include the object names and additional metadata about the objects such as size, checksum, creation timestamp, and the last modification timestamp.

Syntax

```
DBMS_CLOUD.LIST_OBJECTS (
       credential_name      IN VARCHAR2,
       location_uri         IN VARCHAR2)
   RETURN TABLE;
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage.
location_uri	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats.

Usage Notes

- Depending on the capabilities of the object store, DBMS_CLOUD.LIST_OBJECTS does not return values for certain attributes and the return value for the field is NULL in this case.

  All supported Object Stores return values for the OBJECT_NAME, BYTES, and CHECKSUM fields.

The following table shows support for the fields `CREATED` and `LAST_MODIFIED` by Object Store:

Object Store	CREATED	LAST_MODIFIED
Oracle Cloud Infrastructure Native	Returns timestamp	Returns timestamp
Oracle Cloud Infrastructure Swift	Returns `NULL`	Returns timestamp
Oracle Cloud Infrastructure Classic	Returns `NULL`	Returns timestamp
Amazon S3	Returns `NULL`	Returns timestamp
Azure	Returns timestamp	Returns timestamp
GitHub Repository		

- The checksum value is the MD5 checksum. This is a 32-character hexadecimal number that is computed on the object contents. It is expected to have a different checksum value if `OCI$RESOURCE_PRINCIPAL` credential is used.
- This is a pipelined table function with return type as `DBMS_CLOUD_TYPES.list_object_ret_t`.

Example

This is a pipelined function that returns a row for each object. For example, use the following query to use this function:

```
SELECT * FROM DBMS_CLOUD.LIST_OBJECTS('OBJ_STORE_CRED',
    'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/');


OBJECT_NAME   BYTES              CHECKSUM
CREATED         LAST_MODIFIED
------------ ---------- --------------------------------
--------------------- --------------------
cwallet.sso   2965      2339a2731ba24a837b26d344d643dc07
2019-11-23T06:36:54Z
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

46.5.15 PUT_OBJECT Procedure

This procedure is overloaded. In one form the procedure copies a file from Oracle Database to the Cloud Object Storage. In another form the procedure copies a `BLOB` from Oracle Database to the Cloud Object Storage.

Syntax

```
DBMS_CLOUD.PUT_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       directory_name       IN VARCHAR2,
       file_name            IN VARCHAR2);
```

```
DBMS_CLOUD.PUT_OBJECT (
       credential_name      IN VARCHAR2,
       object_uri           IN VARCHAR2,
       contents             IN BLOB,
       file_name            IN VARCHAR2);
```

Parameters

Parameter	Description
credential_name	The name of the credential to access the Cloud Object Storage.
object_uri	Object or file URI. The format of the URI depends on the Cloud Object Storage service you are using, for details see DBMS_CLOUD Package File URI Formats.
directory_name	The name of the directory on the Oracle Database. [1]
file_name	The name of the file in the specified directory.

> **Note:**
>
> To run DBMS_CLOUD.PUT_OBJECT, you need to grant read privileges on the directory to the user. For example, run the following command as ADMIN to grant read privileges to db_user:
>
> ```
> GRANT READ ON DIRECTORY data_pump_dir TO db_user;
> ```

Example

To handle BLOB data after in-database processing and then store the data directly into a file in the object store:

```
DECLARE
      my_blob_data BLOB;
BEGIN
 /* Some processing producing BLOB data and populating my_blob_data */
DBMS_CLOUD.PUT_OBJECT(
     credential_name => 'OBJ_STORE_CRED',
     object_uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/my_new_file',
     contents => my_blob_data));
END;
/
```

Usage Notes

Depending on your Cloud Object Storage, the size of the object you transfer is limited as follows:

Cloud Object Storage Service	Object Transfer Size Limit
Oracle Cloud Infrastructure Object Storage	50 GB

Cloud Object Storage Service	Object Transfer Size Limit
Amazon S3	5 GB
Azure Blob Storage	256 MB

Oracle Cloud Infrastructure object store does not allow writing files into a public bucket without supplying credentials (Oracle Cloud Infrastructure allows users to download objects from public buckets). Thus, you must supply a credential name with valid credentials to store an object in an Oracle Cloud Infrastructure public bucket using `PUT_OBJECT`.

See DBMS_CLOUD Package File URI Formats for more information.

46.5.16 UPDATE_CREDENTIAL Procedure

This procedure updates cloud service credential attributes in Oracle Database.

Use stored credentials for data loading, for querying external data residing in the Cloud, or wherever you use `DBMS_CLOUD` procedures with a `credential_name` parameter. This procedure lets you update an attribute with a new value for a specified `credential_name`.

Syntax

```
DBMS_CLOUD.UPDATE_CREDENTIAL (
    credential_name   IN VARCHAR2,
    attribute         IN VARCHAR2,
    value             IN VARCHAR2);
```

Parameters

Parameter	Description
`credential_name`	The name of the credential to be stored.
`attribute`	Name of attribute to update: `USERNAME` or `PASSWORD`.
`value`	New value for the selected attribute.

Usage Notes

- The user name is case sensitive. It cannot contain double quotes or spaces.
- You only need to create credentials once unless your cloud service credentials change. Once you store the credentials you can then use the same credential name for `DBMS_CLOUD` procedures that require a `credential_name` parameter.

Example

```
BEGIN
  DBMS_CLOUD.UPDATE_CREDENTIAL(
     credential_name => 'OBJ_STORE_CRED',
     attribute => 'PASSWORD',
     value => 'password');
END;
/
```

46.5.17 VALIDATE_EXTERNAL_TABLE Procedure

This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Oracle Database. The overloaded form enables you to use the `operation_id` parameter.

Syntax

```
DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE (
    table_name      IN VARCHAR2,
    schema_name     IN VARCHAR2 DEFAULT,
    rowcount        IN NUMBER DEFAULT,
    stop_on_error   IN BOOLEAN DEFAULT);


DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE(
    table_name      IN VARCHAR2,
    operation_id    OUT NOCOPY NUMBER,
    schema_name     IN VARCHAR2 DEFAULT NULL,
    rowcount        IN NUMBER DEFAULT 0,
    stop_on_error   IN BOOLEAN DEFAULT TRUE);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`operation_id`	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the `USER_LOAD_OPERATIONS` view.
`schema_name`	The name of the schema where the external table resides. The default value is NULL meaning the external table is in the same schema as the user running the procedure.
`rowcount`	Number of rows to be scanned. The default value is NULL meaning all the rows in the source files are scanned.
`stop_on_error`	Determines if the validate should stop when a row is rejected. The default value is `TRUE` meaning the validate stops at the first rejected row. Setting the value to `FALSE` specifies that the validate does not stop at the first rejected row and validates all rows up to the value specified for the `rowcount` parameter. If the external table refers to Avro or Parquet files then the validate stops at the first rejected row. When the external table specifies the `format` parameter `type` set to the value `avro` or `parquet`, the parameter `stop_on_error` effectively always has the value `TRUE`. Thus, the table badfile will always be empty for an external table referring to Avro or Parquet files.

Usage Notes

- `DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE` works with both partitioned external tables and hybrid partitioned tables. This potentially reads data from all external partitions until `rowcount` is reached or `stop_on_error` applies. You do not have control over which partition, or parts of a partition, is read in which order.

46.5.18 VALIDATE_EXTERNAL_PART_TABLE Procedure

This procedure validates the source files for an external partitioned table, generates log information, and stores the rows that do not match the format options specified for the external table in a *badfile* table on Oracle Database. The overloaded form enables you to use the `operation_id` parameter.

Syntax

```
DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
       table_name                 IN VARCHAR2,
       partition_name             IN CLOB DEFAULT,
       schema_name                IN VARCHAR2 DEFAULT,
       rowcount                   IN NUMBER DEFAULT,
       partition_key_validation   IN BOOLEAN DEFAULT,
       stop_on_error              IN BOOLEAN DEFAULT);


DBMS_CLOUD.VALIDATE_EXTERNAL_PART_TABLE (
       table_name                 IN VARCHAR2,
       operation_id               OUT NUMBER,
       partition_name             IN CLOB DEFAULT,
       schema_name                IN VARCHAR2 DEFAULT,
       rowcount                   IN NUMBER DEFAULT,
       partition_key_validation   IN BOOLEAN DEFAULT,
       stop_on_error              IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
`table_name`	The name of the external table.
`operation_id`	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the `USER_LOAD_OPERATIONS` view.
`partition_name`	If defined, then only a specific partition is validated. If not specified then read all partitions sequentially until `rowcount` is reached.
`schema_name`	The name of the schema where the external table resides. The default value is NULL meaning the external table is in the same schema as the user running the procedure.
`rowcount`	Number of rows to be scanned. The default value is NULL meaning all the rows in the source files are scanned.
`partition_key_validation`	For internal use only. Do not use this parameter.

Parameter	Description
stop_on_error	Determines if the validate should stop when a row is rejected. The default value is TRUE meaning the validate stops at the first rejected row. Setting the value to FALSE specifies that the validate does not stop at the first rejected row and validates all rows up to the value specified for the rowcount parameter. If the external table refers to Avro or Parquet files then the validate stops at the first rejected row. When the external table specifies the format parameter type set to the value avro or parquet, the parameter stop_on_error effectively always has the value TRUE. Thus, the table badfile will always be empty for an external table referring to Avro or Parquet files.

46.5.19 VALIDATE_HYBRID_PART_TABLE Procedure

This procedure validates the source files for a hybrid partitioned table, generates log information, and stores the rows that do not match the format options specified for the hybrid table in a *badfile* table on Oracle Database. The overloaded form enables you to use the operation_id parameter.

Syntax

```
DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE (
       table_name                 IN VARCHAR2,
       partition_name             IN CLOB DEFAULT,
       schema_name                IN VARCHAR2 DEFAULT,
       rowcount                   IN NUMBER DEFAULT,
       partition_key_validation   IN BOOLEAN DEFAULT,
       stop_on_error              IN BOOLEAN DEFAULT);


DBMS_CLOUD.VALIDATE_HYBRID_PART_TABLE (
       table_name                 IN VARCHAR2,
       operation_id               OUT NUMBER,
       partition_name             IN CLOB DEFAULT,
       schema_name                IN VARCHAR2 DEFAULT,
       rowcount                   IN NUMBER DEFAULT,
       partition_key_validation   IN BOOLEAN DEFAULT,
       stop_on_error              IN BOOLEAN DEFAULT);
```

Parameters

Parameter	Description
table_name	The name of the external table.
operation_id	Use this parameter to track the progress and final status of the load operation as the corresponding ID in the USER_LOAD_OPERATIONS view.
partition_name	If defined, then only a specific partition is validated. If not specified then read from all external partitions sequentially until rowcount is reached.
schema_name	The name of the schema where the external table resides. The default value is NULL meaning the external table is in the same schema as the user running the procedure.

Parameter	Description
rowcount	Number of rows to be scanned. The default value is NULL meaning all the rows in the source files are scanned.
partition_key_validation	For internal use only. Do not use this parameter.
stop_on_error	Determines if the validate should stop when a row is rejected. The default value is TRUE meaning the validate stops at the first rejected row. Setting the value to FALSE specifies that the validate does not stop at the first rejected row and validates all rows up to the value specified for the rowcount parameter. If the external table refers to Avro or Parquet files then the validate stops at the first rejected row. When the external table specifies the format parameter type set to the value avro or parquet, the parameter stop_on_error effectively always has the value TRUE. Thus, the table badfile will always be empty for an external table referring to Avro or Parquet files.

46.6 DBMS_CLOUD Package File URI Formats

Describes the format of the source file URIs in operations with DBMS_CLOUD. The format depends on the object storage service you are using.

DBMS_CLOUD guarantees secure communication and any URI that you specify must use HTTPS, with https:// as the prefix for the URI.

Topics

- Oracle Cloud Infrastructure Object Storage Native URI Format
- Oracle Cloud Infrastructure Object Storage Swift URI Format
- Oracle Cloud Infrastructure Object Storage URI Format Using Pre-Authenticated Request URL
- URI Format Using Public URL
- Oracle Cloud Infrastructure Object Storage Classic URI Format
- Amazon S3 URI Format
- Azure Blob Storage URI Format

46.6.1 Oracle Cloud Infrastructure Object Storage Native URI Format

If your source files reside on the Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure native URIs, with the format:

https://objectstorage.*region*.oraclecloud.com/n/*namespace-string*/b/*bucket*/o/*filename*

For example, the Native URI for the file `channels.txt` in the *`bucketname`* bucket in the `Phoenix` data center is:

```
https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/channels.txt
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

You can find the URI from the Oracle Cloud Infrastructure Object Storage "Object Details" in the right hand side ellipsis menu in the Object Store:

1. Open the Oracle Cloud Infrastructure Console by clicking the ☰ next to Oracle Cloud.
2. From the Oracle Cloud Infrastructure left navigation menu click **Core Infrastructure**. Under **Object Storage**, click **Object Storage**.
3. Under List Scope, select a **Compartment**.
4. From the **Name** column, select a bucket.
5. In the Objects area, click **View Object Details**.

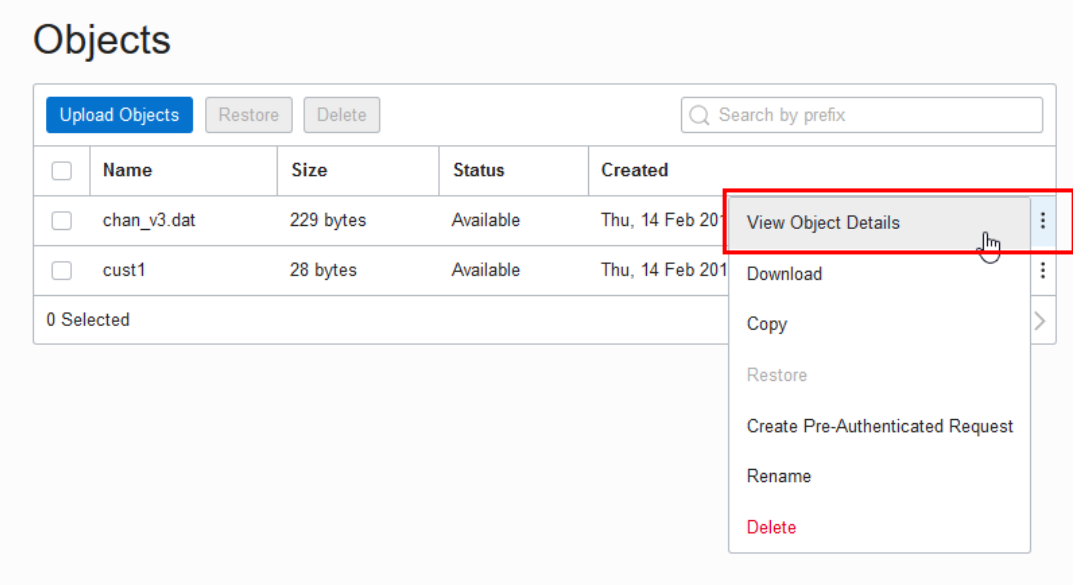

6. On the **Object Details** page, the **URL Path (URI)** field shows the URI to access the object.

> **Note:**
>
> The source files need to be stored in an Object Storage tier bucket. Oracle Database does not support buckets in the Archive Storage tier. See Overview of Object Storage for more information.

46.6.2 Oracle Cloud Infrastructure Object Storage Swift URI Format

If your source files reside on the Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure Swift URIs with the format:

```
https://swiftobjectstorage.region.oraclecloud.com/v1/namespace-string/bucket/
filename
```

For example, the Swift URI for the file `channels.txt` in the *bucketname* bucket in the `Phoenix` data center is:

```
https://swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/namespace-string/
bucketname/channels.txt
```

In this example, *namespace-string* is the Oracle Cloud Infrastructure object storage namespace and *bucketname* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> The source files need to be stored in an Object Storage tier bucket. Oracle Database does not support buckets in the Archive Storage tier. See Overview of Object Storage for more information.

46.6.3 Oracle Cloud Infrastructure Object Storage URI Format Using Pre-Authenticated Request URL

If your source files reside on the Oracle Cloud Infrastructure Object Storage you can use Oracle Cloud Infrastructure pre-authenticated URIs. When you create a pre-authenticated request, a unique URL is generated. You can then provide the unique URL to users in your organization, partners, or third parties to access the Object Storage resource target identified in the pre-authenticated request.

> **Note:**
>
> Carefully assess the business requirement for and the security ramifications of pre-authenticated access. When you create the pre-authenticated request URL, note the **Expiration** and the **Access Type** to make sure they are appropriate for your use. A pre-authenticated request URL gives anyone who has the URL access to the targets identified in the request for as long as the request is active. In addition to considering the operational needs of pre-authenticated access, it is equally important to manage its distribution.

The format for pre-authenticated request URLs is:

```
https://objectstorage.region.oraclecloud.com/p/encrypted_string/n/namespace-string/b/
bucket/o/filename
```

For example, a sample pre-authenticated URI for the file `channels.txt` in the *bucketname* bucket in the `Phoenix` data center is:

```
https://objectstorage.us-phoenix-1.oraclecloud.com/p/2xN-uDtWJNsiD910UCYGue/n/namespace-string/b/bucketname/o/channels.txt
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

You can use a pre-authenticated URL in any `DBMS_CLOUD` procedure that takes a URL to access files in Oracle Cloud Infrastructure object store, without the need to create a credential. You need to either specify the `credential_name` parameter as `NULL` or not supply a `credential_name` parameter.

For example:

```
BEGIN
  DBMS_CLOUD.COPY_DATA(
     table_name =>'CHANNELS',
     file_uri_list =>'https://objectstorage.us-phoenix-1.oraclecloud.com/p/
unique-pre-authenticated-string/n/namespace-string/b/bucketname/o/
channels.txt',
     format => json_object('delimiter' value ',') );
END;
/
```

> **Note:**
>
> A list of mixed URLs is valid. If the URL list contains both pre-authenticated URLs and URLs that require authentication, `DBMS_CLOUD` uses the specified `credential_name` to access the URLs that require authentication and for the pre-authenticated URLs the specified `credential_name` is ignored.

See Using Pre-Authenticated Requests for more information.

46.6.4 URI Format Using Public URL

If your source files reside on an Object Store that provides public URLs, you can use public URLs with `DBMS_CLOUD` procedures. Public means the Object Storage service supports anonymous, unauthenticated access to the Object Store files. See your Cloud Object Storage service for details on how to make an object public in a supported Object Store.

> **Note:**
>
> Carefully assess the business requirement for and the security ramifications of using public URLs. When you use public URLs, due to the file content not being authenticated, make sure this is appropriate for your use.

You can use a public URL in any `DBMS_CLOUD` procedure that takes a URL to access files in your object store, without the need to create a credential. You need to either specify the `credential_name` parameter as `NULL` or not supply a `credential_name` parameter.

For example the following uses `DBMS_CLOUD.COPY_DATA` without a `credential_name`:

```
BEGIN
  DBMS_CLOUD.COPY_DATA(
     table_name =>'CHANNELS',
     file_uri_list =>'https://objectstorage.us-ashburn-1.oraclecloud.com/n/
namespace-string/b/bucketname/o/chan_v3.dat',
     format => json_object('delimiter' value ',') );
END;
/
```

In this example, *`namespace-string`* is the Oracle Cloud Infrastructure object storage namespace and *`bucketname`* is the bucket name. See Understanding Object Storage Namespaces for more information.

> **Note:**
>
> A list of mixed URLs is valid. If the URL list contains both public URLs and URLs that require authentication, `DBMS_CLOUD` uses the specified `credential_name` to access the URLs that require authentication and for the public URLs the specified `credential_name` is ignored.

See Public Buckets for information on using Oracle Cloud Infrastructure public buckets.

46.6.5 Oracle Cloud Infrastructure Object Storage Classic URI Format

If your source files reside in Oracle Cloud Infrastructure Object Storage Classic, see the REST page for a description of the URI format for accessing your files: About REST URLs for Oracle Cloud Infrastructure Object Storage Classic Resources.

46.6.6 Amazon S3 URI Format

If your source files reside in Amazon S3, see the following for a description of the URI format for accessing your files: Accessing a Bucket.

For example the following refers to the file `channels.txt` in the `adb` bucket in the `us-west-2` region.

```
https://s3-us-west-2.amazonaws.com/adb/channels.txt
```

> **Note:**
>
> `DBMS_CLOUD` only supports native S3 storage, and no S3-compliant APIs or storage devices.

46.6.7 Azure Blob Storage URI Format

If your source files reside in Azure Blob Storage, see the following for a description of the URI format for accessing your files: Resource URI Syntax.

For example the following refers to the file `channels.txt` in the `adb` container in the storage account `db_user`:

```
https://db_user.blob.core.windows.net/adb/channels.txt
```

> **Note:**
>
> You can use Shared Access Signatures (SAS) URL in any `DBMS_CLOUD` procedure that takes a URL to access files in Azure Blob Storage, without the need to create a credential. To use a Shared Access Signature (SAS) URL, either specify the `credential_name` parameter as `NULL`, or do not supply a `credential_name` parameter.
> See Grant Limited Access to Azure Storage Resources Using Shared Access Signatures (SAS) for more information.

46.7 DBMS_CLOUD Package Format Options

The format argument in `DBMS_CLOUD` specifies the format of source files.

The two ways to specify the format argument are:

```
format => '{"format_option" : "format_value" }'
```

And:

```
format => json_object('format_option' value 'format_value'))
```

Examples:

```
format => json_object('type' VALUE 'CSV')
```

To specify multiple format options, separate the values with a ",".

For example:

```
format => json_object('ignoremissingcolumns' value 'true', 'removequotes' value 'true',
                      'dateformat' value 'YYYY-MM-DD-HH24-MI-SS', 'blankasnull'
value 'true', 'logretention' value 7)
```

> **Note:**
>
> For Avro or Parquet format options, see DBMS_CLOUD Package Format Options for Avro or Parquet.

Format Option	Description	Syntax
`blankasnull`	When set to `true`, loads fields consisting of spaces as null.	`blankasnull : true` Default value: `False`
`characterset` Valid with `format` JSON and `COPY_DATA`	Specifies the characterset of source files	`characterset`: *string* Default value: Database characterset
`columnpath` Only use with `format` JSON and `COPY_DATA`	Array of JSON path expressions that correspond to the fields that need to be extracted from the JSON records. Each of the JSON path expressions in the array should follow the rules described in SQL/JSON Path Expressions . Only use with format JSON and `DBMS_CLOUD.COPY_DATA`.	JSON Array of json path expressions expressed in string format. For example: 'columnpath' value `'["$.WEATHER_STATION_ID", "$.WEATHER_STATION_NAME"]'`
`compression`	Specifies the compression type of the source file. ZIP archiving format is not supported. Specifying the value `auto` checks for the compression types: `gzip, zlib, bzip2`.	`compression`: `auto|gzip|zlib|bzip2` Default value: Null value meaning no compression.
`conversionerrors`	If a row is rejected because of data type conversion errors, the related columns are stored as null or the row is rejected.	`conversionerrors` : `reject_record | store_null` Default value: `reject_record`

`dateformat`	Specifies the date format in the source file. The format option `AUTO` searches for the following formats:
	`J`
	`MM-DD-YYYYBC`
	`MM-DD-YYYY`
	`YYYYMMDD HHMISS`
	`YYMMDD HHMISS`
	`YYYY.DDD`
	`YYYY-MM-DD`
		`dateformat` : *string*
		Default value: Database date format

`delimiter`	Specifies the field delimiter. To use a special character as the delimiter, specify the HEX value of the ASCII code of the character. For example, the following specifies the TAB character as the delimiter: `format => json_object('delimiter' value 'X''9''')`	`delimiter` : *character* Default value `|` (pipe character)
`enablelogs`	Specifies a boolean value, when set to `TRUE`, logs are generated. When set to `FALSE`, logs are not generated. The value for `enablelogs` must be enclosed in single quotes. For example: `format => JSON_OBJECT('enablelogs' value 'FALSE')`	`enablelogs: false` Default value: `true`

Format Option	Description	Syntax
endquote	Data can be enclosed between two delimiters, specified with quote and endquote. The quote and endquote characters are removed during loading when specified. For example: format => JSON_OBJECT('quote' value '(', 'endquote' value ')')	endquote:character Default value: Null, meaning no endquote.
escape	The character "\" is used as the escape character when specified.	escape : true Default value: false
ignoreblanklines	Blank lines are ignored when set to true.	ignoreblanklines : true Default value: False
ignoremissingcolumns	If there are more columns in the field_list than there are in the source files, the extra columns are stored as null.	ignoremissingcolumns : true Default value False
language	Specifies a language name (for example, FRENCH), from which locale-sensitive information can be derived.	language: *string* Default value: Null See Locale Data in *Oracle Database Globalization Support Guide* for a listing of Oracle-supported languages.
logdir	Specifies a string value that determines the directory object name where the logfile_table or badfile_table files are saved. By default, the logdir is not case-sensitive, but the case is reserved when the specified value is enclosed in double-quotes. For example: format => JSON_OBJECT ('logdir' value 'test_log') The logdir format option specified in the above example saves the logfile_table or badfile_table files in the TEST_LOG directory object. format => JSON_OBJECT ('logdir' value '"test_log"') The logdir format option specified in the above example saves the logfile_table or badfile_table files in the test_log directory object.	logdir: *string* Default value: DATA_PUMP_DIR
logprefix	Specifies a string value that determines the prefix for the logfile_table and badfile_table files. The log table name format is: *logprefix*$*operation_id* By default, the logprefix is in upper case, but the case is reserved when the specified value is enclosed in double-quotes. For example: format => JSON_OBJECT ('logprefix' value 'TEST') Log files then use the TEST prefix, such as: TEST$2_LOG and TEST$2_BAD.	logprefix: *string* Default value: COPY

Format Option	Description	Syntax
`logretention`	Specifies a positive integer duration, in days, for which the `logfile_table` and `badfile_table` files are retained. Valid values: `0` to `99999` For example: `format => JSON_OBJECT ('logretention' value 7)`	`logretention`: *number* Default value: `2`
`numericcharacters`	Specifies the characters to use as the group separator and decimal character. *decimal_character*: The decimal separates the integer portion of a number from the decimal portion. *group_separator*: The group separator separates integer groups (that is, thousands, millions, billions, and so on).	`numericcharacters: 'decimal_character group_separator'` Default value: `"."` See NLS_NUMERIC_CHARACTERS in *Oracle Database Globalization Support Guide* for more information.
`numberformat`	Specifies the number format model. Number format models cause the number to be rounded to the specified number of significant digits. A number format model is composed of one or more number format elements. This is used in combination with `numericcharacters`.	`numberformat: number_format_model` Default value: is derived from the setting of the `NLS_TERRITORY` parameter See Number Format Models in *SQL Language Reference* for more information.

Format Option	Description	Syntax
partition_columns	The format option partition_columns is used with DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE to specify the column names and data types of partition columns when the partition columns are derived from the file path, depending on the type of data file, structured or unstructured: • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE includes the column_list parameter and the data file is unstructured, such as with CSV text files, partition_columns does not include the data type. For example, use a format such as the following for this type of partition_columns specification: '"partition_columns": ["state","zipcode"]' The data type is not required because it is specified in the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE column_list parameter. • When the DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE does not include the column_list parameter and the data files are structured, such as Avro, ORC, or Parquet files, the partition_columns option includes the data type. For example, the following shows a partition_columns specification: '"partition_columns":[{"name":"country", "type":"varchar2(10)"}, {"name":"year", "type":"number"}, {"name":"month", "type":"varchar2(10)"}]' If the data files are unstructured and the type sub-clause is specified with partition_columns, the type sub-clause is ignored. For object names that are not based on hive format, the order of the partition_columns specified columns must match the order as they appear in the object name in the file_uri_list.	
quote	Specifies the quote character for the fields, the quote characters are removed during loading when specified.	quote: *character* Default value: Null meaning no quote

Format Option	Description	Syntax

`recorddelimiter`	Specifies the record delimiter.
	By default, `DBMS_CLOUD` tries to automatically find the correct newline character as the delimiter. It first searches the file for the Windows newline character "`\r\n`". If it finds the Windows newline character, this is used as the record delimiter for all files in the procedure. If a Windows newline character is not found, it searches for the UNIX/Linux newline character "`\n`" and if it finds one it uses "`\n`" as the record delimiter for all files in the procedure.
	Specify this argument explicitly if you want to override the default behavior, for example:
	`format => json_object('recorddelimiter' VALUE '''\r\n''')`
	To indicate that there is no record delimiter you can specify a `recorddelimiter` that does not occur in the input file. For example, to indicate that there is no delimiter, specify the control character 0x01 (SOH) as a value for the `recorddelimiter` and set the `recorddelimiter` value to "`0x''01''`" (this character does not occur in JSON text). For example:
	`format => '{"recorddelimiter" : "0x''01''"}'`
	The `recorddelimiter` is set once per procedure call. If you are using the default value, `detected newline`, then all files use the same record delimiter, if one is detected.
		`recorddelimiter`: *character*
		Default value: detected newline

`rejectlimit`	The operation will error out after specified number of rows are rejected.	`rejectlimit`: *number* Default value: `0`
`removequotes`	Removes any quotes that are around any field in the source file.	`removequotes`: `true` Default value: `False`
`skipheaders`	Specifies how many rows should be skipped from the start of the file.	`skipheaders`: *number* Default value: 0 if not specified, 1 if specified without a value
`territory`	Specifies a territory name to further determine input data characteristics.	`territory`: *string* Default value: Null See Locale Data in *Oracle Database Globalization Support Guide* for a listing of Oracle-supported territories.

`timestampformat`	Specifies the timestamp format in the source file. The format option `AUTO` searches for the following formats:
	`YYYY-MM-DD HH:MI:SS.FF`
	`YYYY-MM-DD HH:MI:SS.FF3`
	`YYYY-MM-DD HH24:MI:SS.FF3`
	`MM/DD/YYYY HH:MI:SS.FF3`
		`timestampformat` : *string*
		Default value: Database timestamp format
		The string can contain wildcard characters such as "$".

Format Option	Description	Syntax

`timestampltzformat`	Specifies the timestamp with local timezone format in the source file. The format option `AUTO` searches for the following formats:
	`DD Mon YYYY HH:MI:SS.FF TZR`
	`MM/DD/YYYY HH:MI:SS.FF TZR`
	`YYYY-MM-DD HH:MI:SS+/-TZR`
	`YYYY-MM-DD HH:MI:SS.FF3`
	`DD.MM.YYYY HH:MI:SS TZR`
		`timestampltzformat` : *string*
		Default value: Database timestamp with local timezone format

`timestamptzformat`	Specifies the timestamp with timezone format in the source file. The format option `AUTO` searches for the following formats:
	`DD Mon YYYY HH:MI:SS.FF TZR`
	`MM/DD/YYYY HH:MI:SS.FF TZR`
	`YYYY-MM-DD HH:MI:SS+/-TZR`
	`YYYY-MM-DD HH:MI:SS.FF3`
	`DD.MM.YYYY HH:MI:SS TZR`
		`timestamptzformat`: *string*
		Default value: Database timestamp with timezone format

`trimspaces`	Specifies how the leading and trailing spaces of the fields are trimmed. See the description of trim_spec.	`trimspaces`: `rtrim| ltrim| notrim| lrtrim| ldrtrim` Default value: `notrim`
`truncatecol`	If the data in the file is too long for a field, then this option will truncate the value of the field rather than reject the row.	`truncatecol`:`true` Default value: `False`
`type`	Specifies the source file type. See the description of `CSV` in field_definitions Clause See DBMS_CLOUD Package Format Options for Avro or Parquet for `type` values `avro` or `parquet`.	`type`: `csv|csv with embedded|csv without embedded` `csv` is the same as `csv without embedded`. Default value: Null

46.8 DBMS_CLOUD Avro and Parquet Support

This section covers the `DBMS_CLOUD` Avro and Parquet support provided with Oracle Database.

Topics

- DBMS_CLOUD Package Format Options for Avro or Parquet
- DBMS_CLOUD Package Avro to Oracle Data Type Mapping
- DBMS_CLOUD Package Parquet to Oracle Data Type Mapping
- DBMS_CLOUD Package Avro and Parquet Complex Types
- DBMS_CLOUD Package Avro and Parquet to Oracle Column Name Mapping

46.8.1 DBMS_CLOUD Package Format Options for Avro or Parquet

The format argument in `DBMS_CLOUD` specifies the format of source files.

The two ways to specify the format argument are:

```
format => '{"format_option" : "format_value" }'
```

And:

```
format => json_object('format_option' value 'format_value'))
```

Examples:

```
format => json_object('type' VALUE 'CSV')
```

To specify multiple format options, separate the values with a ",".

For example:

```
format => json_object('ignoremissingcolumns' value 'true', 'removequotes' value 'true',
'dateformat' value 'YYYY-MM-DD-HH24-MI-SS', 'blankasnull' value 'true')
```

Format Option	Description	Syntax
`type`	Specifies the file type.	`type : avro | parquet`
`schema`	When schema is set to `first` or `all`, the external table columns and data types are automatically derived from the Avro or Parquet file metadata. The column names will match those found in Avro or Parquet. The data types are converted from Avro or Parquet data types to Oracle data types. All columns are added to the table. The value `first` specifies to use the metadata from the first Avro or Parquet file in the `file_uri_list` to auto generate the columns and their data types. Use `first` if all of the files have the same schema. The value `all` specifies to use the metadata from all Avro or Parquet files in the `file_uri_list` to auto generate the columns and their data types. Use `all` (slower) if the files may have different schemas. **Default**: If `column_list` is specified, then the `schema` value, if specified is ignored. If `column_list` is not specified then the `schema` default value is `first`. **Note**: For Avro or Parquet format files the `schema` format option is not available and the `column_list` parameter must be specified for partitioned external tables using the `DBMS_CLOUD.CREATE_EXTERNAL_PART_TABLE` procedure.	`schema : first | all`

46.8.2 DBMS_CLOUD Package Avro to Oracle Data Type Mapping

Describes the mapping of Avro data types to Oracle data types.

> **Note:**
>
> Complex types, such as maps, arrays, and structs are supported starting with Oracle Database 19c. See DBMS_CLOUD Package Avro and Parquet Complex Types for information on using Avro complex types.

Avro Type	Oracle Type
INT	NUMBER(10)
LONG	NUMBER(19)

Avro Type	Oracle Type
BOOL	NUMBER(1)
UTF8 BYTE_ARRAY	RAW(2000)
FLT	BINARY_FLOAT
DBL	BINARY_DOUBLE
DECIMAL(p)	NUMBER(p)
DECIMAL(p,s)	NUMBER(p,s)
DATE	DATE
STRING	VARCHAR2
TIME_MILLIS	VARCHAR2(20 BYTE)
TIME_MICROS	VARCHAR2(20 BYTE)
TIMESTAMP_MILLIS	TIMESTAMP(3)
TIMESTAMP_MICROS	TIMESTAMP(6)
ENUM	VARCHAR2(*n*) Where: "*n*" is the actual maximum length of the AVRO ENUM's possible values
DURATION	RAW(2000)
FIXED	RAW(2000)
NULL	VARCHAR2(1) BYTE

See DBMS_CLOUD Package Avro and Parquet Complex Types for information on using Avro complex types.

46.8.3 DBMS_CLOUD Package Parquet to Oracle Data Type Mapping

Describes the mapping of Parquet data types to Oracle data types.

> **Note:**
>
> Complex types, such as maps, arrays, and structs are supported starting with Oracle Database 19c. See DBMS_CLOUD Package Avro and Parquet Complex Types for information on using Parquet complex types.

Parquet Type	Oracle Type
UINT_64	NUMBER(20)
INT_64	NUMBER(19)
UINT_32	NUMBER(10)
INT_32	NUMBER(10)
UINT_16	NUMBER(5)
INT_16	NUMBER(5)
UINT_8	NUMBER(3)
INT_8	NUMBER(3)
BOOL	NUMBER(1)
UTF8 BYTE_ARRAY	VARCHAR2(4000 BYTE)
FLT	BINARY_FLOAT
DBL	BINARY_DOUBLE
DECIMAL(p)	NUMBER(p)

Parquet Type	Oracle Type
DECIMAL(p,s)	NUMBER(p,s)
DATE	DATE
STRING	VARCHAR2(4000)
TIME_MILLIS	VARCHAR2(20 BYTE)
TIME_MILLIS_UTC	VARCHAR2(20 BYTE)
TIME_MICROS	VARCHAR2(20 BYTE)
TIME_MICROS_UTC	VARCHAR2(20 BYTE)
TIMESTAMP_MILLIS	TIMESTAMP(3)
TIMESTAMP_MILLIS_UTC	TIMESTAMP(3)
TIMESTAMP_MICROS	TIMESTAMP(6)
TIMESTAMP_MICROS_UTC	TIMESTAMP(6)
TIMESTAMP_NANOS	TIMESTAMP(9)

See DBMS_CLOUD Package Avro and Parquet Complex Types for information on using Parquet complex types.

46.8.4 DBMS_CLOUD Package Avro and Parquet Complex Types

Describes the mapping of Avro and Parquet complex data types to Oracle data types.

Oracle Database supports complex data types, including the following complex types:

- struct
- list
- map
- union
- array

When you specify a source file type of Avro or Parquet and the source file includes complex columns, Oracle Database queries return JSON for the complex columns. This simplifies processing of query results; you can use Oracle's powerful JSON parsing features consistently across the file types and data types. The following table shows the format for the complex types in Oracle Database:

> **Note:**
>
> The complex fields map to `VARCHAR2` columns and `VARCHAR2` size limits apply.

Type	Parquet	Avro	Oracle
List: sequence of values	List	Array	`VARCHAR2` (JSON format)
Map: list of objects with single key	Map	Map	`VARCHAR2` (JSON format)
Union: values of different type	Not Available	Union	`VARCHAR2` (JSON format)

Type	Parquet	Avro	Oracle
Object: zero or more key-value pairs	Struct	Record	`VARCHAR2` (JSON format)

46.8.5 DBMS_CLOUD Package Avro and Parquet to Oracle Column Name Mapping

Describes rules for how Avro and Parquet column names are converted to Oracle column names.

The following are supported for Avro and Parquet column names, but may require use of double quotes for Oracle SQL references in external tables. Thus, for ease of use and to avoid having to use double quotes when referencing column names, if possible do not use the following in Avro and Parquet column names:

- Embedded blanks
- Leading numbers
- Leading underscores
- Oracle SQL reserved words

The following table shows various types of Avro and Parquet column names, and rules for using the column names in Oracle column names in external tables.

Avro or Parquet Name	CREATE TABLE Name	Oracle CATALOG	Valid SQL	Notes

part, Part, or PART	part, Part, PART	PART	`select part`
			`select Part`
			`select paRt`
			`select PART`
				Oracle implicitly uppercases unquoted column names

Ord No	"Ord No"	Ord No	`select "Ord No"`	Double quotes are required when there are embedded blanks, which also preserves the character case
__index_key__	"__index_key__"	__index_key__	`select "__index_key__"`	Double quotes are required when there is a leading underscore, which also preserves the character case
6Way	"6Way"	6Way	`select "6Way"`	Double quotes are required when there is a leading numeric digit, which also preserves the character case

Avro or Parquet Name	CREATE TABLE Name	Oracle CATALOG	Valid SQL	Notes
create, Create, or CREATE, and so on. (any case variation) partition, Partition, PARTITION, and so on (for an Oracle Reserved word)	"CREATE" "PARTITION"	CREATE PARTITION	`select "CREATE"` `select "PARTITION"`	Double quotes are required around Oracle SQL Reserved words. These are forced to uppercase, but must always be double-quoted when used anywhere in SQL

rowid, Rowid, ROWid, and so on (for ROWID see notes)	rowid		`select "rowid"`
			`select "Rowid"`
			`select "ROWid"`
			`select "rowid"`
				For ROWID, any mixed or lower-case variation of ROWID preserves the case and must always be double-quoted and use the original case variations. Due to the inherent conflict with Oracle ROWID for the table, if you specify upper-case ROWID, it is automatically stored as lower-case "rowid" and must always be double-quoted when referenced.

Notes:

- In general a column name in an external table can be referenced without double quotes.
- Unless there is an embedded blank, a leading underscore ("_") or leading numeric digit ("0" through "9") in the column name, the original case of the column name is preserved, and it must always be referenced with double quotes and using the original case (upper, lower or mixed-case) of the Avro or Parquet column name.
- After using `DBMS_CLOUD.CREATE_EXTERNAL_TABLE` to create an external table with the format specified as `avro` or `parquet`, use the `DESCRIBE` command in SQL*Plus to view the table's column names.
- When Oracle SQL Reserved Words are used in Avro or Parquet column names, they must always be double-quoted when referenced anywhere in SQL. See Oracle SQL Reserved Words for more information.

47
DBMS_COMPARISON

The DBMS_COMPARISON package provides interfaces to compare and converge database objects at different databases.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Views
- Operational Notes
- Data Structures
- Summary of DBMS_COMPARISON Subprograms

47.1 DBMS_COMPARISON Overview

The DBMS_COMPARISON package is an Oracle-supplied package that you can use to compare database objects at two databases. This package also enables you converge the database objects so that they are consistent at different databases. Typically, this package is used in environments that share a database object at multiple databases. When copies of the same database object exist at multiple databases, the database object is a **shared database object**. Several data dictionary views contain information about comparisons made with the DBMS_COMPARISON package.

Shared database objects might be maintained by data replication. For example, materialized views or Oracle Streams components might replicate the database objects and maintain them at multiple databases. A custom application might also maintain shared database objects. When a database object is shared, it can diverge at the databases that share it. You can use this package to identify differences in the shared database objects. After identifying the differences, you can optionally use this package to synchronize the shared database objects.

To compare a database object that is shared at two different databases, complete the following general steps:

1. Run the CREATE_COMPARISON Procedure in this package to create a **comparison**. The comparison identifies the database objects to compare and specifies parameters for the comparison.
2. Run the COMPARE Function in this package to compare the database object at the two databases and identify differences. This function returns TRUE when no differences are found and FALSE when differences are found. This function also populates data dictionary views with comparison results. Separate comparison results are generated for each execution of the COMPARE function.
3. If you want to examine the comparison results, query the following data dictionary views:
   - DBA_COMPARISON_SCAN
   - USER_COMPARISON_SCAN

- DBA_COMPARISON_SCAN_VALUES
- USER_COMPARISON_SCAN_VALUES
- DBA_COMPARISON_ROW_DIF
- USER_COMPARISON_ROW_DIF

4. If there are differences, and you want to synchronize the database objects at the two databases, then run the CONVERGE procedure in this package.

After you create a comparison with the CREATE_COMPARISON procedure in the DBMS_COMPARISON package, you can run the comparison at any time using the COMPARE function. Each time you run the COMPARE function, it records comparison results in the appropriate data dictionary views. Comparison results might be modified when subprograms in this package are invoked and the scans in the comparison results are specified. For example, comparison results might be modified when you run the RECHECK function.

The comparison results for a single execution of the COMPARE function can include one or more **scans**. A scan checks for differences in some or all of the rows in a shared database object at a single point in time. You can compare database objects multiple times, and a unique scan ID identifies each scan in the comparison results.

A **bucket** is a range of rows in a database object that is being compared. Buckets improve performance by splitting the database object into ranges and comparing the ranges independently. Every comparison divides the rows being compared into an appropriate number of buckets, and each bucket is compared by a scan.

Each time the COMPARE function splits a bucket into smaller buckets, it performs new scans of the smaller buckets. The scan that analyzes a larger bucket is the **parent scan** of each scan that analyzes the smaller buckets into which the larger bucket was split. The **root scan** in the comparison results is the highest level parent scan. The root scan does not have a parent.

You can recheck a scan using the RECHECK function, and you can converge a scan using the CONVERGE procedure. When you want to recheck or converge all of the rows comparison results, specify the root scan ID for the comparison results in the appropriate subprogram. When you want to recheck or converge a portion of the rows in comparison results, specify the scan ID of the scan that contains the differences.

47.2 DBMS_COMPARISON Security Model

Security on this package can be controlled in one of two ways.

- Granting EXECUTE on this package to selected users or roles.
- Granting EXECUTE_CATALOG_ROLE to selected users or roles.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted EXECUTE privilege on the package directly. It cannot be granted through a role.

Each subprogram in the DBMS_COMPARISON package has a comparison_name parameter. The current user must be the owner of the specified comparison to run a subprogram in the DBMS_COMPARISON package.

To run the COMPARE function, RECHECK function, or CONVERGE procedure, the following users must have the SELECT or READ privilege on each copy of the shared database object:

- The comparison owner at the local database

- When a database link is used, the user at the remote database to which the comparison owner connects through a database link

The CONVERGE procedure also requires additional privileges for one of these users at the database where it makes changes to the shared database object. The user must have INSERT, UPDATE, and DELETE privileges on the shared database object at this database.

In addition, when the CONVERGE procedure is run with either the local_converge_tag or remote_converge_tag parameter set to a non-NULL value, then the following additional requirements must be met:

- If the local table "wins," then the user at the remote database to which the invoker of the CONVERGE procedure connects through a database link must be granted either EXECUTE_CATALOG_ROLE or EXECUTE privilege on the DBMS_STREAMS_ADM package.
- If the remote table "wins," then the invoker of the CONVERGE procedure at the local database must be granted either EXECUTE_CATALOG_ROLE or EXECUTE privilege on the DBMS_STREAMS_ADM package.

> **Note:**
>
> The database administrator (DBA) can assume control over some of the DBMS_COMPARISON functions and procedures owned by other users. This control applies to DROP_COMPARISON and PURGE_COMPARISON. This DBA override can be particularly useful in cleanup operations when comparisons created by another user need to be dropped

47.3 DBMS_COMPARISON Constants

The DBMS_COMPARISON package defines several enumerated constants to use specifying parameter values. Enumerated constants must be prefixed with the package name. For example, DBMS_COMPARISON.CMP_SCAN_MODE_FULL.

Table 47-1 lists the parameters and enumerated constants.

Table 47-1 DBMS_COMPARISON Parameters with Enumerated Constants

Parameter	Option	Type	Description
comparison_mode	• CMP_COMPARE_MODE_OBJECT	VARCHAR2(30)	CMP_COMPARE_MODE_OBJECT is a database object. This constant can be specified as 'OBJECT'.

Table 47-1 (Cont.) DBMS_COMPARISON Parameters with Enumerated Constants

Parameter	Option	Type	Description
scan_mode	• CMP_SCAN_MODE_FULL • CMP_SCAN_MODE_RANDOM • CMP_SCAN_MODE_CYCLIC • CMP_SCAN_MODE_CUSTOM	VARCHAR2(30)	CMP_SCAN_MODE_FULL indicates that the entire database object is compared. This constant can be specified as 'FULL'. CMP_SCAN_MODE_RANDOM indicates that a random portion of the database object is compared. This constant can be specified as 'RANDOM'. CMP_SCAN_MODE_CYCLIC indicates that a portion of the database object is compared when you perform a single comparison. When you compare the database object again, another portion of the database object is compared, starting where the last comparison ended. This constant can be specified as 'CYCLIC'. CMP_SCAN_MODE_CUSTOM indicates that the user who runs the subprogram specifies the range to compare in the database object. This constant can be specified as 'CUSTOM'.
converge_options	• CMP_CONVERGE_LOCAL_WINS • CMP_CONVERGE_REMOTE_WINS	VARCHAR2(30)	CMP_CONVERGE_LOCAL_WINS indicates that the column values at the local database replace the column values at the remote database when these column values are different. This constant can be specified as 'LOCAL'. CMP_CONVERGE_REMOTE_WINS indicates that the column values at the remote database replace the column values at the local database when these column values are different. This constant can be specified as 'REMOTE'.
null_value	• CMP_NULL_VALUE_DEF	VARCHAR2(100)	CMP_NULL_VALUE_DEF indicates that ORA$STREAMS$NV is substituted for NULL values in database objects during comparison. This constant can be specified as 'ORA$STREAMS$NV'.
max_num_buckets	• CMP_MAX_NUM_BUCKETS	INTEGER	CMP_MAX_NUM_BUCKETS indicates that the maximum number of buckets is 1,000. This constant can be specified as 1000.
min_rows_in_bucket	• CMP_MIN_ROWS_IN_BUCKET	INTEGER	CMP_MIN_ROWS_IN_BUCKET indicates that the minimum number of rows in a bucket is 10,000. This constant can be specified as 10000.

47.4 DBMS_COMPARISON Views

The DBMS_COMPARISON package uses several views.

These views are listed below:

- DBA_COMPARISON
- USER_COMPARISON
- DBA_COMPARISON_COLUMNS
- USER_COMPARISON_COLUMNS
- DBA_COMPARISON_SCAN
- USER_COMPARISON_SCAN
- DBA_COMPARISON_SCAN_VALUES
- USER_COMPARISON_SCAN_VALUES
- DBA_COMPARISON_ROW_DIF
- USER_COMPARISON_ROW_DIF

> **See Also:**
>
> *Oracle Database Reference*

47.5 DBMS_COMPARISON Operational Notes

The DBMS_COMPARISON package has certain requirements and operational notes.

These include the following:

- Oracle Database Release Requirements for the DBMS_COMPARISON Package
- Database Character Set Requirements for the DBMS_COMPARISON Package
- Database Object Requirements for the DBMS_COMPARISON Package
- Index Column Requirements for the DBMS_COMPARISON Package
- Datatype Requirements for the DBMS_COMPARISON Package
- Only Converge Rows That Are Not Being Updated

Oracle Database Release Requirements for the DBMS_COMPARISON Package

Meet the following Oracle Database release requirements when running the subprograms in the DBMS_COMPARISON package:

- The local database that runs the subprograms in the DBMS_COMPARISON package must be an Oracle Database 11*g* Release 1 (11.1) database.
- The remote database must be an Oracle Database 10*g* Release 1 (10.1) or later database. Oracle databases before this release and non-Oracle databases are not supported.

Database Character Set Requirements for the DBMS_COMPARISON Package

The database character sets must be the same for the databases that contain the database objects being compared.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for information about database character sets

Database Object Requirements for the DBMS_COMPARISON Package

The `DBMS_COMPARISON` package can compare the following types of database objects:

- Tables
- Single-table views
- Materialized views
- Synonyms for tables, single-table views, and materialized views

Database objects of different types can be compared and converged at different databases. For example, a table at one database and a materialized view at another database can be compared and converged with this package.

To run the subprograms in the `DBMS_COMPARISON` package, the specified database objects must have the same shape at each database. Specifically, the database objects must have the same number of columns at each database, and the datatypes of corresponding columns must match.

If a database object being compared contains columns that do not exist in the other database object, then you can compare the database objects by excluding the extra columns during comparison creation. Use the `column_list` parameter in the `CREATE_COMPARISON` procedure to list only the columns that exist in both database objects.

> **See Also:**
>
> CREATE_COMPARISON Procedure

Index Column Requirements for the DBMS_COMPARISON Package

This section discusses number, timestamp, and interval columns. These include the following datatypes:

- Number columns are of the following datatypes: `NUMBER`, `FLOAT`, `BINARY_FLOAT`, and `BINARY_DOUBLE`.
- Timestamp columns are of the following datatypes: `TIMESTAMP`, `TIMESTAMP WITH TIME ZONE`, and `TIMESTAMP WITH LOCAL TIME ZONE`
- Interval columns are of the following datatypes: `INTERVAL YEAR TO MONTH` and `INTERVAL DAY TO SECOND`.

For all scan modes to be supported by the DBMS_COMPARISON package, the database objects must have one of the following types of indexes:

- A single-column index on a number, timestamp, interval, or DATE datatype column
- A composite index that only includes number, timestamp, interval, or DATE datatype columns. Each column in the composite index must either have a NOT NULL constraint or be part of the primary key.

For the scan modes CMP_SCAN_MODE_FULL and CMP_SCAN_MODE_CUSTOM to be supported, the database objects must have one of the following types of indexes:

- A single-column index on a number, timestamp, interval, DATE, VARCHAR2, or CHAR datatype column
- A composite index that only includes number, timestamp, interval, DATE, VARCHAR2, or CHAR columns. Each column in the composite index must either have a NOT NULL constraint or be part of the primary key.

If the database objects do not have one of these types of indexes, then the DBMS_COMPARISON package does not support the database objects. For example, if the database objects only have a single index on an NVARCHAR2 column, then the DBMS_COMPARISON package does not support them. Or, if the database objects have only one index, and it is a composite index that includes a NUMBER column and an NCHAR column, then the DBMS_COMPARISON package does not support them.

You can specify an index when you create a comparison using the index_schema_name and index_name parameters in the CREATE_COMPARISON procedure. If you specify an index, then make sure the columns in the index meet the requirements of the scan mode used for the comparison.

The index columns in a comparison must uniquely identify every row involved in a comparison. The following constraints satisfy this requirement:

- A primary key constraint
- A unique constraint on one or more non-NULL columns

If these constraints are not present on a table, then use the index_schema_name and index_name parameters in the CREATE_COMPARISON procedure to specify an index whose columns satisfy this requirement.

When a single index value identifies both a local row and a remote row, the two rows must be copies of the same row in the replicated tables. In addition, each pair of copies of the same row must always have the same index value.

The DBMS_COMPARISON package can use an index only if all of the columns in the index are included in the column_list parameter when the comparison is created with the CREATE_COMPARISON procedure.

After a comparison is created, you can determine the index column or columns for it by running the following query:

```
SELECT COLUMN_NAME, COLUMN_POSITION FROM DBA_COMPARISON_COLUMNS
  WHERE COMPARISON_NAME = 'COMPARE_CUSTOM' AND
        INDEX_COLUMN    = 'Y';
```

If there are multiple index columns, then the index column with 1 for the COLUMN_POSITION is the lead index column in the composite index.

> **See Also:**
>
> - "Constants" for information about scan modes
> - CREATE_COMPARISON Procedure for information about specifying an index for a comparison

Datatype Requirements for the DBMS_COMPARISON Package

The `DBMS_COMPARISON` package can compare data in columns of the following datatypes:

- `VARCHAR2`
- `NVARCHAR2`
- `NUMBER`
- `FLOAT`
- `DATE`
- `BINARY_FLOAT`
- `BINARY_DOUBLE`
- `TIMESTAMP`
- `TIMESTAMP WITH TIME ZONE`
- `TIMESTAMP WITH LOCAL TIME ZONE`
- `INTERVAL YEAR TO MONTH`
- `INTERVAL DAY TO SECOND`
- `RAW`
- `CHAR`
- `NCHAR`

If a column with datatype `TIMESTAMP WITH LOCAL TIME ZONE` is compared, then the two databases must use the same time zone. Also, if a column with datatype `NVARCHAR2` or `NCHAR` is compared, then the two databases must use the same national character set.

The `DBMS_COMPARISON` package cannot compare data in columns of the following datatypes:

- `LONG`
- `LONG RAW`
- `ROWID`
- `UROWID`
- `CLOB`
- `NCLOB`
- `BLOB`
- `BFILE`
- User-defined types (including object types, `REF`s, varrays, and nested tables)

- Oracle-supplied types (including any types, XML types, spatial types, and media types)

You can compare database objects that contain unsupported columns by excluding the unsupported columns during comparison creation. Use the `column_list` parameter in the `CREATE_COMPARISON` procedure to list only the supported columns in a shared database object.

> **See Also:**
>
> - CREATE_COMPARISON Procedure
> - *Oracle Database SQL Language Reference* for more information about datatypes
> - *Oracle Database Globalization Support Guide* for information about national character sets

Only Converge Rows That Are Not Being Updated

You should only converge rows that are not being updated on either database. For example, if the shared database object is updated by replication components, then only converge rows for which replication changes have been applied and make sure no new changes are in the process of being replicated for these rows. If you compare replicated database objects, then it is typically best to compare them during a time of little or no replication activity to identify persistent differences.

> **Note:**
>
> If a scan identifies that a row is different in the shared database object at two databases, and the row is modified after the scan, then it can result in unexpected data in the row after the `CONVERGE` procedure is run.

47.6 DBMS_COMPARISON Data Structures

The `DBMS_COMPARISON` package defines a `RECORD` type.

Contains information returned by the `COMPARE` function or `CONVERGE` procedure in the `DBMS_COMPARISON` package.

> **Note:**
>
> The `COMPARE` function only returns a value for the `scan_id` field.

47.6.1 COMPARISON_TYPE Record Type

This record type contains information returned by the `COMPARE` function or `CONVERGE` procedure in the `DBMS_COMPARISON` package.

> **Note:**
>
> The `COMPARE` function only returns a value for the `scan_id` field.

Syntax

```
TYPE COMPARISON_TYPE IS RECORD(
  scan_id            NUMBER,
  loc_rows_merged    NUMBER,
  rmt_rows_merged    NUMBER,
  loc_rows_deleted   NUMBER,
  rmt_rows_deleted   NUMBER);
```

Table 47-2 COMPARISON_TYPE Attributes

Field	Description
`scan_id`	The scan ID of the scan
`loc_rows_merged`	The number of rows in the local database object updated with information from the database object at the remote site
`rmt_rows_merged`	The number of rows in the database object updated at the remote site with information from the database object at the local site
`loc_rows_deleted`	The number of rows deleted from the local database object
`rmt_rows_deleted`	The number of rows deleted from the remote database object

47.7 Summary of DBMS_COMPARISON Subprograms

This table lists the `DBMS_COMPARISON` subprograms and briefly describes them.

Table 47-3 DBMS_COMPARISON Package Subprograms

Subprogram	Description
COMPARE Function	Performs the specified comparison
CONVERGE Procedure	Executes data manipulation language (DML) changes to synchronize the portion of the database object that was compared in the specified scan
CREATE_COMPARISON Procedure	Creates a comparison
DROP_COMPARISON Procedure	Drops a comparison
PURGE_COMPARISON Procedure	Purges the comparison results, or a subset of the comparison results, for a comparison
RECHECK Function	Rechecks the differences in a specified scan for a comparison

47.7.1 COMPARE Function

This function performs the specified comparison.

Each time a comparison is performed, it results in at least one new scan, and each scan has a unique scan ID. You can define and name a comparison using the `CREATE_COMPARISON` procedure.

> **See Also:**
>
> - "Overview"
> - CREATE_COMPARISON Procedure

Syntax

```
DBMS_COMPARISON.COMPARE(
   comparison_name  IN   VARCHAR2,
   scan_info        OUT  COMPARISON_TYPE,
   min_value        IN   VARCHAR2   DEFAULT NULL,
   max_value        IN   VARCHAR2   DEFAULT NULL,
   perform_row_dif  IN   BOOLEAN    DEFAULT FALSE)
RETURN BOOLEAN;
```

Parameters

Table 47-4 COMPARE Function Parameters

Parameter	Description
`comparison_name`	The name of the comparison.
`scan_info`	Information about the compare operation returned in the `COMPARISON_TYPE` datatype. See COMPARISON_TYPE Record Type.
`min_value`	When the scan mode for the comparison is set to `CMP_SCAN_MODE_CUSTOM`, specify the minimum index column value for the range of rows that are being compared. To determine the index column for a comparison, query the `DBA_COMPARISON_COLUMNS` data dictionary view. For a composite index, specify a value for the column with `column_position` equal to `1` in the `DBA_COMPARISON_COLUMNS` view. See the index column requirements under DBMS_COMPARISON Operational Notes. If the scan mode is set to a value other than `CMP_SCAN_MODE_CUSTOM`, then this parameter must be set to `NULL`. If `NULL` and the `scan_mode` parameter is set to `CMP_SCAN_MODE_CUSTOM`, then an error is raised. To determine the scan mode for the comparison, query the `DBA_COMPARISON` data dictionary view. See DBMS_COMPARISON Constants for information about scan modes.

Table 47-4 (Cont.) COMPARE Function Parameters

Parameter	Description
max_value	When the scan mode for the comparison is set to CMP_SCAN_MODE_CUSTOM, specify the maximum index column value for the range of rows that are being compared. To determine the index column for a comparison, query the DBA_COMPARISON_COLUMNS data dictionary view. For a composite index, specify a value for the column with column_position equal to 1 in the DBA_COMPARISON_COLUMNS view. See the index column requirements under DBMS_COMPARISON Operational Notes. If the scan mode is set to a value other than CMP_SCAN_MODE_CUSTOM, then this parameter must be set to NULL. If NULL and the scan_mode parameter is set to CMP_SCAN_MODE_CUSTOM, then an error is raised. To determine the scan mode for the comparison, query the DBA_COMPARISON data dictionary view. See DBMS_COMPARISON Constants for information about scan modes.
perform_row_dif	If TRUE, then compares each row individually in the database object being compared after reaching the smallest possible bucket for the comparison. If FALSE, then compares buckets for differences but does not compare each row individually when differences are found in the smallest possible bucket. See DBMS_COMPARISON Overview for information about buckets.

Return Values

This function returns TRUE when no differences are found in the database objects being compared. This function returns FALSE when differences are found in the database objects being compared.

47.7.2 CONVERGE Procedure

This procedure executes data manipulation language (DML) changes to synchronize the portion of the database objects that was compared in the specified scan.

Syntax

```
DBMS_COMPARISON.CONVERGE(
   comparison_name      IN   VARCHAR2,
   scan_id              IN   NUMBER,
   scan_info            OUT  COMPARISON_TYPE,
   converge_options     IN   VARCHAR2  DEFAULT CMP_CONVERGE_LOCAL_WINS,
   perform_commit       IN   BOOLEAN   DEFAULT TRUE,
   local_converge_tag   IN   RAW       DEFAULT NULL,
   remote_converge_tag  IN   RAW       DEFAULT NULL);
```

Parameters

Table 47-5 CONVERGE Procedure Parameters

Parameter	Description
comparison_name	The name of the comparison.

Table 47-5 (Cont.) CONVERGE Procedure Parameters

Parameter	Description
scan_id	The identifier for the scan that contains the differences between the database objects being converged. See "Overview" for more information about specifying a scan ID in this parameter.
scan_info	Information about the converge operation returned in the COMPARISON_TYPE datatype. See COMPARISON_TYPE Record Type.
converge_options	Either the CMP_CONVERGE_LOCAL_WINS constant or the CMP_CONVERGE_REMOTE_WINS constant. See "Constants" for information about these constants.
perform_commit	If TRUE, then performs a COMMIT periodically while making the DML changes. The CONVERGE procedure might perform more than one COMMIT when this parameter is set to TRUE. If FALSE, then does not perform a COMMIT after making DML changes.
local_converge_tag	The Replication tag to set in the session on the local database before performing any changes to converge the data in the database objects being converged. If non-NULL, then this parameter setting takes precedence over the local_converge_tag parameter in the CREATE_COMPARISON procedure that created the comparison. If NULL, then this parameter is ignored, and the local_converge_tag parameter in the CREATE_COMPARISON procedure that created the comparison is used.
remote_converge_tag	The Replication tag to set in the session on the remote database before performing any changes to converge the data in the database objects being converged. If non-NULL, then this parameter setting takes precedence over the remote_converge_tag parameter in the CREATE_COMPARISON procedure that created the comparison. If NULL, then this parameter is ignored, and the remote_converge_tag parameter in the CREATE_COMPARISON procedure that created the comparison is used.

Usage Notes

If one of the database objects being converged is a read-only materialized view, then the converge_options parameter must be set to ensure that the read-only materialized view "wins" in the converge operation. The CONVERGE procedure raises an error if it tries to make changes to a read-only materialized view.

47.7.3 CREATE_COMPARISON Procedure

This procedure creates a comparison.

Syntax

```
DBMS_COMPARISON.CREATE_COMPARISON(
   comparison_name      IN  VARCHAR2,
   schema_name          IN  VARCHAR2,
```

```
  object_name          IN  VARCHAR2,
  dblink_name          IN  VARCHAR2,
  index_schema_name    IN  VARCHAR2  DEFAULT NULL,
  index_name           IN  VARCHAR2  DEFAULT NULL,
  remote_schema_name   IN  VARCHAR2  DEFAULT NULL,
  remote_object_name   IN  VARCHAR2  DEFAULT NULL,
  comparison_mode      IN  VARCHAR2  DEFAULT CMP_COMPARE_MODE_OBJECT,
  column_list          IN  VARCHAR2  DEFAULT '*',
  scan_mode            IN  VARCHAR2  DEFAULT CMP_SCAN_MODE_FULL,
  scan_percent         IN  NUMBER    DEFAULT NULL,
  null_value           IN  VARCHAR2  DEFAULT CMP_NULL_VALUE_DEF,
  local_converge_tag   IN  RAW       DEFAULT NULL,
  remote_converge_tag  IN  RAW       DEFAULT NULL,
  max_num_buckets      IN  NUMBER    DEFAULT CMP_MAX_NUM_BUCKETS,
  min_rows_in_bucket   IN  NUMBER    DEFAULT CMP_MIN_ROWS_IN_BUCKET);
```

Parameters

Table 47-6 CREATE_COMPARISON Procedure Parameters

Parameter	Description
comparison_name	The name of the comparison.
schema_name	The name of the schema that contains the local database object to compare.
object_name	The name of the local database object to compare.
dblink_name	Database link to the remote database. The specified database object in the remote database is compared with the database object in the local database. If NULL, then the comparison is configured to compare two database objects in the local database. In this case, parameters that specify the remote database object apply to the second database object in the comparison and to operations on the second database object. For example, specify the second database object in this procedure by using the remote_schema_name and remote_object_name parameters.
index_schema_name	The name of the schema that contains the index. If NULL, then the schema specified in the schema_name parameter is used.
index_name	The name of the index. If NULL, then the system determines the index columns for the comparison automatically. If the index_schema_name parameter is non-NULL, then the index_name parameter must also be non-NULL. Otherwise, an error is raised. **See Also:** "Usage Notes" for more information about specifying an index
remote_schema_name	The name of the schema that contains the database object at the remote database. Specify a non-NULL value if the schema names are different at the two databases. If NULL, then the schema specified in the schema_name parameter is used.

Table 47-6 (Cont.) CREATE_COMPARISON Procedure Parameters

Parameter	Description
remote_object_name	The name of the database object at the remote database. Specify a non-NULL value if the database object names are different at the two databases. If NULL, then the database object specified in the object_name parameter is used.
comparison_mode	Specify the default value CMP_COMPARE_MODE_OBJECT. Additional modes might be added in future releases.
column_list	Specify '*' to include all of the columns in the database objects being compared. To compare a subset of columns in the database objects, specify a comma-delimited list of the columns to check. Any columns that are not in the list are ignored during a comparison and convergence. See "Usage Notes" for information about columns that are required in the column_list parameter.
scan_mode	Either CMP_SCAN_MODE_FULL, CMP_SCAN_MODE_RANDOM, CMP_SCAN_MODE_CYCLIC, or CMP_SCAN_MODE_CUSTOM. If you specify CMP_SCAN_MODE_CUSTOM, then make sure you specify an index using the index_schema_name and index_name parameters. Specifying an index ensures that you can specify the correct min_value and max_value for the lead index column when you run the COMPARE or RECHECK function. See "Constants" for information about these constants.
scan_percent	The percentage of the database object to scan for comparison when the scan_mode parameter is set to either CMP_SCAN_MODE_RANDOM or CMP_SCAN_MODE_CYCLIC. For these scan_mode settings, a non-NULL value that is greater than 0 (zero) and less than 100 is required. If NULL and the scan_mode parameter is set to CMP_SCAN_MODE_FULL, then the entire database object is scanned for comparison. If NULL and the scan_mode parameter is set to CMP_SCAN_MODE_CUSTOM, then the portion of the database object scanned for comparison is specified when the COMPARE function is run. If non-NULL and the scan_mode parameter is set to either CMP_SCAN_MODE_FULL or CMP_SCAN_MODE_CUSTOM, then the scan_percent parameter is ignored. **Note:** When the scan_percent parameter is non-NULL, and the lead index column for the comparison does not distribute the rows in the database object evenly, the portion of the database object that is compared might be smaller or larger than the specified scan_percent value. See DBMS_COMPARISON Operational Notes for more information about the DBMS_COMPARISON package index requirements for the lead index column.
null_value	The value to substitute for each NULL in the database objects being compared. Specify a value or use the CMP_NULL_VALUE_DEF constant. If a column being compared can contain NULLs, then the value specified for this parameter must be different than any non-NULL value in the column. Otherwise, if the value specified for this parameter can appear in the column, some row differences might not be found. See "Constants" for information about this constant.

Table 47-6 (Cont.) CREATE_COMPARISON Procedure Parameters

Parameter	Description
local_converge_tag	The Oracle Replication tag to set in the session on the local database before performing any changes to converge the data in the database objects being compared. If the local_converge_tag parameter is non-NULL in the CONVERGE procedure when comparison results for this comparison are converged, then the setting in the CONVERGE procedure takes precedence. See CONVERGE Procedure for more information.
remote_converge_tag	The Oracle Replication tag to set in the session on the remote database before performing any changes to converge the data in the database objects being compared. If the remote_converge_tag parameter is non-NULL in the CONVERGE procedure when comparison results for this comparison are converged, then the setting in the CONVERGE procedure takes precedence. See CONVERGE Procedure for more information.
max_num_buckets	Specify the maximum number of buckets to use. Specify a value or use the CMP_MAX_NUM_BUCKETS constant. See "Constants" for information about this constant. See "Overview" for information about buckets. **Note:** If an index column for a comparison is a VARCHAR2 or CHAR column, then the number of buckets might exceed the value specified for the max_num_buckets parameter.
min_rows_in_bucket	Specify the minimum number of rows in each bucket. Specify a value or use the CMP_MIN_ROWS_IN_BUCKET constant. See "Constants" for information about this constant. See "Overview" for information about buckets.

Usage Notes

This section contains usage notes for the CREATE_COMPARISON procedure.

Usage Notes for the index_schema_name and index_name Parameters

When you specify an index for a comparison with the index_schema_name and index_name parameters, the specified index determines the comparison's index columns and their ordering. The order of the columns in the index determines the index column ordering for the comparison. Therefore, the column in column position 1 in the index is the lead column for the comparison.

The index columns and their ordering affect the details of each SQL statement generated and executed for a comparison. For each SQL statement, the optimizer decides whether to use indexes. If the optimizer decides to use indexes, then the optimizer decides which particular indexes to use. An index specified in column_list parameter might or might not be used.

The columns in the specified index must meet the requirements described in DBMS_COMPARISON Operational Notes. If the index columns do not meet these requirements, then an error is raised.

> **Note:**
>
> If you do not specify an index when you create a comparison, then the `CREATE_COMPARISON` procedure selects either the primary key, if it exists, or an existing unique index. The procedure never selects a non-unique index. However, if you specify an index, then the `CREATE_COMPARISON` procedure does not check its uniqueness. Therefore, if you specify a non-unique index, and duplicate index keys exist, then the results might be incorrect when the `CONVERGE` procedure synchronizes data.

Usage Notes for the column_list Parameter

When the `column_list` parameter is set to a value other than `'*'`, the following columns are required in the `column_list` parameter:

- Any columns that are required to meet the index column requirements for the `DBMS_COMPARISON` package. If the `index_name` parameter is non-`NULL`, then the columns in the specified index must be in the column list. If the `index_name` parameter is `NULL`, then see DBMS_COMPARISON Operational Notes for information about the `DBMS_COMPARISON` index requirements.
- If you plan to use the `CONVERGE` procedure to make changes to a database object based on the comparison, then any columns in this database object that have a `NOT NULL` constraint but no default value must be included in the column list. If these columns are not included, then the `CONVERGE` procedure returns an error. See CONVERGE Procedure.

47.7.4 DROP_COMPARISON Procedure

This procedure drops a comparison.

Syntax

```
DBMS_COMPARISON.DROP_COMPARISON(
   comparison_name  IN  VARCHAR2);
```

Parameters

Table 47-7 DROP_COMPARISON Procedure Parameters

Parameter	Description
`comparison_name`	The name of the comparison.

47.7.5 PURGE_COMPARISON Procedure

This procedure purges the comparison results, or a subset of the comparison results, for a comparison.

> **Note:**
>
> At least one of the following parameters must be set to `NULL`: `scan_id` or `purge_time`. If both the `scan_id` and `purge_time` parameters are `NULL`, then this procedure purges all comparison results for the comparison.

Syntax

```
DBMS_COMPARISON.PURGE_COMPARISON(
   comparison_name  IN  VARCHAR2,
   scan_id          IN  NUMBER     DEFAULT NULL,
   purge_time       IN  TIMESTAMP  DEFAULT NULL);
```

Parameters

Table 47-8 PURGE_COMPARISON Procedure Parameters

Parameter	Description
`comparison_name`	The name of the comparison.
`scan_id`	The scan ID of the scan for which results are purged. The scan ID must identify a root scan. If the scan ID does not identify a root scan, then an error is raised. When a root scan ID is specified, it is purged, and all direct and indirect child scans of the specified root scan are purged. If `NULL`, then no scan ID is considered when purging comparison results for the comparison. See "Overview" for information about scans.
`purge_time`	The date before which results are purged. If `NULL`, then no date is considered when purging comparison results for the comparison.

47.7.6 RECHECK Function

This function rechecks the differences in a specified scan for a comparison.

This function performs one of the following actions:

- If the specified scan completed successfully the last time it ran, then this function checks the previously identified differences in the scan.
- If the specified scan completed partially, then this function continues to check the database object from the point where the previous scan ended.

> **Note:**
>
> This function does not compare the shared database object for differences that were not recorded in the specified comparison scan. To check for those differences, run the `COMPARE` function.

> **See Also:**
>
> COMPARE Function

Syntax

```
DBMS_COMPARISON.RECHECK(
   comparison_name  IN  VARCHAR2,
   scan_id          IN  NUMBER,
   perform_row_dif  IN  BOOLEAN  DEFAULT FALSE)
RETURN BOOLEAN;
```

Parameters

Table 47-9 RECHECK Function Parameters

Parameter	Description
`comparison_name`	The name of the comparison.
`scan_id`	The scan ID of the scan to recheck. See "Overview" for more information about specifying a scan ID in this parameter.
`perform_row_dif`	If `TRUE`, then compares each row individually in the database objects being compared after reaching the smallest possible bucket for the comparison. If `FALSE`, then compares buckets for differences but does not compare each row individually when differences are found in the smallest possible bucket. See "Overview" for information about buckets.

Return Values

This function returns `TRUE` when no differences are found in the database objects being compared. This function returns `FALSE` when differences are found in the database objects being compared.

48
DBMS_COMPRESSION

The `DBMS_COMPRESSION` package provides an interface to facilitate choosing the correct compression level for an application.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Data Structures
- Summary of DBMS_COMPRESSION Subprograms

> **See Also:**
>
> - *Oracle Database Administrator's Guide*
> - *Oracle Database Concepts*
> - *Oracle Database SQL Language Reference*
> - *Oracle Database Data Warehousing Guide*
> - *Oracle Database VLDB and Partitioning Guide*
> - *Oracle Database Reference*

48.1 DBMS_COMPRESSION Overview

The `DBMS_COMPRESSION` package gathers compression-related information within a database environment. This includes tools for estimating compressibility of a table for both partitioned and non-partitioned tables, and gathering row-level compression information on previously compressed tables. This gives the user with adequate information to make compression-related decision.

48.2 DBMS_COMPRESSION Security Model

The `DBMS_COMPRESSSION` package is defined with `AUTHID CURRENT USER`, so it executes with the privileges of the current user.

48.3 DBMS_COMPRESSION Constants

The `DBMS_COMPRESSION` package uses constants that can be used for specifying parameter values.

These constants are shown in the following table:

Table 48-1 DBMS_COMPRESSION Constants - Compression Types

Constant	Type	Value	Description
COMP_NOCOMPRESS	NUMBER	1	No compression
COMP_ADVANCED	NUMBER	2	Advanced row compression
COMP_QUERY_HIGH	NUMBER	4	High for query warehouse compression (Hybrid Columnar Compression)
COMP_QUERY_LOW	NUMBER	8	Low for query warehouse compression (Hybrid Columnar Compression)
COMP_ARCHIVE_HIGH	NUMBER	16	High archive compression (Hybrid Columnar Compression)
COMP_ARCHIVE_LOW	NUMBER	32	Low archive compression (Hybrid Columnar Compression)
COMP_BLOCK	NUMBER	64	Compressed block
COMP_LOB_HIGH	NUMBER	128	High compression level for LOB operations
COMP_LOB_MEDIUM	NUMBER	256	Medium compression level for LOB operations
COMP_LOB_LOW	NUMBER	512	Low compression level for LOB operations
COMP_INDEX_ADVANCED_HIGH	NUMBER	1024	High compression level for indexes
COMP_INDEX_ADVANCED_LOW	NUMBER	2048	Low compression level for indexes
COMP_RATIO_LOB_MINROWS	NUMBER	1000	Minimum required number of LOBs in the object for which LOB compression ratio is to be estimated
COMP_BASIC	NUMBER	4096	Basic table compression
COMP_RATIO_LOB_MAXROWS	NUMBER	5000	Maximum number of LOBs used to compute the LOB compression ratio
COMP_INMEMORY_NOCOMPRESS	NUMBER	8192	In-Memory with no compression
COMP_INMEMORY_DML	NUMBER	16384	In-Memory compression level for DML
COMP_INMEMORY_QUERY_LOW	NUMBER	32768	In-Memory compression level optimized for query performance
COMP_INMEMORY_QUERY_HIGH	NUMBER	65536	In-Memory compression level optimized on query performance as well as space saving
COMP_INMEMORY_CAPACITY_LOW	NUMBER	131072	In-Memory low compression level optimizing for capacity
COMP_INMEMORY_CAPACITY_HIGH	NUMBER	262144	In-Memory high compression level optimizing for capacity
COMP_RATIO_MINROWS	NUMBER	1000000	Minimum required number of rows in the object for which HCC ratio is to be estimated

Table 48-1 (Cont.) DBMS_COMPRESSION Constants - Compression Types

Constant	Type	Value	Description
COMP_RATIO_ALLROWS	NUMBER	-1	To indicate the use of all the rows in the object to estimate HCC ratio
OBJTYPE_TABLE	PLS_INTEGER	1	Identifies the object whose compression ratio is estimated as of type table
OBJTYPE_INDEX	PLS_INTEGER	2	Identifies the object whose compression ratio is estimated as of type index

> **Note:**
>
> Hybrid columnar compression is a feature of certain Oracle storage systems. See *Oracle Database Concepts* for more information.

48.4 DBMS_COMPRESSION Data Structures

The DBMS_COMPRESSION package defines a RECORD type and a TABLE type.

RECORD TYPES

COMPREC Record Type

TABLE TYPES

COMPRECLIST Table Type

48.4.1 COMPREC Record Type

The COMPREC record type is a record for calculating an individual index compression ratio on a table.

Syntax

```
TYPE COMPREC IS RECORD(
  ownname           varchar2(255),
  objname           varchar2(255),
  blkcnt_cmp        PLS_INTEGER,
  blkcnt_uncmp      PLS_INTEGER,
  row_cmp           PLS_INTEGER,
  row_uncmp         PLS_INTEGER,
  cmp_ratio         NUMBER,
  objtype           PLS_INTEGER);
```

Fields

Table 48-2 COMPREC Attributes

Field	Description
ownname	Schema of the object owner
objname	Name of the object
blkcnt_cmp	Number of blocks used by the compressed sample of the object
blkcnt_uncmp	Number of blocks used by the uncompressed sample of the object
row_cmp	Number of rows in a block in compressed sample of the object
row_uncmp	Number of rows in a block in uncompressed sample of the object
cmp_ratio	Compression ratio, blkcnt_uncmp divided by blkcnt_cmp
objtype	Type of the object

48.4.2 COMPRECLIST Table Type

COMPRECLIST is a table type of the COMPREC Record Type.

Syntax

```
TYPE compreclist IS TABLE OF comprec;
```

Related Topics

- COMPREC Record Type
  The COMPREC record type is a record for calculating an individual index compression ratio on a table.

48.5 Summary of DBMS_COMPRESSION Subprograms

The DBMS_COMPRESSION package uses the GET_COMPRESSION_RATIO Procedure and GET_COMPRESSION_TYPE Function subprograms.

Table 48-3 DBMS_COMPRESSION Package Subprograms

Subprogram	Description
GET_COMPRESSION_RATIO Procedure	Analyzes the compression ratio of a table, and gives information about compressibility of a table
GET_COMPRESSION_TYPE Function	Returns the compression type for a specified row

48.5.1 GET_COMPRESSION_RATIO Procedure

Use this procedure to estimate the storage space that you can save by enabling the compression feature for an existing SecureFile LOB. It analyzes the compression ratio of a

table or an index and gives information about compressibility of the object. You can provide various parameters to selectively analyze different compression types.

In Oracle Database 23ai, this procedure has been enhanced to estimate the compression ratio faster for LOBs while using less space. Now you can also estimate the compression ratio for BasicFile LOBs. This helps you decide upfront whether you want to compress BasicFile LOBs, before migrating BasicFile LOBs to SecureFile LOBs. You can also estimate the compression ratio at the LOB byte level and the time taken, in hours, to compress the LOB data in the table.

The compression ratio is estimated for the number of rows in the LOB column that you specify. For example, let's consider that the compression ratio is 2.33. It indicates that after you enable the compression feature, you can save around half of the space for the sampled rows in the LOB column.

Disclaimer: The compression ratio is an approximate value, which is calculated based on the sampled rows in the LOB column. The actual space that you save when you enable compression for the complete table may be different.

Syntax

The syntax to get the compression ratio differs for objects, LOBs, IOTs, and indexes on a table.

- Syntax to get the compression ratio for an object (table or index, default is table).

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
   scratchtbsname        IN     VARCHAR2,
   ownname               IN     VARCHAR2,
   objname               IN     VARCHAR2,
   subobjname            IN     VARCHAR2,
   comptype              IN     NUMBER,
   blkcnt_cmp            OUT    PLS_INTEGER,
   blkcnt_uncmp          OUT    PLS_INTEGER,
   row_cmp               OUT    PLS_INTEGER,
   row_uncmp             OUT    PLS_INTEGER,
   cmp_ratio             OUT    NUMBER,
   comptype_str          OUT    VARCHAR2,
   block_compr_ratio     OUT    PLS_INTEGER,
   byte_comp_ratio       OUT    NUMBER,
   subset_numrows        IN     NUMBER  DEFAULT COMP_RATIO_MINROWS,
   objtype               IN     PLS_INTEGER DEFAULT OBJTYPE_TABLE);
```

- Syntax to get compression ratio for BasicFile and SecureFile LOBs:

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
   scratchtbsname        IN     VARCHAR2,
   tabowner              IN     VARCHAR2,
   tabname               IN     VARCHAR2,
   lobname               IN     VARCHAR2,
   partname              IN     VARCHAR2,
   comptype              IN     NUMBER,
   blkcnt_cmp            OUT    PLS_INTEGER,
   blkcnt_uncmp          OUT    PLS_INTEGER,
   lobcnt                OUT    PLS_INTEGER,
   cmp_ratio             OUT    NUMBER,
   comptype_str          OUT    VARCHAR2,
   byte_comp_ratio       OUT    NUMBER,
   total_time            OUT    NUMBER
   subset_numrows        IN     NUMBER DEFAULT COMP_RATIO_LOB_MAXROWS);
```

- Syntax to get the compression ratio for all indexes on a table. The compression ratios are returned as a collection.

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
  scratchtbsname        IN     VARCHAR2,
  ownname               IN     VARCHAR2,
  tabname               IN     VARCHAR2,
  comptype              IN     NUMBER,
  index_cr              OUT    DBMS_COMPRESSION.COMPRECLIST,
  comptype_str          OUT    VARCHAR2,
  subset_numrows        IN     NUMBER DEFAULT COMP_RATIO_INDEX_MINROWS);
```

- Syntax to get the compression ratio for IOTs.

```
DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
     scratchtbsname       IN     VARCHAR2,
     ownname              IN     VARCHAR2,
     objname              IN     VARCHAR2,
     subobjname           IN     VARCHAR2,
     comptype             IN     NUMBER,
     iotcomp_cr           OUT    DBMS_COMPRESSION.COMPRECLIST,
     comptype_str         OUT    VARCHAR2,
     subset_numrows       IN     NUMBER DEFAULT COMP_RATIO_INDEX_MINROWS);
```

Parameters

Table 48-4 GET_COMPRESSION_RATIO Procedure Parameters

Parameter	Description
`scratchtbsname`	Temporary scratch tablespace that can be used for analysis
`ownname` / `tabowner`	Schema of the table to analyze
`tabname`	Name of the table to analyze
`objname`	Name of the object
`subobjname`	Name of the partition or sub-partition of the object
`comptype`	Compression types for which analysis should be performed When the object is an index, only the following compression types are valid: `COMP_INDEX_ADVANCED_HIGH` (value `1024`) and `COMP_INDEX_ADVANCED_LOW` (value `2048`). **Note:** The following compression types cannot be specified in this parameter for any type of object: `COMP_BLOCK` (value `64`) and `COMP_BASIC` (value `4096`).
`blkcnt_cmp`	Number of blocks used by compressed sample of the table
`blkcnt_uncmp`	Number of blocks used by uncompressed sample of the table
`row_cmp`	Number of rows in a block in compressed sample of the table
`row_uncmp`	Number of rows in a block in uncompressed sample of the table
`cmp_ratio`	Compression ratio, `blkcnt_uncmp` divided by `blkcnt_cmp`. It provides the ratio of **blocks** occupied by the uncompressed data to the blocks occupied by the compressed data.
`comptype_str`	String describing the compression type
`subset_numrows`	Number of rows sampled to estimate compression ratio.
`objtype`	Type of the object, either `OBJTYPE_TABLE` or `OBJTYPE_INDEX`
`lobname`	Name of the LOB column
`partname`	In case of partitioned tables, the related partition name
`lobcnt`	Number of lobs actually sampled to estimate compression ratio

Table 48-4 (Cont.) GET_COMPRESSION_RATIO Procedure Parameters

Parameter	Description
byte_comp_ratio	Provides the ratio of bytes of uncompressed data to the bytes of compressed data for LOBs.
index_cr	List of indexes and their estimated compression ratios
iotcomp_cr	Compression ratio for the IOT The first object contains the compression ratio for the whole IOT. The second object contains the compression ratio only for the top index section of the IOT (excludes the overflow segment).
total_time	Provides an estimate of the time taken, in hours, to compress the LOB data in the table.

Example: Estimate the compression ratio for inline and out-of-line LOBs

The following example shows how to estimate the compression ratio for LOBs.

```
SET SERVEROUTPUT ON
DECLARE
    bcmp                    INTEGER;
    buncmp                  INTEGER;
    lobcmp                  INTEGER;
    cr                      NUMBER;
    byte_cr                 NUMBER;
    cstr                    VARCHAR2(2000);
    total_time              NUMBER;
    l_segment_name          VARCHAR2(30);
    l_segment_size_blocks   NUMBER;
    l_segment_size_bytes    NUMBER;
    l_used_blocks           NUMBER;
    l_used_bytes            NUMBER;
    l_expired_blocks        NUMBER;
    l_expired_bytes         NUMBER;
    l_unexpired_blocks      NUMBER;
    l_unexpired_bytes       NUMBER;
BEGIN
   DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
     scratchtbsname        => 'LOBTBSP',
     tabowner              => 'CMPADV',
     tabname               => p_tablename,
     lobname               => 'C',
     partname              => NULL,
     comptype              => 256,
     blkcnt_cmp            => bcmp,
     blkcnt_uncmp          => buncmp,
     lobcnt                => lobcmp,
     cmp_ratio             => cr,
     comptype_str          => cstr,
     subset_numrows        => 1000,
     byte_comp_ratio       => byte_cr,
     total_time            => total_time
);
DBMS_OUTPUT.put_line('Estimated ratio of blocks used by the uncompressed data
```

```
to the compressed data :  ' || cr);
DBMS_OUTPUT.put_line('Estimated ratio of bytes used by the uncompressed data
to the compressed data  :  ' || byte_cr);
END;
/
```

To understand the output of this procedure, let's consider `tab_inline`, an inline table, and `tab_outofline`, an out-of-line table as shown in the following example.

```
CREATE TABLE tab_inline
(
    a NUMBER,
    c CLOB
)
LOB(c) STORE AS SECUREFILE  (ENABLE STORAGE IN ROW CACHE LOGGING);


CREATE TABLE tab_outofline
(
    a NUMBER,
    c CLOB
)
LOB(c) STORE AS SECUREFILE  (DISABLE STORAGE IN ROW CACHE LOGGING);
```

Data is stored in different ways in `tab_inline` and `tab_outofline`. In the `tab_inline` table, if the LOB is less than 4K, then data is stored in the table segment; otherwise, it is stored in the LOB segment. For the `tab_outofline` table, data of all sizes is stored in the LOB segment.

Let's consider that you have inserted 1000 LOBs of 3K each in both the tables, and then calculate the compression ratios. You can use the `dbms_space.space_usage` procedure to calculate the space used by the data that is stored in the LOB segments.

Sample output of compression ratio for inline LOBs.

```
Estimated block compression ratio : 1
Estimated byte compression ratio           : 57.6
Space used(in bytes)                       : 0
space used(in blocks)                      : 0
```

Sample output of compression ratio for out-of-line LOBs.

```
Estimated block compression ratio : 1
Estimated byte compression ratio           : 56.1
Space used(in bytes)                       : 8 MB
space used(in blocks)                      : 1000
```

In this example, even though the estimated byte and block compression ratios are almost the same for inline and out-of-line LOBs, the space that is used is different. In the case of `tab_inline`, LOB segment is not used so the space used is 0. In both cases, the data is approximately 3KB, which is small. Therefore, the data before and after compression uses the same number of blocks (that is 1 block), so the block compression ratio is 1. However, the byte level compression ratio, `byte_comp_ratio`, which compares the actual number of bytes used by the LOBs before and after compression is 57.6 or 56.1.

Example: Estimate the compression ratio for indexes on a table with low compression type

The following example shows how to estimate the compression ratio for advanced index compression (low):

```
SET SERVEROUTPUT ON
DECLARE
  l_blkcnt_cmp       PLS_INTEGER;
  l_blkcnt_uncmp   PLS_INTEGER
  l_row_cmp       PLS_INTEGER;
  l_row_uncmp       PLS_INTEGER;
  l_cmp_ratio       NUMBER;
  l_comptype_str   VARCHAR2(32767);
BEGIN
   DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
    scratchtbsname =>       'USERS' ,
    ownname       =>       'TEST' ,
    objname       =>       'SALES_IDX' ,
    subobjname     =>        NULL ,
    comptype       =>        DBMS_COMPRESSION.COMP_INDEX_ADVANCED_LOW,
    blkcnt_cmp     =>        l_blkcnt_cmp,
    blkcnt_uncmp   =>        l_blkcnt_uncmp,
    row_cmp       =>        l_row_cmp,
    row_uncmp      =>        l_row_uncmp,
    cmp_ratio      =>        l_cmp_ratio,
    comptype_str   =>        l_comptype_str,
    subset_numrows =>        DBMS_COMPRESSION.comp_ratio_minrows,
    objtype       =>        DBMS_COMPRESSION.objtype_index
   );
DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
object   :  ' || l_blkcnt_cmp);
DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the object :  ' || l_blkcnt_uncmp);
DBMS_OUTPUT.put_line( 'Number of rows in a block in compressed sample of the
object   :  ' || l_row_cmp);
DBMS_OUTPUT.put_line( 'Number of rows in a block in uncompressed sample of
the object :  ' || l_row_uncmp);
DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of
Sample                          :  ' || l_cmp_ratio);
DBMS_OUTPUT.put_line( 'Compression Type                                :  ' ||
l_comptype_str);
END;
/
```

Output of compression advisor estimate for advanced index compression (Low):

```
Number of blocks used by the compressed sample of the object     : 243
Number of blocks used by the uncompressed sample of the object   : 539
Number of rows in a block in compressed sample of the object     : 499
Number of rows in a block in uncompressed sample of the object   : 145
Estimated Compression Ratio of Sample                 : 2.2
Compression Type                        : “Compress Advanced Low”
```

Example: Estimate the compression ratio for LOBs with medium compression type

The following example shows how to estimate the compression ratio for advanced LOB compression (medium):

```
SET SERVEROUTPUT ON
DECLARE
  l_blkcnt_cmp     PLS_INTEGER;
  l_blkcnt_uncmp   PLS_INTEGER;
  l_row_cmp       PLS_INTEGER;
  l_lobcnt       PLS_INTEGER;
  l_cmp_ratio      NUMBER;
  l_comptype_str   VARCHAR2(32767);
BEGIN
   DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
    scratchtbsname => 'USERS' ,
    tabowner       => 'TEST' ,
    tabname       => 'PARTS' ,
    lobname       =>  'PART_DESCRIPTION' ,
    partname       =>  NULL ,
    comptype       =>  DBMS_COMPRESSION.COMP_LOB_MEDIUM,
    blkcnt_cmp     => l_blkcnt_cmp,
    blkcnt_uncmp   => l_blkcnt_uncmp,
    row_cmp       => l_row_cmp,
    lobcnt       => l_lobcnt,
    cmp_ratio      => l_cmp_ratio,
    comptype_str   => l_comptype_str,
    subset_numrows => DBMS_COMPRESSION.comp_ratio_lob_maxrows
   );
DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
object    :  ' || l_blkcnt_cmp);
DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the object  :  ' || l_blkcnt_uncmp);
DBMS_OUTPUT.put_line( 'Number of rows in a block in compressed sample of the
object    :  ' || l_row_cmp);
DBMS_OUTPUT.put_line( 'Number of LOBS actually
sampled                     :  ' || l_lobcnt);
DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of
Sample                          :  ' || l_cmp_ratio);
DBMS_OUTPUT.put_line( 'Compression
Type                                                :  ' || l_comptype_str);
END;
/
```

Output of compression advisor estimate for advanced LOB compression (Medium):

```
Number of blocks used by the compressed sample of the object   : 199
Number of blocks used by the uncompressed sample of the object     : 389
Number of rows in a block in compressed sample of the object       : 293
Number of LOBS actually sampled                              : 55
Estimated Compression Ratio of Sample                        : 1.9
Compression Type                                      : “Compress Medium”
```

Example: Estimate the compression ratio for IoTs

The following example shows how to estimate the compression ratio for IOTs:

```
SET SERVEROUTPUT ON
DECLARE
  bcmp       INTEGER;
  buncmp     INTEGER;
  rowcmp     INTEGER;
  rowuncmp   INTEGER;
  cr         NUMBER;
  cstr       VARCHAR2(2000);
  iotcomp_cr DBMS_COMPRESSION.COMPRECLIST;
 BEGIN
    DBMS_COMPRESSION.GET_COMPRESSION_RATIO (
    scratchtbsname        => 'USERS',
    ownname               => 'TEST',
    objname               => 'SALES',
    subobjname            =>  NULL,
    comptype              =>  DBMS_COMPRESSION.COMP_INDEX_ADVANCED_LOW,
    iotcomp_cr            =>  iotcomp_cr,
    comptype_str          =>  cstr,
    subset_numrows        =>  DBMS_COMPRESSION.COMP_RATIO_ALLROWS
    );
 --information about the index and the overflow segment
 DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
IOT table                                    : ' || iotcomp_cr(1).blkcnt_cmp);
 DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the IOT table                                 : ' ||
iotcomp_cr(1).blkcnt_uncmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the compressed
sample of the IOT table                       : ' || iotcomp_cr(1).row_cmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the uncompressed
sample of the IOT table                       : ' || iotcomp_cr(1).row_uncmp);
 DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of the
sample                                                    : ' ||
iotcomp_cr(1).cmp_ratio);
 --information about the index segment
 DBMS_OUTPUT.put_line( 'Number of blocks used by the compressed sample of the
index segment of the IOT table              : ' || iotcomp_cr(2).blkcnt_cmp);
 DBMS_OUTPUT.put_line( 'Number of blocks used by the uncompressed sample of
the index segment of the IOT table            : ' ||
iotcomp_cr(2).blkcnt_uncmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the compressed
sample of the index segment of the IOT table   : ' || iotcomp_cr(2).row_cmp);
 DBMS_OUTPUT.put_line( 'Average number of rows in a block in the uncompressed
sample of the index segment of the IOT table : ' || iotcomp_cr(2).row_uncmp);
 DBMS_OUTPUT.put_line( 'Estimated Compression Ratio of the
sample                                                    : ' ||
iotcomp_cr(2).cmp_ratio);
 END;
 /
```

Output of the compression ratio for IOTs:

```
Number of blocks used by the compressed sample of the IOT
table                                    : 5027
Number of blocks used by the uncompressed sample of the IOT
table                                   : 7950
Average number of rows in a block in the compressed sample of the IOT
table                          : 199
Average number of rows in a block in the uncompressed sample of the IOT
table                        : 126
Estimated Compression Ratio of the
sample                                                          : 1.58
Number of blocks used by the compressed sample of the index segment of the IOT
table                : 3238
Number of blocks used by the uncompressed sample of the index segment of the IOT
table             : 6161
Average number of rows in a block in the compressed sample of the index segment of the
IOT table   : 309
Average number of rows in a block in the uncompressed sample of the index segment of the
IOT table : 162
Estimated Compression Ratio of the
sample                                                          : 1.9
```

Usage Notes

- The procedure creates different tables in the scratch tablespace and runs analysis on these objects. It does not modify anything in the user-specified tables.
- From 23ai onwards, this feature has been enhanced to estimate the compression ratio faster for LOBs while using less space. To get a more accurate result, run the following command to switch to the old method. The older method to calculate the compression ratio takes more time to return the results and uses more space.

  ```
  alter session set "_kdlf_new_compression_adv"= FALSE;
  ```

- To understand the impact of compression, use the value of the byte compression ratio for inline LOBs and for out-of-line LOBs, use the value of the block compression ratio and space used.
- You can get more benefits when you compress large volume of data as compared to small volumes of data. If you want to compress small volumes of data, look at the byte ratio instead of the block ratio to understand the impact of compression.

48.5.2 GET_COMPRESSION_TYPE Function

This function returns the compression type for a specified row. If the row is chained, the function returns the compression type of the head piece only, and does not examine the intermediate or the tail piece since head pieces can be differently compressed.

Syntax

```
DBMS_COMPRESSION.GET_COMPRESSION_TYPE (
   ownname      IN    VARCHAR2,
   tabname      IN    VARCHAR2,
   row_id       IN    ROWID,
   subobjname   IN    VARCHAR2 DEFAULT NULL))
  RETURN NUMBER;
```

Parameters

Table 48-5 GET_COMPRESSION_TYPE Function Parameters

Parameter	Description
ownname	Schema name of the table
tabname	Name of table
rowid	Rowid of the row
subobjname	Name of the table partition or subpartition

Return Values

Flag to indicate the compression type (see Table 48-1).

49
DBMS_CONNECTION_POOL

The DBMS_CONNECTION_POOL package provides an interface to manage the Database Resident Connection Pools (DRCP).

The database initialization parameter ENABLE_PER_PDB_DRCP controls whether DRCP is configured in the CDB ROOT DRCP mode or per-PDB DRCP mode. The default value is FALSE, which will configure DRCP at the CDB level.

- If the value of the database configuration parameter ENABLE_PER_PDB_DRCP is set to TRUE, then:
  - Each PDB administrator can manage their pool configuration using the DBMS_CONNECTION_POOL package. If the ROOT tries to manage the pool configuration using the DBMS_CONNECTION_POOL package, then an error is thrown.
  - The values of the num_cbrok and maxconn_cbrok pool parameters from the cpool$ table are ignored. The PDB administrator cannot modify these parameters using the DBMS_CONNECTION_POOL.ALTER_PARAM() procedure. These parameters can be set using the CONNECTION_BROKERS database parameter. Only the ROOT can alter these parameters dynamically.
- If the value of the ENABLE_PER_PDB_DRCP parameter is set to FALSE, then only the ROOT can manage the pool configuration, using the DBMS_CONNECTION_POOL package. If a PDB administrator tries to manage the pool configuration using the DBMS_CONNECTION_POOL package, then an error is thrown.
- The following DRCP parameters can have values ranging from 0 to SB4MAXVAL (2147483647) - 1:
  - minsize
  - num_cbrok
  - maxconn_cbrok

> **See Also:**
>
> *Oracle Database Concepts* for more information on "Database Resident Connection Pooling"

This chapter contains the following topic:

- Summary of DBMS_CONNECTION_POOL Subprograms

49.1 Summary of DBMS_CONNECTION_POOL Subprograms

This table lists the subprograms of the DBMS_CONNECTION_POOL package in an alphabetical order and describes them briefly.

Table 49-1 DBMS_CONNECTION_POOL Package Subprograms

Subprogram	Description
ADD_POOL Procedure	Adds a new pool to the multiple pool DRCP.
ALTER_PARAM Procedure	Alters a specific configuration parameter as a standalone unit, without affecting the other parameters.
CONFIGURE_POOL Procedure	Configures the pool with advanced options.
REMOVE_POOL Procedure	Removes a pool from the multiple pool DRCP.
RESTORE_DEFAULTS Procedure	Restores the pool to the default settings
START_POOL Procedure	Starts the pool for operations. Only after this procedure is called, the pool can be used by the connection clients for creating sessions .
STOP_POOL Procedure	Stops the pool and makes it unavailable for the registered connection clients.

49.1.1 ADD_POOL Procedure

You can use this procedure to add a new DRCP pool.

Syntax

```
DBMS_CONNECTION_POOL.ADD_POOL (
   pool_name                IN VARCHAR2,
   minsize                  IN PLS_INTEGER   DEFAULT 0,
   maxsize                  IN PLS_INTEGER   DEFAULT 40,
   incrsize                 IN PLS_INTEGER   DEFAULT 2,
   session_cached_cursors   IN PLS_INTEGER   DEFAULT 20,
   inactivity_timeout       IN PLS_INTEGER   DEFAULT 300,
   max_think_time           IN PLS_INTEGER   DEFAULT 120,
   max_use_session          IN PLS_INTEGER   DEFAULT 500000,
   max_lifetime_session     IN PLS_INTEGER   DEFAULT 86400,
   max_txn_think_time       IN PLS_INTEGER   DEFAULT 0);
```

Parameters

> **Note:**
>
> If you are aware of the pool configuration at the time of adding the pool, then you can specify the values for the following parameters. Otherwise, you can call the `CONFIGURE_POOL` subprogram later and specify the new pool name and the configuration options.

Table 49-2 ADD_POOL Procedure Parameters

Parameter	Description
`pool_name`	The name of the pool to be added to the DRCP.
`minsize`	The minimum number of pooled servers in the pool. The default value is 0.
`maxsize`	The maximum number of pooled servers allowed in the pool. The default value is 40.

Table 49-2 (Cont.) ADD_POOL Procedure Parameters

Parameter	Description
incrsize	Pool would increment by this number of pooled server when pooled servers are unavailable at application request time. The default value is 2.
session_cached_cursors	The number of session cursors to cache in each pooled server session. The default value is 20. Turn on SESSION_CACHED_CURSORS for all connections in the pool. This is an existing init.ora parameter.
inactivity_timeout	TTL (Time to live) for an idle session in the pool. This parameter helps to shrink the pool when it is not used to its maximum capacity. If a connection remains in the pool idle for this time, the connection is closed. The default value is 300.
max_think_time	The maximum time of inactivity, in seconds, for a client after it obtains a pooled server from the pool with no open transactions in it. After obtaining a pooled server from the pool, if the client application does not issue a database call for the time specified by MAX_THINK_TIME, the pooled server is freed and the client connection is terminated. The default value is 120.
max_use_session	Maximum number of times a connection can be taken and released to the pool. The default value is 500000.
max_lifetime_session	TTL (Time to live) in seconds for a pooled session. The default value is 86400.
max_txn_think_time	The maximum time of inactivity, in seconds, for a client after it obtains a pooled server from the pool with an open transaction. After obtaining the pooled server from the pool, if the client application does not issue a database call for the time specified by MAX_TXN_THINK_TIME, then the pooled server is freed, and the client connection is terminated. The default value of this parameter is the value of the MAX_THINK_TIME parameter. Applications can set the value of the MAX_TXN_THINK_TIME parameter to a value higher than the MAX_THINK_TIME value to allow more time for the connections with open transactions.

Note:

The pool_name parameter has the following validation checks:

- The pool name must be unique.
- It must begin with an alphabetic character.
- It can contain alphanumeric characters in the second and subsequent positions.
- It can contain an underscore (_) in the second and subsequent positions.
- The maximum length allowed for the pool name is 128.

> **See Also:**
>
> For the list and description of all the database resident connection pooling parameters that can be configured using this procedure, see the *Oracle Database Administrator's Guide*.

Examples

```
exec dbms_connection_pool.add_pool('mypool')
```

49.1.2 ALTER_PARAM Procedure

This procedure alters a specific Database Resident Connection Pool (DRCP) configuration parameter as a standalone unit, and does not affect other parameters.

Syntax

```
DBMS_CONNECTION_POOL.ALTER_PARAM (
   pool_name     IN  VARCHAR2 DEFAULT 'SYS_DEFAULT_CONNECTION_POOL',
   param_name    IN  VARCHAR2,
   param_value   IN  VARCHAR2);
```

Parameters

Table 49-3 ALTER_PARAM Procedure Parameters

Parameter	Description
`pool_name`	Pool to be configured.
`param_name`	Any parameter name from `CONFIGURE_POOL`. For broker parameters: `NUM_CBROK` and `MAXCONN_CBROK`, these rules apply: • If per-PDB DRCP is enabled, then `ALTER_PARAM()` cannot be used to set these parameters. Using the database initialization parameter `connection_brokers` is the only way to set them. • For the CDB ROOT-level DRCP, if these parameters are not set using `connection_brokers`, then the administrator can use the `ALTER_PARAM()` procedure to set them. Using `connection_brokers` is the recommended way to set these parameters.
`param_value`	Parameter value for `param_name`.

> **See Also:**
>
> For the list and description of all the DRCP parameters that can be configured using this procedure, see the *Oracle Database Administrator's Guide*.

Exceptions

Table 49-4 ALTER_PARAM Procedure Exceptions

Exception	Description
ORA-56500	Connection pool not found
ORA-56504	Invalid connection pool configuration parameter name
ORA-56505	Invalid connection pool configuration parameter value
ORA-56507	Connection pool alter configuration failed

Examples

```
DBMS_CONNECTION_POOL.ALTER_PARAM(
   'SYS_DEFAULT_CONNECTION_POOL', 'MAX_LIFETIME_SESSION', '120');
```

49.1.3 CONFIGURE_POOL Procedure

This procedure configures the pool with advanced options.

Syntax

```
DBMS_CONNECTION_POOL.CONFIGURE_POOL (
   pool_name                IN VARCHAR2 DEFAULT 'SYS_DEFAULT_CONNECTION_POOL',
   minsize                  IN NUMBER   DEFAULT 4,
   maxsize                  IN NUMBER   DEFAULT 40,
   incrsize                 IN NUMBER   DEFAULT 2,
   session_cached_cursors   IN NUMBER   DEFAULT 20,
   inactivity_timeout       IN NUMBER   DEFAULT 300,
   max_think_time           IN NUMBER   DEFAULT 120,
   max_use_session          IN NUMBER   DEFAULT 500000,
   max_lifetime_session     IN NUMBER   DEFAULT 86400,
   max_txn_think_time       IN NUMBER);
```

Parameters

Table 49-5 CONFIGURE_POOL Procedure Parameters

Parameter	Description
pool_name	Pool to be configured.
minsize	Minimum number of pooled servers in the pool
maxsize	Maximum allowed pooled servers in the pool
incrsize	Pool would increment by this number of pooled server when pooled server are unavailable at application request time
session_cached_cursors	Turn on SESSION_CACHED_CURSORS for all connections in the pool. This is an existing init.ora parameter
inactivity_timeout	TTL (Time to live) for an idle session in the pool. This parameter helps to shrink the pool when it is not used to its maximum capacity. If a connection remains in the pool idle for this time, it is killed.

Table 49-5 (Cont.) CONFIGURE_POOL Procedure Parameters

Parameter	Description
max_think_time	The maximum time of inactivity, in seconds, for a client after it obtains a pooled server from the pool with no open transactions in it. After obtaining a pooled server from the pool, if the client application does not issue a database call for the time specified by MAX_THINK_TIME, the pooled server is freed and the client connection is terminated.
max_use_session	Maximum number of times a connection can be taken and released to the pool
max_lifetime_session	TTL (Time to live) for a pooled session
max_txn_think_time	The maximum time of inactivity, in seconds, for a client after it obtains a pooled server from the pool with an open transaction. After obtaining the pooled server from the pool, if the client application does not issue a database call for the time specified by MAX_TXN_THINK_TIME, then the pooled server is freed, and the client connection is terminated. The default value of this parameter is the value of the MAX_THINK_TIME parameter. Applications can set the value of the MAX_TXN_THINK_TIME parameter to a value higher than the MAX_THINK_TIME value to allow more time for the connections with open transactions.

Exceptions

Table 49-6 CONFIGURE_POOL Procedure Exceptions

Exception	Description
ORA-56500	Connection pool not found
ORA-56507	Connection pool alter configuration failed

Usage Notes

- All expressions of time are in seconds
- All of the parameters should be set based on statistical request patterns.
- minsize should be set keeping in mind that it puts a lower bound on server resource consumption. This is to prevent the timeout from dragging the pool too low, because of a brief period of inactivity.
- maxsize should be set keeping in mind that it puts an upper bound on concurrency and response-times and also server resource consumption.
- session_cached_cursors is typically set to the number of most frequently used statements. It occupies cursor resource on the server
- In doubt, do not set the increment and inactivity_timeout. The pool will have reasonable defaults.
- max_use_session and max_lifetime_session allow for software rejuvenation or defensive approaches to potential bugs, leaks, accumulations, and like problems, by getting brand new sessions once in a while.

- The connection pool reserves 5% of the pooled servers for authentication, and at least one pooled server is always reserved for authentication. When setting the `maxsize` parameter, ensure that there are enough pooled servers for both authentication and connections.

49.1.4 REMOVE_POOL Procedure

You can use this procedure to remove a pool from DRCP.

Syntax

```
DBMS_CONNECTION_POOL.REMOVE_POOL (
   pool_name      IN  VARCHAR2);
```

Parameters

Table 49-7 REMOVE_POOL Procedure Parameters

Parameter	Description
`pool_name`	The pool to be removed.

> **Note:**
>
> This procedure generates an error if:
>
> - The specified pool does not exist.
> - You try to remove the default pool.
> - You try to remove an active pool without stopping it.

Exceptions

Table 49-8 REMOVE_POOL Procedure Exceptions

Exception	Description
`ORA-56620`	Removing the pool <pool_name> failed because pool doesn't exist.
`ORA-56620`	Removing the pool <pool_name> failed because database is read only.
`ORA-56620`	Removing the pool <pool_name> failed because pool is active.
`ORA-56620`	Removing the pool <pool_name==NULL> failed because pool name cannot be empty.
`ORA-56620`	Removing the pool `SYS_DEFAULT_CONNECTION_POOL` failed because default pool cannot be removed.

> **Note:**
>
> The same error number is used for all exception errors because the error message is a dynamic string.

Examples

```
exec dbms_connection_pool.remove_pool('mypool')
```

49.1.5 RESTORE_DEFAULTS Procedure

This procedure restores the pool to default settings.

Syntax

```
DBMS_CONNECTION_POOL.RESTORE_DEFAULTS (
   pool_name    IN  VARCHAR2 DEFAULT 'SYS_DEFAULT_CONNECTION_POOL');
```

Parameters

Table 49-9 RESTORE_DEFAULTS Procedure Parameters

Parameter	Description
pool_name	Pool to be restored.

Exceptions

Table 49-10 RESTORE_DEFAULTS Procedure Exceptions

Exception	Description
ORA-56500	Connection pool not found
ORA-56507	Connection pool alter configuration failed

49.1.6 START_POOL Procedure

This procedure starts the pool for operations. It is only after this call that the pool could be used by connection classes for creating sessions.

Syntax

```
DBMS_CONNECTION_POOL.START_POOL (
   pool_name  IN  VARCHAR2 DEFAULT 'SYS_DEFAULT_CONNECTION_POOL');
```

Parameters

Table 49-11 START_POOL Procedure Parameters

Parameter	Description
pool_name	Pool to be started.

Exceptions

Table 49-12 START_POOL Procedure Exceptions

Exception	Description
ORA-56500	Connection pool not found
ORA-56501	Connection pool startup failed

Usage Notes

If the instance is restarted (shutdown followed by startup), DRCP pools are automatically started.

49.1.7 STOP_POOL Procedure

This procedure stops the pool and makes it unavailable for the registered connection classes.

Syntax

```
DBMS_CONNECTION_POOL.STOP_POOL (
   pool_name    IN   VARCHAR2 DEFAULT 'SYS_DEFAULT_CONNECTION_POOL',
   draintime    IN   PLS_INTEGER DEFAULT 2147483647);
```

Parameters

Table 49-13 STOP_POOL Procedure Parameters

Parameter	Description
pool_name	Pool to be stopped.
draintime	Allows active DRCP pools to be closed after a specified connection drain time, or be closed immediately without waiting for the connections to be idle. draintime can have the following integer values: • 0 to close the bound pooled servers immediately. • x to close the bound pooled servers after x seconds.

Note:

Do not pass the draintime parameter if you want to retain the old behavior of closing the pools only after they become idle and inactive.

Exceptions

Table 49-14 STOP_POOL Procedure Exceptions

Exception	Description
ORA-56500	Connection pool not found
ORA-56506	Connection pool shutdown failed

Usage Notes

This stops the pool and takes it offline. This does not destroy the persistent data (such as, the pool name and configuration parameters) associated with the pool.

50
DBMS_CQ_NOTIFICATION

The DBMS_CQ_NOTIFICATION package is part of the database change notification feature that provides the functionality to create registration on queries designated by a client application and so to receive notifications in response to DML or DDL changes on the objects associated with the queries. The notifications are published by the database when the DML or DDL transaction commits.

> **See Also:**
>
> *Oracle Database Development Guide* regarding implementing database change notification.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operational Notes
- Examples
- Data Structures
- Summary of DBMS_CQ_NOTIFICATION Subprograms

50.1 DBMS_CQ_NOTIFICATION Overview

The DBMS_CQ_NOTIFICATION package provides PL/SQL based registration interfaces. A client can use this interface to create registrations on queries based on objects of interest and specify a PL/SQL callback handler to receive notifications. In case of object level registration, when a transaction changes any of the objects associated with the registered queries and |COMMIT|s, the PL/SQL callback, specified during registration for those objects, is invoked. The application can define client-specific processing inside the implementation of its PL/SQL callback handler.

The interface lets you define a registration block (using a mechanism similar to a BEGIN-END block). The recipient of notifications namely the name of the PL/SQL callback handler and a few other registration properties like time-outs can be specified during the BEGIN phase. Any queries executed subsequently (inside the registration block) are considered "interesting queries" and objects referenced by those queries during query execution are registered. The registration is completed by ENDing the registration block.The registration block lets you create new registrations or add objects to existing registrations.

When a registration is created through the PL/SQL interface, a unique registration ID is assigned to the registration by the RDBMS. The client application can use the registration ID to keep track of registrations created by it. When a notification is published by the RDBMS, the registration ID will be part of the notification.

Typical Applications

This functionality is useful for example to applications that cache query result sets on mostly read-only objects in the mid-tier to avoid network round trips to the database. Such an application can create a registration on the queries it is interested in caching. On changes to objects referenced inside those queries, the database publishes a notification when the underlying transaction commits. In response to the notification, the mid-tier application can refresh its cache by re-executing the query/queries.

50.2 DBMS_CQ_NOTIFICATION Security Model

The `DBMS_CQ_NOTIFICATION` package requires that the user have the `CHANGE NOTIFICATION` system privilege in order to receive notifications, and be granted `EXECUTE` privilege on the `DBMS_CQ_NOTIFICATION` package.

In addition the user is required to have `SELECT` or `READ` privileges on all objects to be registered. Note that if the `SELECT` or `READ` privilege on an object was granted at the time of registration creation but lost subsequently (due to a revoke), then the registration will be purged and a notification to that effect will be published.

50.3 DBMS_CQ_NOTIFICATION Constants

DBMS_CQ_NOTIFICATION constants are used as flag parameters either during registration or when received during the notification.

The `DBMS_CQ_NOTIFICATION` package has sets of constants:

- `EVENT_STARTUP`, `EVENT_SHUTDOWN`, `EVENT_SHUTDOWN_ANY`, `EVENT_DEREG` describe the type of the notification published by the database.
- `INSERTOP`, `DELETEOP`, `UPDATEOP`, `ALTEROP`, `DROPOP` and `UNKNOWNOP` describe the type of operation on a table (during a notification published by the database).
- `QOS_RELIABLE`, `QOS_DEREG_NFY`, `QOS_ROWIDs` describe registration Quality of Service properties that the client requires. These are specified during registration.

The constants are shown in the following table.

Table 50-1 DBMS_CQ_NOTIFICATION Constants

Name	Type	Value	Description
ALL_OPERATIONS	BINARY_INTEGER	0	Interested in being notified on all operations, specified as a parameter during registration
ALL_ROWS	BINARY_INTEGER	1	All rows within the table may have been potentially modified
EVENT_STARTUP	BINARY_INTEGER	1	Instance startup notification
EVENT_SHUTDOWN	BINARY_INTEGER	2	Instance shutdown notification
EVENT_SHUTDOWN_ANY	BINARY_INTEGER	3	Any instance shutdown when running Oracle Real Application Clusters (Oracle RAC)
EVENT_DEREG	BINARY_INTEGER	5	Registration has been removed
EVENT_OBJCHANGE	BINARY_INTEGER	6	Notification for object change

Table 50-1 (Cont.) DBMS_CQ_NOTIFICATION Constants

Name	Type	Value	Description
EVENT_QUERYCHANGE	BINARY_INTEGER	7	Notification for query result set change
INSERTOP	BINARY_INTEGER	2	Insert operation
UPDATEOP	BINARY_INTEGER	4	Update operation
DELETEOP	BINARY_INTEGER	8	Delete operation
ALTEROP	BINARY_INTEGER	16	Table altered
DROPOP	BINARY_INTEGER	32	Table dropped
UNKNOWNOP	BINARY_INTEGER	64	Unknown operation
QOS_RELIABLE	BINARY_INTEGER	1	Reliable or persistent notification. Also implies that the notifications will be inserted into the persistent storage atomically with the committing transaction that results in an object change.
QOS_DEREG_NFY	BINARY_INTEGER	2	Purge registration on first notification
QOS_ROWIDS	BINARY_INTEGER	4	Require rowids of modified rows
QOS_QUERY	BINARY_INTEGER	8	Register at query granularity
QOS_BEST_EFFORT	BINARY_INTEGER	16	Best effort evaluation
NTFN_GROUPING_CLASS_TIME	BINARY_INTEGER	1	Group notifications by time
NTFN_GROUPING_TYPE_SUMMARY	BINARY_INTEGER	1	Summary grouping of notifications
NTFN_GROUPING_TYPE_LAST	BINARY_INTEGER	2	Last notification in the group
NTFN_GROUPING_FOREVER	BINARY_INTEGER	-1	Repeat notifications forever

50.4 DBMS_CQ_NOTIFICATION Operational Notes

The following are DBMS_CQ_NOTIFICATION operational notes.

Object Level Registration

- The notifications are published by the database when a transaction changes the registered objects and COMMITs.
- All objects referenced in the queries executed inside the registration block starting from the previous NEW_REG_START or ENABLE_REG to REG_END are considered interesting objects and added to the registration.

Query Result Change Registration

- The notifications are published by the database when a transaction changes the result set of the registered query and COMMITs.

Troubleshooting

If you have created a registration and seem to not receive notifications when the underlying tables are changed, please check the following.

- Is the `job_queue_processes` parameter set to a nonzero value? This parameter needs to be configured to a nonzero value in order to receive PL/SQL notifications through the handler.
- Are the registrations being created as a non-`SYS` user?
- If you are attempting DML changes on the registered object, are you `COMMIT`ing the transaction? Please note that the notifications are transactional and will be generated when the transaction `COMMIT`s.
- It maybe possible that there are run-time errors during the execution of the PL/SQL callback due to implementation errors. If so, they would be logged to the trace file of the `JOBQ` process that attempts to execute the procedure. The trace file would be usually named <`ORACLE_SID`>_j*_<`PID`>.`trc`. '

  For example, if the ORACLE_SID is 'dbs1' and the process is 12483, the trace file might be named '`dbs1_j000_12483.trc`.

  Suppose a registration is created with '`chnf_callback` as the notification handler and with `registration_id` 100. Let us suppose the user forgets to define the `chnf_callback` procedure. Then the `JOBQ` trace file might contain a message of the following form.

```
Runtime error during execution of PL/SQL cbk chnf_callback for reg CHNF100
   Error in PLSQL notification of msgid:
   Queue :
   Consumer Name :
   PLSQL function :chnf_callback
   Exception Occured, Error msg:
   ORA-00604: error occurred at recursive SQL level 2
   ORA-06550: line 1, column 7:
   PLS-00201: identifier 'CHNF_CALLBACK' must be declared
   ORA-06550: line 1, column 7:
   PL/SQL: Statement ignored
```

> **See Also:**
>
> For more information about troubleshooting Database Change Notification, see *Oracle Database Development Guide*.

50.5 DBMS_CQ_NOTIFICATION Examples

Suppose that a mid-tier application has a lot of queries on the `HR.EMPLOYEES` table. If the `EMPLOYEES` table is infrequently updated, it can obtain better performance by caching rows from the table because that would avoid a round-trip to the backend database server and server side execution latency. Let us assume that the application has implemented a mid-tier HTTP listener that listens for notifications and updates the mid-tier cache in response to a notification.

Object Change Registration Example

The `DBMS_CQ_NOTIFICATION` package can be utilized in this scenario to send notifications about changes to the table by means of the following steps:

1. Implement a mid-tier listener component of the cache management system (for example, using HTTP) that listens to notification messages sent from the database and refreshes the mid-tier cache in response to the notification.
2. Create a server side stored procedure to process notifications

```
CONNECT system;
Enter password: password
GRANT CHANGE NOTIFICATION TO hr;
GRANT EXECUTE ON DBMS_CQ_NOTIFICATION TO hr;

Rem Enable job queue processes to receive notifications.
ALTER SYSTEM SET "job_queue_processes"=2;

CONNECT hr;
Enter password: password
Rem Create a table to record notification events
CREATE TABLE nfevents(regid number, event_type number);

Rem create a table to record changes to registered tables
CREATE TABLE nftablechanges(regid number, table_name varchar2(100),
                            table_operation number);

Rem create a table to record rowids of changed rows.
CREATE TABLE nfrowchanges(regid number, table_name varchar2(100),
                          row_id varchar2(30));

Rem Create a PL/SQL callback handler to process notifications.
CREATE OR REPLACE PROCEDURE chnf_callback(ntfnds IN SYS.CHNF$_DESC) IS
   regid           NUMBER;
   tbname          VARCHAR2(60);
   event_type      NUMBER;
   numtables       NUMBER;
   operation_type  NUMBER;
   numrows         NUMBER;
   row_id          VARCHAR2(20);
 BEGIN
     regid      := ntfnds.registration_id;
     numtables  := ntfnds.numtables;
     event_type := ntfnds.event_type;

  INSERT INTO nfevents VALUES(regid, event_type);
  IF (event_type = DBMS_CQ_NOTIFICATION.EVENT_OBJCHANGE) THEN
    FOR i IN 1..numtables LOOP
      tbname          := ntfnds.table_desc_array(i).table_name;
      operation_type  := ntfnds.table_desc_array(I). Opflags;
      INSERT INTO nftablechanges VALUES(regid, tbname, operation_type);
      /* Send the table name and operation_type to client side listener using
UTL_HTTP */
      /* If interested in the rowids, obtain them as follows */
      IF (bitand(operation_type, DBMS_CQ_NOTIFICATION.ALL_ROWS) = 0) THEN
        numrows := ntfnds.table_desc_array(i).numrows;
      ELSE
        numrows :=0;   /* ROWID INFO NOT AVAILABLE */
      END IF;

      /* The body of the loop is not executed when numrows is ZERO */
      FOR j IN 1..numrows LOOP
          Row_id := ntfnds.table_desc_array(i).row_desc_array(j).row_id;
          INSERT INTO nfrowchanges VALUES(regid, tbname, Row_id);
         /* optionally Send out row_ids to client side listener using UTL_HTTP; */
      END LOOP;

    END LOOP;
  END IF;
  COMMIT;
END;
/
```

In Step 2 we can send as much information about the invalidation as the mid-tier application needs based on the information obtained from the notification descriptor.

Notes

a. In the above example, a registration was created on the EMPLOYEES table with 'chnf_callback' as the PL/SQL handler for notifications. During registration, the client specified reliable notifications (QOS_RELIABLE) and rowid notifications (QOS_ROWIDS)

b. The handler accesses the table descriptor array from the notification descriptor only if the notification type is of EVENT_OBJCHANGE. In all other cases (e.g EVENT_DEREG, EVENT_SHUTDOWN), the table descriptor array should not be accessed.

c. The handler accesses the row descriptor array from the table notification descriptor only if the ALL_ROWS bit is not set in the table operation flag. If the ALL_ROWS bit is set in the table operation flag, then it means that all rows within the table may have been potentially modified. In addition to operations like TRUNCATE that affect all rows in the tables, this bit may also be set if individual rowids have been rolled up into a FULL table invalidation.

   This can occur if too many rows were modified on a given table in a single transaction (more than 80) or the total shared memory consumption due to rowids on the RDBMS is determined too large (exceeds 1% of the dynamic shared pool size). In this case, the recipient must conservatively assume that the entire table has been invalidated and the callback/application must be able to handle this condition.

   Also note that the implementation of the user defined callback is up to the developer. In the above example, the callback was used to record event details into database tables. The application can additionally send the notification details to a mid-tier HTTP listener of its cache management system (as in the example) using UTL_HTTP. The listener could then refresh its cache by querying from the back-end database.

3. Create a registrations on the tables that we wish to be notified about. We pass in the previously defined procedure name (chnf_callback) as the name of the server side PL/SQL procedure to be executed when a notification is generated.

```
Rem Create a REGISTRATION on the EMPLOYEES TABLE
DECLARE
  REGDS      SYS.CHNF$_REG_INFO;
  regid      NUMBER;
  mgr_id     NUMBER;
  dept_id    NUMBER;
  qosflags   NUMBER;
BEGIN
 qosflags := DBMS_CQ_NOTIFICATION.QOS_RELIABLE +
         DBMS_CQ_NOTIFICATION.QOS_ROWIDS;
REGDS := SYS.CHNF$_REG_INFO ('chnf_callback', qosflags, 0,0,0);
regid := DBMS_CQ_NOTIFICATION.NEW_REG_START (REGDS);
SELECT manager_id INTO mgr_id FROM EMPLOYEES WHERE employee_id = 200;
DBMS_CQ_NOTIFICATION.REG_END;
END;
/
```

Once the registration is created in Step 3 above, the server side PL/SQL procedure defined in Step 2 is executed in response to any COMMITted changes to the HR.EMPLOYEES table. As an example, let us assume that the following update is performed on the employees table.

```
UPDATE employees SET salary=salary*1.05 WHERE employee_id=203;COMMIT;
```

Once the notification is processed, you will find rows which might look like the following in the `nfevents`, `nftablechanges` and `nfrowchanges` tables.

```
SQL> SELECT * FROM nfevents;

     REGID     EVENT_TYPE
--------------------------
     20045          6


SQL> SELECT * FROM nftablechanges;

  REGID      TABLE_NAME         TABLE_OPERATION
-----------------------------------------
  20045      HR.EMPLOYEES            4


SQL> select * from nfrowchanges;

    REGID      TABLE_NAME          ROW_ID
-----------------------------------------------------
    20045     HR.EMPLOYEES    AAAKB/AABAAAJ8zAAF
```

Query Result Change Registration Example

1. Creating a Callback

```
CONNECT system;
Enter password: password
GRANT CHANGE NOTIFICATION TO hr;
GRANT EXECUTE ON DBMS_CQ_NOTIFICATION TO hr;
CONNECT hr;
Enter password: password
Rem Create a table to record notification events
CREATE TABLE nfevents(regid NUMBER, event_type NUMBER);

Rem Create a table to record notification queries
CREATE TABLE nfqueries (qid NUMBER, qop NUMBER);

Rem Create a table to record changes to registered tables
CREATE TABLE nftablechanges(
       qid                NUMBER,
       table_name         VARCHAR2(100),
       table_operation    NUMBER);

Rem Create a table to record rowids of changed rows.
CREATE TABLE nfrowchanges(
       qid                NUMBER,
       table_name         VARCHAR2(100),
       row_id             VARCHAR2(2000));


CREATE OR REPLACE PROCEDURE chnf_callback
  (ntfnds IN CQ_NOTIFICATION$_DESCRIPTOR)
IS
  regid           NUMBER;
  tbname          VARCHAR2(60);
  event_type      NUMBER;
  numtables       NUMBER;
  operation_type  NUMBER;
  numrows         NUMBER;
  row_id          VARCHAR2(2000);
```

```
  numqueries      NUMBER;
  qid NUMBER;
  qop NUMBER;

BEGIN
  regid := ntfnds.registration_id;
  event_type := ntfnds.event_type;
  INSERT INTO nfevents VALUES(regid, event_type);
  numqueries :=0;
  IF (event_type = DBMS_CQ_NOTIFICATION.EVENT_QUERYCHANGE) THEN
      numqueries := ntfnds.query_desc_array.count;
      FOR i in 1..numqueries LOOP
         qid := ntfnds.QUERY_DESC_ARRAY(i).queryid;
         qop := ntfnds.QUERY_DESC_ARRAY(i).queryop;
         INSERT INTO nfqueries VALUES(qid, qop);
         numtables := 0;
         numtables := ntfnds.QUERY_DESC_ARRAY(i).table_desc_array.count;
         FOR j IN 1..numtables LOOP
           tbname := ntfnds.QUERY_DESC_ARRAY(i).table_desc_array(j).table_name;
           operation_type := ntfnds.QUERY_DESC_ARRAY(i).table_desc_array(j).Opflags;
           INSERT INTO nftablechanges VALUES(qid, tbname, operation_type);
           IF (bitand(operation_type, DBMS_CQ_NOTIFICATION.ALL_ROWS) = 0)
           THEN
             numrows := ntfnds.query_desc_array(i).table_desc_array(j).numrows;
           ELSE
             numrows :=0;   /* ROWID INFO NOT AVAILABLE */
           END IF;

            /* The body of the loop is not executed when numrows is ZERO */
            FOR k IN 1..numrows LOOP
               Row_id :=
ntfnds.query_desc_array(i).table_desc_array(j).row_desc_array(k).row_id;
               INSERT INTO nfrowchanges VALUES(qid, tbname, Row_id);

           END LOOP;  /* loop over rows */
        END LOOP;     /* loop over tables */
      END LOOP;        /* loop over queries */
  END IF;
  COMMIT;
END;
/
```

2. Creates a query registration

```
DECLARE
  reginfo    cq_notification$_reg_info;
  mgr_id     NUMBER;
  dept_id    NUMBER;
  v_cursor   SYS_REFCURSOR;
  regid      NUMBER;
  qosflags   NUMBER;

BEGIN
  /* Register two queries for result-set-change notifications: */

  /* 1. Construct registration information.
        'chnf_callback' is name of notification handler.
        QOS_QUERY specifies result-set-change notifications. */

   qosflags := DBMS_CQ_NOTIFICATION.QOS_QUERY +
               DBMS_CQ_NOTIFICATION.QOS_ROWIDS;
```

```
   reginfo := cq_notification$_reg_info('chnf_callback', qosflags,0, 0, 0);

  /* 2. Create registration */

        regid := DBMS_CQ_NOTIFICATION.NEW_REG_START(reginfo);

        OPEN v_cursor FOR
           SELECT DBMS_CQ_NOTIFICATION.CQ_NOTIFICATION_QUERYID, manager_id
           FROM HR.employees
           WHERE employee_id = 7902;
        CLOSE v_cursor;

       OPEN v_cursor for
        SELECT DBMS_CQ_NOTIFICATION.CQ_NOTIFICATION_QUERYID, department_id
        FROM HR.departments
        WHERE department_name = 'IT';
       CLOSE v_cursor;

  DBMS_CQ_NOTIFICATION.REG_END;
END;
/
```

3. After creating the query registrations, the output from `USER_CQ_NOTIFICATION_QUERIES` would appear as follows.

```
SQL> SELECT queryid, regid, to_char(querytext)
     FROM user_cq_notification_queries;

   QUERYID      REGID
---------- ----------
TO_CHAR(QUERYTEXT)
--------------------------------------------------------------------------------
        22         41
 SELECT HR.DEPARTMENTS.DEPARTMENT_ID  FROM HR.DEPARTMENTS  WHERE HR.DEPARTMENTS.
DEPARTMENT_NAME  = 'IT'

        21         41
 SELECT HR.EMPLOYEES.MANAGER_ID  FROM HR.EMPLOYEES  WHERE HR.EMPLOYEES.EMPLOYEE_
ID  = 7902


Now, let us perform an UPDATE that changes the result of the query with queryid 22
by renaming the department with name 'IT' to FINANCE.

SQL>  update departments set department_name = 'FINANCE' where department_name =
'IT';

1 row updated.

SQL> commit;

Commit complete.

Now we can query the notifications that we recorded in the callback.

SQL> select * from nfevents;

     REGID EVENT_TYPE
---------- ----------
        61          7


Event type 7 corresponds to EVENT_QUERYCHANGE
```

```
SQL> select * from nfqueries;

       QID        QOP
---------- ----------
        42          7

Event type 7 corresponds to EVENT_QUERYCHANGE

SQL> select * from nftablechanges;
SQL> select * from nftablechanges;

     REGID
----------
TABLE_NAME
--------------------------------------------------------------------------------
TABLE_OPERATION
---------------
        42
HR.DEPARTMENTS
              4


TABLE_OPERATION 4 corresponds to UPDATEOP

SQL> select * from nfrowchanges;
     REGID
----------
TABLE_NAME
--------------------------------------------------------------------------------
ROW_ID
--------------------------------------------------------------------------------
        61
HR.DEPARTMENTS
AAANkdAABAAALinAAF
```

50.6 DBMS_CQ_NOTIFICATION Data Structures

The `DBMS_CQ_NOTIFICATION` package defines several `OBJECT` types.

OBJECT Types

- CQ_NOTIFICATION$_DESCRIPTOR Object Type
- CQ_NOTIFICATION$_QUERY Object Type
- CQ_NOTIFICATION$_QUERY_ARRAY Object (Array) Type
- CQ_NOTIFICATION$_TABLE Object Type
- CQ_NOTIFICATION$_TABLE_ARRAY Object (Array) Type
- CQ_NOTIFICATION$_ROW Object Type
- CQ_NOTIFICATION$_ROW_ARRAY Object (Array) Type
- CQ_NOTIFICATION$_REG_INFO Object Type

50.6.1 CQ_NOTIFICATION$_DESCRIPTOR Object Type

This is the top level change notification descriptor type. It is a synonym for the `SYS.CHNF$_DESC` type.

Syntax

```
TYPE SYS.CHNF$_DESC IS OBJECT(
   registration_id    NUMBER,
   transaction_id     RAW(8),
   dbname             VARCHAR2(30),
   event_type         NUMBER,
   numtables          NUMBER,
   table_desc_array   CQ_NOTIFICATION$_TABLE_ARRAY,
   query_desc_array   CQ_NOTIFICATION$_QUERY_ARRAY);
```

Attributes

Table 50-2 CQ_NOTIFICATION$_DESCRIPTOR Object Type

Attribute	Description
`registration_id`	Registration ID returned during registration
`transaction_id`	Transaction ID. `transaction_id` of the transaction that made the change. Will be `NULL` unless the event_type is `EVENT_OBJCHANGE` or `EVENT_QUERYCHANGE`.
`dbname`	Name of database
`event_type`	Database event associated with the notification. Can be one of `EVENT_OBJCHANGE` (change to a registered object), `EVENT_STARTUP`, or `EVENT_QUERYCHANGE`, `EVENT_SHUTDOWN` or `EVENT_DEREG` (registration has been removed due to a timeout or other reason)
`numtables`	Number of modified tables. Will be `NULL` unless the `event_type` is `EVENT_OBJCHANGE`.
`table_desc_array`	Array of table descriptors. Will be `NULL` unless the `event_type` is `EVENT_OBJCHANGE`.
`query_desc_array`	Array of queries changed. This will be `NULL` unless event_type is `EVENT_QUERYCHANGE`

50.6.2 CQ_NOTIFICATION$_QUERY Object Type

The object type describes the changes to a query result caused by an event such as a transaction commit.

An array of `CQ_NOTIFICATION$_QUERY` descriptors is embedded inside the top level notification descriptor (`CQ_NOTIFICATION$_DESCRIPTOR`) for events of type `EVENT_QUERYCHANGE`. The array corresponds to the `SET` of queryids which were invalidated as a result of the event.

This is a synonym for the base type `SYS.CHNF$_QDESC`.

Syntax

```
TYPE SYS.CHNF$_QDESC IS OBJECT (
  queryid            NUMBER,
```

```
  queryop            NUMBER,
  table_desc_array   CQ_NOTIFICATION$_TABLE_ARRAY);
```

Attributes

Table 50-3 TYPE SYS.CQ_NOTIFICATION$_QUERY Object Type

Attribute	Description
queryid	QueryId of the changed query
queryop	Operation describing change to the query
table_desc_array	Array of table changes which contributed to the query Result Set change

50.6.3 CQ_NOTIFICATION$_QUERY_ARRAY Object (Array) Type

This type corresponds to an array of `CQ_NOTIFICATION$_QUERY` objects. It is a synonym for the `SYS.CHNF$_QUERY_ARRAY` type.

Syntax

```
TYPE CQ_NOTIFICATION$_TABLE_ARRAY IS VARRAY (1073741824) OF CQ_NOTIFICATION$_TABLE;
```

50.6.4 CQ_NOTIFICATION$_TABLE Object Type

This descriptor type describes a change to a table and is embedded inside the top level change notification descriptor type for `EVENT_OBJCHANGE` event types. For query result set changes (event type will be set to `EVENT_QUERYCHANGE`), the array of table descriptors is embedded inside each query change descriptor.

Note that this is a synonym for the type `previously named SYS.CHNF$_TDESC`.

Syntax

```
TYPE SYS.CHNF$_TDESC IS OBJECT (
  opflags            NUMBER,
  table_name         VARCHAR2(2*M_IDEN+1),
  numrows            NUMBER,
  row_desc_array     CQ_NOTIFICATION$_ROW_ARRAY)
```

Attributes

Table 50-4 TYPE SYS.CQ_NOTIFICATION$_TABLE Object Type

Attribute	Description
opflags	Table level operation flags. This is a flag field (bit-vector) that describes the operations that occurred on the table. It can be an `OR` of the following bit fields - `INSERTOP`, `UPDATEOP`, `DELETEOP`, `DROPOP`, `ALTEROP`, `ALL_ROWS`. If the `ALL_ROWS` (0x1) bit is set it means that either the entire table is modified (for example, `DELETE * FROM t`) or row level granularity of information is not requested or not available in the notification and the receiver has to conservatively assume that the entire table has been invalidated.
table_name	Name of modified table

Table 50-4 (Cont.) TYPE SYS.CQ_NOTIFICATION$_TABLE Object Type

Attribute	Description
numrows	Number of modified rows within the table. numrows will be NULL and hence should not be accessed if the ALL_ROWS bit is set in the table change descriptor.
row_desc_array	Array of row descriptors. This field will be NULL if the ALL_ROWS bit is set in opflags.

50.6.5 CQ_NOTIFICATION$_TABLE_ARRAY Object (Array) Type

This type corresponds to an array of CQ_NOTIFICATION$_TABLE objects. It is a synonym for the SYS.CHNF$_TDESC_ARRAY type.

Syntax

```
TYPE CQ_NOTIFICATION$_TABLE_ARRAY IS VARRAY (1073741824) OF CQ_NOTIFICATION$_TABLE;
```

50.6.6 CQ_NOTIFICATION$_ROW Object Type

An array of CQ_NOTIFICATION$_ROW is embedded inside a CQ_NOTIFICATION$_TABLE (table change descriptor), if the QOS_ROWIDS option was chosen at registration and the ALL_ROWS bit is not set in the opflags field of the table change descriptor.

Note that this is a synonym for the type previously named SYS.CHNF$_RDESC.

Syntax

```
TYPE SYS.CHNF$_RDESC IS OBJECT (
  opflags            NUMBER,
  row_id             VARCAHR2 (2000));
```

Attributes

Table 50-5 TYPE SYS.CQ_NOTIFICATION$_ROW Object Type

Attribute	Description
opflags	Row level operation flags. The flag field (bit vector) describes the operations in the row (could be INSERTOP, UPDATEOP or DELETEOP).
row_id	The rowid of the modified row

50.6.7 CQ_NOTIFICATION$_ROW_ARRAY Object (Array) Type

This object type corresponds to an array of CQ_NOTIFICATION$_ROW objects. It is embedded inside the CQ_NOTIFICATION$_TABLE if QOS_ROWIDS was specified during registration and the ALL_ROWS bit is not set in the opflags field of the table change descriptor.

This type is a synonym for the SYS.CHNF$_RDESC_ARRAY type.

Syntax

```
TYPE CQ_NOTIFICATION$_ROW_ARRAY IS VARRAY (1073741824) OF CQ_NOTIFICATION$_ROW;
```

50.6.8 CQ_NOTIFICATION$_REG_INFO Object Type

The object type describes the attributes associated with creating a new registration. It is a synonym for the type previously named `SYS.CHNF$_REG_INFO`.

Syntax

```
TYPE SYS.CHNF$_REG_INFO IS OBJECT (
  callback                      VARCHAR2(20),
  qosflags                      NUMBER,
  timeout                       NUMBER,
  operations_filter             NUMBER,
  transaction_lag               NUMBER,
  ntfn_grouping_class           NUMBER,
  ntfn_grouping_value           NUMBER,
  ntfn_grouping_type            NUMBER,
  ntfn_grouping_start_time      TIMESTAMP WITH TIME ZONE,
  ntfn_grouping_repeat_count    NUMBER);
```

Attributes

Table 50-6 TYPE CQ_NOTIFICATION$_REG_INFO Object Type

Attribute	Description
`callback`	Name of the server side PL/SQL procedure to be executed on a notification. Prototype is `<call_backname>(ntfnds IN SYS.chnf$_desc)`
`qosflags`	Quality of service flags. Can be set to an `OR` of the following values: • `QOS_RELIABLE` (`0x1`): Notifications are reliable (persistent) and survive instance death. This means that on an instance death in an Oracle RAC cluster, surviving instances will be able to deliver any queued invalidations. Similarly, pending invalidations can be delivered on instance restart, in a single instance configuration. The disadvantage is that there is a CPU cost/ latency involved in inserting the invalidation message to a persistent store. If this parameter is false, then server side CPU and latency are minimized, because invalidations are buffered into an in memory queue but the client could lose invalidation messages on an instance shutdown. • `QOS_DEREG_NFY` (`0x2`): The registration will be expunged on the first notification • `QOS_ROWIDS` (`0x4`): The notification needs to include information about the rowids that were modified • `QOS_QUERY (0x8)`: specifies query result change notification as opposed to object change notification • `QOS_BEST_EFFORT (0x10)` or `QOS_BEST_EFFORT (0x16)`: can register simplified versions of queries and minimizes evaluation with some false positives.

Table 50-6 (Cont.) TYPE CQ_NOTIFICATION$_REG_INFO Object Type

Attribute	Description
timeout	If set to a nonzero value, specifies the time in seconds after which the registration is automatically expunged by the database. If zero / NULL, the registration lives until explicitly deregistered. Note that the timeout option can be combined with the purge on notification (QOS_DEREG_NFY) option as well.
operations_filter	if nonzero, specifies a filter to be selectively notified on certain operations. These flags can be used to filter based on specific operation types: • 0: Notify on all operations (DBMS_CQ_NOTIFICATION.ALL_OPERATIONS) • 0x2: Notify on every INSERT (DBMS_CQ_NOTIFICATION.INSERTOP) • 0x4: Notify on every UPDATE (DBMS_CQ_NOTIFICATION.UPDATEOP) • 0x8: Notify on every DELETE (DBMS_CQ_NOTIFICATION.DELETEOP) A combination of operations can be specified by using a bitwise OR. Caution: This parameter will be honored for object level registrations but ignored for query result change registrations. To implement notification flow control in 11g, the applications can use the "GROUPING notification" option.
transaction_lag	Lag between consecutive notifications in units of transactions. Can be used to specify the number of transactions/database changes, by which the client is willing to lag behind the database. If 0, it means that the client needs to receive an invalidation message as soon as it is generated. **Caution**: This parameter will be honored for object level registrations but ignored for query result change notification registrations.
ntfn_grouping_class	When grouping notifications, the class based on which the group is derived. Currently, the only allowed value is DBMS_CQ_NOTIFICATION.NTFN_GROUPING_CLASS_TIME by which notifications are grouped by time.
ntfn_grouping_value	The grouping value. This describes the time interval that defines the group in seconds. For example, if this were set to 900, it would mean that notifications that were generated in each 15 minute interval would be grouped together.
ntfn_grouping_type	The type of grouping desired. It can be one of two allowed values • DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_SUMMARY - all notifications in the group are summarized into a single notification • DBMS_CQ_NOTIFICATION.NTFN_GROUPING_TYPE_LAST - only the last notification in the group is published and the earlier ones discarded
ntfn_grouping_start_time	When to start generating notifications. If specified as NULL, it defaults to the current system generated time.

Table 50-6 (Cont.) TYPE CQ_NOTIFICATION$_REG_INFO Object Type

Attribute	Description
ntfn_grouping_repeat_count	How many times the notification should be repeated. Set this to DBMS_CQ_NOTIFICATION.NTFN_GROUPING_FOREVER to receive notifications for the life time of the registration. Set to a nonzero value if only a certain number of notifications are desired for the life time of the registration.

Usage Notes

- The type declaration incorporates three other alternative constructors. In the first case all other parameters default to their default values.

```
TYPE CQ_NOTIFICATION$_REG_INFO IS OBJECT (
  callback                      VARCHAR2(20),
  qosflags                      NUMBER,
  timeout                       NUMBER);
```

  The second option applies to the type constructor defined in a previous release, and which is retained for backward compatibility:

```
TYPE CQ_NOTIFICATION$_REG_INFO IS OBJECT (
  callback                      VARCHAR2(20),
  qosflags                      NUMBER,
  timeout                       NUMBER,
  operations_filter             NUMBER,
  transaction_lag               NUMBER);
```

  The third definition contains all the members of the type except `transaction_lag` which is being deprecated:

```
TYPE CQ_NOTIFICATION$_REG_INFO IS OBJECT (
  callback                      VARCHAR2(20),
  qosflags                      NUMBER,
  timeout                       NUMBER,
  operations_filter             NUMBER,
  ntfn_grouping_class           NUMBER,
  ntfn_grouping_value           NUMBER,
  ntfn_grouping_type            NUMBER,
  ntfn_grouping_start_time      TIMESTAMP WITH TIME ZONE,
  ntfn_grouping_repeat_count    NUMBER);
```

- In response to a database change, the server side PL/SQL procedure specified by "callback" is executed. The PL/SQL procedure name has to be specified in the format `schema_name.procedure_name`. The procedure must have the following signature:

```
PROCEDURE <procedure_name>(ntfnds IN SYS.chnf$_desc)
```

  `CHNF$_DESC` describes the change notification descriptor.

- The init.ora parameter `job_queue_processes` must be set to a nonzero value to receive PL/SQL notifications, because the specified procedure is executed inside a job queue process when a notification is generated.

50.7 Summary of DBMS_CQ_NOTIFICATION Subprograms

This table lists the `DBMS_CQ_NOTIFICATION` subprograms and briefly describes them.

Table 50-7 DBMS_CQ_NOTIFICATION Package Subprograms

Subprogram	Description
CQ_NOTIFICATION_QUERYID Function	Returns the queryid of the most recent query that was attempted to be registered in a registration block
DEREGISTER Procedure	De-subscribes the client with the supplied registration identifier (ID)
ENABLE_REG Procedure	Begins a registration block using an existing registration identifier (ID)
NEW_REG_START Function	Begins a new registration block
REG_END Procedure	Ends the registration boundary
SET_ROWID_THRESHOLD Procedure	Configures the maximum number of rows of a table published in a change notification if the rows of the table are modified in a transaction

50.7.1 CQ_NOTIFICATION_QUERYID Function

This function returns the queryid of the most recent query that was attempted to be registered in a registration block.

Syntax

```
DBMS_CQ_NOTIFICATION.CQ_NOTIFICATION_QUERYID
 RETURN NUMBER;
```

Return Values

Returns the queryid of the most recently registered query.

50.7.2 DEREGISTER Procedure

This procedure describes the client with the specified registration identifier (ID).

Syntax

```
DBMS_CQ_NOTIFICATION.DEREGISTER (
  regid IN NUMBER);
```

Parameters

Table 50-8 DEREGISTER Procedure Parameters

Parameter	Description
`regid`	Client registration ID

Usage Notes

Only the user that created the registration (or the `SYS` user) will be able to desubscribe the registration.

50.7.3 ENABLE_REG Procedure

This procedure adds objects to an existing registration identifier (ID).

It is similar to the interface for creating a new registration, except that it takes an existing `regid` to which to add objects.Subsequent execution of queries causes the objects referenced in the queries to be added to the specified `regid`, and the registration is completed on invoking the REG_END Procedure.

Syntax

```
DBMS_CQ_NOTIFICATION.ENABLE_REG (
  regid IN NUMBER);
```

Parameters

Table 50-9 ENABLE_REG Procedure Parameters

Parameter	Description
`regid`	Client registration ID

Usage Notes

Only the user that created the registration will be able to add further objects to the registration.

50.7.4 NEW_REG_START Function

This procedure begins a new registration block.

Any objects referenced by queries executed within the registration block are considered interesting objects and added to the registration. The registration block ends upon calling the `REG_END` procedure.

Syntax

```
DBMS_CQ_NOTIFICATION.NEW_REG_START (
  regds IN sys.chnf$_reg_info)
 RETURN NUMBER;
```

Parameters

Table 50-10 NEW_REG_START Function Parameters

Parameter	Description
`sys.chnf$_reg_info`	Registration descriptor describing the notification handler and other properties of the registration

Return Values

The procedure returns a registration-id which is a unique integer assigned by the database to this registration. The registration-id will be echoed back in every notification received for this registration.

Usage Notes

- The only operations permitted inside a registration block are queries (the ones the user wishes to register). DML and DDL operations are not permitted.
- The registration block is a session property and implicitly terminates upon exiting the session. While the registration block is a session property, the registration itself is a persistent database entity. Once created, the registration survives until explicitly deregistered by the client application or timed-out or removed by the database for some other reason (such as loss of privileges).
- The user must have the `CHANGE NOTIFICATION` system privilege and `SELECT` or `READ` privileges on any objects to be registered.
- The `SYS` user will not be permitted to create new registrations.
- Nesting of registration block is not permitted.

50.7.5 REG_END Procedure

This procedure marks the end of the registration block. No newly executed queries are tracked.

Syntax

```
DBMS_CQ_NOTIFICATION.REG_END;
```

50.7.6 SET_ROWID_THRESHOLD Procedure

This procedure configures the maximum number of rows of a table published in a change notification if the rows of the table are modified in a transaction.

Syntax

```
DBMS_CQ_NOTIFICATION.SET_ROWID_THRESHOLD (
  tbname     IN  VARCHAR2,
  threshold  IN  NUMBER);
```

Parameters

Table 50-11 SET_ROWID_THRESHOLD Procedure Parameters

Parameter	Description
`tbname`	Table name qualified by the schema name in the form `schemaname.tablename`
`threshold`	Maximum number of modified rows of the table to be published in the change notification

Usage Notes

- The table needs to be registered for change notification either at object change granularity or at query result set granularity.
- The threshold set by means of this subprogram applies to that instance only and does not persist across instance startup/shutdown.

51
DBMS_CREDENTIAL

The DBMS_CREDENTIAL package provides an interface for authenticating and impersonating EXTPROC callout functions, as well as external jobs, remote jobs and file watchers from the SCHEDULER.

> **See Also:**
>
> - *Oracle Database Administrator's Guide* regarding Specifying Job Credentials
> - *Oracle Database Security Guide* regarding Guidelines for Securing External Processes

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Summary of DBMS_CREDENTIAL Subprograms

51.1 DBMS_CREDENTIAL Overview

Credentials are database objects that hold a username/password pair for authenticating and impersonating EXTPROC callout functions, as well as remote jobs, external jobs and file watchers from the SCHEDULER.

They are created using the CREATE_CREDENTIAL Procedure. The procedure also allows you to specify the Windows domain for remote external jobs executed against a Windows server.

51.2 DBMS_CREDENTIAL Security Model

Every Oracle credential has a unique credential name and you can associate a credential through its unique credential name with EXTPROC by means of a PL/SQL alias library.

In order to associate a credential with a PL/SQL alias library and external procedure, you must have the CREATE AND/OR REPLACE LIBRARY privilege or CREATE AND/OR REPLACE FUNCTION / PROCEDURE privilege and read permission of the DLL or shared object that the alias library to be associated with so that you can create and/or replace function or procedure to make use of the alias library.

Once authenticated, EXTPROC must act on behalf of the client based on client's identity defined in the supplied user credential. If not authenticated, EXTPROC must return an error message.

In order to create or alter a credential, you must have the CREATE CREDENTIAL privilege. If you are attempting to create or alter a credential in a schema other than your own, you must have the CREATE ANY CREDENTIAL privilege.

51.3 DBMS_CREDENTIAL Operational Notes

As the existing `CREATE OR REPLACE LIBRARY` statement and `CREATE OR REPLACE FUNCTION/PROCEDURE` do not support a `CREDENTIAL` clause, this model requires syntax and semantic changes in `CREATE OR REPLACE LIBRARY` and `CREATE OR REPLACE FUNCTION/PROCEDURE` statement.

For example:

```
CREATE OR REPLACE LIBRARY test
  AS '$ORACLE_HOME/bin/test.so' CREDENTIAL ricky_cred;
CREATE OR REPLACE FUNCTION ftest1
  (x VARCHAR2, y BINARY_INTEGER)
RETURN BINARY_INTEGER
AS LANGUAGE C
LIBRARY test
NAME "negative"
PARAMETERS(x STRING, y INT);
```

The credential name defined in the `CREDENTIAL` clause is a name of a database object. Therefore, do not enclose the credential name with single or double quotes.

An example of a credential being used on an external job:

```
BEGIN
   DBMS_SCHEDULER.CREATE_JOB(
   job_name            => 'example_job',
   job_type            => 'EXECUTABLE',
   job_action          => '/bin/ls',
   credential_name     => 'ricky_cred');
END;
/
```

51.4 Summary of DBMS_CREDENTIAL Subprograms

This table lists the `DBMS_CREDENTIAL` subprograms and briefly describes them.

Table 51-1 DBMS_CREDENTIAL Package Subprograms

Subprogram	Description
CREATE_CREDENTIAL Procedure	Creates a stored username/password pair in a database object called an Oracle credential
DISABLE_CREDENTIAL Procedure	Disables an Oracle credential
DROP_CREDENTIAL Procedure	Drops an Oracle credential
ENABLE_CREDENTIAL Procedure	Enables an Oracle credential
UPDATE_CREDENTIAL Procedure	Updates an existing Oracle credential

51.4.1 CREATE_CREDENTIAL Procedure

This procedure creates a stored username/password pair in a database object called an Oracle credential. You can also use this procedure to manage the credentials used for accessing files stored in cloud object storage.

Syntax

```
DBMS_CREDENTIAL.CREATE_CREDENTIAL (
   credential_name   IN  VARCHAR2,
   username          IN  VARCHAR2,
   password          IN  VARCHAR2,
   database_role     IN  VARCHAR2 DEFAULT NULL
   windows_domain    IN  VARCHAR2 DEFAULT NULL,
   comments          IN  VARCHAR2 DEFAULT NULL,
   enabled           IN  BOOLEAN DEFAULT TRUE);
```

Parameters

Table 51-2 ***CREATE_CREDENTIAL Procedure Parameters***

Parameter	Description
credential_name	Name of the credential. It can optionally be prefixed with a schema. This cannot be set to NULL. It is converted to upper case unless enclosed in double quotes.
username	User name to login to the operating system or remote database to run a job if this credential is chosen. This cannot be set to NULL.
password	Password to login to the remote operating system to run a job if this credential is chosen. It is case sensitive.
database_role	Whether a database job using this credential should attempt to log in with administrative privileges. Values: SYSDBA, SYSDG, SYSADMIN or SYSBACKUP.
windows_domain	For a Windows remote executable target, this is the domain that the specified user belongs to. The domain will be converted to uppercase automatically.
comments	A text string that can be used to describe the credential to the user. The Scheduler does not use this field.
enabled	Determines whether the credential is enabled or not

Usage Notes

- Credentials reside in a particular schema and can be created by any user with the CREATE CREDENTIAL or CREATE ANY CREDENTIAL system privilege. To create a credential in a schema other than your own, you must have the CREATE CREDENTIAL or CREATE ANY CREDENTIAL privilege.
- The user name is case sensitive. It cannot contain double quotes or spaces.
- Attempting to create a credential with an existing credential name returns an error. To alter an existing credential, users must drop the existing credential first using the DROP_CREDENTIAL Procedure.
- Attempting to drop an existing credential, which is already referenced by alias libraries, returns an error. To drop an existing credential without any checking, users must set the force parameter of DROP_CREDENTIAL Procedure to TRUE.

- You may also alter a credential, by means of the UPDATE_CREDENTIAL Procedure.

Examples

Create a Basic Credential

```
CONN scott
Enter password: password

BEGIN
-- Basic credential.
   DBMS_CREDENTIAL.CREATE_CREDENTIAL(
      credential_name  =>  'JAMES_SMITH',
      username         =>  'james_smith',
      password         =>  'password');
END
```

Create a Windows Credential

```
CONN scott
Enter password: password

-- Credential including Windows domain
BEGIN
  DBMS_CREDENTIAL.CREATE_CREDENTIAL(
     credential_name   =>  'JAMES_SMITH_WIN_CREDENTIAL',
     username          =>  'james_smith',
     password          =>  'password',
     windows_domain    =>  'localdomain');
END
```

Display Information about Credentials

Information about credentials is displayed using the `[DBA|ALL|USER] _CREDENTIALS` views.

```
COLUMN credential_name FORMAT A25
COLUMN username FORMAT A20
COLUMN windows_domain FORMAT A20
SELECT credential_name,
       username,
       windows_domain
FROM   user_credentials
ORDER BY credential_name;

CREDENTIAL_NAME           USERNAME             WINDOWS_DOMAIN
------------------------- -------------------- --------------------
JAMES_SMITH_CREDENTIAL       james_smith
JAMES_SMITH_WIN_CREDENTIAL   james_smith             LOCALDOMAIN
2 rows selected.

SQL>
```

51.4.2 DISABLE_CREDENTIAL Procedure

This procedure disables an Oracle credential.

Syntax

```
DBMS_CREDENTIAL.DISABLE_CREDENTIAL (
   credential_name   IN  VARCHAR2,
   force             IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 51-3 *DISABLE_CREDENTIAL Procedure Parameters*

Parameter	Description
credential_name	Name of the credential. It can optionally be prefixed with a schema. This cannot be set to NULL. It is converted to upper case unless enclosed in double quotes.
force	If FALSE, the credential is not disabled provided it has no dependency on any existing scheduler job or PL/SQL library. An error is returned if the dependency is observed.If TRUE, the credential is disabled whether or not there is any scheduler job or PL/SQL library referencing it.

Usage Notes

- Credentials reside in a particular schema and can be disabled by any user with the CREATE CREDENTIAL or CREATE ANY CREDENTIAL system privilege. To disable a credential in a schema other than your own, you must have the CREATE ANY CREDENTIAL privilege.
- A credential for an OS user can be viewed as an entry point into an operating system as a particular user. Allowing a credential to be disabled lets an administrator (or credential owner) to quickly, easily and reversibly disallow all logins from the database to the OS as a particular user of external jobs, database jobs, file transfers, external procedures, and file watching. To enable an existing disabled credential, you need to use the ENABLE_CREDENTIAL Procedure.
- A library can become invalid if the properties of the credential – windows domain, username, password, its enable/disable bit – are changed.

51.4.3 DROP_CREDENTIAL Procedure

This procedure drops an Oracle credential.

Syntax

```
DBMS_CREDENTIAL.DROP_CREDENTIAL (
   credential_name   IN  VARCHAR2,
   force             IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 51-4 *DROP_CREDENTIAL Procedure Parameters*

Parameter	Description
credential_name	Name of the credential. It can optionally be prefixed with a schema. This cannot be set to NULL.
force	If set to FALSE, the credential must not be referenced by any EXTPROC alias library or an error is raised. If set to TRUE, the credential is dropped whether or not there are extproc alias libraries referencing it. EXTPROC alias libraries that reference the dropped credential become invalid.

Usage Notes

Only the owner of a credential or a user with the `CREATE ANY CREDENTIAL` system privilege may drop the credential.

Examples

```
EXEC DBMS_CREDENTIAL.DROP_CREDENTIAL('JAMES_SMITH_CREDENTIAL', FALSE);
EXEC DBMS_CREDENTIAL.DROP_CREDENTIAL('JAMES_SMITH_WIN_CREDENTIAL', FALSE);
```

51.4.4 ENABLE_CREDENTIAL Procedure

This procedure enables an Oracle credential.

Syntax

```
DBMS_CREDENTIAL.ENABLE_CREDENTIAL (
   credential_name   IN  VARCHAR2);
```

Parameters

Table 51-5 ***ENABLE_CREDENTIAL Procedure Parameters***

Parameter	Description
`credential_name`	Name of the credential. It can optionally be prefixed with a schema. This cannot be set to `NULL`. It is converted to upper case unless enclosed in double quotes.

Usage Notes

- Credentials reside in a particular schema and can be disabled by any user with the `CREATE CREDENTIAL OR CREATE ANY CREDENTIAL` system privilege. To disable a credential in a schema other than your own, you must have the `CREATE CREDENTIAL OR CREATE ANY CREDENTIAL` privilege.
- A credential for an OS user can be viewed as an entry point into an operating system as a particular user. Allowing a credential to be disabled would allow an administrator (or credential owner) to quickly, easily and reversibly disallow all logins from the database to the OS as a particular user (external jobs, file transfers, external procedures, file watching). To disable an existing credential, you need to use the DISABLE_CREDENTIAL Procedure.
- A library can become invalid if the properties of the credential – windows domain, username, password, its enable/disable bit – are changed.

51.4.5 UPDATE_CREDENTIAL Procedure

This procedure updates an existing Oracle credential.

Syntax

```
DBMS_CREDENTIAL.UPDATE_CREDENTIAL (
   credential_name   IN   VARCHAR2,
   attribute         IN   VARCHAR2,
   value             IN   VARCHAR2);
```

Parameters

Table 51-6 *UPDATE_CREDENTIAL Procedure Parameters*

Parameter	Description
credential_name	Name of the credential. It can optionally be prefixed with a schema. This cannot be set to NULL. It is converted to upper case unless enclosed in double quotation marks.
attribute	Name of attribute to update: USERNAME, PASSWORD, WINDOWS_DOMAIN, DATABASE_ROLE or COMMENTS
value	New value for the selected attribute

Usage Notes

- Credentials reside in a particular schema and can be created by any user with the CREATE CREDENTIAL or CREATE ANY CREDENTIAL system privilege. To create a credential in a schema other than your own, you must have the CREATE ANY CREDENTIAL privilege.
- The user name is case sensitive. It cannot contain double quotes or spaces.
- EXTPROC alias libraries that reference the updated credential will become invalid. A library becomes invalid if the properties of the credential – windows domain, username, password, its enable/disable bit – are changed.

Examples

Update a Basic Credential

```
CONN scott
Enter password: password

BEGIN
-- Basic credential.
   DBMS_CREDENTIAL.UPDATE_CREDENTIAL (
      credential_name  =>  'JAMES_SMITH_CREDENTIAL',
      attribute        =>  'password',
      value            =>  'password2');

   DBMS_CREDENTIAL.UPDATE_CREDENTIAL (
      credential_name  =>  'JAMES_SMITH_CREDENTIAL',
      attribute        =>  'username',
      value            => 'james_smith');
END;
```

Update a Windows Credential

```
CONN scott
Enter password: password

-- Credential including Windows domain
BEGIN
  DBMS_CREDENTIAL.UPDATE_CREDENTIAL(
     credential_name   =>  'JAMES_SMITH_WIN_CREDENTIAL',
     username          =>  'james_smith',
     password          =>  'password',
     windows_domain    =>  'localdomain');
END
```

Display Information about Credentials

Information about credentials is displayed using the `[DBA|ALL|USER] _CREDENTIALS` views.

```
COLUMN credential_name FORMAT A25
COLUMN username FORMAT A20
COLUMN windows_domain FORMAT A20
SELECT credential_name,
       username,
       windows_domain
FROM   all_credentials
ORDER BY credential_name;

CREDENTIAL_NAME           USERNAME             WINDOWS_DOMAIN
------------------------- -------------------- --------------------
JAMES_SMITH_CREDENTIAL       james_smith
JAMES_SMITH_WIN_CREDENTIAL   james_smith             LOCALDOMAIN
2 rows selected.

SQL>
```

52
DBMS_CRYPTO

`DBMS_CRYPTO` provides an interface to encrypt and decrypt stored data, and can be used in conjunction with PL/SQL programs running network communications. It provides support for several industry-standard encryption and hashing algorithms, including the Advanced Encryption Standard (AES) encryption algorithm. AES has been approved by the National Institute of Standards and Technology (NIST) to replace the Data Encryption Standard (DES).

Note:

- Starting from Oracle Database 23ai Release, the use of the `MD4` algorithm is desupported.
- Starting from Oracle Database 21c Release, the use of the `MD5` and `MD4` algorithms are deprecated.
- Starting from Oracle Database 21c Release, the use of the `SHA-1` algorithm for `SQLNET` and `DBMS_CRYPTO` is deprecated.

Oracle recommends that you do not use deprecated values in new applications. Support for deprecated features is for backward compatibility only.

This chapter contains the following topics:

- DBMS_CRYPTO Overview
- DBMS_CRYPTO Security Model
- DBMS_CRYPTO Constants
- DBMS_CRYPTO Datatypes
- DBMS_CRYPTO Algorithms
- DBMS_CRYPTO Restrictions
- DBMS_CRYPTO Exceptions
- DBMS_CRYPTO Operational Notes
- DBMS_CRYPTO Examples
- Summary of DBMS_CRYPTO Subprograms

See Also:

- *Oracle Database Security Guide* for more information about using this package and about encrypting data in general.

52.1 DBMS_CRYPTO Overview

DBMS_CRYPTO contains basic cryptographic functions and procedures. To use this package correctly and securely, a general level of security expertise is assumed.

The DBMS_CRYPTO package enables encryption and decryption for common Oracle datatypes, including RAW and large objects (LOBs), such as images and sound. Specifically, it supports BLOBs and CLOBs. In addition, it provides Globalization Support for encrypting data across different database character sets. It also supports asymmetric key functions.

The following cryptographic algorithms are supported:

- Data Encryption Standard (DES), Triple DES (3DES, 2-key and 3-key) (deprecated)
- Advanced Encryption Standard (AES)
- MD5 (deprecated), SHA-1 (deprecated), and SHA-2 and SHA-3 cryptographic hashes
- MD5, SHA-1, and SHA-2 and SHA-3 Message Authentication Code (MAC), SHA-3 KMACXOF
- SM2, SM3, SM4 (SM stands for ShangMi)

Block cipher modifiers are also provided with DBMS_CRYPTO. You can choose from several padding options, including PKCS (Public Key Cryptographic Standard) #5, and from six block cipher chaining modes, including Cipher Block Chaining (CBC).

Table 52-1 summarizes the DBMS_CRYPTO package features.

Table 52-1 DBMS_CRYPTO Features

Package Feature	DBMS_CRYPTO
Cryptographic algorithms	DES, 3DES, AES, 3DES_2KEY, SM4
PKENCRYPT and PKDECRYPT algorithm	PKENCRYPT_RSA_PKCS1_OAEP_SHA2, SM4
Padding forms	PAD_PKCS5 (PKCS5 padding), PAD_ZERO (zeroes padding) PAD_NONE (no padding), PAD_ORCL (Oracle padding)
Block cipher chaining modes	CBC, CFB, ECB, OFB, GCM, CCM, XTS
Cryptographic hash algorithms	MD5, SHA-1, SHA-2 (SHA-256, SHA-384, SHA-512), SM3, (MD5 and SHA-1 are deprecated; MD4 is desupported) HASH_SHAKE128, HASH_SHAKE256
Keyed hash (MAC) algorithms	HMAC_MD5, HMAC_SH1, HMAC_SH256, HMAC_SH384, HMAC_SH512, HMAC_SHA3_224, HMAC_SHA3_256, HMAC_SHA3_384, HMAC_SHA3_512
Keccak MAC algorithms	KMACXOF_128, KMACXOF_256
Cryptographic pseudo-random number generator	RAW, NUMBER, BINARY_INTEGER
Database types	RAW, CLOB, BLOB

The DBMS_CRYPTO package replaces DBMS_OBFUSCATION_TOOLKIT, providing greater ease of use and support for a range of algorithms to accommodate new and existing systems. Specifically, 3DES_2KEY are provided for backward compatibility. It is not recommended that you use these algorithms because they do not provide the same level of security as provided by 3DES, AES, MD5, SHA-1, or SHA-2.

52.2 DBMS_CRYPTO Security Model

Oracle Database installs this package in the `SYS` schema. You can then grant package access to existing users and roles as needed.

52.3 DBMS_CRYPTO Constants

The `DBMS_CRYPTO` package uses the constants listed and described in this topic.

Table 52-2 DBMS_CRYPTO Constants

Name	Type	Value	Description
`AES_CCM_NONE`	`PLS_INTEGER`	`ENCRYPT_AES + CHAIN_CCM + PAD_NONE`	If you are using `DBMS_CRYPTO` to encrypt a plaintext with AES algorithm in CCM mode, then set the `typ` argument of `DBMS_CRYPTO.ENCRYPT` to the value `AES_CCM_NONE` to ensure that padding is disabled, because CCM mode cannot be used together with padding. Provide one extra input AAD (Additional Authenticated Data, optional) and one extra output TAG (non-optional). For example: `ciphertext := dbms_crypto.encrypt (src => plaintext, typ => DBMS_CRYPTO.AES_CCM_NONE , key => key_information, iv => iv_information, aad => aad_information, tag); plaintext := dbms_crypto.decrypt (src => plaintext, typ => DBMS_CRYPTO.AES_CCM_NONE , key => key_information, iv => iv_information, aad => aad_information, tag=> tag_information);`

Table 52-2 (Cont.) DBMS_CRYPTO Constants

Name	Type	Value	Description
AES_GCM_NONE	PLS_INTEGER	ENCRYPT_AES + CHAIN_GCM + PAD_NONE	If you are using DBMS_CRYPTO to encrypt a plaintext with AES algorithm in GCM mode, then set the typ argument of DBMS_CRYPTO.ENCRYPT to the value AES_GCM_NONE to ensure that padding is disabled, because GCM mode cannot be used together with padding. Provide one extra input AAD (Additional Authenticated Data, optional) and one extra output TAG (non-optional). For example: `ciphertext := dbms_crypto.encrypt (src => plaintext, typ => DBMS_CRYPTO.AES_GCM_NONE , key => key_information, iv => iv_information, aad => aad_information, tag); plaintext := dbms_crypto.decrypt (src => plaintext, typ => DBMS_CRYPTO.AES_GCM_NONE , key => key_information, iv => iv_information, aad => aad_information, tag=> tag_information);`
AES_XTS_NONE	PLS_INTEGER	ENCRYPT_AES + CHAIN_XTS + PAD_NONE	Constant for AES encryption algorithm in XTS mode. You can use this constant with DBMS_CRYPTO.ENCRYPT. To use AES with XTS mode, the input must be at least one full block (16 bytes).
DES3_CBC_NONE	PLS_INTEGER	ENCRYPT_3DES + CHAIN_CBC + PAD_NONE	If you are using DBMS_CRYPTO to decipher a triple-DES ciphertext that you created in the past using the desupported DBMS_OBFUSCATION_TOOLKIT, then set the typ argument of DBMS_CRYPTO.decrypt to the value DBMS_CRYPTO.DES3_CBC_NONE to ensure that the PKCS#5 padding is disabled. For example: `plaintext := DBMS_CRYPTO.decrypt (src => ciphertext_from_legacy_DES3Encrypt ,typ => DBMS_CRYPTO.DES3_CBC_NONE ,key => key_information ,iv => hextoraw(DBMS_CRYPTO.LEGACY_DEFAULT_IV) );`
AES_CBC_PKCS5	PLS_INTEGER	ENCRYPT_AES + CHAIN_CBC + PAD_PKCS5	Constant for AES encryption algorithm in CBC mode and PKCS5 padding. You can use this constant with DBMS_CRYPTO.ENCRYPT.

Table 52-2 (Cont.) DBMS_CRYPTO Constants

Name	Type	Value	Description
LEGACY_DEFAULT_IV	VARCHAR2(16)	0123456789ABCDEF	If you did not provide the IV argument when creating a triple-DES ciphertext using the desupported DBMS_OBFUSCATION_TOOLKIT, then provide IV as hextoraw(DBMS_CRYPTO.LEGACY_DEFAULT_IV) when invoking DBMS_CRYPTO to decrypt the triple-DES ciphertext. For example: `plaintext := DBMS_CRYPTO.decrypt (src => ciphertext_from_legacy_DES3Encrypt ,typ => DBMS_CRYPTO.DES3_CBC_NONE ,key => key_information ,iv => hextoraw(dbms_crypto.LEGACY_DEFAULT_IV) );`
SM4_CFB_NONE	PLS_INTEGER	ENCRYPT_SM4 + CHAIN_CFB + PAD_NONE	Cipher suite for the SM4 encryption algorithm in Ciphertext Feedback (CFB) mode. You can only use CFB mode with no padding mode PAD_NONE.
SM4_OFB_NONE	PLS_INTEGER	ENCRYPT_SM4 + CHAIN_OFB + PAD_NONE	Cipher suite for the SM4 encryption algorithm in Output Feedback (OFB) mode. You can only use OFB mode with no padding mode PAD_NONE.

52.4 DBMS_CRYPTO Datatypes

Parameters for the DBMS_CRYPTO subprograms use these datatypes.

Table 52-3 DBMS_CRYPTO Datatypes

Type	Description
BLOB	A source or destination binary LOB
CLOB	A source or destination character LOB (excluding NCLOB)
PLS_INTEGER	Specifies a cryptographic algorithm type (used with BLOB, CLOB, and RAW datatypes)
RAW	A source or destination RAW buffer

52.5 DBMS_CRYPTO Algorithms

The DBMS_CRYPTO package contains predefined cryptographic algorithms, modifiers, and cipher suites.

These are shown in the following tables.

Table 52-4 DBMS_CRYPTO Cryptographic Hash Functions

Name	Description
HASH_MD5	Produces a 128-bit hash, but is more complex than MD4 (which is desupported). Note that MD5 is deprecated.
HASH_SH1	Secure Hash Algorithm (SHA-1) (deprecated). Produces a 160-bit hash.
HASH_SH256	SHA-2, produces a 256-bit hash.
HASH_SH384	SHA-2, produces a 384-bit hash.
HASH_SH512	SHA-2, produces a 512-bit hash.
HASH_SHA3_224	SHA-3, produces a 224-bit hash.
HASH_SHA3_256	SHA-3, produces a 256-bit hash.
HASH_SHA3_384	SHA-3, produces a 384-bit hash.
HASH_SHA3_512	SHA-3, produces a 512-bit hash.
HASH_SHAKE128	Produces variable length hash with 128-bit security level; used for the DBMS_CRYPTO.HASH_LEN only.
HASH_SHAKE256	Produces variable length hash with 256-bit security level; used for the DBMS_CRYPTO.HASH_LEN only.
HASH_SM3	Produces a 256-bit hash

Table 52-5 DBMS_CRYPTO MAC (Message Authentication Code) Functions

Name	Description
HMAC_MD5 (deprecated)	Same as MD5 (deprecated) hash function, except it requires a secret key to verify the hash value.
HMAC_SH1 (deprecated)	Same as SHA hash function, except it requires a secret key to verify the hash value. Complies with IETF RFC 2104 standard.
HMAC_SH256	Same as SHA-2 256-bit hash function, except it requires a secret key to verify the hash value.
HMAC_SH384	Same as SHA-2 384-bit hash function, except it requires a secret key to verify the hash value.
HMAC_SH512	Same as SHA-2 512-bit hash function, except it requires a secret key to verify the hash value.
HMAC_SHA3_224	Same as SHA-3 224-bit hash function, except it requires a secret key to verify the hash value.
HMAC_SHA3_256	Same as SHA-3 256-bit hash function, except it requires a secret key to verify the hash value.
HMAC_SHA3_384	Same as SHA-3 384-bit hash function, except it requires a secret key to verify the hash value.
HMAC_SHA3_512	Same as SHA-3 512-bit hash function, except it requires a secret key to verify the hash value.

Table 52-6 DBMS_CRYPTO KMACXOF Functions

Name	Description
KMACXOF_128	Variable length KMAC with 128-bit security level. KMAC stands for KECCAK Message Authentication Code.

Table 52-6 (Cont.) DBMS_CRYPTO KMACXOF Functions

Name	Description
KMACXOF_256	Variable length KMAC with 256-bit security level.

Table 52-7 DBMS_CRYPTO Encryption Algorithms

Name	Description
ENCRYPT_DES (deprecated)	Data Encryption Standard. Block cipher. Uses key length of 56 bits.
ENCRYPT_3DES_2KEY (deprecated)	Data Encryption Standard. Block cipher. Operates on a block 3 times with 2 keys. Effective key length of 112 bits.
ENCRYPT_3DES (deprecated)	Data Encryption Standard. Block cipher. Operates on a block 3 times.
ENCRYPT_AES128	Advanced Encryption Standard. Block cipher. Uses 128-bit key size.
ENCRYPT_AES192	Advanced Encryption Standard. Block cipher. Uses 192-bit key size.
ENCRYPT_AES256	Advanced Encryption Standard. Block cipher. Uses 256-bit key size.
PKENCRYPT_RSA_PKCS1_OAEP (deprecated)	Public key encryption algorithm; only allowed for decryption.
PKENCRYPT_RSA_PKCS1_OAEP_SHA2	Public key encryption algorithm.
PKENCRYPT_SM2	Provides Chinese SM2 signature and encryption algorithm support
ENCRYPT_SM4	Block cipher used in the Chinese National Standard for Wireless LAN WAPI (WLAN Authentication and Privacy Infrastructure) and also used with Transport Layer Security

Table 52-8 DBMS_CRYPTO Block Cipher Suites

Name	Description
AES_CBC_PKCS5	ENCRYPT_AES + CHAIN_CBC + PAD_PKCS5
AES_CCM_NONE	ENCRYPT_AES + CHAIN_GCM + PAD_NONE
AES_GCM_NONE	ENCRYPT_AES + CHAIN_CCM + PAD_NONE
AES_XTS_NONE	ENCRYPT_AES + CHAIN_XTS + PAD_NONE
DES_CBC_PKCS5	ENCRYPT_DES[1] + CHAIN_CBC[2]+ PAD_PKCS5[3]
DES3_CBC_PKCS5	ENCRYPT_3DES[1] + CHAIN_CBC[2] + PAD_PKCS5[3]

1 See Table 52-7
2 See Table 52-9
3 See Table 52-10

Table 52-9 DBMS_CRYPTO Block Cipher Chaining Modifiers

Name	Description
CHAIN_CBC	Cipher Block Chaining. Plaintext is XORed with the previous ciphertext block before it is encrypted.

Table 52-9 (Cont.) DBMS_CRYPTO Block Cipher Chaining Modifiers

Name	Description
CHAIN_CCM	The Counter with CBC-MAC (CCM) mode. It is a generic authenticated encryption block cipher mode. It can be used with AES (AES_CCM_NONE). The encryption of AES_CCM_NONE takes one extra input AAD (Additional Authenticated Data, optional) and one extra output TAG (non-optional). The decryption of AES_CCM_NONE takes two extra inputs AAD and TAG.
CHAIN_CFB	Cipher-Feedback. Enables encrypting units of data smaller than the block size.
CHAIN_ECB	Electronic Codebook. Encrypts each plaintext block independently.
CHAIN_GCM	The Galois/Counter Mode. It is a generic authenticated encryption block cipher mode. It can be used with AES (AES_GCM_NONE). The encryption of AES_GCM_NONE takes one extra input AAD (Additional Authenticated Data, optional) and one extra output TAG (non-optional). The decryption of AES_GCM_NONE takes two extra inputs AAD and TAG.
CHAIN_OFB	Output-Feedback. Enables running a block cipher as a synchronous stream cipher. Similar to CFB, except that *n* bits of the previous output block are moved into the right-most positions of the data queue waiting to be encrypted.
CHAIN_XTS	XEX-based modified-codebook mode with ciphertext stealing (XTS), a chain mode for AES. It can be used with the AES algorithm in DBMS_CRYPTO and PAD_NONE

Table 52-10 DBMS_CRYPTO Block Cipher Padding Modifiers

Name	Description
PAD_NONE	Provides option to specify no padding. Caller must ensure that blocksize is correct, else the package returns an error.
PAD_ORCL	Provides padding that adds a single byte containing the pad length to the end of the crypto vector.
PAD_ZERO	Provides padding consisting of zeroes

52.6 DBMS_CRYPTO Restrictions

The VARCHAR2 datatype is not directly supported by DBMS_CRYPTO. Before you can perform cryptographic operations on data of the type VARCHAR2, you must convert it to the uniform database character set AL32UTF8, and then convert it to the RAW datatype. After performing these conversions, you can then encrypt it with the DBMS_CRYPTO package.

See Also:

DBMS_CRYPTO Operational Notes for information about the conversion rules for converting datatypes.

DBMS_CRYPTO Examples for examples of using the call to UTL_I18N.STRING_TO_RAW to allow VARCHAR2 to be encrypted.

52.7 DBMS_CRYPTO Exceptions

The following table lists exceptions that have been defined for DBMS_CRYPTO.

Table 52-11 DBMS_CRYPTO Exceptions

Exception	Code	Description
CipherSuiteInvalid	28827	The specified cipher suite is not defined.
CipherSuiteNull	28829	No value has been specified for the cipher suite to be used.
KeyNull	28239	The encryption key has not been specified or contains a NULL value.
KeyBadSize	28234	DES keys: Specified key size is too short. DES keys must be at least 8 bytes (64 bits). AES keys: Specified key size is not supported. AES keys must be 128, 192, or 256 bits in length.
DoubleEncryption	28233	Source data was previously encrypted.

52.8 DBMS_CRYPTO Operational Notes

This section describes several DBMS_CRYPTO operational notes.

- When to Use Encrypt and Decrypt Procedures or Functions
- When to Use Hash or Message Authentication Code (MAC) Functions
- About Generating and Storing Encryption Keys
- Conversion Rules

When to Use Encrypt and Decrypt Procedures or Functions

This package includes both ENCRYPT and DECRYPT procedures and functions. The procedures are used to encrypt or decrypt LOB datatypes (overloaded for CLOB and BLOB datatypes). In contrast, the ENCRYPT and DECRYPT functions are used to encrypt and decrypt RAW datatypes. Data of type VARCHAR2 must be converted to RAW before you can use DBMS_CRYPTO functions to encrypt it.

When to Use Hash or Message Authentication Code (MAC) Functions

This package includes two different types of one-way hash functions: the HASH function and the MAC function. Hash functions operate on an arbitrary-length input message, and return a fixed-length hash value. One-way hash functions work in one direction only. It is easy to compute a hash value from an input message, but it is extremely difficult to generate an input message that hashes to a particular value. Note that hash values should be at least 256 bits in length to be considered secure.

You can use hash values to verify whether data has been altered. For example, before storing data, the user runs DBMS_CRYPTO.HASH against the stored data to create a hash value. On returning the stored data, the user can again run the hash function against it, using the same algorithm. If the second hash value is identical to the first one, then the data has not been altered. Hash values are similar to "file fingerprints" and are used to ensure data integrity.

The HASH function included with DBMS_CRYPTO, is a one-way hash function that you can use to generate a hash value from either RAW or LOB data. The MAC function is also a one-way hash function, but with the addition of a secret key. It works the same way as the DBMS_CRYPTO.HASH function, except only someone with the key can verify the hash value.

MACs can be used to authenticate files between users. They can also be used by a single user to determine if her files have been altered, perhaps by a virus. A user could compute the MAC of his files and store that value in a table. If the user did not use a MAC function, then the virus could compute the new hash value after infection and replace the table entry. A virus cannot do that with a MAC because the virus does not know the key.

About Generating and Storing Encryption Keys

The DBMS_CRYPTO package can generate random material for encryption keys, but it does not provide a mechanism for maintaining them. Application developers must take care to ensure that the encryption keys used with this package are securely generated and stored. Also note that the encryption and decryption operations performed by DBMS_CRYPTO occur on the server, not on the client. Consequently, if the key is sent over the connection between the client and the server, the connection must be protected by using network encryption. Otherwise, the key is vulnerable to capture over the wire.

Although DBMS_CRYPTO cannot generate keys on its own, it does provide tools you can use to aid in key generation. For example, you can use the RANDOMBYTES function to generate random material for keys.

When generating encryption keys for DES, it is important to remember that some numbers are considered weak and semiweak keys. Keys are considered weak or semiweak when the pattern of the algorithm combines with the pattern of the initial key value to produce ciphertext that is more susceptible to cryptanalysis. To avoid this, filter out the known weak DES keys. Lists of the known weak and semiweak DES keys are available on several public Internet sites.

See Also:

- *Oracle Database Security Guide* for information about configuring network encryption and SSL
- RANDOMBYTES Function

Conversion Rules

- To convert VARCHAR2 to RAW, use the UTL_I18N.STRING_TO_RAW function to perform the following steps:
  1. Convert VARCHAR2 in the current database character set to VARCHAR2 in the AL32UTF8 database character.
  2. Convert VARCHAR2 in the AL32UTF8 database character set to RAW.

  Syntax example:

  ```
  UTL_I18N.STRING_TO_RAW (string, 'AL32UTF8');
  ```
- To convert RAW to VARCHAR2, use the UTL_I18N.RAW_TO_CHAR function to perform the following steps:
  1. Convert RAW to VARCHAR2 in the AL32UTF8 database character set.

2. Convert VARCHAR2 in the AL32UTF8 database character set to VARCHAR2 in the database character set you wish to use.

   Syntax example:

   ```
   UTL_I18N.RAW_TO_CHAR (data, 'AL32UTF8');
   ```

> **See Also:**
>
> UTL_I18N for information about using the UTL_I18N PL/SQL package.

- If you want to store encrypted data of the RAW datatype in a VARCHAR2 database column, then use RAWTOHEX or UTL_ENCODE.BASE64_ENCODE to make it suitable for VARCHAR2 storage. These functions expand data size by 2 and 4/3, respectively.

52.9 DBMS_CRYPTO Examples

The examples in this section demonstrate different types of PL/SQL coding using DBMS_CRYPTO functions.

Example 1: AES Encryption with Cipher Block Chaining and PKCS#5 Compliant Padding

This example shows PL/SQL block encrypting and decrypting pre-defined input_string using 256-bit AES algorithm with Cipher Block Chaining and PKCS#5 compliant padding.

```
DECLARE
   input_string       VARCHAR2 (200) :=  'Secret Message';
   output_string      VARCHAR2 (200);
   encrypted_raw      RAW (2000);             -- stores encrypted binary text
   decrypted_raw      RAW (2000);             -- stores decrypted binary text
   num_key_bytes      NUMBER := 256/8;        -- key length 256 bits (32 bytes)
   key_bytes_raw      RAW (32);               -- stores 256-bit encryption key
   encryption_type    PLS_INTEGER :=          -- total encryption type

                        DBMS_CRYPTO.ENCRYPT_AES256
                      + DBMS_CRYPTO.CHAIN_CBC
                      + DBMS_CRYPTO.PAD_PKCS5;

   iv_raw             RAW (16);

BEGIN
   DBMS_OUTPUT.PUT_LINE ( 'Original string: ' || input_string);
   key_bytes_raw := DBMS_CRYPTO.RANDOMBYTES (num_key_bytes);
   iv_raw        := DBMS_CRYPTO.RANDOMBYTES (16);
   encrypted_raw := DBMS_CRYPTO.ENCRYPT
      (
         src => UTL_I18N.STRING_TO_RAW (input_string,  'AL32UTF8'),
         typ => encryption_type,
         key => key_bytes_raw,
         iv  => iv_raw
      );

    -- The encrypted value "encrypted_raw" can be used here

    decrypted_raw := DBMS_CRYPTO.DECRYPT
      (
         src => encrypted_raw,
         typ => encryption_type,
```

```
        key => key_bytes_raw,
        iv  => iv_raw
      );

   output_string := UTL_I18N.RAW_TO_CHAR (decrypted_raw, 'AL32UTF8');

   DBMS_OUTPUT.PUT_LINE ('Decrypted string: ' || output_string);
END;
```

Example 2: PKENCRYPTION and PKDECRYPTION Functions

```
DECLARE
  ip_str     VARCHAR (200) := 'Secret Message';
  op_str     VARCHAR (200);
  -- Use OpenSSL to generate the private and public key (2048 bit RSA key)
  -- openssl genrsa -out private.pem 2048
  -- openssl rsa -in private.pem -outform PEM -pubout -out public.pem
  pubkey    VARCHAR (2000) := 'public_key';
  prvkey    VARCHAR (2000) := 'private_key';
  enc_raw    RAW (2000);
  dec_raw    RAW (2000);
  eType      PLS_INTEGER := DBMS_CRYPTO.PKENCRYPT_RSA_PKCS1_OAEP_SHA2;
  kType      PLS_INTEGER := DBMS_CRYPTO.KEY_TYPE_RSA;

BEGIN

  DBMS_OUTPUT.PUT_LINE('-------------------------------------------------');
  DBMS_OUTPUT.PUT_LINE('Original String := ' || ip_str);
  DBMS_OUTPUT.PUT_LINE('-------------------------------------------------');
  enc_raw:= DBMS_CRYPTO.PKENCRYPT
  (
    src        => UTL_I18N.STRING_TO_RAW(ip_str,'AL32UTF8'),
    pub_key    => UTL_I18N.STRING_TO_RAW( pubkey, 'AL32UTF8'),
    pubkey_alg => kType,
    enc_alg    => eType
  );

  dec_raw := DBMS_CRYPTO.PKDECRYPT
  (
    src        => enc_raw,
    prv_key    => UTL_I18N.STRING_TO_RAW( prvkey, 'AL32UTF8'),
    pubkey_alg => kType,
    enc_alg    => eType
  );
  op_str := UTL_I18N.RAW_TO_CHAR(dec_raw,'AL32UTF8');
  dbms_output.put_line('-------------------------------------------------');
  dbms_output.put_line('Decrypted String := ' || op_str);
  dbms_output.put_line('-------------------------------------------------');
end;
/
```

Example 3: SIGN and VERIFY Functions

```
DECLARE
  ip_str     VARCHAR2 (200) := 'Secret Message';
  -- Use OpenSSL to generate the private and public key (2048 bit RSA key)
  -- openssl genrsa -out private.pem 2048
  -- openssl rsa -in private.pem -outform PEM -pubout -out public.pem
  pubkey    VARCHAR (2000) := 'public_key';
  prvkey    VARCHAR (2000) := 'private_key';
  sign_raw   RAW (2000);
  returnval  BOOLEAN := false;
```

```
  sType      PLS_INTEGER := DBMS_CRYPTO.SIGN_SHA224_RSA;
  kType      PLS_INTEGER := DBMS_CRYPTO.KEY_TYPE_RSA;
BEGIN
  sign_raw := DBMS_CRYPTO.SIGN
  (
   src        => UTL_I18N.STRING_TO_RAW(ip_str,'AL32UTF8'),
   prv_key    => UTL_I18N.STRING_TO_RAW( prvkey, 'AL32UTF8'),
   pubkey_alg => kType,
   sign_alg   => sType
  );
  returnval := DBMS_CRYPTO.VERIFY
  (
   src        => UTL_I18N.STRING_TO_RAW( ip_str,'AL32UTF8'),
   sign       => sign_raw,
   pub_key    => UTL_I18N.STRING_TO_RAW( pubkey, 'AL32UTF8'),
   pubkey_alg => kType,
   sign_alg   => sType
  );
  DBMS_OUTPUT.PUT_LINE('-------------------------------------------------');
  IF returnval THEN
    DBMS_OUTPUT.PUT_LINE('True');
  ELSE
    DBMS_OUTPUT.PUT_LINE('False');
  END IF;
  DBMS_OUTPUT.PUT_LINE('-------------------------------------------------');
END;
/
```

Example 4: HASH_SHA3_256 and HASH_SHAKE256 Hash Functions

```
DECLARE
  ip_str VARCHAR2 (200) := 'Secret Message';
  output_string VARCHAR2 (2000);
  source_text_raw RAW (2000);
  hash_raw RAW (2000);
BEGIN
  source_text_raw := UTL_I18N.STRING_TO_RAW(ip_str,'AL32UTF8');
  hash_raw := dbms_crypto.hash(source_text_raw, dbms_crypto.hash_sha3_256);

  output_string := RAWTOHEX(hash_raw);

  DBMS_OUTPUT.PUT_LINE ('Hash string: ' || output_string);
  DBMS_OUTPUT.PUT_LINE (utl_raw.length(hash_raw);
END;
/


HASH_SHAKE256:

DECLARE
  ip_str VARCHAR2 (200) := 'Secret Message';
  output_string VARCHAR2 (2000);
  source_text_raw RAW (2000);
  hash_raw RAW (2000);
  len PLS_INTEGER := 30;
BEGIN
  source_text_raw := UTL_I18N.STRING_TO_RAW(ip_str,'AL32UTF8');
  hash_raw := dbms_crypto.hash_len(source_text_raw, dbms_crypto.hash_shake256, len);

  output_string := RAWTOHEX(hash_raw);

  DBMS_OUTPUT.PUT_LINE ('Hash string: ' || output_string);
```

```
END;
/
```

Example 5: HMAC_SHA3_256 and SIGN_SHA3_256_RSA Functions

```
DECLARE
  ip_str VARCHAR2 (200) := 'Secret Message';
  output_string VARCHAR2 (2000);
  source_text_raw RAW (2000);
  hmac_raw RAW (2000);
  testkey RAW (2000);
BEGIN
  source_text_raw := UTL_I18N.STRING_TO_RAW(ip_str,'AL32UTF8');
  -- Use your test key
  testkey := HEXTORAW('');
  hmac_raw := dbms_crypto.mac(source_text_raw, dbms_crypto.HMAC_SHA3_256, testkey);

  output_string := RAWTOHEX(hmac_raw);

  DBMS_OUTPUT.PUT_LINE ('Hmac string: ' || output_string);
END;
/

SIGN_SHA3_256_RSA:

DECLARE
  ip_str VARCHAR2 (200) := 'Secret Message';
  -- Use OpenSSL to generate the private and public key (2048 bit RSA key)
  -- openssl genrsa -out private.pem 2048
  -- openssl rsa -in private.pem -outform PEM -pubout -out public.pem
  pubkey VARCHAR (2000) := 'public_key';
  prvkey VARCHAR (2000) := 'private_key';
  sign_raw RAW (2000);
  returnval BOOLEAN := false;
  sType PLS_INTEGER := DBMS_CRYPTO.SIGN_SHA3_256_RSA;
  kType PLS_INTEGER := DBMS_CRYPTO.KEY_TYPE_RSA;
BEGIN
  sign_raw := DBMS_CRYPTO.SIGN
  (
   src => UTL_I18N.STRING_TO_RAW(ip_str,'AL32UTF8'),
   prv_key => UTL_I18N.STRING_TO_RAW( prvkey, 'AL32UTF8'),
   pubkey_alg => kType,
   sign_alg => sType
  );
  returnval := DBMS_CRYPTO.VERIFY
  (
   src => UTL_I18N.STRING_TO_RAW( ip_str,'AL32UTF8'),
   sign => sign_raw,
   pub_key => UTL_I18N.STRING_TO_RAW( pubkey, 'AL32UTF8'),
   pubkey_alg => kType,
   sign_alg => sType
  );
  DBMS_OUTPUT.PUT_LINE('-------------------------------------------------');
  IF returnval THEN
    DBMS_OUTPUT.PUT_LINE('True');
  ELSE
    DBMS_OUTPUT.PUT_LINE('False');
  END IF;
  DBMS_OUTPUT.PUT_LINE('-------------------------------------------------');
END;
/
```

52.10 Summary of DBMS_CRYPTO Subprograms

This table lists the `DBMS_CRYPTO` subprograms in alphabetical order and briefly describes them.

Table 52-12 ***DBMS_CRYPTO Package Subprograms***

Subprogram	Description
DECRYPT Function	Decrypts `RAW` data using a stream or block cipher with a user supplied key and optional IV (initialization vector)
DECRYPT Procedures	Decrypts `LOB` data using a stream or block cipher with a user supplied key and optional IV
ECDHDERIVE_SHAREDSECRET Function	Derives shared secret using private key of local application and public key from the remote application
ECDH_GENKEYPAIR Function	Generates an EC public/private key pair
ENCRYPT Function	Encrypts `RAW` data using a stream or block cipher with a user supplied key and optional IV
ENCRYPT Procedures	Encrypts `LOB` data using a stream or block cipher with a user supplied key and optional IV
HASH Function	Applies one of the supported cryptographic hash algorithms (MD5, SHA-1, or SHA-2) to data.
	Note: Starting from Oracle Database 20c Release, the use of MD5 and SHA-1 are deprecated. MD4 is desupported. Oracle recommends that you do not use deprecated values in new applications. Support for deprecated features is for backward compatibility only.
HASH_LEN Function	Similar to the `HASH` function, except that it includes an extra input length that uses the `PLS_INTEGER` type
MAC Function	Applies Message Authentication Code algorithms (MD5, SHA-1, or SHA-2) to data to provide keyed message protection.
	Note: Starting from Oracle Database 20c Release, the use of MD5 and SHA-1 are deprecated. Oracle recommends that you do not use deprecated values in new applications. Support for deprecated features is for backward compatibility only.
KMACXOF Function	Similar to the `MAC` function, except that is includes the `length` and `custStr` fields.

Table 52-12 (Cont.) *DBMS_CRYPTO Package Subprograms*

Subprogram	Description
PKDECRYPT Function	Decrypts `RAW` data using a private key assisted with key algorithm and encryption algorithm and returns decrypted data.
PKENCRYPT Function	Encrypts `RAW` data using a public key assisted with key algorithm and encryption algorithm and returns encrypted data
RANDOMBYTES Function	Returns a `RAW` value containing a cryptographically secure pseudo-random sequence of bytes, and can be used to generate random material for encryption keys
RANDOMINTEGER Function	Returns a random `BINARY_INTEGER`
RANDOMNUMBER Function	Returns a random 128-bit integer of the `NUMBER` datatype
SIGN Function	Signs `RAW` data using a private key assisted with key algorithm and sign algorithm, and returns a signature
VERIFY Function	Verifies `RAW` data using the signature, public key assisted with key algorithm, and sign algorithm. It returns `TRUE` if the signature was verified

52.10.1 DECRYPT Functions

These functions decrypt `RAW` data using a stream or block cipher with a user-supplied key and optional IV (initialization vector).

Syntax

```
DBMS_CRYPTO.DECRYPT(
   src IN RAW,
   typ IN PLS_INTEGER,
   key IN RAW,
   iv  IN RAW DEFAULT NULL)
 RETURN RAW;

DBMS_CRYPTO.DECRYPT (
   src IN RAW,
   typ IN PLS_INTEGER,
   key IN RAW,
   iv IN RAW DEFAULT NULL,
   aad IN RAW DEFAULT NULL,
   tag IN RAW)
 RETURN RAW;
```

Pragmas

```
pragma restrict_references(decrypt,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-13 DECRYPT Function Parameters

Parameter Name	Description
`src`	`RAW` data to be decrypted.

Table 52-13 (Cont.) DECRYPT Function Parameters

Parameter Name	Description
typ	Stream or block cipher type and modifiers to be used.
key	Key to be used for decryption.
iv	Optional initialization vector for block ciphers. Default is NULL.
aad	Additional authenticated data, which is any string that you pass to an Oracle Cloud key management service as part of the request.
tag	Authentication tag that is used for the authentication check.

Usage Notes

- To retrieve original plaintext data, DECRYPT must be called with the same cipher, modifiers, key, and IV that was used to encrypt the data originally.

> **See Also:**
>
> "Usage Notes" for the ENCRYPT function for additional information about the ciphers and modifiers available with this package.

- If VARCHAR2 data is converted to RAW before encryption, then it must be converted back to the appropriate database character set by using the UTL_I18N package.

> **See Also:**
>
> DBMS_CRYPTO Operational Notes for a discussion of the VARCHAR2 to RAW conversion rules

52.10.2 DECRYPT Procedures

These procedures decrypt LOB data using a stream or block cipher with a user supplied key and optional IV (initialization vector).

Syntax

```
DBMS_CRYPTO.DECRYPT(
   dst IN OUT NOCOPY BLOB,
   src IN            BLOB,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW          DEFAULT NULL);

DBMS_CRYPT.DECRYPT(
   dst IN OUT NOCOPY CLOB         CHARACTER SET ANY_CS,
   src IN            BLOB,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW          DEFAULT NULL);

DBMS_CRYPTO.DECRYPT (
```

```
   dst IN OUT NOCOPY BLOB,
   src IN            BLOB,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW DEFAULT NULL,
   aad IN            RAW DEFAULT NULL,
   tag IN            RAW);

DBMS_CRYPTO.DECRYPT (
  dst IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
  src IN            BLOB,
  typ IN            PLS_INTEGER,
  key IN            RAW,
  iv  IN            RAW DEFAULT NULL,
  aad IN            RAW DEFAULT NULL,
  tag IN            RAW);
```

Pragmas

```
pragma restrict_references(decrypt,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-14 DECRYPT Procedure Parameters

Parameter Name	Description
dst	LOB locator of output data. The value in the output LOB <dst> will be overwritten.
src	LOB locator of input data.
typ	Stream or block cipher type and modifiers to be used.
key	Key to be used for decryption.
iv	Optional initialization vector for block ciphers. Default is all zeroes.
aad	Additional authenticated data, which is any string that you pass to an Oracle Cloud key management service as part of the request.
tag	Authentication tag that is used for the authentication check.

52.10.3 ECDHDERIVE_SHAREDSECRET Function

This function derives shared secret using private key of local application and public key from the remote application.

Syntax

```
DBMS_CRYPTO.
DBMS_CRYPTO.ECDHDERIVE_SHAREDSECRET (
   curveid       IN  BINARY_INTEGER,
   pubkey        IN  RAW,
   privkey       IN  RAW,
   sharedsecret  OUT RAW);
```

Parameters

Table 52-15 HASH Function Parameters

Parameter Name	Description
curveid	Constants to denote the curve name that the algorithm uses. Example: SECP_256_R1
pubkey	Public key from the other side
privkey	Private key
sharedsecret	Shared secret generated from private key and public key from the other side

Usage Note

The supported curve id SECP_256_R1 denotes NIST Recommended Curve secp256r1.

52.10.4 ECDH_GENKEYPAIR Function

This function generates an EC public/private key pair.

Syntax

```
DBMS_CRYPTO.
DBMS_CRYPTO.ECDH_GENKEYPAIR (
   curveid       IN  BINARY_INTEGER,
   pubkey        OUT  RAW,
   privkey       OUT  RAW);
```

Parameters

Table 52-16 HASH Function Parameters

Parameter Name	Description
curveid	Constants to denote the curve name that the algorithm uses. Example: SECP_256_R1
pubkey	Public key
privkey	Private key

Usage Note

The supported curve id SECP_256_R1 denotes NIST Recommended Curve secp256r1.

52.10.5 ENCRYPT Functions

These functions encrypt RAW data using a stream or block cipher with a user supplied key and optional IV (initialization vector).

Syntax

```
DBMS_CRYPTO.ENCRYPT(
   src IN RAW,
   typ IN PLS_INTEGER,
   key IN RAW,
```

```
   iv  IN RAW DEFAULT NULL)
 RETURN RAW;

DBMS_CRYPTO.ENCRYPT (
   src IN RAW,
   typ IN PLS_INTEGER,
   key IN RAW,
   iv  IN RAW DEFAULT NULL,
   aad IN RAW DEFAULT NULL,
   tag OUT RAW)
 RETURN RAW;
```

Pragmas

```
pragma restrict_references(encrypt,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-17 ENCRYPT Function Parameters

Parameter Name	Description
src	RAW data to be encrypted.
typ	Stream or block cipher type and modifiers to be used.
key	Encryption key to be used for encrypting data.
iv	Optional initialization vector for block ciphers. Default is NULL.
aad	Additional authenticated data, which is any string that you pass to an Oracle Cloud key management service as part of the request.
tag	Authentication tag that is used for the authentication check.

Usage Notes

- Block ciphers may be modified with chaining and padding type modifiers. The chaining and padding type modifiers are added to the block cipher to produce a cipher suite. Cipher Block Chaining (CBC) is the most commonly used chaining type, and PKCS #5 is the recommended padding type. See Table 52-9 and Table 52-10 for block cipher chaining and padding modifier constants that have been defined for this package.
- To improve readability, you can define your own package-level constants to represent the cipher suites you use for encryption and decryption. For example, the following example defines a cipher suite that uses DES, cipher block chaining mode, and no padding:

  ```
  DES_CBC_NONE CONSTANT PLS_INTEGER := DBMS_CRYPTO.ENCRYPT_DES
                                     + DBMS_CRYPTO.CHAIN_CBC
                                     + DBMS_CRYPTO.PAD_NONE;
  ```

  See Table 52-8 for the block cipher suites already defined as constants for this package.
- To encrypt VARCHAR2 data, it should first be converted to the AL32UTF8 character set.

> **See Also:**
>
> The discussion of conversion rules under DBMS_CRYPTO Operational Notes

52.10.6 ENCRYPT Procedures

These procedures encrypt LOB data using a stream or block cipher with a user supplied key and optional IV (initialization vector).

Syntax

```
DBMS_CRYPTO.ENCRYPT(
   dst IN OUT NOCOPY BLOB,
   src IN            BLOB,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW          DEFAULT NULL);

DBMS_CRYPTO.ENCRYPT(
   dst IN OUT NOCOPY BLOB,
   src IN            CLOB         CHARACTER SET ANY_CS,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW          DEFAULT NULL);

DBMS_CRYPTO.ENCRYPT (
   dst IN OUT NOCOPY BLOB,
   src IN            BLOB,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW DEFAULT NULL,
   aad IN            RAW DEFAULT NULL,
   tag OUT           RAW);

DBMS_CRYPTO.ENCRYPT(
   dst IN OUT        NOCOPY BLOB,
   src IN            CLOB CHARACTER SET ANY_CS,
   typ IN            PLS_INTEGER,
   key IN            RAW,
   iv  IN            RAW DEFAULT NULL,
   aad IN            RAW DEFAULT NULL,
   tag OUT           RAW);
```

Pragmas

```
pragma restrict_references(encrypt,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-18 ENCRYPT Procedure Parameters

Parameter Name	Description
dst	LOB locator of output data. The value in the output LOB <dst> will be overwritten.
src	LOB locator of input data.
typ	Stream or block cipher type and modifiers to be used.
key	Encryption key to be used for encrypting data.
iv	Optional initialization vector for block ciphers. Default is NULL.

Table 52-18 (Cont.) ENCRYPT Procedure Parameters

Parameter Name	Description
aad	Additional authenticated data, which is any string that you pass to an Oracle Cloud key management service as part of the request.
tag	Authentication tag that is used for the authentication check.

Usage Notes

- ENCRYPT can use the ENCRYPT_SM4 constant for SM4 encryption algorithm.
- See DBMS_DEBUG Operational Notes for more information about the conversion rules for the ENCRYPT procedure.

52.10.7 HASH Function

A one-way hash function takes a variable-length input string, the data, and converts it to a fixed-length (generally smaller) output string called a *hash value*. The hash value serves as a unique identifier (like a fingerprint) of the input data. You can use the hash value to verify whether data has been changed or not.

Note that a one-way hash function is a hash function that works in one direction. It is easy to compute a hash value from the input data, but it is hard to generate data that hashes to a particular value. Consequently, one-way hash functions work well to ensure data integrity. Refer to “When to Use Hash or Message Authentication Code (MAC) Functions” in DBMS_CRYPTO Operational Notes for more information about using one-way hash functions.

This function applies to data one of the supported cryptographic hash algorithms listed in Table 52-4.

Syntax

```
DBMS_CRYPTO.HASH (
   src IN RAW,
   typ IN PLS_INTEGER)
 RETURN RAW;

DBMS_CRYPTO.HASH (
   src IN BLOB,
   typ IN PLS_INTEGER)
 RETURN RAW;

DBMS_CRYPTO.HASH (
   src IN CLOB CHARACTER SET ANY_CS,
   typ IN PLS_INTEGER)
 RETURN RAW;
```

Pragmas

```
pragma restrict_references(hash,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-19 HASH Function Parameters

Parameter Name	Description
src	The source data to be hashed.
typ	The hash algorithm to be used.

Usage Notes

- Oracle recommends that you use SHA-2 (SHA-256, SHA-384, SHA-512). SHA-1 (HASH_SH1) is deprecated.
- The HASH function can use the HASH_SM3 constant for the SM3 hash algorithm

52.10.8 HASH_LEN Function

HASH_LEN is an extension of the HASH function that can generate variable length hash output.

HASH_LEN includes an extra input length in PLS_INTEGER type, which is hash length. HASH_LEN only supports two types of hash: SHAKE128 and SHAKE256.

Syntax

```
DBMS_CRYPTO.HASH_LEN (
   src IN RAW,
   typ IN PLS_INTEGER,
   length IN PLS_INTEGER)
 RETURN RAW DETERMINISTIC;

DBMS_CRYPTO.HASH_LEN (
   src IN BLOB,
   typ IN PLS_INTEGER,
   length IN PLS_INTEGER)
 RETURN RAW DETERMINISTIC;

DBMS_CRYPTO.HASH_LEN (
   src IN CLOB CHARACTER SET ANY_CS,
   typ IN PLS_INTEGER,
   length IN PLS_INTEGER)
 RETURN RAW DETERMINISTIC;
```

Parameters

Table 52-20 HASH Function Parameters

Parameter Name	Description
src	The source data to be hashed.
typ	The hash algorithm to be used.
len	The variable length for SHAKE128 and SHAKE256 hash algorithms.

Usage Notes

`HASH_LEN` can only use hash types `HASH_SHAKE128` and `HASH_SHAKE256`. Other hash types are invalid for this function.

52.10.9 MAC Function

This function applies Message Authentication Code (MAC) algorithms to data to provide keyed message protection.

A MAC is a key-dependent one-way hash function. MACs have the same properties as the one-way hash function described in HASH Function, but they also include a key. Only someone with the identical key can verify the hash. Also refer to “When to Use Hash or Message Authentication Code (MAC) Functions” in DBMS_CRYPTO Operational Notesfor more information about using MACs.

See Table 52-5 for a list of MAC algorithms that have been defined for this package.

Syntax

```
DBMS_CRYPTO.MAC (
   src IN RAW,
   typ IN PLS_INTEGER,
   key IN RAW)
 RETURN RAW;

DBMS_CRYPTO.MAC (
   src IN BLOB,
   typ IN PLS_INTEGER
   key IN RAW)
 RETURN RAW;

DBMS_CRYPTO.MAC (
   src IN CLOB CHARACTER SET ANY_CS,
   typ IN PLS_INTEGER
   key IN RAW)
 RETURN RAW;
```

Pragmas

```
pragma restrict_references(mac,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-21 MAC Function Parameters

Parameter Name	Description
`src`	Source data to which MAC algorithms are to be applied.
`typ`	MAC algorithm to be used.
`key`	Key to be used for MAC algorithm.

52.10.10 KMACXOF Function

KMAC stands for KECCAK Message Authentication Code. This function is similar to the `MAC` function except that is includes the `length` and `custStr` fields.

Syntax

```
DBMS_CRYPTO.KMACXOF(
   src IN RAW,
   typ IN PLS_INTEGER,
   key IN RAW,
   length IN PLS_INTEGER,
   custStr IN RAW)
 RETURN RAW;

DBMS_CRYPTO.KMACXOF (
   src IN BLOB,
   typ IN PLS_INTEGER,
   key IN RAW,
   length IN PLS_INTEGER,
   custStr IN RAW)
 RETURN RAW;

DBMS_CRYPTO.KMACXOF (
   src IN CLOB CHARACTER SET ANY_CS,
   typ IN PLS_INTEGER,
   key IN RAW,
   length IN PLS_INTEGER,
   custStr IN RAW)
 RETURN RAW;
```

Parameters

Table 52-22 KMACXOF Function Parameters

Parameter Name	Description
`src`	Source data to which KMACXOF algorithms are to be applied.
`typ`	KMACXOF algorithm to be used.
`key`	Key to be used for KMACXOF algorithm.
`length`	Length of KMACXOF output in bytes.
`custStr`	Custom string for KMACXOF.

52.10.11 PKDECRYPT Function

This function decrypts `RAW` data using a private key assisted with key algorithm and encryption algorithm and returns decrypted data.

Syntax

```
DBMS_CRYPTO.PKDECRYPT(
   src IN RAW,
   prv_key IN RAW,
   pubkey_alg IN BINARY_INTEGER,
   enc_alg  IN BINARY_INTEGER)
 RETURN RAW;
```

Parameters

Table 52-23 PKDECRYPT Function Parameters

Parameter Name	Description
src	RAW data to be decrypted.
prv_key	Private key.
pubkey_alg	Specify the KEY_TYPE_RSA RSA key type.
enc_alg	Specify the algorithm PKENCRYPT_RSA_PKCS1_OAEP, for RSA Public Key Cryptosystem with PKCS1 and OAEP padding.

52.10.12 PKENCRYPT Function

This function encrypts RAW data using a public key assisted with key algorithm and encryption algorithm and returns encrypted data.

Syntax

```
DBMS_CRYPTO.PKENCRYPT(
   src IN RAW,
   pub_key IN RAW,
   pubkey_alg IN BINARY_INTEGER,
   enc_alg  IN BINARY_INTEGER)
 RETURN RAW;
```

Parameters

Table 52-24 PKENCRYPT Function Parameters

Parameter Name	Description
src	RAW data to be encrypted.
pub_key	Public key.
pubkey_alg	Specify the KEY_TYPE_RSA RSA key type.
enc_alg	Specify the algorithm PKENCRYPT_RSA_PKCS1_OAEP, for RSA Public Key Cryptosystem with PKCS1 and OAEP padding.

Usage Notes

You can use the PKENCRYPT_SM2 constant for the SM2 public key encryption algorithm and the KEY_TYPE_SM2 constant for SM2 key type.

52.10.13 RANDOMBYTES Function

This function returns a RAW value containing a cryptographically secure pseudo-random sequence of bytes, which can be used to generate random material for encryption keys.

Syntax

```
DBMS_CRYPTO.RANDOMBYTES (
   number_bytes IN POSITIVE)
 RETURN RAW;
```

Pragmas

```
pragma restrict_references(randombytes,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 52-25 RANDOMBYTES Function Parameter

Parameter Name	Description
number_bytes	The number of pseudo-random bytes to be generated.

Usage Note

The number_bytes value should not exceed the maximum length of a RAW variable.

52.10.14 RANDOMINTEGER Function

This function returns an integer in the complete range available for the Oracle BINARY_INTEGER datatype.

Syntax

```
DBMS_CRYPTO.RANDOMINTEGER
 RETURN BINARY_INTEGER;
```

Pragmas

```
pragma restrict_references(randominteger,WNDS,RNDS,WNPS,RNPS);
```

52.10.15 RANDOMNUMBER Function

This function returns an integer in the Oracle NUMBER datatype in the range of [0..2**128-1].

Syntax

```
DBMS_CRYPTO.RANDOMNUMBER
 RETURN NUMBER;
```

Pragmas

```
pragma restrict_references(randomnumber,WNDS,RNDS,WNPS,RNPS);
```

52.10.16 SIGN Function

This function signs RAW data using a private key assisted with key algorithm and sign algorithm, and returns a signature.

Syntax

```
DBMS_CRYPTO.SIGN(
   src IN RAW,
   prv_key IN RAW,
   pubkey_alg IN BINARY_INTEGER,
   sign_alg IN BINARY_INTEGER)
 RETURN RAW;
```

Parameters

Table 52-26 SIGN Function Parameters

Parameter Name	Description
src	RAW data to be signed.
prv_key	Private key.
pubkey_alg	Specify the KEY_TYPE_RSA RSA key type for RSA algorithms and KEY_TYPE_ECDSA ECDSA key type for ECDSA algorithms.
sign_alg	Specify one of the algorithms that are listed in the Usage Notes.

Usage Notes

Table 52-27 Signature Type Algorithms

Hash Algorithm	Description
SIGN_SHA1_RSA	SHA1 hash function with RSA
SIGN_SHA1_RSA_X931	SHA1 hash function with RSA and X931 padding
SIGN_SHA224_ECDSA	SHA 224-bit hash function with ECDSA
SIGN_SHA224_RSA	SHA 224-bit hash function with RSA
SIGN_SHA256_ECDSA	SHA 256-bit hash function with ECDSA
SIGN_SHA256_RSA	SHA 256-bit hash function with RSA
SIGN_SHA256_RSA_X931	SHA 256-bit hash function with RSA and X931 padding
SIGN_SHA3_224_ECDSA	SHA-3 224-bit hash function with ECDSA
SIGN_SHA3_224_RSA	SHA-3 234-bit hash function with RSA
SIGN_SHA3_256_ECDSA	SHA-3 256-bit hash function with ECDSA
SIGN_SHA3_256_RSA	SHA-3 256-bit hash function with RSA
SIGN_SHA3_384_ECDSA	SHA-3 384-bit hash function with ECDSA
SIGN_SHA3_384_RSA	SHA-3 384-bit hash function with RSA
SIGN_SHA3_512_ECDSA	SHA-3 512-bit hash function with ECDSA
SIGN_SHA3_512_RSA	SHA-3 512-bit hash function with RSA
SIGN_SHA384_ECDSA	SHA 384-bit hash function with ECDSA
SIGN_SHA384_RSA	SHA 384-bit hash function with RSA
SIGN_SHA384_RSA_X931	SHA 384-bit hash function with RSA and X931 padding
SIGN_SHA512_ECDSA	SHA 512-bit hash function with ECDSA
SIGN_SHA512_RSA	SHA 512-bit hash function with RSA
SIGN_SHA512_RSA_X931	SHA 512-bit hash function with RSA and X931 padding
SIGN_SM3_SM2	SM3 256-bit hash function with SM2

Usage Notes

You can use the SIGN_SM3_SM2 constant for the SM3_SM2 encryption algorithm

52.10.17 VERIFY Function

This function verifies `RAW` data using the signature, public key assisted with key algorithm, and sign algorithm. It returns `TRUE` if the signature was verified.

Syntax

```
DBMS_CRYPTO.VERIFY(
   src IN RAW,
   sign IN RAW,
   pub_key IN RAW,
   pubkey_alg IN BINARY_INTEGER,
   sign_alg  IN BINARY_INTEGER)
 RETURN BOOLEAN;
```

Parameters

Table 52-28 VERIFY Function Parameters

Parameter Name	Description
`src`	`RAW` data to be verified.
`sign`	Message signature.
`pub_key`	Public key.
`pubkey_alg`	Specify the KEY_TYPE_RSA RSA key type for RSA algorithms and KEY_TYPE_ECDSA ECDSA key type for ECDSA algorithms.
`sign_alg`	Specify one of the algorithms that are listed the Usage Notes.

Usage Notes

Table 52-29 Verify Type Algorithms

Hash Algorithm	Description
`SIGN_SHA1_RSA`	SHA hash function with RSA
`SIGN_SHA1_RSA_X931`	SHA hash function with RSA and X931 padding
`SIGN_SHA224_ECDSA`	SHA 224-bit hash function with ECDSA
`SIGN_SHA224_RSA`	SHA 224-bit hash function with RSA
`SIGN_SHA256_ECDSA`	SHA 256-bit hash function with ECDSA
`SIGN_SHA256_RSA`	SHA 256-bit hash function with RSA
`SIGN_SHA256_RSA_X931`	SHA 256-bit hash function with RSA and X931 padding
`SIGN_SHA3_224_ECDSA`	SHA-3 224-bit hash function with ECDSA
`SIGN_SHA3_224_RSA`	SHA-3 234-bit hash function with RSA
`SIGN_SHA3_256_ECDSA`	SHA-3 256-bit hash function with ECDSA
`SIGN_SHA3_256_RSA`	SHA-3 256-bit hash function with RSA
`SIGN_SHA3_384_ECDSA`	SHA-3 384-bit hash function with ECDSA
`SIGN_SHA3_384_RSA`	SHA-3 384-bit hash function with RSA

Table 52-29 (Cont.) Verify Type Algorithms

Hash Algorithm	Description
SIGN_SHA3_512_ECDSA	SHA-3 512-bit hash function with ECDSA
SIGN_SHA3_512_RSA	SHA-3 512-bit hash function with RSA
SIGN_SHA384_ECDSA	SHA 384-bit hash function with ECDSA
SIGN_SHA384_RSA	SHA 384-bit hash function with RSA
SIGN_SHA384_RSA_X931	SHA 384-bit hash function with RSA and X931 padding
SIGN_SHA512_ECDSA	SHA 512-bit hash function with ECDSA
SIGN_SHA512_RSA	SHA 512-bit hash function with RSA
SIGN_SHA512_RSA_X931	SHA 512-bit hash function with RSA and X931 padding
SIGN_SM3_SM2	SM3 256-bit hash function with SM2

Usage Notes

You can use the SIGN_SM3_SM2 constant for the SM3_SM2 encryption algorithm

53
DBMS_CSX_ADMIN

The DBMS_CSX_ADMIN package provides an interface to customize the setup when transporting a tablespace containing binary XML data.

The chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_CSX_ADMIN

53.1 DBMS_CSX_ADMIN Overview

This package can be used by DBAs to customize the setup when transporting a tablespace containing binary XML data. The use of the package is not required in order for a transportable tablespace job to run.

By default, all binary XML tables will use the default token table set, which will be replicated during transport on the target database. To avoid the cost of transporting a potentially large token table set, the DBA may opt for registering a new set of token tables for a given tablespace. The package provides routines for token table set registration and lookup.

53.2 DBMS_CSX_ADMIN Security Model

Owned by XDB, the DBMS_CSX_ADMIN package must be created by SYS or XDB. The EXECUTE privilege is granted to SYS or XDB or DBA. Subprograms in this package are executed using the privileges of the current user.

53.3 DBMS_CSX_ADMIN Constants

The DBMS_CSX_ADMIN package uses several constants.

These are described in the following table.

Table 53-1 DBMS_CSX_ADMIN Constants

Name	Type	Value	Description
DEFAULT_LEVEL	BINARY_INTEGER	0	Default token table
TAB_LEVEL	BINARY_INTEGER	1	Token table set associated with tables, not tablespaces
TBS_LEVEL	BINARY_INTEGER	2	Token table set associated with a tablespace
NO_CREATE	BINARY_INTEGER	0	Token tables already exist, associate them with the given table/tablespace

Table 53-1 (Cont.) DBMS_CSX_ADMIN Constants

Name	Type	Value	Description
NO_INDEXES	BINARY_INTEGER	1	Do not create indexes on the new set of token tables
WITH_INDEXES	BINARY_INTEGER	2	Create indexes on the token tables
DEFAULT_TOKS	BINARY_INTEGER	0	Prepopulate the token tables with default token mappings
NO_DEFAULT_TOKS	BINARY_INTEGER	1	Do not prepopulate the token tables with default token mappings

53.4 Summary of DBMS_CSX_ADMIN Subprograms

This table lists and describes the DBMS_CSX_ADMIN Package subprograms.

Table 53-2 DBMS_CSX_ADMIN Package Subprograms

Subprogram	Description
GETTOKENTABLEINFO Procedure & Function	Returns the GUID of the token table set where token mappings for this table
GETTOKENTABLEINFOBYTABLESPACE Procedure	Returns the GUID and the token table names for this tablespace
NAMESPACEIDTABLE Function	Returns default namespace-ID token table
PATHIDTABLE Function	Returns the default path-ID token table
QNAMEIDTABLE Function	Returns the default qname-ID token table.

53.4.1 GETTOKENTABLEINFO Procedure & Function

This procedure is overloaded. The specific forms of functionality are described along with the syntax declarations.

Given the table name and the owner, the first overload of the procedure returns the globally unique identifier (GUID) of the token table set where token mappings for this table can be found. The procedure returns also the names of the token tables, and whether the token table set is the default one.

Given the object number of a table, the second overload of the procedure returns the GUID of the token table set used by the table, and whether this is the default token table set.

Syntax

```
DBMS_CSX_ADMIN.GETTOKENTABLEINFO  (
   ownername       IN   VARCHAR2,
   tablename       IN   VARCHAR2,
   guid            OUT  RAW,
   qnametable      OUT  VARCHAR2,
   nmspctable      OUT  VARCHAR2,
   level           OUT  NUMBER,
   tabno           OUT  NUMBER);

DBMS_CSX_ADMIN.GETTOKENTABLEINFO  (
   tabno           IN   NUMBER,
```

```
  guid           OUT  RAW);
RETURN BOOLEAN;
```

Parameters

Table 53-3 GETTOKENTABLEINFO Procedure & Function Parameters

Parameter	Description
ownername	Owner of the table
tablename	Name of the table
guid	GUID of the token table set used by the given table
qnametable	Name of the qname-ID table in the new set
nmspctable	Name of the namespace-ID table in the new set
level	DEFAULT_LEVEL if default token table set, TBS_LEVEL if same token table set is used by all tables in the same tablespace as the given table, TAB_LEVEL otherwise
tabno	Table object number

53.4.2 GETTOKENTABLEINFOBYTABLESPACE Procedure

Given a tablespace number, this procedure returns the GUID and the token table names for this tablespace.

Syntax

```
DBMS_CSX_ADMIN.GETTOKENTABLEINFOBYTABLESPACE  (
   tsname          IN   VARCHAR2,
   tablespaceno    IN   NUMBER,
   guid            OUT  RAW,
   qnametable      OUT  VARCHAR2,
   nmspctable      OUT  VARCHAR2,
   isdefault       OUT  BOOLEAN,
   containTokTab   OUT  BOOLEAN);
```

Parameters

Table 53-4 GETTOKENTABLEINFOBYTABLESPACE Procedure Parameters

Parameter	Description
tsname	Tablespace name
tablespaceno	Tablespace number
guid	GUID of the token table set associated with this tablespace (if any)
qnametable	Name of the qname-ID table
nmspctable	Name of the namespace-ID table
isdefault	TRUE if the token table is the default one
containTokTab	TRUE if the tablespace contains its own token table set

53.4.3 NAMESPACEIDTABLE Function

This procedure returns default namespace-ID token table.

Syntax

```
DBMS_CSX_ADMIN.NAMESPACEIDTABLE
  RETURN VARCHAR2;
```

53.4.4 PATHIDTABLE Function

This procedure returns the default path-ID token table. This is used for granting permissions on the default path-ID token table for a user before executing `EXPLAIN PLAN` for a query on an XML table with an XML index.

Syntax

```
DBMS_CSX_ADMIN.PATHIDTABLE
  RETURN VARCHAR2;
```

53.4.5 QNAMEIDTABLE Function

This procedure returns the default qname-ID token table.

Syntax

```
DBMS_CSX_ADMIN.QNAMEIDTABLE
  RETURN VARCHAR2;
```

54
DBMS_CUBE

DBMS_CUBE contains subprograms that create OLAP cubes and dimensions, and that load and process the data for querying.

> **See Also:**
>
> OLAP Technology in the Oracle Database in *Oracle OLAP User's Guide* regarding use of the OLAP option to support business intelligence and analytical applications.

This chapter contains the following topics:

- Using DBMS_CUBE
- Using SQL Aggregation Management
- Upgrading 10g Analytic Workspaces
- Summary of DBMS_CUBE Subprograms

54.1 Using DBMS_CUBE

Cubes and cube dimensions are first class data objects that support multidimensional analytics. They are stored in a container called an analytic workspace. Multidimensional objects and analytics are available with the OLAP option to Oracle Database.

Cubes can be enabled as cube materialized views for automatic refresh of the cubes and dimensions, and for query rewrite. Several DBMS_CUBE subprograms support the creation and maintenance of cube materialized views as a replacement for relational materialized views. These subprograms are discussed in "Using SQL Aggregation Management".

The metadata for cubes and dimensions is defined in XML documents, called *templates*, which you can derive from relational materialized views using the CREATE_CUBE or DERIVE_FROM_MVIEW functions. Using a graphical tool named Analytic Workspace Manager, you can enhance the cube with analytic content or create the metadata for new cubes and cube dimensions from scratch.

Several other DBMS_CUBE subprograms provide a SQL alternative to Analytic Workspace Manager for creating an analytic workspace from an XML template and for refreshing the data stored in cubes and dimensions. The IMPORT_XML procedure creates an analytic workspace with its cubes and cube dimensions from an XML template. The BUILD procedure loads data into the cubes and dimensions from their data sources and performs whatever processing steps are needed to prepare the data for querying.

54.1.1 DBMS_CUBE Security Model

Certain roles and system privileges are required to use the DBMS_CUBE package.

To create dimensional objects in the user's own schema:

- `OLAP_USER` role
- `CREATE SESSION` privilege

To create dimensional objects in different schemas:

- `OLAP_DBA` role
- `CREATE SESSION` privilege

To create cube materialized views in the user's own schema:

- `CREATE MATERIALIZED VIEW` privilege
- `CREATE DIMENSION` privilege
- `ADVISOR` privilege

To create cube materialized views in different schemas:

- `CREATE ANY MATERIALIZED VIEW` privilege
- `CREATE ANY DIMENSION` privilege
- `ADVISOR` privilege

If the source tables are in a different schema, then the owner of the dimensional objects needs `SELECT` object privileges on those tables.

54.2 Using SQL Aggregation Management

SQL Aggregation Management is a group of PL/SQL subprograms in `DBMS_CUBE` that supports the rapid deployment of cube materialized views from existing relational materialized views.

Cube materialized views are cubes that have been enhanced to use the automatic refresh and query rewrite features of Oracle Database. A single cube materialized view can replace many of the relational materialized views of summaries on a fact table, providing uniform response time to all summary data.

Cube materialized views bring the fast update and fast query capabilities of the OLAP option to applications that query summaries of detail relational tables. The summary data is generated and stored in a cube, and query rewrite automatically redirects queries to the cube materialized views. Applications experience excellent querying performance.

In the process of creating the cube materialized views, `DBMS_CUBE` also creates a fully functional analytic workspace including a cube and the cube dimensions. The cube stores the data for a cube materialized view instead of the table that stores the data for a relational materialized view. A cube can also support a wide range of analytic functions that enhance the database with information-rich content.

Cube materialized views are registered in the data dictionary along with all other materialized views. A `CB$` prefix identifies a cube materialized view.

The `DBMS_CUBE` subprograms also support life-cycle management of cube materialized views.

> **See Also:**
>
> Adding Materialized View Capability to a Cube in *Oracle OLAP User's Guide* for more information about cube materialized views and enhanced OLAP analytics.

54.2.1 Subprograms in SQL Aggregation Management

SQL Aggregation Management includes four subprograms.

- CREATE_MVIEW Function
- DERIVE_FROM_MVIEW Function
- DROP_MVIEW Procedure
- REFRESH_MVIEW Procedure

54.2.2 Requirements for the Relational Materialized View

SQL Aggregation Management uses an existing relational materialized view to derive all the information needed to generate a cube materialized view. The relational materialized view determines the detail level of data that is stored in the cube materialized view. The related relational dimension objects determine the scope of the aggregates, from the lowest level specified in the GROUP BY clause of the materialized view subquery, to the highest level of the dimension hierarchy.

The relational materialized view must conform to these requirements:

- Explicit `GROUP BY` clause for one or more columns.
- No expressions in the select list or `GROUP BY` clause.
- At least one of these numeric aggregation methods: `SUM`, `MIN`, `MAX`, or `AVG`.
- No outer joins.
- Summary keys with at least one simple column associated with a relational dimension.

  or

  Summary keys with at least one simple column and no hierarchies or levels.
- Numeric datatype of any type for the fact columns. All facts are converted to `NUMBER`.
- Eligible for rewrite. `REWRITE_CAPABILITY` should be `GENERAL`; it cannot be `NONE`. Refer to the `ALL_MVIEWS` entry in the *Oracle Database Reference*.
- Cannot use the `DISTINCT` or `UNIQUE` keywords with an aggregate function in the defining query. For example, `AVG(DISTINCT units)` causes an error in `STRICT` mode and is ignored in `LOOSE` mode.

You can choose between two modes when rendering the cube materialized view, `LOOSE` and `STRICT`. In `STRICT` mode, any deviation from the requirements raises an exception and prevents the materialized view from being created. In `LOOSE` mode (the default), some deviations are allowed, but they affect the content of the materialized view. These elements in the relational materialized view generate warning messages:

- Complex expressions in the defining query are ignored and do not appear in the cube materialized view.

- The AVG function is changed to SUM and COUNT.
- The COUNT function without a SUM, MIN, MAX, or AVG function is ignored.
- The STDDEV and VARIANCE functions are ignored.

You can also choose how conditions in the WHERE clause are filtered. When filtering is turned off, the conditions are ignored. When turned on, valid conditions are rendered in the cube materialized view, but asymmetric conditions among dimension levels raise an exception.

54.2.3 Permissions for Managing and Querying Cube Materialized Views

Certain permissions are required to manage and query cube materialized views.

To create cube materialized views, you must have these privileges:

- CREATE [ANY] MATERIALIZED VIEW privilege
- CREATE [ANY] DIMENSION privilege
- ADVISOR privilege

To access cube materialized views from another schema using query rewrite, you must have these privileges:

- GLOBAL QUERY REWRITE privilege
- SELECT or READ privilege on the relational source tables
- SELECT or READ privilege on the analytic workspace (AW$*name*) that supports the cube materialized view
- SELECT or READ privilege on the cube
- SELECT or READ privilege on the dimensions of the cube

Note that you need SELECT or READ privileges on the database objects that *support* the cube materialized views, but not on the cube materialized views.

54.2.4 Example of SQL Aggregation Management

Six examples of SQL Aggregate Management are given. All these examples use the sample Sales History schema, which is installed in Oracle Database with two relational materialized views: CAL_MONTH_SALES_MV and FWEEK_PSCAT_SALES_MV.

- About Relational Materialized View CAL_MONTH_SALES_MV
- Creating the Cube Materialized View
- Disabling the Relational Materialized Views
- Creating Execution Plans for Cube Materialized Views
- Maintaining Cube Materialized Views
- New Database Objects

About Relational Materialized View CAL_MONTH_SALES_MV

This example uses CAL_MONTH_SALES_MV as the basis for creating a cube materialized view. The following query was used to create CAL_MONTH_SALES_MV. CAL_MONTH_SALES_MV summarizes the daily sales data stored in the SALES table by month.

```
SELECT query FROM user_mviews
     WHERE mview_name='CAL_MONTH_SALES_MV';

QUERY
-----------------------------------------
SELECT   t.calendar_month_desc
  ,        sum(s.amount_sold) AS dollars
  FROM     sales s
  ,        times t
  WHERE    s.time_id = t.time_id
  GROUP BY t.calendar_month_desc
```

DBMS_CUBE uses relational dimensions to derive levels and hierarchies for the cube materialized view. The SH schema has relational dimensions for most dimension tables in the schema, as shown by the following query.

```
SELECT dimension_name FROM user_dimensions;

DIMENSION_NAME
-----------------------------------
CUSTOMERS_DIM
PRODUCTS_DIM
TIMES_DIM
CHANNELS_DIM
PROMOTIONS_DIM
```

Creating the Cube Materialized View

This PL/SQL script uses the CREATE_MVIEW function to create a cube materialized view from CAL_MONTH_SALES_MV. CREATE_MVIEW sets the optional BUILD parameter to refresh the cube materialized view immediately.

```
SET serverout ON format wrapped

DECLARE
     salesaw  varchar2(30);

BEGIN
     salesaw := dbms_cube.create_mview('SH', 'CAL_MONTH_SALES_MV',
                'build=immediate');
END;
/
```

These messages confirm that the script created and refreshed CB$CAL_MONTH_SALES successfully:

```
Completed refresh of cube mview "SH"."CB$CAL_MONTH_SALES" at 20130212 08:42:58.0
03.
Created cube organized materialized view "CB$CAL_MONTH_SALES" for rewrite at 200
130212 08:42:58.004.
```

The following query lists the materialized views in the SH schema:

```
SELECT mview_name FROM user_mviews;

MVIEW_NAME
------------------------------
CB$CAL_MONTH_SALES
CB$TIMES_DIM_D1_CAL_ROLLUP
CAL_MONTH_SALES_MV
FWEEK_PSCAT_SALES_MV
```

Two new materialized views are registered in the data dictionary:

- CB$CAL_MONTH_SALES: Cube materialized view
- CB$TIMES_DIM_D1_CAL_ROLLUP: Cube dimension materialized view for the TIME_DIM Calendar Rollup hierarchy

Cube dimension materialized views support refresh of the cube materialized view. You do not directly administer dimension materialized views.

Disabling the Relational Materialized Views

After creating a cube materialized view, disable query rewrite on all relational materialized views for the facts now supported by the cube materialized view. You can drop them when you are sure that you created the cube materialized view with the optimal parameters.

```
ALTER MATERIALIZED VIEW cal_month_sales_mv DISABLE QUERY REWRITE;

Materialized view altered.
```

You can also use the DISABLEQRW parameter in the CREATE_MVIEW function, which disables query rewrite on the source materialized view as described in Table 54-7.

Creating Execution Plans for Cube Materialized Views

You can create execution plans for cube materialized views the same as for relational materialized views. The following command generates an execution plan for a query against the SALES table, which contains data at the day level. The answer set requires data summarized by quarter. Query rewrite would not use the original relational materialized view for this query, because its data is summarized by month. However, query rewrite can use the new cube materialized view for summary data for months, quarters, years, and all years.

```
EXPLAIN PLAN FOR SELECT
           t.calendar_quarter_desc,
           sum(s.amount_sold) AS dollars
  FROM     sales s,
           times t
  WHERE    s.time_id = t.time_id
  AND      t.calendar_quarter_desc LIKE '2001%'
  GROUP BY t.calendar_quarter_desc
  ORDER BY t.calendar_quarter_desc;
```

The query returns these results:

```
CALENDAR_QUARTER_DESC    DOLLARS
-------------------- ----------
2001-01               6547097.44
2001-02               6922468.39
2001-03               7195998.63
2001-04               7470897.52
```

The execution plan shows that query rewrite returned the summary data from the cube materialized view, CB$CAL_MONTH_SALES, instead of recalculating it from the SALES table.

```
SELECT plan_table_output FROM TABLE(dbms_xplan.display());

PLAN_TABLE_OUTPUT
-------------------------------------------------------------------------------------------------------------------
Plan hash value: 2999729407

-------------------------------------------------------------------------------------------------
| Id  | Operation                     | Name              | Rows  | Bytes | Cost (%CPU)| Time     |
-------------------------------------------------------------------------------------------------
```

```
|   0 | SELECT STATEMENT               |                    |     1 |    30 |     3  (34)| 00:00:01 |
|   1 |  SORT GROUP BY                 |                    |     1 |    30 |     3  (34)| 00:00:01 |
|*  2 |   MAT_VIEW REWRITE CUBE ACCESS | CB$CAL_MONTH_SALES |     1 |    30 |     2   (0)| 00:00:01 |
-------------------------------------------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------------

   2 - filter("CB$CAL_MONTH_SALES"."D1_CALENDAR_QUARTER_DESC" LIKE '2001%' AND
              "CB$CAL_MONTH_SALES"."SYS_GID"=63)

15 rows selected.
```

Maintaining Cube Materialized Views

You can create a cube materialized view that refreshes automatically. However, you can force a refresh at any time using the REFRESH_MVIEW Procedure:

```
BEGIN
     dbms_cube.refresh_mview('SH', 'CB$CAL_MONTH_SALES');
END;
/

Completed refresh of cube mview "SH"."CB$CAL_MONTH_SALES" at 20130212
14:30:59.534.
```

If you want to drop a cube materialized view, use the DROP_MVIEW Procedure so that all supporting database objects (analytic workspace, cube, cube dimensions, and so forth) are also dropped:

```
BEGIN
     dbms_cube.drop_mview('SH', 'CB$CAL_MONTH_SALES');
END;
/

Dropped cube organized materialized view "SH"."CAL_MONTH_SALES" including
container analytic workspace "SH"."CAL_MONTH_SALES_AW" at 20130212 13:38:47.878.
```

New Database Objects

The `CREATE_MVIEW` function creates several first class database objects in addition to the cube materialized views. You can explore these objects through the data dictionary by querying views such as `ALL_CUBES` and `ALL_CUBE_DIMENSIONS`.

This example created the following supporting objects:

- Analytic workspace `CAL_MONTH_SALES_AW` (`AW$CAL_MONTH_SALES_AW` table)
- Cube `CAL_MONTH_SALES`
- Cube dimension `TIMES_DIM_D1`
- Dimension hierarchy `CAL_ROLLUP`
- Dimension levels `ALL_TIMES_DIM`, `YEAR`, `QUARTER`, and `MONTH`
- Numerous attributes for levels in the `CAL_ROLLUP` hierarchy

54.3 Upgrading 10*g* Analytic Workspaces

You can upgrade an OLAP 10*g* analytic workspace to OLAP 12*c* by saving the OLAP 10*g* objects as an XML template and importing the XML into a different schema. The original analytic workspace remains accessible and unchanged by the upgrade process.

Oracle OLAP metadata is the same in OLAP 11*g* and OLAP 12*c* so you do not need to upgrade an OLAP 11*g* analytic workspace to OLAP 12*c*. This topic describes upgrading an Oracle OLAP 10*g* analytic workspace to OLAP 12*c*.

> **Tip:**
>
> Oracle recommends using Analytic Workspace Manager for performing upgrades. See Upgrading Metadata From Oracle OLAP 10*g* in *Oracle OLAP User's Guide*.

These subprograms in `DBMS_CUBE` support the upgrade process:

- CREATE_EXPORT_OPTIONS Procedure
- CREATE_IMPORT_OPTIONS Procedure
- EXPORT_XML Procedure
- EXPORT_XML_TO_FILE Procedure
- IMPORT_XML Procedure
- INITIALIZE_CUBE_UPGRADE Procedure
- UPGRADE_AW Procedure

Prerequisites:

- The OLAP 10*g* analytic workspace can use OLAP standard form metadata.
- Customizations to the OLAP 10*g* analytic workspace may not be exported to the XML template. You must re-create them in OLAP 12*c*.
- The original relational source data must be available to load into the new analytic workspace. If the data is in a different schema or the table names are different, then you must remap the dimensional objects to the new relational sources after the upgrade.
- You can create the OLAP 12*c* analytic workspace in the same schema as the OLAP 10*g* analytic workspace. However, if you prefer to create it in a different schema, then create a new user with the following privileges:
  - `SELECT` or `READ` privileges on the OLAP 10*g* analytic workspace (`GRANT SELECT ON` *`schema`*`.AW$`*`analytic_workspace`*).
  - `SELECT` or `READ` privileges on all database tables and views that contain the source data for the OLAP 10*g* analytic workspace.
  - Appropriate privileges for an OLAP administrator.
  - Same default tablespace as the Oracle 10*g* user.

  See the *Oracle OLAP User's Guide*.

Correcting Naming Conflicts

The namespaces are different in OLAP 10*g* than those in OLAP 12*c*. For a successful upgrade, you must identify any 10*g* object names that are used multiple times under the 12*c* naming rules and provide unique names for them.

The following namespaces control the uniqueness of OLAP object names in Oracle 12*c*:

- **Schema**: The names of cubes, dimensions, and measure folders must be unique within a schema. They cannot conflict with the names of tables, views, indexes, relational dimensions, or any other first class objects. However, these OLAP 12*c* object names do not need to be distinct from 10*g* object names, because they are in different namespaces.
- **Cube**: The names of measures must be unique within a cube.
- **Dimension**: The names of hierarchies, levels, and attributes must be unique within a dimension. For example, a dimension cannot have a hierarchy named Customers and a level named Customers.

You can use an initialization table and a rename table to rename objects in the upgraded 12*c* analytic workspace.

Initialization Table

The `INITIALIZE_CUBE_UPGRADE` procedure identifies ambiguous names under the OLAP 12*c* naming rules. For example, a 10*g* dimension might have a hierarchy and a level with the same name. Because hierarchies and levels are in the same 12*c* namespace, the name is not unique in 12*c*; to a 12*c* client, the hierarchy and the level cannot be differentiated by name.

`INITIALIZE_CUBE_UPGRADE` creates and populates a table named `CUBE_UPGRADE_INFO` with unique names for these levels, hierarchies, and attributes. By using the unique names provided in the table, a 12*c* client can browse the OLAP 12*c* metadata. You cannot attach an OLAP 12*c* client to the analytic workspace or perform an upgrade without a `CUBE_UPGRADE_INFO` table, if the 10*g* metadata contains ambiguous names.

You can edit `CUBE_UPGRADE_INFO` to change the default unique names to names of your choosing. You can also add rows to change the names of any other objects. When using an 12*c* client, you see the new object names. When using an 10*g* client, you see the original names. However, the `INITIALIZE_CUBE_UPGRADE` procedure overwrites this table, so you may prefer to enter customizations in a rename table.

During an upgrade from OLAP 10*g,* the unique object names in `CUBE_UPGRADE_INFO` are used as the names of 12*c* objects in the new analytic workspace. However, `INITIALIZE_CUBE_UPGRADE` does not automatically provide unique names for cubes, dimensions, and measure folders. To complete an upgrade, you must assure that these objects have unique names within the 12*c* namespace. You can provide these objects with new names in the `CUBE_UPGRADE_INFO` table or in a rename table.

OLAP 12*c* clients automatically use `CUBE_UPGRADE_INFO` when it exists in the same schema as the OLAP 10*g* analytic workspace.

> **See Also:**
>
> "INITIALIZE_CUBE_UPGRADE Procedure"

Rename Table

You can create a rename table that contains new object names for an OLAP 12*c* analytic workspace. You can then use the rename table in the `CREATE_IMPORT_OPTIONS` and `UPGRADE_AW` procedures.

When upgrading within the same schema, you must provide a unique name for the 12*c* analytic workspace. The `UPGRADE_AW` procedure provides a parameter for this purpose; otherwise, you must provide the new name in the rename table. The duplication of cube names does not create ambiguity because the 12*c* cubes are created in a different namespace than the 10*g* cubes.

The names provided in a rename table are used only during an upgrade and overwrite any names entered in the `CUBE_UPGRADE_INFO` table.

To create a rename table:

1. Open SQL*Plus or another SQL client, and connect to Oracle Database as the owner of the 10*g* analytic workspace.
2. Issue a command like the following:

```
CREATE TABLE table_name (
          source_id    VARCHAR2(300),
          new_name     VARCHAR2(30),
          object_type  VARCHAR2(30));
```

3. Populate the rename table with the appropriate values, as follows.

`table_name` is the name of the rename table.

`source_id` is the identifier for an object described in the XML document supplied to IMPORT_XML. The identifier must have this format:

schema_name.object_name[*.subobject_name*]

`new_name` is the object name given during the import to the object specified by `source_id`.

`object_type` is the object type as described in the XML, such as StandardDimension or DerivedMeasure.

For example, these SQL statements populate the table with new names for the analytic workspace, a cube, and four dimensions:

```
INSERT INTO my_object_map VALUES('GLOBAL_AW.GLOBAL10.AW', 'GLOBAL12', 'AW');
INSERT INTO my_object_map VALUES('GLOBAL_AW.UNITS_CUBE', 'UNIT_SALES_CUBE', 'Cube');
INSERT INTO my_object_map VALUES('GLOBAL_AW.CUSTOMER', 'CUSTOMERS', 'StandardDimension');
INSERT INTO my_object_map VALUES('GLOBAL_AW.CHANNEL', 'CHANNELS', 'StandardDimension');
INSERT INTO my_object_map VALUES('GLOBAL_AW.PRODUCT', 'PRODUCTS', 'StandardDimension');
INSERT INTO my_object_map VALUES('GLOBAL_AW.TIME', 'TIME_PERIODS', 'TimeDimension');
```

> **See Also:**
>
> "CREATE_IMPORT_OPTIONS Procedure"

Simple Upgrade

A simple upgrade creates an OLAP 12*c* analytic workspace from an OLAP 10*g* analytic workspace.

To perform a simple upgrade of an Oracle OLAP 10*g* analytic workspace:

1. Open SQL*Plus or a similar SQL command-line interface and connect to Oracle Database 12*c* as the schema owner for the OLAP 12*c* analytic workspace.
2. To rename any objects in the 12*c* analytic workspace, create a rename table as described in the Rename Table section. (Optional)
3. Perform the upgrade, as described in "UPGRADE_AW Procedure".
4. Use the `DBMS_CUBE.BUILD` procedure to load data into the cube.

Example 54-1 Performing a Simple Upgrade to the GLOBAL Analytic Workspace

This example creates an OLAP 12*c* analytic workspace named `GLOBAL12` from an OLAP 10*g* analytic workspace named `GLOBAL10`. `GLOBAL10` contains no naming conflicts between cubes, dimensions, measure folders, or tables in the schema, so a rename table is not needed in this example.

```
BEGIN

  -- Upgrade the analytic workspace
  dbms_cube.upgrade_aw(sourceaw =>'GLOBAL10', destaw => 'GLOBAL12');

  -- Load and aggregate the data
  dbms_cube.build(script => 'UNITS_CUBE, PRICE_AND_COST_CUBE');

END;
/
```

Custom Upgrade

A custom upgrade enables you to set the export and import options.

To perform a custom upgrade of an Oracle OLAP 10*g* analytic workspace:

1. Open SQL*Plus or a similar SQL command-line interface and connect to Oracle Database 12*c* as the schema owner of the OLAP 12*c* analytic workspace.
2. Generate an initialization table, as described in the Initialization Table section. Review the new, default object names and modify them as desired.
3. Create a rename table, as described in the Rename Table section. If you are upgrading in the same schema, you must use a rename table to provide a unique name for the 12*c* analytic workspace. Otherwise, a rename table is needed only if names are duplicated among the cubes, dimensions, and measure folders of the analytic workspace, or between those names and the existing cubes, dimensions, measure folders, or tables of the destination schema.
4. Create a SQL script that does the following:
   a. Create an XML document for the export options, as described in "CREATE_EXPORT_OPTIONS Procedure". The `SUPPRESS_NAMESPACE` option must be set to `TRUE` for the upgrade to occur.
   b. Create an XML document for the import options, as described in "CREATE_IMPORT_OPTIONS Procedure".

c. Create an XML template in OLAP 12*c* format, as described in "EXPORT_XML Procedure".

d. Create an OLAP 12*c* analytic workspace from the XML template, as described in "IMPORT_XML Procedure".

5. Load and aggregate the data in the new analytic workspace, as described in "BUILD Procedure".

Example 54-2 Performing a Custom Upgrade to the GLOBAL Analytic Workspace

This example upgrades the `GLOBAL10` analytic workspace from OLAP 10*g* metadata to OLAP 12*c* metadata in the `GLOBAL_AW` schema.

The rename table provides the new name of the analytic workspace. These commands define the rename table.

```
CREATE TABLE my_object_map(
       source_id    VARCHAR2(300),
       new_name     VARCHAR2(30),
       object_type  VARCHAR2(30));

INSERT INTO my_object_map VALUES('GLOBAL_AW.GLOBAL10.AW',  'GLOBAL12', 'AW');
COMMIT;
```

Following is the script for performing the upgrade.

```
set serverout on

DECLARE
  importClob    clob;
  exportClob    clob;
  exportOptClob clob;
  importOptClob clob;

BEGIN

  -- Create table of reconciled names
  dbms_cube.initialize_cube_upgrade;

  -- Create a CLOB containing the export options
  dbms_lob.createtemporary(exportOptClob, TRUE);
  dbms_cube.create_export_options(out_options_xml=>exportOptClob,
suppress_namespace=>TRUE, preserve_table_owners=>TRUE);

  -- Create a CLOB containing the import options
  dbms_lob.createtemporary(importOptClob, TRUE);
  dbms_cube.create_import_options(out_options_xml=>importOptClob, rename_table =>
'MY_OBJECT_MAP');

   -- Create CLOBs for the metadata
  dbms_lob.createtemporary(importClob, TRUE);
  dbms_lob.createtemporary(exportClob, TRUE);

  -- Export metadata from a 10g analytic workspace to a CLOB
  dbms_cube.export_xml(object_ids=>'GLOBAL_AW', options_xml=>exportOptClob,
out_xml=>exportClob);

  -- Import metadata from the CLOB
  dbms_cube.import_xml(in_xml => exportClob, options_xml=>importOptClob,
out_xml=>importClob);

  -- Load and aggregate the data
```

```
  dbms_cube.build('UNITS_CUBE, PRICE_AND_COST_CUBE');

END;
/
```

54.4 Summary of DBMS_CUBE Subprograms

This table lists and describes the DBMS_CUBE procedure subprograms.

Table 54-1 DBMS_CUBE Subprograms

Subprogram	Description
BUILD Procedure	Loads data into one or more cubes and dimensions, and prepares the data for querying.
CREATE_EXPORT_OPTIONS Procedure	Creates an input XML document of processing options for the EXPORT_XML procedure.
CREATE_IMPORT_OPTIONS Procedure	Creates an input XML document of processing options for the IMPORT_XML procedure.
CREATE_MVIEW Function	Creates a cube materialized view from the definition of a relational materialized view.
DERIVE_FROM_MVIEW Function	Creates an XML template for a cube materialized view from the definition of a relational materialized view.
DROP_MVIEW Procedure	Drops a cube materialized view.
EXPORT_XML Procedure	Exports the XML of an analytic workspace to a CLOB.
EXPORT_XML_TO_FILE Procedure	Exports the XML of an analytic workspace to a file.
IMPORT_XML Procedure	Creates, modifies, or drops an analytic workspace by using an XML template
INITIALIZE_CUBE_UPGRADE Procedure	Processes Oracle OLAP 10*g* objects with naming conflicts to enable Oracle 12*c* clients to access them.
REFRESH_MVIEW Procedure	Refreshes a cube materialized view.
UPGRADE_AW Procedure	Upgrades an analytic workspace from Oracle OLAP 10*g* to 12*c*.
VALIDATE_XML Procedure	Checks the XML to assure that it is valid, without committing the results to the database.

54.4.1 BUILD Procedure

This procedure loads data into one or more cubes and dimensions, and generates aggregate values in the cubes. The results are automatically committed to the database.

Syntax

```
DBMS_CUBE.BUILD (
       script                IN  VARCHAR2,
       method                IN  VARCHAR2        DEFAULT NULL,
       refresh_after_errors  IN  BOOLEAN         DEFAULT FALSE,
       parallelism           IN  BINARY_INTEGER  DEFAULT 0,
       atomic_refresh        IN  BOOLEAN         DEFAULT FALSE,
       automatic_order       IN  BOOLEAN         DEFAULT TRUE,
       add_dimensions        IN  BOOLEAN         DEFAULT TRUE,
       scheduler_job         IN  VARCHAR2        DEFAULT NULL,
       master_build_id       IN  BINARY_INTEGER  DEFAULT 0,
```

```
nested                IN  BOOLEAN         DEFAULT FALSE);
job_class             IN  VARCHAR2        DEFAULT 'DEFAULT_JOB_CLASS'
```

Parameters

Table 54-2 BUILD Procedure Parameters

Parameter	Description
script	A list of cubes and dimensions and their build options (see "SCRIPT Parameter").
method	A full or a fast (partial) refresh. In a fast refresh, only changed rows are inserted in the cube and the affected areas of the cube are re-aggregated. You can specify a method for each cube and dimension in sequential order, or a single method to apply to all cubes and dimensions. If you list more objects than methods, then the last method applies to the additional objects. • C: Complete refresh clears all dimension values before loading. (Default) • F: Fast refresh of a cube materialized view, which performs an incremental refresh and re-aggregation of only changed rows in the source table. • ?: Fast refresh if possible, and otherwise a complete refresh. • P: Recomputes rows in a cube materialized view that are affected by changed partitions in the detail tables. • S: Fast solve of a compressed cube. A fast solve reloads all the detail data and re-aggregates only the changed values. See the "Usage Notes" for additional details. Methods do not apply to dimensions.
refresh_after_errors	TRUE to roll back just the cube or dimension with errors, and then continue building the other objects. FALSE to roll back all objects in the build.
parallelism	Number of parallel processes to allocate to this job (see Usage Notes).
atomic_refresh	TRUE prevents users from accessing intermediate results during a build. It freezes the current state of an analytic workspace at the beginning of the build to provide current sessions with consistent data. This option thaws the analytic workspace at the end of the build to give new sessions access to the refreshed data. If an error occurs during the build, then all objects are rolled back to the frozen state. FALSE enables users to access intermediate results during an build.
automatic_order	TRUE enables optimization of the build order. Dimensions are loaded before cubes. FALSE builds objects in the order you list them in the script.
add_dimensions	TRUE automatically includes all the dimensions of the cubes in the build, whether or not you list them in the script. If a cube materialized view with a particular dimension is fresh, then that dimension is not reloaded. You can list a cube once in the script. FALSE includes only dimensions specifically listed in the script.
scheduler_job	Any text identifier for the job, which will appear in the log table. The string does not need to be unique.
master_build_id	A unique name for the build.

Table 54-2 (Cont.) BUILD Procedure Parameters

Parameter	Description
nested	TRUE performs nested refresh operations for the specified set of cube materialized views. Nested refresh operations refresh all the depending materialized views and the specified set of materialized views based on a dependency order to ensure the nested materialized views are truly fresh with respect to the underlying base tables. All objects must reside in a single analytic workspace.
job_class	The class this job is associated with.

SCRIPT Parameter

The SCRIPT parameter identifies the objects to include in the build, and specifies the type of processing to perform on each one. The parameter has this syntax:

```
[VALIDATE | NO COMMIT] objects [ USING ( commands ) ][,...]
```

Where:

VALIDATE checks all steps of the build and sends the planned steps to CUBE_BUILD_LOG without executing the steps. You can view all generated SQL in the OUTPUT column of the log table.

NO COMMIT builds the objects in the current attach mode (or Read Only when the analytic workspace is not attached) but does not commit the changes. This option supports what-if analysis, since it enables you to change data values temporarily. See "SCRIPT Parameter: USING Clause: SET command".

objects is the qualified name of one or more cubes or dimensions, separated by commas, in the form *[aw_name.]object*, such as UNITS_CUBE or GLOBAL.UNITS_CUBE.

SCRIPT Parameter: USING Clause

The USING clause specifies the processing options. It consists of one or more commands separated by commas.

> **Note:**
>
> A cube with a rewrite materialized view cannot have a USING clause, except for the ANALYZE command. It uses the default build options.

The commands can be any of the following.

- AGGREGATE USING [MEASURE]

  Generates aggregate values using the syntax described in "SCRIPT Parameter: USING Clause: AGGREGATE command".

- ANALYZE

  Runs DBMS_AW_STATS.ANALYZE, which generates and stores optimizer statistics for cubes and dimensions.

- CLEAR [VALUES | LEAVES | AGGREGATES] [SERIAL | PARALLEL]

Prepares the cube for a data refresh. It can also be used on dimensions, but CLEAR removes all dimension keys, and thus deletes all data values for cubes that use the dimension.

These optional arguments control the refresh method. If you omit the argument, then the behavior of CLEAR depends on the refresh method. The 'C' (complete) refresh method runs CLEAR VALUES, and all other refresh methods run CLEAR LEAVES.

- VALUES: Clears all data in the cube. All facts must be reloaded and all aggregates must be recomputed. This option supports the COMPLETE refresh method. (Default for the C and F methods)
- LEAVES: Clears the detail data and retains the aggregates. All facts must be reloaded, and the aggregates for any new or changed facts must be computed. This option supports the FAST refresh method. (Default for the ? method)
- AGGREGATES: Retains the detail data and clears the aggregates. All aggregates must be recomputed.

These optional arguments control the load method, and can be combined with any of the refresh options:

- PARALLEL: Each partition is cleared separately. (Default)
- SERIAL: All partitions are cleared together.

If you omit the CLEAR command, DBMS_CUBE loads new and updated facts, but does not delete any old detail data. This is equivalent to a LOAD NO SYNC for dimensions.

- COMPILE [SORT | NO SORT | SORT ONLY]

  Creates the supporting structures for the dimension. (Dimensions only)

  These options control the use of a sort order attribute:

  - SORT: The user-defined sort order attribute populates the sort column in the embedded-total (ET) view. (Default)
  - NO SORT: Any sort order attribute is ignored. This option is for very large dimensions where sorting could consume too many resources.
  - SORT ONLY: The compile step only runs the sort.

- EXECUTE PLSQL string

  Executes a PL/SQL command or script in the database.

- EXECUTE OLAP DML string [PARALLEL | SERIAL]

  Executes an OLAP DML command or program in the analytic workspace. The options control execution of the command or program:

  - PARALLEL: Execute the command or program once for each partition. This option can be used to provide a performance boost to complex DML operations, such as forecasts and models.
  - SERIAL: Execute the command or program once for the entire cube. (Default)

- [INSERT | MERGE] INTO [ALL HIERARCHIES | HIERARCHIES (*dimension.hierarchy*)] VALUES (*dim_key, parent, level_name*)

  Adds a dimension member to one or more hierarchies. INSERT throws an error if the member already exists, while MERGE does not. See "Dimension Maintenance Example".

  dimension.hierarchy: The name of a hierarchy the new member belongs to. Enclose each part of the name in double quotes, for example, "PRODUCT"."PRIMARY".

dim_key: The DIM_KEY value of the dimension member.

parent: The parent of the dimension key.

level_name: The level of the dimension key.

- UPDATE [ALL HIERARCHIES | HIERARCHIES (*dimension.hierarchy*)] SET PARENT = *parent*, LEVEL=*level_name* WHERE MEMBER = *dim_key*

  Alters the level or parent of an existing dimension member. See INSERT for a description of the options. Also see "Dimension Maintenance Example".

- DELETE FROM DIMENSION WHERE MEMBER=*dim_key*

  Deletes a dimension member. See "Dimension Maintenance Example".

  dim_key: The DIM_KEY value of the dimension member to be deleted.

- SET *dimension.attribute[qdr]* = CAST('*attribute_value*' AS VARCHAR2))

  Sets the value of an attribute for a dimension member. See "Dimension Maintenance Example".

  dimension.attribute: The name of the attribute. Enclose each part of the name in double quotes, for example, "PRODUCT"."LONG_DESCRIPTION".

  qdr: The dimension member being given an attribute value in the form of a qualified data reference, such as "PRODUCT"='OPT MOUSE'.

  attribute_value: The value of the attribute, such as 'Optical Mouse'.

- FOR *dimension_clause measure_clause* BUILD (*commands*)

  Restricts the build to particular measures and dimension values, using the following arguments. See "FOR Clause Example".

  – *dimension_clause*:

    dimension ALL | NONE | WHERE *condition* | LEVELS (*level* [, *level*...])

    dimension is the name of a dimension of the cube.

    ALL sets the dimension status to all members before executing the list of commands.

    NONE loads values for no dimension members.

    WHERE loads values for those dimension members that match the condition.

    LEVELS loads values for dimension members in the named levels.

    level is a level of the named dimension.

  – *measure_clause*:

    MEASURES (*measure* [, *measure*...])

    measure is the name of a measure in the cube.

  – *commands*: Any of the other USING commands.

- LOAD [SYNCH | NO SYNCH | RETAIN] [PRUNE | PARALLEL | SERIAL] [WHERE *condition*]

  Loads data into the dimension or cube.

  – WHERE limits the load to those values in the mapped relational table that match *condition*.

  – *condition* is a valid predicate based on the columns of the mapped table. See the "Examples".

These optional arguments apply only to dimensions:

- `SYNCH` matches the dimension keys to the relational data source. (Default)
- `NO SYNCH` loads new dimension keys but does not delete old keys.

  If the parent of a dimension key has changed in the relational data source, this option allows the load to change the parent/child relation in the analytic workspace.
- `RETAIN` loads new dimension keys but does not delete old keys.

  This option does not allow the parent of a dimension key to change. If the parent has changed, the load rejects the record. The rejection generates an error in the rejected records log, if the log is enabled.

These optional arguments apply only to cubes:

- `PRUNE`: Runs a full table scan on the fact table to determine which partitions to load. For example, if a cube is partitioned by month and the fact table has values only for the last two months, then jobs are only started to load the partitions for the last two months.
- `PARALLEL`: Each partition is loaded separately. (Default)
- `SERIAL`: All partitions are loaded in one `SELECT` statement.

- `MODEL model_name [PARALLEL | SERIAL]`

  Executes a model previously created for the cube. It accepts these arguments:
  - `PARALLEL`: The model runs separately on each partition.
  - `SERIAL`: The model runs on all cubes at the same time. (Default)
- `SET`

  Supports write-back to the cube using the syntax described in "SCRIPT Parameter: USING Clause: SET command". (Cubes only)
- `SOLVE [PARALLEL | SERIAL]`

  Aggregates the cube using the rules defined for the cube, including the aggregation operator and the precompute specifications. (Cubes only)

  It accepts these arguments:
  - `PARALLEL`: Each partition is solved separately. (Default)
  - `SERIAL`: All partitions are solved at the same time.

SCRIPT Parameter: USING Clause: AGGREGATE command

The `AGGREGATE` command in a script specifies the aggregation rules for one or more measures.

> **Note:**
>
> The `AGGREGATE` command is available only for uncompressed cubes.

`AGGREGATE` has the following syntax:

```
{ AGGREGATE USING MEASURE
     WHEN measure1 THEN operator1
     WHEN measure2 THEN operator2...
```

```
          ELSE default_operator
|
 [AGGREGATE USING] operator_clause }
processing_options
OVER { ALL | dimension | dimension HIERARCHIES (hierarchy)}
```

USING MEASURE Clause

This clause enables you to specify different aggregation operators for different measures in the cube.

Operator Clause

The `operator_clause` has this syntax:

```
operator(WEIGHTBY expression | SCALEBY expression)
```

`WEIGHTBY` multiplies each data value by an expression before aggregation.

`SCALEBY` adds the value of an expression to each data value before aggregation.

Table 54-3 Aggregation Operators

Operator	Option	Description
AVG	WEIGHTBY	Adds data values, then divides the sum by the number of data values that were added together.
FIRST	WEIGHTBY	The first real data value.
HIER_AVG	WEIGHTBY	Adds data values, then divides the sum by the number of the children in the dimension hierarchy. Unlike AVERAGE, which counts only non-NA children, HAVERAGE counts all of the logical children of a parent, regardless of whether each child does or does not have a value.
HIER_FIRST	WEIGHTBY	The first data value in the hierarchy, even when that value is NA.
HIER_LAST	WEIGHTBY	The last data value in the hierarchy, even when that value is NA.
LAST	WEIGHTBY	The last real data value.
MAX	WEIGHTBY	The largest data value among the children of each parent.
MIN	WEIGHTBY	The smallest data value among the children of each parent.
NO AGGREGATION	No option	Do not aggregate the values of the dimension or dimensions. Leave all aggregated values as NA.
SUM	SCALEBY | WEIGHTBY	Adds data values. (Default)

Processing Options

You can specify these processing options for aggregation:

- `(ALLOW | DISALLOW) OVERFLOW`

  Specifies whether to allow decimal overflow, which occurs when the result of a calculation is very large and can no longer be represented by the exponent portion of the numerical representation.

  – `ALLOW`: A calculation that generates overflow executes without error and produces null results. (Default)

  - DISALLOW: A calculation involving overflow stops executing and generates an error message.

- (ALLOW | DISALLOW) DIVISION BY ZERO

  Specifies whether to allow division by zero.

  - ALLOW: A calculation involving division by zero executes without error but returns a null value. (Default)
  - DISALLOW: A calculation involving division by zero stops executing and generates an error message.

- (CONSIDER | IGNORE) NULLS

  Specifies whether nulls are included in the calculations.

  - CONSIDER: Nulls are included in the calculations. A calculation that includes a null value returns a null value.
  - IGNORE: Only actual data values are used in calculations. Nulls are treated as if they do not exist. (Default)

- MAINTAIN COUNT

  Stores an up-to-date count of the number of dimension members for use in calculating averages. Omit this option to count the members on the fly.

SCRIPT Parameter: USING Clause: SET command

The SET command in a script assigns values to one or more cells in a stored measure. It has this syntax:

```
SET target = expression
```

Where:

target is a a measure or a qualified data reference.

expression returns values of the appropriate datatype for *target*.

Qualified Data References

Qualified data references (QDRs) limit a dimensional object to a single member in one or more dimensions for the duration of a query.

A QDR has the following syntax:

```
expression [ { dimension = member }[ , { dimension = member } ...] ]
```

Where:

expression is a dimensional expression, typically the name of a measure.

dimension is a primary dimension of expression.

member is a value of dimension.

The outside square brackets shown in bold are literal syntax elements; they do not indicate an optional argument. The inside square brackets shown in regular text delimit an optional argument and are not syntax elements.

This example returns Sales values for calendar year 2007:

```
global.sales[global.time = 'CY2007'
]
```

The next example returns Sales values only for the United States in calendar year 2007:

```
sales[customer = 'US', time = 'CY2007'
]
```

See the Examples for qualified data references in SET commands.

Usage Notes

Build Methods

The `C`, `S`, and `?` methods always succeed and can be used on any cube.

The `F` and `P` methods require that the cube have a materialized view that was created as a fast or a rewrite materialized view.

Parallelism

Partitioned cubes can be loaded and aggregated in parallel processes. For example, a cube with five partitions can use up to five processes. Dimensions are always loaded serially.

The number of parallel processes actually allocated by a build is controlled by the smallest of these factors:

- Number of cubes in the build and the number of partitions in each cube.
- Setting of the `PARALLELISM` argument of the `BUILD` procedure.
- Setting of the `JOB_QUEUE_PROCESSES` database initialization parameter.

Suppose `UNITS_CUBE` has 12 partitions, `PARALLELISM` is set to 10, and `JOB_QUEUE_PROCESSES` is set to 4. OLAP uses four processes, which appear as slave processes in the build log.

The SQL engine may allocate additional processes when the PARALLEL_DEGREE_POLICY database initialization parameter is set to AUTO or LIMITED. For example, if OLAP allocates four processes, the SQL engine might determine that two of those processes should be done by four processes instead, for a total of six processes.

Build Logs

OLAP generates three logs that provide diagnostic information about builds:

- Cube build log
- Rejected records log
- Cube dimension compile log

Analytic Workspace Manager creates these logs automatically as tables in the same schema as the analytic workspace. If you do not use Analytic Workspace Manager, you can create and manage the logs in PL/SQL using the `DBMS_CUBE_LOG` package.

You can also create the cube log file by running `$ORACLE_HOME/olap/admin/utlolaplog.sql`. This script creates three additional views:

- `CUBE_BUILD_LATEST`: Returns rows only from the last build.
- `CUBE_BUILD_REPORT`: Returns one row for each command with elapsed times.
- `CUBE_BUILD_REPORT_LATEST`: Returns a report like `CUBE_BUILD_REPORT` only from the last build.

This report shows a successfully completed build of the objects in the GLOBAL analytic workspace, which has four dimensions and two cubes.

```
SELECT command, status, build_object, build_object_type type
     FROM cube_build_report_latest;

COMMAND                   STATUS     BUILD_OBJECT                   TYPE
------------------------- ---------- ------------------------------ ----------
BUILD                     COMPLETED                                 BUILD
FREEZE                    COMPLETED                                 BUILD
LOAD NO SYNCH             COMPLETED  CHANNEL                        DIMENSION
COMPILE                   COMPLETED  CHANNEL                        DIMENSION
UPDATE/COMMIT             COMPLETED  CHANNEL                        DIMENSION
LOAD NO SYNCH             COMPLETED  CUSTOMER                       DIMENSION
COMPILE                   COMPLETED  CUSTOMER                       DIMENSION
UPDATE/COMMIT             COMPLETED  CUSTOMER                       DIMENSION
LOAD NO SYNCH             COMPLETED  PRODUCT                        DIMENSION
COMPILE                   COMPLETED  PRODUCT                        DIMENSION
UPDATE/COMMIT             COMPLETED  PRODUCT                        DIMENSION
LOAD NO SYNCH             COMPLETED  TIME                           DIMENSION
COMPILE                   COMPLETED  TIME                           DIMENSION
UPDATE/COMMIT             COMPLETED  TIME                           DIMENSION
COMPILE AGGMAP            COMPLETED  PRICE_CUBE                     CUBE
UPDATE/COMMIT             COMPLETED  PRICE_CUBE                     CUBE
COMPILE AGGMAP            COMPLETED  UNITS_CUBE                     CUBE
UPDATE/COMMIT             COMPLETED  UNITS_CUBE                     CUBE
DBMS_SCHEDULER.CREATE_JOB COMPLETED  PRICE_CUBE                     CUBE
DBMS_SCHEDULER.CREATE_JOB COMPLETED  UNITS_CUBE                     CUBE
BUILD                     COMPLETED                                 BUILD
LOAD                      COMPLETED  PRICE_CUBE                     CUBE
SOLVE                     COMPLETED  PRICE_CUBE                     CUBE
UPDATE/COMMIT             COMPLETED  PRICE_CUBE                     CUBE
BUILD                     COMPLETED                                 BUILD
LOAD                      COMPLETED  UNITS_CUBE                     CUBE
SOLVE                     COMPLETED  UNITS_CUBE                     CUBE
UPDATE/COMMIT             COMPLETED  UNITS_CUBE                     CUBE
ANALYZE                   COMPLETED  PRICE_CUBE                     CUBE
ANALYZE                   COMPLETED  UNITS_CUBE                     CUBE
THAW                      COMPLETED                                 BUILD

31 rows selected.
```

Examples

This example uses the default parameters to build UNITS_CUBE.

```
EXECUTE DBMS_CUBE.BUILD('GLOBAL.UNITS_CUBE');
```

The next example builds UNITS_CUBE and explicitly builds two of its dimensions, TIME and CHANNEL. The dimensions use the complete (C) method, and the cube uses the fast solve (S) method.

```
BEGIN
  DBMS_CUBE.BUILD(
    script=>'GLOBAL."TIME", GLOBAL.CHANNEL, GLOBAL.UNITS_CUBE',
    method=>'CCS',
    parallelism=>2);
END;
/
```

The following example loads only the selection of data identified by the WHERE clause:

```
BEGIN
  DBMS_CUBE.BUILD(q'!
  GLOBAL."TIME",
  GLOBAL.CHANNEL,
  GLOBAL.CUSTOMER,
  GLOBAL.PRODUCT,
  GLOBAL.UNITS_CUBE USING (LOAD NO SYNCH
     WHERE UNITS_FACT.MONTH_ID LIKE '2006%'
     AND UNITS_FACT.SALES > 5000)!');
END;
/
```

FOR Clause Example

In this example, the Time dimension is partitioned by calendar year, and `DBMS_CUBE` builds only the partition identified by `CY2006`. The `HIER_ANCESTOR` is an analytic function in the OLAP expression syntax.

```
BEGIN
   dbms_cube.build(q'!
   UNITS_CUBE USING
   (
   FOR "TIME"
     WHERE HIER_ANCESTOR(WITHIN "TIME".CALENDAR LEVEL "TIME".CALENDAR_YEAR) = 'CY2006'
     BUILD (LOAD, SOLVE)
   )!',
   parallelism=>1);
END;
/
```

The next example uses a `FOR` clause to limit the build to the `SALES` measure in 2006. All objects are built using the complete (`C`) method.

```
BEGIN
  DBMS_CUBE.BUILD(
  script => '
  GLOBAL."TIME",
  GLOBAL.CHANNEL,
  GLOBAL.CUSTOMER,
  GLOBAL.PRODUCT,
  GLOBAL.UNITS_CUBE USING
  (
    FOR MEASURES(GLOBAL.UNITS_CUBE.SALES)
      BUILD(LOAD NO SYNCH WHERE GLOBAL.UNITS_FACT.MONTH_ID LIKE ''2006%'')
   )',
  method => 'C',
  parallelism => 2);
END;
/
```

Write-Back Examples

The following examples show various use of the `SET` command in a `USING` clause.

This example sets Sales Target to Sales increased by 5%:

```
DBMS_CUBE.BUILD('UNITS_CUBE USING(
   SET UNITS_CUBE.SALES_TARGET = UNITS_CUBE.SALES * 1.05, SOLVE)');
```

This example sets the price of the Deluxe Mouse in May 2007 to $29.99:

```
DBMS_CUBE.BUILD('PRICE_CUBE USING(
  SET PRICE_CUBE.UNIT_PRICE["TIME"=''2007.05'', "PRODUCT"=''DLX MOUSE'']
  = 29.99, SOLVE)');
```

The next example contains two SET commands, but does not reaggregate the cube:

```
DBMS_CUBE.BUILD('PRICE_CUBE USING(
   SET PRICE_CUBE.UNIT_PRICE["TIME"=''2006.12'', "PRODUCT"=''DLX MOUSE'']
   = 29.49,
   SET PRICE_CUBE.UNIT_PRICE["TIME"=''2007.05'', "PRODUCT"=''DLX MOUSE'']
   = 29.99)');
```

Dimension Maintenance Example

This script shows dimension maintenance. It adds a new dimension member named `OPT MOUSE` to all hierarchies, alters its position in the Primary hierarchy, assigns it a long description, then deletes it from the dimension.

```
BEGIN
dbms_output.put_line('Add optical mouse');
dbms_cube.build(q'!
   "PRODUCT" using (MERGE INTO ALL HIERARCHIES
   VALUES ('ITEM_OPT MOUSE', 'CLASS_SFT', "PRODUCT"."FAMILY"))
!');

dbms_output.put_line('Alter optical mouse');
dbms_cube.build(q'!
   "PRODUCT" using (UPDATE HIERARCHIES ("PRODUCT"."PRIMARY")
   SET PARENT = 'FAMILY_ACC', LEVEL = "PRODUCT"."ITEM"
   WHERE MEMBER = 'ITEM_OPT MOUSE')
!');

dbms_output.put_line('Provide attributes to optical mouse');
dbms_cube.build(q'!
   "PRODUCT" USING (SET "PRODUCT"."LONG_DESCRIPTION"["PRODUCT" = 'ITEM_OPT MOUSE']
    = CAST('Optical Mouse' AS VARCHAR2))
!');

dbms_output.put_line('Delete optical mouse');
dbms_cube.build(q'!
   "PRODUCT" USING (DELETE FROM DIMENSION WHERE MEMBER='ITEM_OPT MOUSE')
!');

END;
/
```

OLAP DML Example

This example uses the OLAP DML to add comments to the cube build log:

```
BEGIN
 DBMS_CUBE.BUILD(q'!
  global.units_cube USING (
   EXECUTE OLAP DML 'SHOW STATLEN(units_cube_prt_list)' PARALLEL,
   EXECUTE OLAP DML 'SHOW LIMIT(units_cube_prt_list KEEP ALL)' PARALLEL,
   EXECUTE OLAP DML 'SHOW STATLEN(time)' parallel,
   EXECUTE OLAP DML 'SHOW LIMIT(time KEEP time_levelrel ''CALENDAR_YEAR'')' parallel)!',
  parallelism=>2,
  add_dimensions=>false);
END;
/
```

This query shows the comments in the cube build log:

```
SELECT partition, slave_number, TO_CHAR(output) output
   FROM cube_build_log
   WHERE command = 'OLAP DML'
   AND status = 'COMPLETED'
   ORDER BY slave_number, time;

PARTITION    SLAVE_NUMBER OUTPUT
------------ ------------ ------------------------------------------------------
P10:CY2007              1 <OLAPDMLExpression
                            Expression="TO_CHAR(statlen(units_cube_prt_list))"
                            Value="1"/>

P10:CY2007              1 <OLAPDMLExpression
                            Expression="TO_CHAR(limit(units_cube_prt_list keep al
                          l))"
                            Value="P10"/>

P10:CY2007              1 <OLAPDMLExpression
                            Expression="TO_CHAR(statlen(time))"
                            Value="17"/>

P10:CY2007              1 <OLAPDMLExpression
                            Expression="TO_CHAR(limit(time keep time_levelrel &ap
                          os;CALENDAR_YEAR&apos;))"
                            Value="CALENDAR_YEAR_CY2007"/>

P9:CY2006               2 <OLAPDMLExpression
                            Expression="TO_CHAR(statlen(units_cube_prt_list))"
                            Value="1"/>

P9:CY2006               2 <OLAPDMLExpression
                            Expression="TO_CHAR(limit(units_cube_prt_list keep al
                          l))"
                            Value="P9"/>

P9:CY2006               2 <OLAPDMLExpression
                            Expression="TO_CHAR(statlen(time))"
                            Value="17"/>
               .
               .
               .
```

54.4.2 CREATE_EXPORT_OPTIONS Procedure

This procedure creates an input XML document that describes processing options for the EXPORT_XML Procedure and the EXPORT_XML_TO_FILE Procedure.

Syntax

```
DBMS_CUBE.CREATE_EXPORT_OPTIONS (
          out_options_xml        IN/OUT  CLOB,
          target_version         IN      VARCHAR2  DEFAULT NULL,
          suppress_owner         IN      BOOLEAN   DEFAULT FALSE,
          suppress_namespace     IN      BOOLEAN   DEFAULT FALSE,
          preserve_table_owners  IN      BOOLEAN   DEFAULT FALSE,
          metadata_changes       IN      CLOB      DEFAULT NULL);
```

Parameters

Table 54-4 CREATE_EXPORT_OPTIONS Procedure Parameters

Parameter	Description
`out_options_xml`	Contains the generated XML document, which can be passed into the `options_xml` parameter of the EXPORT_XML Procedure.
`target_version`	Specifies the version of Oracle Database in which the XML document generated by EXPORT_XML or EXPORT_XML_TO_FILE will be imported. You can specify two to five digits, such as 12.1 or 12.1.0.1.0. This parameter defaults to the current database version, and so can typically be omitted.
`suppress_owner`	Controls the use of the Owner attribute in XML elements and the owner qualifier in object names. Enter `True` to drop the owner from the XML, or enter `False` to retain it. Enter `True` if you plan to import the exported metadata into a different schema.
`suppress_namespace`	Controls the use of Namespace attributes in XML elements and the namespace qualifier in object names. Enter `True` to drop the namespace from the XML, or enter `False` to retain it (default). Enter `True` when upgrading to Oracle OLAP 12*c* metadata. Namespaces allow objects created in Oracle 10*g* to coexist with objects created in Oracle 12*c*. You cannot set or change namespaces.
`preserve_table_owners`	Controls the use of the owner in qualifying table names in the mapping elements, such as GLOBAL.UNITS_HISTORY_FACT instead of UNITS_HISTORY_FACT. Enter `True` to retain the table owner, or enter `False` to default to the current schema for table mappings. If you plan to import the exported metadata to a different schema, you must set this option to `True` to load data from tables and views in the original schema, unless the destination schema has its own copies of the tables and views.
`metadata_changes`	Contains an 12*c* XML description of an object that overwrites the exported object description. The XML document must contain all parent XML elements of the modified element with the attributes needed to uniquely identify them. Use the Name attribute if it exists. See the Examples.

Examples

The following example generates an XML document of export options:

```
DECLARE
   optionsClob  CLOB;

BEGIN
   dbms_lob.createtemporary(optionsClob, false, dbms_lob.CALL);
   dbms_cube.create_export_options(out_options_xml=>optionsClob,
suppress_namespace=>TRUE);
   dbms_output.put_line(optionsClob);
END;
/
```

The `DBMS_OUTPUT.PUT_LINE` procedure displays this XML document (formatted for readability:

```
<?xml version="1.0"?>
<Export TargetVersion="12.1.0.1">
  <ExportOptions>
```

```
    <Option Name="SuppressOwner" Value="FALSE"/>
    <Option Name="SuppressNamespace" Value="TRUE"/>
    <Option Name="PreserveTableOwners" Value="FALSE"/>
  </ExportOptions>
</Export>
```

The next example generates an XML document with a metadata change to the mapping of the American long description attribute of the `CHANNEL` dimension.

```
DECLARE
  importClob         clob;
  exportClob         clob;
  overClob           clob;
  exportOptClob      clob;
  importOptClob      clob;

BEGIN
  dbms_lob.createtemporary(overClob, TRUE);
  dbms_lob.open(overClob, DBMS_LOB.LOB_READWRITE);
  dbms_lob.writeappend(overClob,58, '<Metadata Version="1.3"
MinimumDatabaseVersion="12.1.0.1">');
  dbms_lob.writeappend(overClob,34, '<StandardDimension Name="CHANNEL">');
  dbms_lob.writeappend(overClob,75, '<Description Type="Description" Language="AMERICAN"
Value="Sales Channel"/>');
  dbms_lob.writeappend(overClob,20, '</StandardDimension>');
  dbms_lob.writeappend(overClob,11, '</Metadata>');
  dbms_lob.close(overClob);

  -- Enable Oracle Database 12c Release 1 (12.1) clients to access 10g metadata
  dbms_cube.initialize_cube_upgrade;

  -- Create a CLOB containing the export options
  dbms_lob.createtemporary(exportOptClob, TRUE);
  dbms_cube.create_export_options(out_options_xml=>exportOptClob,
suppress_namespace=>TRUE, metadata_changes=>overClob);

  -- Create a CLOB containing the import options
  dbms_lob.createtemporary(importOptClob, TRUE);
  dbms_cube.create_import_options(out_options_xml=>importOptClob, rename_table =>
'MY_OBJECT_MAP');

   -- Create CLOBs for the metadata
  dbms_lob.createtemporary(importClob, TRUE);
  dbms_lob.createtemporary(exportClob, TRUE);

  -- Export metadata from a 10g analytic workspace to a CLOB
  dbms_cube.export_xml(object_ids=>'GLOBAL_AW', options_xml=>exportOptClob,
out_xml=>exportClob);

  -- Import metadata from the CLOB
  dbms_cube.import_xml(in_xml => exportClob, options_xml=>importOptClob,
out_xml=>importClob);

  -- Load and aggregate the data
  dbms_cube.build(script=>'UNITS_CUBE, PRICE_AND_COST_CUBE');

END;
/
```

The following is the content of `exportClob` (formatting added for readability). The XML document changes the description of Channel to Sales Channel.

```
<Metadata Version="1.3" MinimumDatabaseVersion="12.1.0.1">
  <StandardDimension Name="CHANNEL">
    <Description Type="Description" Language="AMERICAN" Value="Sales Channel"/>
  </StandardDimension>
</Metadata>
```

Related Topics

- EXPORT_XML Procedure
  This procedure writes OLAP metadata to a CLOB.
- EXPORT_XML_TO_FILE Procedure
  This procedure exports OLAP metadata to a file. This file can be imported into a new or existing analytic workspace using the `IMPORT_XML` procedure. In this way, you can create a copy of the analytic workspace in another schema or database.

54.4.3 CREATE_IMPORT_OPTIONS Procedure

This procedure creates an input XML document that describes processing options for the IMPORT_XML Procedure.

Syntax

```
DBMS_CUBE.CREATE_IMPORT_OPTIONS (
          out_options_xml  IN/OUT  CLOB,
          validate_only    IN      BOOLEAN   DEFAULT FALSE,
          rename_table     IN      VARCHAR2  DEFAULT NULL);
```

Parameters

Table 54-5 CREATE_IMPORT_OPTIONS Procedure Parameters

Parameter	Description
`out_options_xml`	Contains the generated XML document, which can be passed to the `options_xml` parameter of the IMPORT_XML Procedure.
`validate_only`	`TRUE` causes the `IMPORT_XML` procedure to validate the metadata described in the input file or the `in_xml` parameter, without committing the changes to the metadata.
`rename_table`	The name of a table identifying new names for the imported objects, in the form [*schema_name*.]*table_name*. The `IMPORT_XML` procedure creates objects using the names specified in the table instead of the ones specified in the XML document. See the Usage Notes for the format of the rename table.

Usage Notes

See the information about using a rename table in DBMS_CUBE - Upgrading 10g Analytic Workspaces.

Examples

This example specifies validation only and a rename table. For an example of the import CLOB being used in an import, see "IMPORT_XML Procedure".

```
DECLARE
importClob  clob;

BEGIN
```

```
  dbms_lob.createtemporary(importClob, TRUE);

  dbms_cube.create_import_options(out_options_xml => importClob, rename_table =>
'MY_OBJECT_MAP', validate_only => TRUE);

  dbms_output.put_line(importClob);
END;
/
```

It generates the following XML document:

```
<?xml version="1.0"?>
<Import>
  <ImportOptions>
    <Option Name="ValidateOnly" Value="TRUE"/>
    <Option Name="RenameTable" Value="MY_OBJECT_MAP"/>
  </ImportOptions>
</Import>
```

Related Topics

- IMPORT_XML Procedure
  This procedure creates, modifies, or drops an analytic workspace by using an XML template.

54.4.4 CREATE_MVIEW Function

This function creates a cube materialized view from the definition of a relational materialized view.

Syntax

```
DBMS_CUBE.CREATE_MVIEW (
          mvowner        IN  VARCHAR2,
          mvname         IN  VARCHAR2,
          sam_parameters IN  CLOB  DEFAULT NULL)
     RETURN VARCHAR2;
```

Parameters

Table 54-6 CREATE_MVIEW Function Parameters

Parameter	Description
mvowner	Owner of the relational materialized view.
mvname	Name of the relational materialized view. For restrictions, see "Requirements for the Relational Materialized View". A single cube materialized view can replace many of the relational materialized views for a table. Choose the materialized view that has the lowest levels of the dimension hierarchies that you want represented in the cube materialized view.
sam_parameters	Parameters in the form '*parameter1=value1*, *parameter2=value2*,...'. See "SQL Aggregation Management Parameters".

SQL Aggregation Management Parameters

The CREATE_MVIEW and DERIVE_FROM_MVIEW functions use the SQL aggregation management (SAM) parameters described in Table 54-7. Some parameters support the development of cubes with advanced analytics. Other parameters support the development of Java

applications. The default settings are appropriate for cube materialized views that are direct replacements for relational materialized views.

Table 54-7 SQL Aggregation Management Parameters

Parameter	Description
ADDTOPS	Adds a top level and a level member to every dimension hierarchy in the cube. If the associated relational dimension has no hierarchy, then a dimension hierarchy is created. TRUE: Creates levels named ALL_*dimension* with level members All_*dimension*. (Default) FALSE: Creates only the hierarchies and levels identified by the relational dimensions.
ADDUNIQUEKEYPREFIX	Controls the creation of dimension keys. TRUE: Creates cube dimension keys by concatenating the level name with the relational dimension key. This practice assures that the dimension keys are unique across all levels, such as CITY_NEW_YORK and STATE_NEW_YORK. (Default) FALSE: Uses the relational dimension keys as cube dimension keys.
ATRMAPTYPE	Specifies whether attributes are mapped by hierarchy levels, dimension levels, or both. HIER_LEVEL: Maps attributes to the levels of a particular dimension hierarchy. (Default) DIM_LEVEL: Maps attributes to the levels of the dimension regardless of hierarchy. BOTH: Maps attributes to both dimension and hierarchy levels. AUTO: Maps attributes to the levels of the dimension for a star schema and to the levels of a particular dimension hierarchy for a snowflake schema.
AWNAME	Provides the name of the analytic workspace that owns the cube. Choose a simple database object name of 1 to 30 bytes. The default name is *fact_tablename*_AW*n*.
BUILD	Specifies whether a data refresh will immediately follow creation of the cube materialized view. IMMEDIATE: Refreshes immediately. DEFERRED: Does not perform a data refresh. (Default) **Note**: Only the CREATE_MVIEW function uses this parameter.

Table 54-7 (Cont.) SQL Aggregation Management Parameters

Parameter	Description
CUBEMVOPTION	Controls validation and creation of a cube materialized view. Regardless of this setting, the function creates an analytic workspace containing a cube and its related cube dimensions. COMPLETE_REFRESH: Creates a complete refresh cube materialized view (full update). FAST_REFRESH: Creates a fast refresh materialized view (incremental update). REWRITE_READY: Runs validation checks for a rewrite cube materialized view, but does not create it. REWRITE: Creates a rewrite cube materialized view. REWRITE_WITH_ATTRIBUTES: Creates a rewrite cube materialized view that includes columns with dimension attributes, resulting in faster query response times. (Default) **Note**: *The following settings do not create a cube materialized view.* Use Analytic Workspace Manager to drop an analytic workspace that does not have a cube materialized view. You can use the DROP_MVIEW procedure to delete an analytic workspace only when it supports a cube materialized view. NONE: Does not create a cube materialized view. COMPLETE_REFRESH_READY: Runs validation checks for a complete refresh cube materialized view, but does not create it. FAST_REFRESH_READY: Runs validation checks for fast refresh, but does not create the cube materialized view.
CUBENAME	Provides the name of the cube derived from the relational materialized view. Choose simple database object name of 1 to 30 bytes. The default name is *fact_tablename_Cn*.
DIMJAVABINDVARS	Supports access by Java programs to the XML document. TRUE: Generates an XML template that uses Java bind variable notation for the names of dimensions. No XML validation is performed. You cannot use the IMPORT_XML procedure to create a cube using this template. FALSE: Generates an XML template that does not support Java bind variables. (Default)
DISABLEQRW	Controls disabling of query rewrite on the source relational materialized view. TRUE: Issues an ALTER MATERIALIZED VIEW *mview_name* DISABLE QUERY REWRITE command. FALSE: No action. **Note**: Only the CREATE_MVIEW function with BUILD=IMMEDIATE uses this parameter.
EXPORTXML	Exports the XML that defines the dimensional objects to a file, which you specify as *dir/filename*. Both the directory and the file name are case sensitive. *dir*: Name of a database directory. *filename*: The name of the file, typically given an XML filename extension.

Table 54-7 (Cont.) SQL Aggregation Management Parameters

Parameter	Description
FILTERPARTITIONANCESTORLEVELS	Controls the generation of aggregate values above the partitioning level of a partitioned cube. TRUE: Removes levels above the partitioning level from the cube. Requests for summary values above the partitioning level are solved by SQL. FALSE: All levels are retained in the cube. Requests for summary values are solved by OLAP. (Default)
LOGDEST	Directs and stores log messages. By default, the messages are not available. SERVEROUT: Sends messages to server output (typically the screen), which is suitable when working interactively such as in SQL*Plus or SQL Developer. TRACEFILE: Sends messages to the session trace file.
PARTITIONOPTION	Controls partitioning of the cube. NONE: Prevents partitioning. DEFAULT: Allows the Sparsity Advisor to determine whether partitioning is needed and how to partition the cube. (Default) FORCE: Partitions the cube even when the Sparsity Advisor recommends against it. The Sparsity Advisor identifies the best dimension, hierarchy, and level to use for partitioning. *dimension.hierarchy.level*: Partitions the cube using the specified dimension, hierarchy, and level.
POPULATELINEAGE	Controls the appearance of attributes in a cube materialized view. TRUE: Includes all dimension attributes in the cube materialized view. (Default) FALSE: Omits all dimension attributes from the cube materialized view.
PRECOMPUTE	Identifies a percentage of the data that is aggregated and stored. The remaining values are calculated as required by queries during the session. *precompute_percentage*[:*precompute_top_percentage*] Specify the top percentage for partitioned cubes. The default value is 35:0, which specifies precomputing 35% of the bottom partition and 0% of the top partition. If the cube is not partitioned, then the second number is ignored.
REMAPCOMPOSITEKEYS	Controls how multicolumn keys are rendered in the cube. TRUE: Creates a unique key attribute whose values are concatenated string expressions with an underscore between the column values. For example, the value BOSTON_MA_USA might be an expression produced from a multicolumn key composed of CITY, STATE, and COUNTRY columns. In addition, an attribute is created for each individual column to store the relational keys. (Default) FALSE: Creates a unique key attribute for each column.

Table 54-7 (Cont.) SQL Aggregation Management Parameters

Parameter	Description
RENDERINGMODE	Controls whether a loss in fidelity between the relational materialized view and the cube materialized view results in a warning message or an exception. See "Requirements for the Relational Materialized View". LOOSE: Losses are noted in the optional logs generated by the CREATE_MVIEW Function and the DERIVE_FROM_MVIEW Function. No exceptions are raised. (Default) STRICT: Any loss in fidelity raises an exception so that no XML template is created.
SEEFILTERS	Controls whether conditions in the WHERE clause of the relational materialized view's defining query are retained or ignored. TRUE: Renders valid conditions in the XML template. (Default) FALSE: Ignores all conditions.
UNIQUENAMES	Controls whether top level dimensional objects have unique names. Cross namespace conflicts may occur because dimensional objects have different namespaces than relational objects. TRUE: Modifies all relational names when they are rendered in the cube.(Default) FALSE: Duplicates relational names in the cube unless a naming conflict is detected. In that case, a unique name is created.
UNKNOWNKEYASDIM	Controls handling of simple columns with no levels or hierarchies in the GROUP BY clause of the relational materialized view's defining query. TRUE: Renders a simple column without a relational dimension as a cube dimension with no levels or hierarchies. FALSE: Raises an exception when no relational dimension is found for the column. (Default)
VALIDATEXML	Controls whether the generated XML document is validated. TRUE: Validates the template using the VALIDATE_XML procedure. (Default) FALSE: No validation is done.

Returns

The name of the cube materialized view created by the function.

Usage Notes

See "Using SQL Aggregation Management"

Examples

All examples for the SQL Aggregate Management subprograms use the sample Sales History schema, which is installed in Oracle Database with two relational materialized views: CAL_MONTH_SALES_MV and FWEEK_PSCAT_SALES_MV.

The following script creates a cube materialized view using CAL_MONTH_SALES_MV as the relational materialized view. It uses all default options.

```
SET serverout ON format wrapped

DECLARE
     salesaw  varchar2(30);
```

```
BEGIN
     salesaw := dbms_cube.create_mview('SH', 'CAL_MONTH_SALES_MV');
END;
/
```

The next example sets several parameters for creating a cube materialized view from FWEEK_PSCAT_SALES_MV. These parameters change the cube materialized view in the following ways:

- ADDTOPS: Adds a top level consisting of a single value to the hierarchies. All of the dimensions in Sales History have a top level already.
- PRECOMPUTE: Changes the percentage of materialized aggregates from 35:0 to 40:10.
- EXPORTXML: Creates a text file for the XML document.
- BUILD: Performs a data refresh.

```
DECLARE
     salescubemv   varchar2(30);
     sam_param     clob := 'ADDTOPS=FALSE,
                            PRECOMPUTE=40:10,
                            EXPORTXML=WORK_DIR/sales.xml,
                            BUILD=IMMEDIATE';

BEGIN
     salescubemv := dbms_cube.create_mview('SH', 'FWEEK_PSCAT_SALES_MV',
                    sam_param);
END;
/
```

54.4.5 DERIVE_FROM_MVIEW Function

This function generates an XML template that defines a cube with materialized view capabilities, using the information derived from an existing relational materialized view.

Syntax

```
DBMS_CUBE.DERIVE_FROM_MVIEW (
          mvowner        IN  VARCHAR2,
          mvname         IN  VARCHAR2,
          sam_parameters IN  CLOB  DEFAULT NULL)
     RETURN CLOB;
```

Parameters

Table 54-8 DERIVE_FROM_MVIEW Function Parameters

Parameter	Description
mvowner	Owner of the relational materialized view.
mvname	Name of the relational materialized view. For restrictions, see "Requirements for the Relational Materialized View". A single cube materialized view can replace many of the relational materialized views for a table. Choose the materialized view that has the lowest levels of the dimension hierarchies that you want represented in the cube materialized view.

Table 54-8 (Cont.) DERIVE_FROM_MVIEW Function Parameters

Parameter	Description
sam_parameters	Optional list of parameters in the form '*parameter1=value1*, *parameter2=value2*,...'. See "SQL Aggregation Management Parameters".

Returns

An XML template that defines an analytic workspace containing a cube enabled as a materialized view.

Usage Notes

To create a cube materialized view from an XML template, use the IMPORT_XML procedure. Then use the REFRESH_MVIEW procedure to refresh the cube materialized view with data.

See "Using SQL Aggregation Management".

Examples

The following example generates an XML template named sales_cube.xml from the CAL_MONTH_SALES_MV relational materialized view in the SH schema.

```
DECLARE
     salescubexml  clob := null;
     sam_param     clob := 'exportXML=WORK_DIR/sales_cube.xml';

BEGIN
     salescubexml := dbms_cube.derive_from_mview('SH', 'CAL_MONTH_SALES_MV',
     sam_param);
END;
/
```

54.4.6 DROP_MVIEW Procedure

This procedure drops a cube materialized view and all associated objects from the database. These objects include the dimension materialized views, cubes, cube dimensions, levels, hierarchies, and the analytic workspace.

Syntax

```
DBMS_CUBE.DROP_MVIEW (
          mvowner        IN  VARCHAR2,
          mvname         IN  VARCHAR2,
          sam_parameters IN  CLOB  DEFAULT NULL);
```

Parameters

Table 54-9 DROP_MVIEW Procedure Parameters

Parameter	Description
mvowner	Owner of the cube materialized view
mvname	Name of the cube materialized view

Table 54-9 (Cont.) DROP_MVIEW Procedure Parameters

Parameter	Description
sam_parameters	EXPORTXML: Exports the XML that drops the dimensional objects to a file, which you specify as *dir/filename*. Both the directory and the file name are case sensitive. *dir*: Name of a database directory. *filename*: The name of the file, typically given an XML filename extension.

Usage Notes

Use this procedure to drop a cube materialized view that you created using the CREATE_MVIEW and DERIVE_FROM_MVIEW functions. If you make modifications to the cubes or dimensions, then DROP_MVIEW may not be able to drop the cube materialized view.

Some of the CUBEMVOPTION parameters used by the CREATE_MVIEW and DERIVE_FROM_MVIEW functions do not create a materialized view. Use Analytic Workspace Manager to drop the analytic workspace, cubes, and cube dimensions.

If you use the EXPORTXML parameter, then you can use the XML document to drop the cube materialized view, after you re-create it. Use the IMPORT_XML procedure.

See "Using SQL Aggregation Management".

Examples

The current schema has four materialized views. CB$CAL_MONTH_SALES is a cube materialized view for the SALES table. CB$TIMES_DIM_D1_CAL_ROLLUP is a cube dimension materialized view for the TIMES_DIM dimension on the TIMES dimension table. The others are relational materialized views.

```
SELECT mview_name FROM user_mviews;

MVIEW_NAME
------------------------------
CB$CAL_MONTH_SALES
CB$TIMES_DIM_D1_CAL_ROLLUP
CAL_MONTH_SALES_MV
FWEEK_PSCAT_SALES_MV
```

The following command drops both CB$CAL_MONTH_SALES and CB$TIMES_DIM_D1_CAL_ROLLUP.

```
EXECUTE dbms_cube.drop_mview('SH', 'CB$CAL_MONTH_SALES');

Dropped cube organized materialized view "SH"."CAL_MONTH_SALES"
including container analytic workspace "SH"."CAL_MONTH_SALES_AW"
at 20130213 16:31:40.056.
```

This query against the data dictionary confirms that the materialized views have been dropped.

```
SELECT mview_name FROM user_mviews;

MVIEW_NAME
------------------------------
CAL_MONTH_SALES_MV
FWEEK_PSCAT_SALES_MV
```

54.4.7 EXPORT_XML Procedure

This procedure writes OLAP metadata to a CLOB.

Syntax

```
DBMS_CUBE.EXPORT_XML
       (object_ids            IN      VARCHAR2,
        out_xml               IN/OUT  CLOB;

DBMS_CUBE.EXPORT_XML
       (object_ids            IN      VARCHAR2,
        options_xml           IN      CLOB,
        out_xml               IN/OUT  CLOB;

DBMS_CUBE.EXPORT_XML
       (object_ids            IN      VARCHAR2,
        options_dirname       IN      VARCHAR2,
        options_filename      IN      VARCHAR2,
        out_xml               IN/OUT  CLOB;
```

Parameters

Table 54-10 EXPORT_XML Procedure Parameters

Parameter	Description
object_ids	Any of these identifiers. • The name of a schema, such as GLOBAL. • The fully qualified name of an analytic workspace in the form *owner.aw_name*.AW, such as GLOBAL.GLOBAL.AW. • Cube • Dimension • Named build process • Measure folder You can specify multiple objects by separating the names with commas. **Note**: When exporting an individual object, be sure to export any objects required to reconstruct it. For example, when exporting a cube, you must also export the dimensions of the cube.
options_dirname	The case-sensitive name of a database directory that contains *options_filename*.
options_filename	A file containing an XML document of export options.
options_xml	A CLOB variable that contains an XML document of export options. Use the CREATE_EXPORT_OPTIONS Procedure to generate this document.
out_xml	A CLOB variable that will store the XML document of OLAP metadata for the objects listed in *object_ids*.

Export Options

The default settings for the export options are appropriate in many cases, so you can omit the *options_xml* parameter or the *options_dirname* and *options_filename* parameters. However, when upgrading Oracle OLAP 10*g* metadata to OLAP 12*c*, you must specify an XML document

that changes the default settings. This example changes all of the parameters from False to True; set them appropriately for your schema.

```
<?xml version="1.0"?>
<Export>
  <ExportOptions>
    <Option Name="SuppressNamespace" Value="True"/>
    <Option Name="SuppressOwner" Value="True"/>
    <Option Name="PreserveTableOwners" Value="True"/>
  </ExportOptions>
</Export>
```

You can create this XML document manually or by using the CREATE_EXPORT_OPTIONS Procedure.

Usage Notes

See "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

Example

For an example of using `EXPORT_XML` in an upgrade to the same schema, see "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

The following PL/SQL script copies an OLAP 12*c* analytic workspace named `GLOBAL12` from the `GLOBAL_AW` schema to the `GLOBAL` schema. No upgrade is performed.

To upgrade into a different schema, change the example as follows:

- Call the `INITIALIZE_CUBE_UPGRADE` procedure.
- Call the `CREATE_EXPORT_OPTIONS` procedure with the additional parameter setting `SUPPRESS_NAMESPACE=>TRUE`.

The PL/SQL client must be connected to the database as `GLOBAL`. The `GLOBAL` user must have `SELECT` permissions on `GLOBAL_AW.AW$GLOBAL` and on all relational data sources.

```
BEGIN
  -- Create a CLOB for the export options
  dbms_lob.createtemporary(optionsClob, TRUE);
  dbms_cube.create_export_options(out_options_xml=>optionsClob, suppress_owner=>TRUE,
preserve_table_owners=>TRUE);

  -- Create a CLOB for the XML template
  dbms_lob.createtemporary(exportClob, TRUE);

  -- Export metadata from an analytic workspace to a CLOB
  dbms_cube.export_xml(object_ids=>'GLOBAL_AW.GLOBAL12.AW', options_xml=>optionsClob,
out_xml=>exportClob);

  -- Import metadata from the CLOB
  dbms_cube.import_xml(in_xml=>exportClob);

  -- Load and aggregate the data
  dbms_cube.build(script=>'GLOBAL.UNITS_CUBE, GLOBAL.PRICE_AND_COST_CUBE');

END;
/
```

54.4.8 EXPORT_XML_TO_FILE Procedure

This procedure exports OLAP metadata to a file. This file can be imported into a new or existing analytic workspace using the `IMPORT_XML` procedure. In this way, you can create a copy of the analytic workspace in another schema or database.

This procedure can also be used as part of the process for upgrading OLAP standard form metadata contained in an Oracle OLAP 10*g* analytic workspace to OLAP 12*c* format.

Syntax

```
DBMS_CUBE.EXPORT_XML_TO_FILE
       (object_ids             IN      VARCHAR2,
        output_dirname         IN      VARCHAR2,
        output_filename        IN      VARCHAR2;

DBMS_CUBE.EXPORT_XML_TO_FILE
       (object_ids             IN      VARCHAR2,
        options_dirname        IN      VARCHAR2,
        options_filename       IN      VARCHAR2,
        output_dirname         IN      VARCHAR2,
        output_filename        IN      VARCHAR2;
```

Parameters

Table 54-11 EXPORT_XML_TO_FILE Procedure Parameters

Parameter	Description
`object_ids`	Any of these identifiers. • The name of a schema, such as `GLOBAL`. • The fully qualified name of an analytic workspace in the form *`owner.aw_name`*`.AW`, such as `GLOBAL.GLOBAL.AW`. • Cube • Dimension • Named build process • Measure folder You can specify multiple objects by separating the names with commas. **Note**: When exporting an individual object, be sure to export any objects required to reconstruct it. For example, when you export a cube, you must also export the dimensions of the cube.
`options_dirname`	The case-sensitive name of a database directory that contains `options_filename`. See "Export Options".
`options_filename`	The name of a file containing an XML document of export options. See "Export Options".
`output_dirname`	The case-sensitive name of a database directory where *`output_filename`* is created.
`output_filename`	The name of the template file created by the procedure.

Export Options

The default settings for the export options are appropriate in most cases, and you can omit the *`options_dirname`* and *`options_filename`* parameters. However, when upgrading Oracle OLAP

10*g* metadata to OLAP 12*c*, you must specify an XML document that changes the default settings, like the following:

```
<?xml version="2.0"?>
<Export>
  <ExportOptions>
    <Option Name="SuppressNamespace" Value="True"/>
    <Option Name="SuppressOwner" Value="True"/>
    <Option Name="PreserveTableOwners" Value="True"/>
  </ExportOptions>
</Export>
```

You can create this XML document manually or by using the CREATE_EXPORT_OPTIONS Procedure.

Usage Notes

See "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

Examples

The following example generates an XML file named global.xml in OLAP 12*c* format using the default export settings. The metadata is derived from all analytic workspaces and CWM metadata in the `GLOBAL_AW` schema. The output file is generated in the `WORK_DIR` database directory.

```
execute dbms_cube.export_xml_to_file('GLOBAL_AW', 'WORK_DIR', 'global.xml');
```

The next example also generates an XML file named global.xml in OLAP 12*c* format using the export options set in options.xml. The metadata is derived from the GLOBAL analytic workspace in the `GLOBAL_AW` schema. Both the options file and the output file are in the `WORK_DIR` database directory.

```
execute dbms_cube.export_xml_to_file('GLOBAL_AW.GLOBAL.AW', 'WORK_DIR', 'options.xml',
'WORK_DIR', 'global.xml');
```

54.4.9 IMPORT_XML Procedure

This procedure creates, modifies, or drops an analytic workspace by using an XML template.

Syntax

```
DBMS_CUBE.IMPORT_XML
       (dirname               IN      VARCHAR2,
        filename              IN      VARCHAR2 );

DBMS_CUBE.IMPORT_XML
       (dirname               IN      VARCHAR2,
        filename              IN      VARCHAR2,
        out_xml               IN/OUT  CLOB );

DBMS_CUBE.IMPORT_XML
       (input_dirname         IN      VARCHAR2,
        input_filename        IN      VARCHAR2
        options_dirname       IN      VARCHAR2,
        options_filename      IN      VARCHAR2,
        out_xml               IN/OUT  CLOB );

DBMS_CUBE.IMPORT_XML
       (in_xml                IN      CLOB );
```

```
DBMS_CUBE.IMPORT_XML
       (in_xml                IN      CLOB,
        out_xml               IN/OUT  CLOB );

DBMS_CUBE.IMPORT_XML

       (in_xml                IN      CLOB,
        options_xml           IN      CLOB,
        out_xml               IN/OUT  CLOB );
```

Parameters

Table 54-12 IMPORT_XML Procedure Parameters

Parameter	Description
dirname	The case-sensitive name of a database directory containing the XML document describing an analytic workspace.
filename	A file containing an XML document describing an analytic workspace.
in_xml	A CLOB containing an XML document describing an analytic workspace.
input_dirname	The case-sensitive name of a database directory containing the XML document describing an analytic workspace.
input_filename	A file containing an XML document describing an analytic workspace.
options_dirname	The case-sensitive name of a database directory containing a file of import options.
options_filename	A file of import options.
options_xml	An XML document describing the import options. Use the CREATE_IMPORT_OPTIONS Procedure to generate this document.
out_xml	An XML document that either describes the analytic workspace or, for validation only, describes any errors. It may contain changes that DBMS_CUBE made to the imported XML, such as setting default values or making minor corrections to the XML.

Usage Notes

The XML can define, modify, or drop an entire analytic workspace, or one or more cubes or dimensions. When defining just cubes or dimensions, you must do so within an existing analytic workspace.

You can also use IMPORT_XML to drop an analytic workspace by using the XML document generated by the DROP_MVIEW procedure with the EXPORTXML parameter.

See "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

Example

This example loads an XML template from a file named GLOBAL.XML and located in a database directory named XML_DIR.

```
EXECUTE dbms_cube.import_xml('XML_DIR', 'GLOBAL.XML');
```

The next example exports an OLAP 10*g* template and uses IMPORT_XML to validate it before an upgrade to 12*c*.

```
DECLARE
```

```
  exportOptClob clob;
  importOptClob clob;
  importClob    clob;
  exportClob    clob;

BEGIN

  -- Create a CLOB for the export options
  dbms_lob.createtemporary(exportOptClob, TRUE);
  dbms_cube.create_export_options(out_options_xml=>exportOptClob,
suppress_namespace=>TRUE, preserve_table_owners=>TRUE);

  -- Create a CLOB for the XML template
  dbms_lob.createtemporary(exportClob, TRUE);

  -- Create a CLOB for import options
  dbms_lob.createtemporary(importOptClob, TRUE);
  dbms_cube.create_import_options(out_options_xml=>importOptClob, validate_only=>TRUE);

  -- Create a CLOB for the change log
  dbms_lob.createtemporary(importClob, TRUE);

  -- Enable Oracle Database 12c Release 1 (12.1) clients to access 10g metadata
  dbms_cube.initialize_cube_upgrade;

  -- Export metadata from an analytic workspace to a CLOB
  dbms_cube.export_xml(object_ids=>'GLOBAL_AW', options_xml=>exportOptClob,
out_xml=>exportClob);

  /* Import metadata from the CLOB. No objects are committed to the database
     because the validate_only parameter of CREATE_IMPORT_OPTIONS is set to
     TRUE.
  */

  dbms_cube.import_xml(in_xml=>exportClob, options_xml=>importOptClob,
out_xml=>importClob);

   -- Output the metadata changes
  dbms_output.put_line('This is the validation log:');
  dbms_output.put_line(importClob);

END;
/
```

The contents of `importClob` show that the XML is valid. Otherwise, error messages appear in the `<RootCommitResult>` element.

```
This is the validation log:
<?xml version="1.0" encoding="UTF-16"?>
<RootCommitResult>

</RootCommitResult>
```

For an example of `IMPORT_XML` within the context of an upgrade from 10*g* to 12*c* metadata, see the Custom Upgrade section of "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

54.4.10 INITIALIZE_CUBE_UPGRADE Procedure

This procedure processes analytic workspaces created in Oracle OLAP 10*g* so they can be used by Oracle OLAP 12*c* clients. It processes all analytic workspaces in the current schema. Run this procedure once for each schema in which there are 10*g* analytic workspaces.

Without this processing step, 12*c* clients cannot connect to a database containing a 10*g* analytic workspace with subobjects of a dimension or cube having the same name. Additionally, some `DBMS_CUBE` procedures and functions, such as `EXPORT_XML` and `EXPORT_XML_TO_FILE`, do not work on the 10*g* metadata.

After processing, OLAP 12*c* clients can connect and use the alternate names provided by `INITIALIZE_CUBE_UPGRADE` for the conflicting subobjects. OLAP 10*g* clients continue to use the original names.

`INITIALIZE_CUBE_UPGRADE` does not upgrade any OLAP 10*g* objects to OLAP 12*c* format.

See "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

Syntax

```
DBMS_CUBE.INITIALIZE_CUBE_UPGRADE;
```

Usage Notes

This procedure creates and populates a table named `CUBE_UPGRADE_INFO`. If it already exists, the table is truncated and repopulated.

While the 10*g* namespace allowed subobjects with the same name in the same dimension or cube, the 12*c* namespace does not. When `INITIALIZE_CUBE_UPGRADE` detects a name conflict among subobjects such as levels, hierarchies, and dimension attributes, it creates a row in `CUBE_UPGRADE_INFO` providing a new, unique name for each one. Rows may also be created for objects that do not require renaming; these rows are distinguished by a value of 0 or null in the `CONFLICT` column. Top-level objects, such as dimensions and cubes, are not listed.

You can edit the table using SQL `INSERT` and `UPDATE` if you want to customize the names of OLAP 10*g* objects on OLAP 12*c* clients.

The `UPGRADE_AW`, `EXPORT_XML` and `EXPORT_XML_TO_FILE` procedures use the names specified in the `NEW_NAME` column of the table to identify objects in CWM or OLAP standard form (AWXML) analytic workspaces, rather than the original names.

The following table describes the columns of `CUBE_UPGRADE_INFO`.

Column	Datatype	NULL	Description
`OWNER`	`VARCHAR2`	`NOT NULL`	Owner of the analytic workspace.
`AW`	`VARCHAR2`	`NOT NULL`	Name of the analytic workspace.
`AWXML_ID`	`VARCHAR2`	`NOT NULL`	Full logical name of the object requiring modification, in the form *`simple_name.[subtype_name].object_type`*. For example, `TIME.DIMENSION` and `PRODUCT.COLOR.ATTRIBUTE`.
`NEW_NAME`	`VARCHAR2`	`NOT NULL`	The name the object will have in Oracle 12*c* after the upgrade.
`OBJECT_CLASS`	`VARCHAR2`	--	`DerivedMeasure` for calculated measures, or empty for all other object types.

Column	Datatype	NULL	Description
CONFLICT	NUMBER	--	Indicates the reason that the row was added to CUBE_UPGRADE_INFO: • 0: The object does not have a naming conflict but appears in the table for other reasons. • 1: Two objects have the same name and would create a conflict in the OLAP 12*c* namespace. The object type (such as level or hierarchy) will be added to the names.

Examples

The following command creates and populates the CUBE_UPGRADE_INFO table:

```
EXECUTE dbms_cube.initialize_cube_upgrade;
```

The table shows that the OLAP 10*g* analytic workspace has a hierarchy and a level named MARKET_SEGMENT, which will be renamed. The table also contains rows for calculated measures, but these objects do not require renaming: The value of CONFLICT is 0.

```
SELECT awxml_id, new_name, conflict FROM cube_upgrade_info;

AWXML_ID                                   NEW_NAME                    CONFLICT
------------------------------------------ ------------------------- ----------
CUSTOMER.MARKET_SEGMENT.HIERARCHY          MARKET_SEGMENT_HIERARCHY           1
CUSTOMER.MARKET_SEGMENT.LEVEL              MARKET_SEGMENT_LEVEL               1
UNITS_CUBE.EXTENDED_COST.MEASURE           EXTENDED_COST                      0
UNITS_CUBE.EXTENDED_MARGIN.MEASURE         EXTENDED_MARGIN                    0
UNITS_CUBE.CHG_SALES_PP.MEASURE            CHG_SALES_PP                       0
UNITS_CUBE.CHG_SALES_PY.MEASURE            CHG_SALES_PY                       0
UNITS_CUBE.PCTCHG_SALES_PP.MEASURE         PCTCHG_SALES_PP                    0
UNITS_CUBE.PCTCHG_SALES_PY.MEASURE         PCTCHG_SALES_PY                    0
UNITS_CUBE.PRODUCT_SHARE.MEASURE           PRODUCT_SHARE                      0
UNITS_CUBE.CHANNEL_SHARE.MEASURE           CHANNEL_SHARE                      0
UNITS_CUBE.MARKET_SHARE.MEASURE            MARKET_SHARE                       0
UNITS_CUBE.CHG_EXTMRGN_PP.MEASURE          CHG_EXTMRGN_PP                     0
UNITS_CUBE.CHG_EXTMRGN_PY.MEASURE          CHG_EXTMRGN_PY                     0
UNITS_CUBE.PCTCHG_EXTMRGN_PP.MEASURE       PCTCHG_EXTMRGN_PP                  0
UNITS_CUBE.PCTCHG_EXTMRGN_PY.MEASURE       PCTCHG_EXTMRGN_PY                  0
UNITS_CUBE.CHG_UNITS_PP.MEASURE            CHG_UNITS_PP                       0
UNITS_CUBE.EXTMRGN_PER_UNIT.MEASURE        EXTMRGN_PER_UNIT                   0
UNITS_CUBE.SALES_YTD.MEASURE               SALES_YTD                          0
UNITS_CUBE.SALES_YTD_PY.MEASURE            SALES_YTD_PY                       0
UNITS_CUBE.PCTCHG_SALES_YTD_PY.MEASURE     PCTCHG_SALES_YTD_PY                0
UNITS_CUBE.SALES_QTD.MEASURE               SALES_QTD                          0
UNITS_CUBE.CHG_UNITS_PY.MEASURE            CHG_UNITS_PY                       0
```

54.4.11 REFRESH_MVIEW Procedure

This procedure refreshes the data in a cube materialized view.

Syntax

```
DBMS_CUBE.REFRESH_MVIEW (
          mvowner              IN  VARCHAR2,
          mvname               IN  VARCHAR2,
          method               IN  VARCHAR2       DEFAULT NULL,
          refresh_after_errors IN  BOOLEAN        DEFAULT FALSE,
          parallelism          IN  BINARY_INTEGER DEFAULT 0,
```

```
atomic_refresh       IN  BOOLEAN        DEFAULT FALSE,
scheduler_job        IN  VARCHAR2       DEFAULT NULL,
sam_parameters       IN  CLOB           DEFAULT NULL,
nested               IN  BOOLEAN        DEFAULT FALSE );
```

Parameters

Table 54-13 REFRESH_MVIEW Procedure Parameters

Parameter	Description
mvowner	Owner of the cube materialized view.
mvname	Name of the cube materialized view.
method	A full or a fast (partial) refresh. In a fast refresh, only changed rows are inserted in the cube and the affected areas of the cube are re-aggregated. You can specify a method for each cube in sequential order, or a single method to apply to all cubes. If you list more cubes than methods, then the last method applies to the additional cubes. • C: Complete refresh clears all dimension values before loading. (Default) • F: Fast refresh of a cube materialized view, which performs an incremental refresh and re-aggregation of only changed rows in the source table. • ?: Fast refresh if possible, and otherwise a complete refresh. • P: Recomputes rows in a cube materialized view that are affected by changed partitions in the detail tables. • S: Fast solve of a compressed cube. A fast solve reloads all the detail data and re-aggregates only the changed values. See the "Usage Notes" for the BUILD procedure for additional details.
refresh_after_errors	TRUE to roll back just the cube or dimension with errors, and then continue building the other objects. FALSE to roll back all objects in the build.
parallelism	Number of parallel processes to allocate to this job. See the "Usage Notes" for the BUILD procedure for additional details.
atomic_refresh	TRUE prevents users from accessing intermediate results during a build. It freezes the current state of an analytic workspace at the beginning of the build to provide current sessions with consistent data. This option thaws the analytic workspace at the end of the build to give new sessions access to the refreshed data. If an error occurs during the build, then all objects are rolled back to the frozen state. FALSE enables users to access intermediate results during an build.
scheduler_job	Any text identifier for the job, which will appear in the log table. The string does not need to be unique.
sam_parameters	None.
nested	TRUE performs nested refresh operations for the specified set of cube materialized views. Nested refresh operations refresh all the depending materialized views and the specified set of materialized views based on a dependency order to ensure the nested materialized views are truly fresh with respect to the underlying base tables. All objects must reside in a single analytic workspace.

Usage Notes

REFRESH_MVIEW changes *mvname* to the name of the cube, then passes the cube name and all parameters to the BUILD procedure. Thus, you can use the BUILD procedure to refresh a cube materialized view. See the "BUILD Procedure" for additional information about the parameters.

Examples

The following example uses the default settings to refresh a cube materialized view named CB$FWEEK_PSCAT_SALES.

```
SET serverout ON format wrapped

EXECUTE dbms_cube.refresh_mview('SH', 'CB$FWEEK_PSCAT_SALES');
```

The next example changes the refresh method to use fast refresh if possible, continue refreshing after an error, and use two parallel processes.

```
EXECUTE dbms_cube.refresh_mview('SH', 'CB$FWEEK_PSCAT_SALES', '?', TRUE, 2);
```

After successfully refreshing the cube materialized view, REFRESH_MVIEW returns a message like the following:

```
Completed refresh of cube mview "SH"."CB$FWEEK_PSCAT_SALES" at 20130212 15:04:46.370.
```

54.4.12 UPGRADE_AW Procedure

This procedure creates an Oracle OLAP 12*c* analytic workspace from a copy of the metadata contained in an OLAP 10*g* analytic workspace. The original OLAP 10*g* analytic workspace is not affected and can exist at the same time and in the same schema as the OLAP 12*c* analytic workspace.

UPGRADE_AW automatically runs INITIALIZE_CUBE_UPGRADE if the CUBE_UPGRADE_INFO table does not exist. If it does exist, then UPGRADE_AW does not overwrite it, thus preserving any changes you made to the table.

See "DBMS_CUBE - Upgrading 10g Analytic Workspaces".

Syntax

```
DBMS_CUBE.UPGRADE_AW
       (sourceaw              IN  VARCHAR2,
        destaw                IN  VARCHAR2,
        upgoptions            IN  CLOB DEFAULT NULL);
```

Parameters

Table 54-14 UPGRADE_AW Procedure Parameters

Parameter	Description
sourceaw	The name of a 10*g* analytic workspace.
destaw	A new name for the generated 12*c* analytic workspace. It cannot be the same as *sourceaw*.

Table 54-14 (Cont.) UPGRADE_AW Procedure Parameters

Parameter	Description
upgoptions	One or more of these upgrade options, as a string in the form '*OPTION=VALUE*'. Separate multiple options with commas. • PRESERVE_TABLE_OWNERS: YES preserves the original source table mappings. Use this option when creating an OLAP 12*c* analytic workspace in a different schema from the 10*g* analytic workspace, and you want the new objects mapped to tables in the original schema. (Default) NO removes the schema owner from the source table mappings. Use this option when creating an OLAP 12*c* analytic workspace in a different schema from the 10*g* analytic workspace, and you want the new objects mapped to tables in the destination schema. • RENAME_TABLE: The name of a table that specifies new names for objects as they are created in OLAP 12*c* format. These changes are in addition to those specified by the INITIALIZE_CUBE_UPGRADE procedure. See "CREATE_IMPORT_OPTIONS Procedure" for information about creating a rename table. • TARGET_VERSION: The version of the upgrade, specified by a 2- to 5-part number, such as 11.2 or 11.2.0.2.0. If you enter an unsupported version number, then the closest version below it is used.

Examples

This example upgrades an OLAP 10*g* analytic workspace named GLOBAL10 to an OLAP 12*c* analytic workspace named GLOBAL12, using a rename table named MY_OBJECT_MAP:

```
BEGIN

  -- Upgrade the analytic workspace
  dbms_cube.upgrade_aw(sourceaw =>'GLOBAL10', destaw => 'GLOBAL12', upgoptions =>
'RENAME_TABLE=MY_OBJECT_MAP');

  -- Load and aggregate the data
  dbms_cube.build(script=>'UNITS_CUBE, PRICE_AND_COST_CUBE');

END;
/
```

54.4.13 VALIDATE_XML Procedure

This procedure checks the XML to assure that it is valid without committing the results to the database. It does not create an analytic workspace.

Syntax

```
DBMS_CUBE.VALIDATE_XML
       (dirname               IN  VARCHAR2,
        filename              IN  VARCHAR2 );

DBMS_CUBE.VALIDATE_XML
       (in_xml                IN  CLOB );
```

Parameters

Table 54-15 VALIDATE_XML Procedure Parameters

Parameter	Description
dirname	The case-sensitive name of a database directory.
filename	The name of a file containing an XML template.
IN_XML	The name of a CLOB containing an XML template.

Usage Notes

You should always load a template into the same version and release of Oracle Database as the one used to generate the template. The XML may not be valid if it was generated by a different release of the software.

Example

This example reports a problem in the schema:

```
EXECUTE dbms_cube.validate_xml('UPGRADE_DIR', 'MYGLOBAL.XML');
BEGIN dbms_cube.validate_xml('UPGRADE_DIR', 'MYGLOBAL.XML'); END;

*
ERROR at line 1:
ORA-37162: OLAP error
'GLOBAL.PRICE_CUBE.$AW_ORGANIZATION': XOQ-01950: The AWCubeOrganization for
cube "GLOBAL.PRICE_CUBE" contains multiple BuildSpecifications with the same
name.
'GLOBAL.UNITS_CUBE.$AW_ORGANIZATION': XOQ-01950: The AWCubeOrganization for
cube "GLOBAL.UNITS_CUBE" contains multiple BuildSpecifications with the same
name.
XOQ-01400: invalid metadata objects
ORA-06512: at "SYS.DBMS_CUBE", line 411
ORA-06512: at "SYS.DBMS_CUBE", line 441
ORA-06512: at "SYS.DBMS_CUBE", line 501
ORA-06512: at "SYS.DBMS_CUBE", line 520
ORA-06512: at line 1
```

After the problems are corrected, the procedure reports no errors:

```
EXECUTE dbms_cube.validate_xml('UPGRADE_DIR', 'MYGLOBAL.XML');

PL/SQL procedure successfully completed.
```

This example loads an XML template into a temporary CLOB, then validates it. The script is named `GLOBAL.XML`, and it is located in a database directory named `XML_DIR`.

```
DEFINE xml_file = 'GLOBAL.XML';

SET ECHO ON;
SET SERVEROUT ON;


DECLARE
     xml_file     BFILE := bfilename('XML_DIR', '&xml_file');
     in_xml       CLOB;
     out_xml      CLOB := NULL;
     dest_offset  INTEGER := 1;
```

```
     src_offset   INTEGER := 1;
     lang_context INTEGER := 0;
     warning      INTEGER;
BEGIN
     -- Setup the clob from a file
     DBMS_LOB.CREATETEMPORARY(in_xml, TRUE);
     DBMS_LOB.OPEN(in_xml, DBMS_LOB.LOB_READWRITE);
     DBMS_LOB.OPEN(xml_file, DBMS_LOB.FILE_READONLY);
     DBMS_LOB.LOADCLOBFROMFILE(in_xml, xml_file, DBMS_LOB.LOBMAXSIZE,
        dest_offset, src_offset, 0, lang_context, warning);

     -- Validate the xml
     DBMS_CUBE.VALIDATE_XML(in_xml);
END;
/
```

55
DBMS_CUBE_ADVISE

DBMS_CUBE_ADVISE contains subprograms for evaluating cube materialized views to support log-based fast refresh and query rewrite.

This chapter contains the following topics:

- DBMS_CUBE_ADVISE Security Model
- Summary of DBMS_CUBE_ADVISE Subprograms

> **See Also:**
>
> *Oracle OLAP User's Guide* for information about cube materialized views

55.1 DBMS_CUBE_ADVISE Security Model

The MV_CUBE_ADVICE function requires the ADVISOR privilege.

55.2 Summary of DBMS_CUBE_ADVISE Subprograms

This table lists and describes the DBMS_CUBE_ADVISE subprograms.

Table 55-1 Summary of DBMS_CUBE_ADVISE Subprograms

Subprogram	Description
MV_CUBE_ADVICE Function	Evaluates the metadata of a cube materialized view and generates recommendations for constraints, SQL dimension objects, and materialized view logs to support a broad range of query rewrite and fast refresh opportunities.
SET_CNS_EXCEPTION_LOG Procedure	Identifies the name of an exception log used in validated constraints generated by MV_CUBE_ADVICE.
TRACE Procedure	Displays or suppresses diagnostic messages for MV_CUBE_ADVICE.

55.2.1 MV_CUBE_ADVICE Function

This table function evaluates the metadata for a specified cube materialized view. It generates recommendations and returns them as a SQL result set. These SQL statements can be used to create constraints, SQL dimension objects, and materialized view logs that allow the broadest range of query rewrite transformations and log-based fast refresh of the cube materialized view.

Syntax

```
DBMS_CUBE_ADVISE.MV_CUBE_ADVICE (
          owner         IN  VARCHAR2  DEFAULT USER,
```

```
        mvname       IN  VARCHAR2,
        reqtype      IN  VARCHAR2  DEFAULT '0',
        validate     IN  NUMBER    DEFAULT 0)
    RETURN COAD_ADVICE_T  PIPELINED;
```

Parameters

Table 55-2 MV_CUBE_ADVICE Function Parameters

Parameter	Description
owner	Owner of the cube materialized view
mvname	Name of the cube, such as UNITS_CUBE, or the cube materialized view, such as CB$UNITS_CUBE
reqtype	Type of advice to generate: • 0: All applicable advice types • 1: Column NOT NULL constraints • 2: Primary key constraints • 3: Foreign key constraints • 4: Relational dimension objects • 5: Cube materialized view logs with primary key
validate	Validation option: • 0: Validate the constraints • 1: Do not validate the constraints

Returns

A table of type COAD_ADVICE_T, consisting of a set of rows of type COAD_ADVICE_REC. Table 55-3 describes the columns.

Table 55-3 MV_CUBE_ADVICE Return Values

Column	Datatype	Description
OWNER	VARCHAR2(30)	Owner of the dimensional object identified in APIOBJECT.
APIOBJECT	VARCHAR2(30)	Name of a cube enhanced with materialized view capabilities, such as UNITS_CUBE.
SQLOBJOWN	VARCHAR2(30)	Owner of the relational object identified in SQLOBJECT.
SQLOBJECT	VARCHAR2(65)	Name of the master table, such as UNITS_FACT, or the cube materialized view, such as CB$UNITS_CUBE.
ADVICETYPE	NUMBER(38,0)	Type of recommendation: • 1: Create NOT NULL constraints on the foreign key columns • 2: Create primary key constraints on the master table • 3: Create primary key constraints on the master view • 4: Create foreign key constraints on the master table • 5: Create foreign key constraints on the master view • 6: Create relational dimensions on the master dimension tables • 7: Create a materialized view log • 8: Compile the materialized view

Table 55-3 (Cont.) MV_CUBE_ADVICE Return Values

Column	Datatype	Description
DISPOSITION	CLOB	Pre-existing conditions that conflict with the recommendations and should be resolved before SQLTEXT can be executed.
SQLTEXT	CLOB	SQL statement that implements the recommendation.
DROPTEXT	CLOB	SQL statement that reverses SQLTEXT. Pre-existing conditions may prevent these statements from restoring the schema to its previous state.

Usage Notes

This function is available in Analytic Workspace Manager as the Materialized View Advisor, which will generate a SQL script with the recommendations.

You can query the returned rows the same as any other table, as shown in the example.

MV_CUBE_ADVICE generates unique object names each time it is called. You should execute the function once, capture the results, and work with those SQL statements.

Take care when dropping database objects. If a table already has a materialized view log, it will have the same name used in the SQL DROP MATERIALIZED VIEW LOG statement in the DROPTEXT column. You should avoid inadvertently dropping materialized view logs, especially when they may be used for remote data replication.

Examples

The following query displays the SQL statements recommended by MV_CUBE_ADVICE. UNITS_FACT is the master table for UNITS_CUBE, and MV_CUBE_ADVICE generates an ALTER TABLE command to add primary key constraints.

It also generates an ALTER MATERIALIZED VIEW command to compile the CB$UNITS_CUBE cube materialized view.

```
SQL> SELECT apiobject, sqlobject, sqltext
      FROM TABLE(dbms_cube_advise.mv_cube_advice('GLOBAL', 'CB$UNITS_CUBE'));

APIOBJECT    SQLOBJECT       SQLTEXT
------------ --------------- --------------------------------------------
UNITS_CUBE   UNITS_FACT      alter table "GLOBAL"."UNITS_FACT" add constra
                             int "COAD_PK000208" PRIMARY KEY ("CHANNEL_ID"
                             , "ITEM_ID", "SHIP_TO_ID", "MONTH_ID") rely d
                             isable novalidate

UNITS_CUBE   CB$UNITS_CUBE   alter materialized view "GLOBAL"."CB$UNITS_CU
                             BE" compile
```

55.2.2 SET_CNS_EXCEPTION_LOG Procedure

This procedure identifies the name of an exception log used in validated constraints generated by `MV_CUBE_ADVICE`.

Syntax

```
DBMS_CUBE_ADVISE.SET_CNS_EXCEPTION_LOG (
       exceptlogtab     IN   VARCHAR2  DEFAULT user.EXCEPTIONS);
```

Parameters

Table 55-4 SET_CNS_EXCEPTION_LOG Procedure Parameters

Parameter	Description
`exceptlogtab`	The name of an existing exception log.

Usage Notes

To create an exception log, use the `utlexcpt.sql` or the `utlexpt1.sql` script before executing `SET_CNS_EXCEPTION_LOG`.

The `validate` parameter of `MV_CUBE_ADVICE` must be set to `1`.

Examples

The `utlexcpt.sql` script creates a table named `EXCEPTIONS`, and the `SET_CNS_EXCEPTION_LOG` procedure identifies it as the exception log for `MV_CUBE_ADVICE`. The `ALTER TABLE` statement now includes the clause `VALIDATE EXCEPTIONS INTO "GLOBAL"."EXCEPTIONS"`.

```
SQL> @utlexcpt
Table created.

SQL> EXECUTE dbms_cube_advise.set_cns_exception_log;
PL/SQL procedure successfully completed.

SQL> SELECT apiobject, sqlobject, advicetype type, sqltext
     FROM TABLE(
     dbms_cube_advise.mv_cube_advice('GLOBAL', 'CB$UNITS_CUBE', '2', 1));

APIOBJECT    SQLOBJECT       TYPE SQLTEXT
------------ --------------- ---- ----------------------------------------------
UNITS_CUBE   UNITS_FACT         2 alter table "GLOBAL"."UNITS_FACT" add constrai
                                  nt "COAD_PK000219" PRIMARY KEY ("CHANNEL_ID",
                                  "ITEM_ID", "SHIP_TO_ID", "MONTH_ID") norely en
                                  able validate exceptions into "GLOBAL"."EXCEPT
                                  IONS"

UNITS_CUBE   CB$UNITS_CUBE      8 alter materialized view "GLOBAL"."CB$UNITS_CUB
                                  E" compile
```

55.2.3 TRACE Procedure

This procedure turns on and off diagnostic messages to server output for the `MV_CUBE_ADVICE` function.

Syntax

```
DBMS_CUBE_ADVISE.TRACE (
     diaglevel       IN  BINARY_INTEGER DEFAULT 0);
```

Parameters

Table 55-5 TRACE Procedure Parameters

Parameter	Description
`diaglevel`	`0` to turn tracing off, or `1` to turn tracing on.

Examples

The following example directs the diagnostic messages to server output. The SQL*Plus `SERVEROUTPUT` setting displays the messages.

```
SQL> SET SERVEROUT ON FORMAT WRAPPED
SQL> EXECUTE dbms_cube_advise.trace(1);
DBMS_COAD_DIAG: Changing diagLevel from [0] to [1]

PL/SQL procedure successfully completed.

SQL> SELECT sqlobject, sqltext, droptext
     FROM TABLE(
     dbms_cube_advise.mv_cube_advice('GLOBAL', 'CB$UNITS_CUBE'))
     WHERE apiobject='UNITS_CUBE';

SQLOBJECT       SQLTEXT                                  DROPTEXT
--------------- ---------------------------------------- ----------------------------------------
UNITS_FACT      alter table "GLOBAL"."UNITS_FACT" add co alter table "GLOBAL"."UNITS_FACT" drop c
                nstraint "COAD_PK000222" PRIMARY KEY ("C onstraint "COAD_PK000222" cascade
                HANNEL_ID", "ITEM_ID", "SHIP_TO_ID", "MO
                NTH_ID") rely disable novalidate

CB$UNITS_CUBE   alter materialized view "GLOBAL"."CB$UNI alter materialized view "GLOBAL"."CB$UNI
                TS_CUBE" compile                         TS_CUBE" compile


20070706 07:25:27.462780000 DBMS_COAD_DIAG NOTE: Parameter mvOwner  : GLOBAL
20070706 07:25:27.462922000 DBMS_COAD_DIAG NOTE: Parameter mvName   : CB$UNITS_CUBE
20070706 07:25:27.462967000 DBMS_COAD_DIAG NOTE: Parameter factTab  : .
20070706 07:25:27.463011000 DBMS_COAD_DIAG NOTE: Parameter cubeName : UNITS_CUBE
20070706 07:25:27.463053000 DBMS_COAD_DIAG NOTE: Parameter cnsState : rely disable novalidate
20070706 07:25:27.463094000 DBMS_COAD_DIAG NOTE: Parameter NNState  : disable novalidate
20070706 07:25:27.462368000 DBMS_COAD_DIAG NOTE: Begin NN:
20070706 07:25:27.833530000 DBMS_COAD_DIAG NOTE: End   NN:
20070706 07:25:27.833620000 DBMS_COAD_DIAG NOTE: Begin PK:
20070706 07:25:28.853418000 DBMS_COAD_DIAG NOTE: End   PK:
20070706 07:25:28.853550000 DBMS_COAD_DIAG NOTE: Begin FK:
20070706 07:25:28.853282000 DBMS_COAD_DIAG NOTE: End   FK:
20070706 07:25:28.853359000 DBMS_COAD_DIAG NOTE: Begin RD:
20070706 07:25:29.660471000 DBMS_COAD_DIAG NOTE: End   RD:
```

```
20070706 07:25:29.661363000 DBMS_COAD_DIAG NOTE: Begin CM:
20070706 07:25:29.665106000 DBMS_COAD_DIAG NOTE: End   CM:

SQL> EXECUTE dbms_cube_advise.trace(0);
DBMS_COAD_DIAG: Changing diagLevel from [1] to [0]

PL/SQL procedure successfully completed.
```

56
DBMS_CUBE_LOG

DBMS_CUBE_LOG contains subprograms for creating and managing logs for cubes and cube dimensions.

> **See Also:**
>
> OLAP Technology in the Oracle Database in*Oracle OLAP User's Guide* regarding use of the OLAP option to support business intelligence and analytical applications.

This chapter contains the following topics:

- Using DBMS_CUBE_LOG
- Summary of DBMS_CUBE_LOG Subprograms

56.1 Using DBMS_CUBE_LOG

DBMS_CUBE_LOG manages several logs that enable you to track the progress of long running processes, then use the results to profile performance characteristics.

They provide information to help you diagnose and remedy problems that may occur during development and maintenance of a cube, such as hierarchies that are improperly structured in the relational source tables, records that fail to load, or data refreshes that take too long to complete. They also help diagnose performance problems in querying cubes.

Analytic Workspace Manager creates the logs automatically using the default names and types. It also disables the logs when Analytic Workspace Manager is closed. To use the same logs outside of Analytic Workspace Manager, you must first enable them. Alternatively, you can create and manage different logs for use outside of Analytic Workspace Manager.

This section contains the following topics:

- Logging Types
- Logging Targets
- Verbosity Levels
- Security Model
- Creating Cube Logs
- Cube Build Log
- Cube Dimension Compile Log
- Cube Operations Log
- Cube Rejected Records Log

56.1.1 DBMS_CUBE_LOG — Logging Types

Several logs are available, each one dedicated to storing messages of a particular type. You may use all of them or only those that you find particularly valuable.

The logs and their contents are described later in this topic.

- Cube Build Log
- Cube Dimension Compile Log
- Cube Operations Log
- Cube Rejected Records Log

`DBMS_CUBE_LOG` provides functions that return the binary integer for each log type. You can produce more readable code by using these functions instead of integers for the argument values of other `DBMS_CUBE_LOG` procedures and functions. Refer to these descriptions:

- TYPE_BUILD Function
- TYPE_DIMENSION_COMPILE Function
- TYPE_OPERATIONS Function
- TYPE_REJECTED_RECORDS Function

56.1.2 DBMS_CUBE_LOG — Logging Targets

The `TABLE_CREATE` procedure creates database tables for storing the logs. Using the `ENABLE` procedure, you can create additional targets with changes in the destination or logging level. For example, you might target the Cube Operations log to both a table and a disk file.

These are the available targets:

- Disk file
- LOB
- Database table
- Trace file

See ENABLE Procedure for more information about creating multiple targets.

`DBMS_CUBE_LOG` provides functions that return the binary integer for each target type. You can produce more readable code by using these functions instead of integers for the argument values of other `DBMS_CUBE_LOG` procedures and functions. Refer to these descriptions:

- TARGET_FILE Function
- TARGET_LOB Function
- TARGET_TABLE Function
- TARGET_TRACE Function

56.1.3 DBMS_CUBE_LOG — Verbosity Levels

You can decide how much information is recorded in a log. You may want fewer details when leaving a job to run overnight than when you are monitoring the success of a new build. You can choose from these verbosity levels. Each level adds to the preceding level.

- `LOWEST`: Logs the status of each command used to build the cube dimensions and cubes, the use of slave processes, and summary records. This is the basic logging level.
- `LOW`: Logs messages from the OLAP engine, such as start and finish records for SQL Import, Aggregate, and Update.
- `MEDIUM`: Logs messages at the level used by Analytic Workspace Manager.
- `HIGH`: Logs messages that provide tuning information, such as composite lengths, partitioning details, object sizes, and aggregation work lists. This level is intended for use by Oracle Field Services.
- `HIGHEST`: Logs debugging messages and other information typically sent to a trace file. This level is intended for use by Oracle Support Services.

`DBMS_CUBE_LOG` provides functions that return the binary integer for each verbosity level. You can produce more readable code by using these functions instead of integers for the argument values of other `DBMS_CUBE_LOG` procedures and functions. Refer to these descriptions:

- LEVEL_LOWEST Function
- LEVEL_LOW Function
- LEVEL_MEDIUM Function
- LEVEL_HIGH Function
- LEVEL_HIGHEST Function

56.1.4 DBMS_CUBE_LOG Security Model

The `TABLE_CREATE` procedure requires the `CREATE TABLE` privilege.

56.1.5 DBMS_CUBE_LOG — Creating Cube Logs

To store logging information in a database table, you must create that table using the TABLE_CREATE procedure. Cube Build logs are always stored in tables. The ENABLE procedure creates the other target types for the other logs.

To create a Cube Build log:

- Execute the `TABLE_CREATE` procedure.

The following command creates a Cube Build log with the default name of `CUBE_BUILD_LOG`:

```
EXECUTE dbms_cube_log.table_create(dbms_cube_log.type_build);
```

To create a Cube Dimension Compile log, Cube Operations log, or Cube Rejected Records log with a database table target:

1. Execute the `TABLE_CREATE` procedure to create the table.
2. Execute the `ENABLE` procedure to begin logging.

These commands create and enable a Cube Operations table with the default name of CUBE_OPERATIONS_LOG and the default verbosity level:

```
EXECUTE dbms_cube_log.table_create(dbms_cube_log.type_operations);
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_operations);
```

To create a Cube Dimension Compile log, Cube Operations log, or Cube Rejected Records log with a trace file, disk file, or LOB target:

- Execute the ENABLE procedure.

This command enables the Cube Rejected Records log, sets verbosity to the lowest level, and directs the output to a disk file named rejects.log in the WORK_DIR database directory:

```
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_rejected_records, -
     dbms_cube_log.target_file, dbms_cube_log.level_lowest, -
     'WORK_DIR/rejects.log');
```

56.1.6 DBMS_CUBE_LOG — Cube Build Log

The Cube Build log provides information about what happened during a build. Use this log to determine whether the build produced the results you were expecting, and if not, why not.

The contents of the Cube Build log are refreshed continuously during a build. You can query the log at any time to evaluate the progress of the build and to estimate the time to completion.

The default name of the Cube Build log is CUBE_BUILD_LOG. The following table describes its contents.

> **Note:**
>
> To manage a Cube Build log, use only the TABLE_CREATE and VERSION procedures.

Column	Datatype	NULL	Description
BUILD_ID	NUMBER	--	A unique sequence number for the build. The same number is used for slave processes in a parallel build.
SLAVE_NUMBER	NUMBER	--	A counter for slave processes in a parallel build: 0 is the master process, 1 is the first slave, 2 is the second slave, and so forth.
STATUS	VARCHAR2(10)	--	The current status of the command: STARTED, COMPLETED, FAILED, or WORKING.
COMMAND	VARCHAR2(25)	--	The name of the command being executed, such as BUILD, LOAD, and SOLVE.
BUILD_OBJECT	VARCHAR2(500)	--	The name of the cube or cube dimension being processed.
BUILD_OBJECT_TYPE	VARCHAR2(10)	--	The type of object: CUBE, DIMENSION, or BUILD.
OUTPUT	CLOB	--	Information structured like an XML document about the command, or NULL when there is no additional information, such as for a STARTED row.

Column	Datatype	NULL	Description
AW	VARCHAR2(30)	--	The name of the analytic workspace that contains the objects of the build.
OWNER	VARCHAR2(30)		The owner of the analytic workspace and all the objects of the build.
PARTITION	VARCHAR2(10)	--	The name of the partition being processed, or NULL when the current operation does not correspond to a partition.
SCHEDULER_JOB	VARCHAR2(100)	--	A user-specified string to identify the build.
TIME	TIMESTAMP(6)	--	The time the row is added to the table.
BUILD_SCRIPT	CLOB	--	The cube build script. Populated only in rows where COMMAND is BUILD.
BUILD_TYPE	VARCHAR2(22)	--	The origin of the build: DBMS_CUBE, DBMS_MVIEW, JAVA, or SLAVE.
COMMAND_DEPTH	NUMBER(2)	--	The nesting depth of the command. For example, COMPILE HIERARCHIES is a component step of COMPILE, so if COMPILE has a depth of 1, then COMPILE HIERARCHIES has a depth of 2.
BUILD_SUB_OBJECT	VARCHAR2(30)	--	The name of a subobject being processed, such as a measure that does not inherit the aggregation rules of the cube.
REFRESH_METHOD	VARCHAR2(1)	--	The refresh method, such as C or F, that is associated with the current command. The refresh method is important only for the CLEAR step.
SEQ_NUMBER	NUMBER	--	Not currently used.
COMMAND_NUMBER	NUMBER	--	The sequence number of the command in the current process, which can be used to distinguish the same command on different objects. For example, a LOAD on PRODUCT and a LOAD on TIME.
IN_BRANCH	NUMBER(1)	--	Not currently used.
COMMAND_STATUS_NUMBER	NUMBER	--	Identifies the sequence number of all rows for a particular command. For example, a particular command might be represented by four rows: The first row has a status of STARTED and the last row has a status of COMPLETED. This column is used for sorting.

56.1.7 DBMS_CUBE_LOG — Cube Dimension Compile Log

When solving a cube, OLAP checks the dimension hierarchies to make sure they are valid. Errors that occur during this validation are written to the Cube Dimension Compile log.

The checks include:

- Circularity: Hierarchies are defined by parent-child relations among dimension members. Circularity occurs when a dimension member is specified as its own ancestor or descendant.

- Hierarchy type: Hierarchies can be level based or value based. You can define a cube so that only level-based hierarchies are valid, such as a cube materialized view.
- Level options: Level-based hierarchies can be regular, ragged, or skip level. You can define a dimension so that only regular hierarchies are valid, such as a Time dimension.

The default name of the Cube Dimension Compile log is `CUBE_DIMENSION_COMPILE`. The following table describes its contents.

Column	Datatype	NULL	Description
`ID`	`NUMBER`	--	Current operation identifier
`SEQ_NUMBER`	`NUMBER`	--	Sequence number in the Cube Build log
`ERROR#`	`NUMBER(8)`	`NOT NULL`	Number of the error being reported
`ERROR_MESSAGE`	`VARCHAR2(2000)`	--	Error message associated with the error
`DIMENSION`	`VARCHAR2(100)`	--	Name of the dimension being compiled
`DIMENSION_MEMBER`	`VARCHAR2(100)`	--	Faulty dimension member
`MEMBER_ANCESTOR`	`VARCHAR2(100)`	--	Parent of `DIMENSION_MEMBER`
`HIERARCHY1`	`VARCHAR2(100)`	--	First hierarchy involved in the error
`HIERARCHY2`	`VARCHAR2(100)`	--	Second hierarchy involved in the error
`ERROR_CONTEXT`	`CLOB`	--	Additional information about the error

56.1.8 DBMS_CUBE_LOG — Cube Operations Log

The Cube Operations log contains messages and debugging information for all OLAP engine events. You can track current operations at a very detailed level. Using the `SQL_ID` column, you can join the Cube Operations log to dynamic performance views such as `V$SQL`, `V$SESSION`, and `V$SESSION_LONGOPS` to see cube operations in the context of other database operations such as I/O Wait and CPU.

The default name of the Cube Operations log is `CUBE_OPERATIONS_LOG`. The following table describes its contents.

Column	Datatype	NULL	Description
`INST_ID`	`NUMBER`	`NOT NULL`	Instance identifier
`SID`	`NUMBER`	`NOT NULL`	Session identifier
`SERIAL#`	`NUMBER`	`NOT NULL`	Session serial number
`USER#`	`NUMBER`	`NOT NULL`	User identifier
`SQL_ID`	`VARCHAR2(13)`	--	Executing SQL statement identifier
`JOB`	`NUMBER`	--	Job identifier
`ID`	`NUMBER`	--	Current operation identifier
`PARENT_ID`	`NUMBER`	--	Parent operation identifier
`SEQ_NUMBER`	`NUMBER`	--	Sequence number in the Cube Build log
`TIME`	`TIMESTAMP(6) WITH TIME ZONE`	`NOT NULL`	Time the record was added to the Cube Operations log
`LOG_LEVEL`	`NUMBER(4)`	`NOT NULL`	Verbosity level of the record, as specified by the `DBMS_CUBE_LOG.ENABLE` procedure.

Column	Datatype	NULL	Description
DEPTH	NUMBER(4)	--	Nesting depth of the record. For example, a level of 0 indicates that the operation and suboperation are not nested within other operations and suboperations.
OPERATION	VARCHAR2(15)	NOT NULL	Current operation, such as AGGREGATE, ROWSOURCE, or SQLIMPORT.
SUBOPERATION	VARCHAR2(20)	--	Current suboperation, such as Loading or Import
STATUS	VARCHAR2(10)	NOT NULL	Current status of the operation, such as START, TRACE, COMPLETED, or Failed.
NAME	VARCHAR2(20)	NOT NULL	Name of the record, such as ROWS LOADED, AVE_ROW_LEN, and PAGEPOOLSIZE
VALUE	VARCHAR2(4000)	--	Value of NAME
DETAILS	CLOB	--	Additional information about NAME.

56.1.9 DBMS_CUBE_LOG — Cube Rejected Records Log

The Cube Rejected Records log contains a summary of the loader job and any records that were rejected because they did not meet the expected format.

A single row in the source table may have errors in more than one field. Each field generates an error in the log, resulting in multiple rows with the same rowid4

in the SOURCE_ROW column.

The default name of the Cube Rejected Records log is CUBE_REJECTED_RECORDS. The following table describes its contents.

Column	Datatype	NULL	Description
ID	NUMBER	--	Current operation identifier
SEQ_NUMBER	NUMBER	--	Sequence number in the Cube Build log
ERROR#	NUMBER(8)	NOT NULL	Number of the error triggered by the record
ERROR_MESSAGE	VARCHAR2	--	Error message associated with the error
RECORD#	NUMBER(38)	--	Input record number
SOURCE_ROW	ROWID	--	Rowid of the row in the source table; null when the source is a view or a query

56.2 Summary of DBMS_CUBE_LOG Subprograms

This table lists and describes the DBMS_CUBE_LOG subprograms.

Table 56-1 DBMS_CUBE_LOG Subprograms

Subprogram	Description
DEFAULT_NAME Function	Returns the default table names of the various log types.
DISABLE Procedure	Turns logging off for the duration of a session.
ENABLE Procedure	Turns on logging for the duration of a session, redirects logging to additional output types, and changes the verbosity level in the logs.

Table 56-1 (Cont.) DBMS_CUBE_LOG Subprograms

Subprogram	Description
FLUSH Procedure	Forces all buffered messages to be written to the logs.
GET_LOG Procedure	Returns the current settings for the level and location of a particular log.
GET_LOG_SPEC Function	Retrieves a description of all active logs.
GET_PARAMETER Function	Returns the current values of the options that control various aspects of logging.
LEVEL_HIGH Function	Returns the integer value of the high verbosity level.
LEVEL_HIGHEST Function	Returns the integer value of the highest verbosity level.
LEVEL_LOW Function	Returns the integer value of the low verbosity level.
LEVEL_LOWEST Function	Returns the integer value of the lowest verbosity level.
LEVEL_MEDIUM Function	Returns the integer value of the medium verbosity level.
SET_LOG_SPEC Procedure	Sets all logging to the values specified in the input string.
SET_PARAMETER Procedure	Sets options that control various aspects of logging.
TABLE_CREATE Procedure	Creates the table targets for the OLAP logs.
TARGET_FILE Function	Returns the integer value of a disk file target.
TARGET_LOB Function	Returns the integer value of a LOB target.
TARGET_TABLE Function	Returns the integer value of a database table target
TARGET_TRACE Function	Returns the integer value of a trace file target.
TYPE_BUILD Function	Returns the integer value of the Cube Build log.
TYPE_DIMENSION_COMPILE Function	Returns the integer value of the Cube Dimension Compile log.
TYPE_OPERATIONS Function	Returns the integer value of the Cube Operations log.
TYPE_REJECTED_RECORDS Function	Returns the integer value of the Cube Rejected Records log.
VERSION Function	Returns the version number of a specific log table or the current version number of a specific log type.

56.2.1 DEFAULT_NAME Function

This function returns the default table names of the various log types.

Syntax

```
DBMS_CUBE_LOG.DEFAULT_NAME (
          LOG_TYPE       IN   BINARY_INTEGER  DEFAULT TYPE_OPERATIONS)
    RETURN VARCHAR2;
```

Parameters

Table 56-2 DEFAULT_NAME Function Parameters

Parameter	Description
log_type	One of the following log types: • 1: TYPE_OPERATIONS • 2: TYPE_REJECTED_RECORDS • 3: TYPE_DIMENSION_COMPILE • 4: TYPE_BUILD See "Logging Types".

Returns

The default table name of the specified log type.

Examples

This example returns the default name of the Cube Operations log:

```
SELECT dbms_cube_log.default_name FROM dual;

DEFAULT_NAME
------------------------------
CUBE_OPERATIONS_LOG
```

The next example returns the default name of the Cube Rejected Records log:

```
select dbms_cube_log.default_name(dbms_cube_log.type_rejected_records) -
     "Default Name" from dual;

Default Name
-------------------------
CUBE_REJECTED_RECORDS
```

56.2.2 DISABLE Procedure

This procedure turns logging off for the duration of a session, unless logging is explicitly turned on again with the ENABLE procedure.

Syntax

```
DBMS_CUBE_LOG.DISABLE (
          LOG_TYPE       IN   BINARY_INTEGER  DEFAULT,
          LOG_TARGET     IN   BINARY_INTEGER  DEFAULT);
```

Parameters

Table 56-3 DISABLE Procedure Parameters

Parameter	Description
log_type	One of the following log types: • 1: TYPE_OPERATIONS • 2: TYPE_REJECTED_RECORDS • 3: TYPE_DIMENSION_COMPILE **Note**: You cannot disable the Cube Build log with this procedure. See "Logging Types".
log_target	One of the following destinations for the logging records. The logs are sent to a table unless you previously specified a different target using the ENABLE procedure. • 1: TARGET_TABLE • 2: TARGET_TRACE • 3: TARGET_FILE • 4: TARGET_LOB See "Logging Targets"

Example

This command disables the dimension compilation error log table:

```
EXECUTE dbms_cube_log.disable(dbms_cube_log.type_dimension_compile);
```

56.2.3 ENABLE Procedure

This procedure turns on logging for the duration of a session or until it is turned off using the DISABLE procedure.

The ENABLE procedure also allows you to direct logging to additional output types and to change the amount of detail in the logs. You can enable a log type to each of the log targets. For example, you can enable the Cube Operations log to the trace file, a table, and a file at different verbosity levels, but you cannot enable the Cube Operations log to two files at the same time.

This procedure also checks the format of the logs and updates them if necessary.

Syntax

```
DBMS_CUBE_LOG.ENABLE (
          LOG_TYPE       IN     BINARY_INTEGER  DEFAULT NULL,
          LOG_TARGET     IN     BINARY_INTEGER  DEFAULT NULL,
          LOG_LEVEL      IN     BINARY_INTEGER  DEFAULT NULL);

DBMS_CUBE_LOG.ENABLE (
          LOG_TYPE       IN     BINARY_INTEGER  DEFAULT NULL,
          LOG_TARGET     IN     BINARY_INTEGER  DEFAULT NULL,
          LOG_LEVEL      IN     BINARY_INTEGER  DEFAULT NULL,
          LOG_LOCATION   IN     VARCHAR2        DEFAULT NULL);

DBMS_CUBE_LOG.ENABLE (
          LOG_TYPE       IN     BINARY_INTEGER  DEFAULT NULL,
          LOG_TARGET     IN     BINARY_INTEGER  DEFAULT NULL,
```

```
LOG_LEVEL      IN     BINARY_INTEGER  DEFAULT NULL,
LOG_LOCATION   IN/OUT CLOB );
```

Parameters

Table 56-4 ENABLE Procedure Parameters

Parameter	Description
log_type	One of the following log types: • 1: TYPE_OPERATIONS • 2: TYPE_REJECTED_RECORDS • 3: TYPE_DIMENSION_COMPILE **Note**: You cannot disable the Cube Build log with this procedure. See "Logging Types".
log_target	One of the following destinations for the logging records. The logs are sent to a table unless you previously specified a different target using the ENABLE procedure. • 1: TARGET_TABLE • 2: TARGET_TRACE • 3: TARGET_FILE • 4: TARGET_LOB See "Logging Targets"
log_level	One of the following log verbosity levels. Each level adds new types of messages to the previous level. • 1: LEVEL_LOWEST • 2: LEVEL_LOW • 3: LEVEL_MEDIUM • 4: LEVEL_HIGH • 5: LEVEL_HIGHEST See "Verbosity Levels".
log_location	The full identity of the log, such as *owner.table_name* when log_target is a table.

Examples

The following command enables all cube logs:

```
EXECUTE dbms_cube_log.enable;
```

The following PL/SQL procedure sets the log level to LEVEL_LOWEST:

```
BEGIN
     dbms_cube_log.disable(dbms_cube_log.type_rejected_records);
     dbms_cube_log.enable(dbms_cube_log.type_rejected_records,
          dbms_cube_log.target_table, dbms_cube_log.level_lowest);
END;
/
```

56.2.4 FLUSH Procedure

This procedure forces all buffered messages to be written to the logs.

The buffers are flushed automatically throughout a session, but manually flushing them before viewing the logs assures that you can view all of the messages.

Syntax

```
DBMS_CUBE_LOG.FLUSH ( );
```

Example

The following example flushes the buffers for all of the logs:

```
EXECUTE dbms_cube_log.flush;
```

56.2.5 GET_LOG Procedure

This procedure returns the current settings for the level and location of a particular log.

Syntax

```
DBMS_CUBE_LOG.GET_LOG (
          LOG_TYPE       IN   BINARY_INTEGER  DEFAULT TYPE_OPERATIONS,
          LOG_TARGET     IN   BINARY_INTEGER  DEFAULT TARGET_TABLE,
          LOG_LEVEL      OUT  BINARY_INTEGER,
          LOG_LOCATION   OUT  VARCHAR2 );
```

Parameters

Table 56-5 GET_LOG Procedure Parameters

Parameter	Description
log_type	One of the following log types: • 1: TYPE_OPERATIONS • 2: TYPE_REJECTED_RECORDS • 3: TYPE_DIMENSION_COMPILE See "Logging Types".
log_target	One of the following destinations for the logging records. The logs are sent to a table unless you previously specified a different target using the ENABLE procedure. • 1: TARGET_TABLE • 2: TARGET_TRACE • 3: TARGET_FILE • 4: TARGET_LOB See "Logging Targets"
log_level	One of the following log verbosity levels. Each level adds new types of messages to the previous level. • 1: LEVEL_LOWEST • 2: LEVEL_LOW • 3: LEVEL_MEDIUM • 4: LEVEL_HIGH • 5: LEVEL_HIGHEST See "Verbosity Levels".
log_location	The full identity of the log, such as *owner.table_name* when log_target is a table.

Usage Notes

If `log_type` is not active, then `log_level` and `log_location` are null. Use `DBMS_CUBE_LOG.ENABLE` to activate a log.

Examples

This PL/SQL procedure provides information about the Cube Rejected Records log:

```
SET serverout ON format wrapped

DECLARE
     myloglevel  binary_integer;
     mylogtarget varchar2(128);

BEGIN
     dbms_cube_log.get_log(dbms_cube_log.type_rejected_records,
     dbms_cube_log.target_table, myloglevel, mylogtarget);

     dbms_output.put_line('Log Level: ' || myloglevel);
     dbms_output.put_line('Log Target: ' || mylogtarget);
END;
/
```

The procedure generates results like the following:

```
Log Level: 5

Log Target: GLOBAL.CUBE_REJECTED_RECORDS
```

56.2.6 GET_LOG_SPEC Function

This function retrieves a description of all active Cube Operations logs, Cube Rejected Records logs, and Cube Dimension Compile logs.

Syntax

```
DBMS_CUBE_LOG.GET_LOG_SPEC ( )
     RETURN VARCHAR2;
```

Returns

The type and target of all active logs.

Usage Notes

You can use the output from this function as the input to `SET_LOG_SPEC`.

Examples

The following example shows that the Cube Operations log, Cube Rejected Records log, and Cube Dimension Compile log are active. The Cube Operations log is stored in the session trace file and the other logs are stored in tables.

```
SELECT dbms_cube_log.get_log_spec FROM dual;

GET_LOG_SPEC
--------------------------------------------------------------------------------
OPERATIONS(TABLE, TRACE) REJECTED_RECORDS(TABLE[DEBUG])
```

56.2.7 GET_PARAMETER Function

This function returns the current values of the options that control various aspects of logging. To set these options, use the `SET_PARAMETER` function.

Syntax

```
DBMS_CUBE_LOG.GET_PARAMETER (
          LOG_TYPE       IN   BINARY_INTEGER,
          LOG_PARAMETER  IN   BINARY_INTEGER )
     RETURN BINARY_INTEGER;
```

Parameters

Table 56-6 GET_PARAMETER Function Parameters

Parameter	Description
`log_type`	One of the following log types: • 1: `TYPE_OPERATIONS` • 2: `TYPE_REJECTED_RECORDS` • 3: `TYPE_DIMENSION_COMPILE` See "Logging Types".
`log_parameter`	One of the following options: • 1: `MAX_ERRORS` • 2: `FLUSH_INTERVAL` • 3: `LOG_FULL_RECORD` • 4: `LOG_EVERY_N` • 5: `ALLOW_ERRORS` See "SET_PARAMETER Procedure".

Returns

The value of the specified *`log_parameter`*.

Examples

This example shows the current maximum number of errors in the Cube Rejected Records log before logging stops. This parameter was previously set with the `SET_PARAMETER` procedure.

```
SELECT dbms_cube_log.get_parameter(dbms_cube_log.type_rejected_records, 1) -
     "Maximum Records" FROM dual;

Maximum Records
---------------
            100
```

56.2.8 LEVEL_HIGH Function

This function returns the integer value of the high verbosity level.

Syntax

```
DBMS_CUBE_LOG.LEVEL_HIGH ()
     RETURN BINARY_INTEGER;
```

Returns

4

Usage Notes

Use this function instead of its binary integer equivalent for the *LOG_LEVEL* parameter in DBMS_CUBE_LOG subprograms. See "Verbosity Levels".

Example

This command sets the verbosity level of the cube operations table to high:

```
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_operations, -
     dbms_cube_log.target_table, dbms_cube_log.level_high);
```

56.2.9 LEVEL_HIGHEST Function

This function returns the integer value of the highest verbosity level.

Syntax

```
DBMS_CUBE_LOG.LEVEL_HIGHEST ()
     RETURN BINARY_INTEGER;
```

Returns

5

Usage Notes

Use this function instead of its binary integer equivalent for the LOG_LEVEL parameter in DBMS_CUBE_LOG subprograms. See "Verbosity Levels".

Example

This command sets the verbosity level of the cube operations table to highest:

```
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_operations, -
     dbms_cube_log.target_table, dbms_cube_log.level_highest);
```

56.2.10 LEVEL_LOW Function

This function returns the integer value of the low verbosity level.

Syntax

```
DBMS_CUBE_LOG.LEVEL_LOW ()
     RETURN BINARY_INTEGER;
```

Returns

2

Usage Notes

Use this function instead of its binary integer equivalent for the LOG_LEVEL parameter in DBMS_CUBE_LOG subprograms. See "Verbosity Levels".

Example

This command sets the verbosity level of the cube operations table to low:

```
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_operations, -
     dbms_cube_log.target_table, dbms_cube_log.level_low);
```

56.2.11 LEVEL_LOWEST Function

This function returns the integer value of the lowest verbosity level. This level logs the status of each command used to build the cube dimensions and cubes, the use of slave processes, and summary records.

Syntax

```
DBMS_CUBE_LOG.LEVEL_LOWEST ()
     RETURN BINARY_INTEGER;
```

Returns

```
1
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_LEVEL` parameter in `DBMS_CUBE_LOG` subprograms. See "Verbosity Levels".

Example

This command sets the verbosity level of the cube operations table to lowest:

```
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_operations, -
     dbms_cube_log.target_table, dbms_cube_log.level_lowest);
```

56.2.12 LEVEL_MEDIUM Function

This function returns the integer value of the medium verbosity level.

Syntax

```
DBMS_CUBE_LOG.LEVEL_MEDIUM ()
     RETURN BINARY_INTEGER;
```

Returns

```
3
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_LEVEL` parameter in `DBMS_CUBE_LOG` subprograms. See "Verbosity Levels".

Example

This command sets the verbosity level of the cube operations table to medium:

```
EXECUTE dbms_cube_log.enable(dbms_cube_log.type_operations, -
     dbms_cube_log.target_table, dbms_cube_log.level_medium);
```

56.2.13 SET_LOG_SPEC Procedure

This procedure sets all logging to the values specified in the input string.

Syntax

```
DBMS_CUBE_LOG.SET_LOG_SPEC (
          LOG_SPEC       IN   VARCHAR2 );
```

Parameters

Table 56-7 SET_LOG_SPEC Procedure Parameters

Parameter	Description
log_spec	A string consisting of *type*(*target*) pairs. *Type* can be: • OPERATIONS • REJECTED_RECORDS • DIMENSION_COMPILE *Target* can be: • TABLE • TRACE • FILE • LOB

Usage Notes

The GET_LOG_SPEC function returns a properly formatted string for SET_LOG_SPEC.

Examples

This PL/SQL procedure disables all logs, verifies that they are disabled, then activates the Cube Operations log and the Cube Rejected Records log.

```
BEGIN
     dbms_cube_log.disable;
     dbms_output.put_line('Cube Logs: ' || dbms_cube_log.get_log_spec);

     dbms_cube_log.set_log_spec('OPERATIONS(TRACE) REJECTED_RECORDS(TABLE)');
     dbms_output.put_line('Cube Logs: ' || dbms_cube_log.get_log_spec);
END;
/
```

The output from the procedure verifies that the DISABLE function de-activated all logs, and the SET_LOG_SPEC function activated two logs:

```
Cube Logs:

Cube Logs: OPERATIONS(TRACE) REJECTED_RECORDS(TABLE)
```

56.2.14 SET_PARAMETER Procedure

This procedure sets options that control various aspects of logging.

To obtain the current value of these options, use the GET_PARAMETER function.

Syntax

```
DBMS_CUBE_LOG.SET_PARAMETER (
          LOG_TYPE       IN   BINARY_INTEGER,
          LOG_PARAMETER  IN   BINARY_INTEGER,
          VALUE          IN   BINARY_INTEGER );
```

Parameters

Table 56-8 SET_PARAMETER Procedure Parameters

Parameter	Description
log_type	One of the following log types: • 1: TYPE_OPERATIONS • 2: TYPE_REJECTED_RECORDS • 3: TYPE_DIMENSION_COMPILE • 4: TYPE_BUILD See "Logging Types".
log_parameter	One of the following parameters: • 1: MAX_ERRORS Maximum number of records before signalling an end to logging, such as the number of rejected records in the Cube Rejected Records log or the number of compilation errors in the dimension compilation error log. • 2: FLUSH_INTERVAL The number of seconds to buffer the records before writing them to a log. When this parameter is 0, the records are written directly to the logs without buffering. • 3: LOG_FULL_RECORD Controls logging of rejected records. Set this parameter to one of the following constants: 0: FULL_RECORD_AUTO: Log the full record when no row ID is available. 1: FULL_RECORD_ALWAYS: Always log the full record. 2: FULL_RECORD_NEVER: Never log the full record. • 4: LOG_EVERY_N Enters a progress message every *n* rows during data maintenance. • 5: ALLOW_ERRORS: Displays logging errors, which are initially turned off to allow processing to proceed.
value	The new value of *log_parameter*.

Examples

This PL/SQL procedure sets the two parameters, then uses the GET_PARAMETER function to show the settings:

```
BEGIN
     dbms_cube_log.set_parameter(dbms_cube_log.type_rejected_records, 1, 150);
     dbms_cube_log.set_parameter(dbms_cube_log.type_rejected_records, 2, 5);

     dbms_output.put_line('Max rejected records: ' ||
        dbms_cube_log.get_parameter(dbms_cube_log.type_rejected_records, 1));
```

```
     dbms_output.put_line('Buffer time: ' ||
        dbms_cube_log.get_parameter(dbms_cube_log.type_rejected_records, 2) ||
          ' seconds');
END;
/
```

The procedure displays this information:

```
Max rejected records: 150

Buffer time: 5 seconds
```

56.2.15 TABLE_CREATE Procedure

This procedure creates the table targets for the OLAP logs. You must have the `CREATE TABLE` privilege to use this procedure.

`TABLE_CREATE` also upgrades existing log tables to the current version while preserving the data.

Syntax

```
DBMS_CUBE_LOG.TABLE_CREATE (
          LOG_TYPE       IN   BINARY_INTEGER  DEFAULT,
          TBLNAME        IN   VARCHAR2        DEFAULT );
```

Parameters

Table 56-9 TABLE_CREATE Procedure Parameters

Parameter	Description
`log_type`	One of the following log types: • 1: `TYPE_OPERATIONS` • 2: `TYPE_REJECTED_RECORDS` • 3: `TYPE_DIMENSION_COMPILE` • 4: `TYPE_BUILD` See "Logging Types".
`tblname`	A table name for the log. These are the default names: • `CUBE_OPERATIONS_LOG` • `CUBE_REJECTED_RECORDS` • `CUBE_DIMENSION_COMPILE` • `CUBE_BUILD_LOG`

Examples

This example creates a Cube Operations log table named `CUBE_OPERATIONS_LOG`:

```
EXECUTE dbms_cube_log.table_create;
```

This example creates a Cube Rejected Records log table named `CUBE_REJECTED_RECORDS`:

```
EXECUTE dbms_cube_log.table_create(dbms_cube_log.type_rejected_records);
```

The next example creates a Cube Build log table named `MY_BUILD_LOG`:

```
EXECUTE dbms_cube_log.table_create -
     (dbms_cube_log.type_build, 'MY_BUILD_LOG');
```

56.2.16 TARGET_FILE Function

This function returns the integer value of a file target in `DBMS_CUBE_LOG` subprograms.

Syntax

```
DBMS_CUBE_LOG.TARGET_FILE ()
     RETURN BINARY_INTEGER;
```

Returns

```
3
```

Usage Notes

Use this function instead of its binary integer equivalent for the *`LOG_LEVEL`* parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Targets".

Example

This command disables the Cube Operations log file:

```
EXECUTE dbms_cube_log.disable -
     (dbms_cube_log.type_operations, dbms_cube_log.target_file);
```

56.2.17 TARGET_LOB Function

This function returns the integer value of a LOB target.

Syntax

```
DBMS_CUBE_LOG.TARGET_LOB ()
     RETURN BINARY_INTEGER;
```

Returns

```
4
```

Usage Notes

Use this function instead of its binary integer equivalent for the *`LOG_LEVEL`* parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Targets".

Example

This command disables the Cube Operations log LOB:

```
EXECUTE dbms_cube_log.disable -
     (dbms_cube_log.type_operations, dbms_cube_log.target_lob);
```

56.2.18 TARGET_TABLE Function

This function returns the integer value of a table target.

Syntax

```
DBMS_CUBE_LOG.TARGET_TABLE ()
     RETURN BINARY_INTEGER;
```

Returns

1

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_TARGET` parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Targets".

Example

This command disables the Cube Operations log table:

```
EXECUTE dbms_cube_log.disable -
     (dbms_cube_log.type_operations, dbms_cube_log.target_table);
```

56.2.19 TARGET_TRACE Function

This function returns the integer value of a trace file target.

Syntax

```
DBMS_CUBE_LOG.TARGET_TRACE ()
     RETURN BINARY_INTEGER;
```

Returns

```
2
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_TARGET` parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Targets".

Example

This command disables the Cube Operations log trace file:

```
EXECUTE dbms_cube_log.disable -
     (dbms_cube_log.type_operations, dbms_cube_log.target_trace);
```

56.2.20 TYPE_BUILD Function

This function returns the integer value of the Cube Build log.

Syntax

```
DBMS_CUBE_LOG.TYPE_BUILD ()
     RETURN BINARY_INTEGER;
```

Returns

```
4
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_TYPE` parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Types".

Example

This query returns the default name of a Cube Build log:

```
SELECT dbms_cube_log.default_name(dbms_cube_log.type_build) "Log Name" -
     FROM dual;

Log Name
-------------------------
CUBE_BUILD_LOG
```

56.2.21 TYPE_DIMENSION_COMPILE Function

This function returns the integer value of the Cube Dimension Compile log.

Syntax

```
DBMS_CUBE_LOG.TYPE_DIMENSION_COMPILE ()
     RETURN BINARY_INTEGER;
```

Returns

```
3
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_TYPE` parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Types".

Example

This query returns the default name of a Cube Dimension Compile log:

```
SELECT dbms_cube_log.default_name(dbms_cube_log.type_dimension_compile) -
     "Log Name" FROM dual;

Log Name
-------------------------
CUBE_DIMENSION_COMPILE
```

56.2.22 TYPE_OPERATIONS Function

This function returns the integer value of the Cube Operations log.

Syntax

```
DBMS_CUBE_LOG.TYPE_OPERATIONS ()
     RETURN BINARY_INTEGER;
```

Returns

```
1
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_TYPE` parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Types".

Example

This query returns the default name of a Cube Dimension Compile log:

```
SELECT dbms_cube_log.default_name(dbms_cube_log.type_operations) "Log Name" -
     FROM dual;

Log Name
-------------------------
CUBE_OPERATIONS_LOG
```

56.2.23 TYPE_REJECTED_RECORDS Function

This function returns the integer value of the cube Cube Rejected Records log.

Syntax

```
DBMS_CUBE_LOG.TYPE_REJECTED_RECORDS ()
     RETURN BINARY_INTEGER;
```

Returns

```
2
```

Usage Notes

Use this function instead of its binary integer equivalent for the `LOG_TYPE` parameter in `DBMS_CUBE_LOG` subprograms. See "Logging Types".

Example

This query returns the default name of a Cube Rejected Records log:

```
SELECT dbms_cube_log.default_name(dbms_cube_log.type_rejected_records) -
     "Log Name" FROM dual;

Log Name
-------------------------
CUBE_REJECTED_RECORDS
```

56.2.24 VERSION Function

This function returns the version number of a specific log table or the current version number of a specific log type.

Syntax

```
DBMS_CUBE_LOG.VERSION (
          LOG_TYPE       IN   BINARY_INTEGER  DEFAULT 1,
          TBLNAME        IN   VARCHAR2        DEFAULT NULL)
     RETURN BINARY_INTEGER;
```

Parameters

Table 56-10 VERSION Function Parameters

Parameter	Description
log_type	One of the following log types: • 1: TYPE_OPERATIONS • 2: TYPE_REJECTED_RECORDS • 3: TYPE_DIMENSION_COMPILE • 4: TYPE_BUILD See "Logging Types".
tblname	The name of the log table whose version is returned.

Returns

A version number

Examples

This example returns the current version of the Cube Operations log:

```
SELECT dbms_cube_log.version FROM dual;

   VERSION
----------
         2
```

This example returns the version number of an existing Cube Rejected Records log named CUBE_REJECTED_RECORDS.

```
SELECT dbms_cube_log.version(dbms_cube_log.type_rejected_records, -
     'CUBE_REJECTED_RECORDS') version FROM dual;

   VERSION
----------
         2
```

57
DBMS_DATA_MINING

The `DBMS_DATA_MINING` package is the application programming interface for creating, evaluating, and querying Oracle Machine Learning for SQL models.

In Oracle Database Release 21c, Oracle Data Mining has been rebranded to Oracle Machine Learning for SQL (Oracle Machine Learning for SQL). The PL/SQL package name, however, has not changed and remains `DBMS_DATA_MINING`.

This chapter contains the following topics:

- Overview
- Security Model
- Mining Functions
- Model Settings
- Algorithm Specific Settings
- Solver Settings
- Datatypes
- Summary of DBMS_DATA_MINING Subprograms

> **See Also:**
>
> - *Oracle Machine Learning for SQL Concepts*
> - *Oracle Machine Learning for SQL User's Guide*
> - DBMS_DATA_MINING_TRANSFORM
> - DBMS_PREDICTIVE_ANALYTICS

57.1 DBMS_DATA_MINING Overview

Oracle Machine Learning for SQL supports both supervised and unsupervised machine learning. Supervised machine learning predicts a target value based on historical data. Unsupervised machine learning discovers natural groupings and does not use a target. You can use Oracle Machine Learning for SQL procedures on structured data and unstructured text.

Supervised machine learning techniques include:

- Classification
- Regression
- Feature Selection (Attribute Importance)
- Time Series

Unsupervised machine learning techniques include:

- Clustering
- Association
- Feature Extraction
- Anomaly Detection

The steps you use to build and apply a machine learning model depend on the machine learning technique and the algorithm being used. The algorithms supported by Oracle Machine Learning for SQL are listed in the following table.

Table 57-1 Oracle Machine Learning for SQL Algorithms

Algorithm	Abbreviation	Function
Apriori	AR	Association
CUR Matrix Decomposition	CUR	Attribute importance
Decision Tree	DT	Classification
Expectation Maximization	EM	Clustering
Explicit Semantic Analysis	ESA	Feature extraction, classification
Exponential Smoothing	ESM	Time series
Generalized Linear Models	GLM	Classification, regression
k-Means	KM	Clustering
Minimum Descriptor Length	MDL	Attribute importance
Multivariate State Estimation Technique - Sequential Probability Ratio Test	MSET-SPRT	Anomaly detection, classification
Naive Bayes	NB	Classification
Neural Network	NN	Classification, regression
Non-Negative Matrix Factorization	NMF	Feature extraction
Orthogonal Partitioning Clustering	O-Cluster	Clustering
Random Forest	RF	Classification
Singular Value Decomposition and Principal Component Analysis	SVD and PCA	Feature extraction
Support Vector Machine	SVM	Classification, regression, anomaly detection
XGBoost	XGBoost	Classification, regression

Oracle Machine Learning for SQL supports more than one algorithm for the classification, regression, clustering, and feature extraction machine learning techniques. Each of these machine learning techniques has a default algorithm, as shown in the following table.

Table 57-2 Oracle Machine Learning for SQL Default Algorithms

Mining Function	Default Algorithm
Classification	Naive Bayes
Clustering	*k*-Means
Feature Extraction	Non-Negative Matrix Factorization
Feature Selection	Minimum Descriptor Length

Table 57-2 (Cont.) Oracle Machine Learning for SQL Default Algorithms

Mining Function	Default Algorithm
Regression	Support Vector Machine
Time Series	Exponential Smoothing

57.2 DBMS_DATA_MINING Security Model

The `DBMS_DATA_MINING` package is owned by user `SYS` and is installed as part of database installation. Execution privilege on the package is granted to public. The routines in the package are run with invokers' rights (run with the privileges of the current user).

The `DBMS_DATA_MINING` package exposes APIs that are leveraged by the Oracle Machine Learning for SQL. Users who wish to create machine learning models in their own schema require the `CREATE MINING MODEL` system privilege. Users who wish to create machine learning models in other schemas require the `CREATE ANY MINING MODEL` system privilege.

Users have full control over managing models that exist within their own schema. Additional system privileges necessary for managing machine learning models in other schemas include `ALTER ANY MINING MODEL`, `DROP ANY MINING MODEL`, `SELECT ANY MINING MODEL`, `COMMENT ANY MINING MODEL`, and `AUDIT ANY`.

Individual object privileges on machine learning models, `ALTER MINING MODEL` and `SELECT MINING MODEL`, can be used to selectively grant privileges on a model to a different user.

> **See Also:**
>
> *Oracle Data Mining User's Guide* for more information about the security features of Oracle Machine Learning for SQL

57.3 DBMS_DATA_MINING — Machine Learning Functions

A machine learning **function** refers to the methods for solving a given class of machine learning problems.

The machine learning function must be specified when a model is created. You specify a machine learning function with the `mining_function` parameter of the CREATE_MODEL Procedure or the CREATE_MODEL2 Procedure.

Table 57-3 Machine Learning Functions

Value	Description
`ASSOCIATION`	Association is a descriptive machine learning function. An association model identifies relationships and the probability of their occurrence within a data set. Association models use the Apriori algorithm.

Table 57-3 (Cont.) Machine Learning Functions

Value	Description
ATTRIBUTE_IMPORTANCE	Attribute importance is a predictive machine learning function, also known as feature selection. An attribute importance model identifies the relative importance of an attribute in predicting a given outcome. Attribute importance models can use Minimum Description Length (MDL) or CUR Matrix Decomposition. MDL is the default.
CLASSIFICATION	Classification is a predictive machine learning function. A classification model uses historical data to predict a categorical target. Classification models can use: Decision Tree, logistic regression, Multivariate State Estimation Technique - Sequential Probability Ratio Test, Naive Bayes, Support Vector Machine (SVM), or XGBoost. The default is Naive Bayes. The classification function can also be used for **anomaly detection**. For anomaly detection, you can use the Multivariate State Estimation Technique - Sequential Probability Ratio Test algorithm or the SVM algorithm with a null target (One-Class SVM), or the EM algorithm with a null target (EM Anomaly).
CLUSTERING	Clustering is a descriptive machine learning function. A clustering model identifies natural groupings within a data set. Clustering models can use *k*-Means, O-Cluster, or Expectation Maximization. The default is *k*-Means.
FEATURE_EXTRACTION	Feature extraction is a descriptive machine learning function. A feature extraction model creates an optimized data set on which to base a model. Feature extraction models can use Explicit Semantic Analysis, Non-Negative Matrix Factorization, Singular Value Decomposition, or Principal Component Analysis. Non-Negative Matrix Factorization is the default.
REGRESSION	Regression is a predictive machine learning function. A regression model uses historical data to predict a numerical target. Regression models can use linear regression, Support Vector Machine, or XGBoost. The default is Support Vector Machine.
TIME_SERIES	Time series is a predictive machine learning function. A time series model forecasts the future values of a time-ordered series of historical numeric data over a user-specified time window. Time series models use the Exponential Smoothing algorithm.

See Also:

Oracle Machine Learning for SQL Concepts for more information about mining functions

57.4 DBMS_DATA_MINING — Model Settings

Oracle Machine Learning for SQL uses settings to specify the algorithm and other characteristics of a model. Some settings are general, some are specific to a machine learning function, and some are specific to an algorithm.

All settings have default values. If you want to override one or more of the settings for a model, then you must create a settings table. The settings table must have the column names and data types shown in the following table.

Table 57-4 Required Columns in the Model Settings Table

Column Name	Data Type
SETTING_NAME	VARCHAR2(30)
SETTING_VALUE	VARCHAR2(4000)

The information you provide in the settings table is used by the model at build time. The name of the settings table is an optional argument to the CREATE_MODEL Procedure. You can also provide these settings through the CREATE_MODEL2 Procedure.

The settings used by a model can be found by querying the data dictionary view ALL_MINING_MODEL_SETTINGS. This view displays the model settings used by the machine learning models to which you have access. All of the default and user-specified setting values are included in the view.

See Also:

- ALL_MINING_MODEL_SETTINGS in *Oracle Database Reference*
- *Oracle Machine Learning for SQL User's Guide* for information about specifying model settings

57.4.1 DBMS_DATA_MINING — Algorithm Names

The ALGO_NAME setting specifies the model algorithm.

The values for the ALGO_NAME setting are listed in the following table.

Table 57-5 Algorithm Names

ALGO_NAME Value	Description	Machine Learning Function
ALGO_AI_MDL	Minimum Description Length	Attribute importance
ALGO_APRIORI_ASSOCIATION_RULES	Apriori	Association rules
ALGO_CUR_DECOMPOSITION	CUR Matrix Decomposition	Attribute importance
ALGO_DECISION_TREE	Decision Tree	Classification
ALGO_EXPECTATION_MAXIMIZATION	Expectation Maximization	Clustering, Classification

Table 57-5 (Cont.) Algorithm Names

ALGO_NAME Value	Description	Machine Learning Function
ALGO_EXPLICIT_SEMANTIC_ANALYS	Explicit Semantic Analysis	Feature extraction Classification
ALGO_EXPONENTIAL_SMOOTHING	Exponential Smoothing	Time series
ALGO_EXTENSIBLE_LANG	Language used for extensible algorithm	All mining functions supported
ALGO_GENERALIZED_LINEAR_MODEL	Generalized Linear Model	Classification, regression; also feature selection and generation
ALGO_KMEANS	Enhanced *k*-Means	Clustering
ALGO_MSET_SPRT	Multivariate State Estimation Technique - Sequential Probability Ratio Test	Classification
ALGO_NAIVE_BAYES	Naive Bayes	Classification
ALGO_NEURAL_NETWORK	Neural Network	Classification
ALGO_NONNEGATIVE_MATRIX_FACTOR	Non-Negative Matrix Factorization	Feature extraction
ALGO_O_CLUSTER	O-Cluster	Clustering
ALGO_RANDOM_FOREST	Random Forest	Classification
ALGO_SINGULAR_VALUE_DECOMP	Singular Value Decomposition	Feature extraction
ALGO_SUPPORT_VECTOR_MACHINES	Support Vector Machine	Classification and regression
ALGO_XGBOOST	XGBoost	Classification and regression

See Also:

Oracle Machine Learning for SQL Concepts for information about algorithms

57.4.2 DBMS_DATA_MINING — Automatic Data Preparation

Oracle Machine Learning for SQL supports fully Automatic Data Preparation (ADP), user-directed general data preparation, and user-specified embedded data preparation. The `PREP_*` settings enable the user to request fully automated or user-directed general data preparation. By default, fully Automatic Data Preparation (`PREP_AUTO_ON`) is enabled.

When you enable ADP, the model uses heuristics to transform the build data according to the requirements of the algorithm. Instead of fully ADP, the user can request that the data be shifted and/or scaled with the `PREP_SCALE*` and `PREP_SHIFT*` settings. The transformation instructions are stored with the model and reused whenever the model is applied. The model settings can be viewed in `USER_MINING_MODEL_SETTINGS`.

You can choose to supplement Automatic Data Preparations by specifying additional transformations in the `xform_list` parameter when you build the model. See "CREATE_MODEL Procedure" and "CREATE_MODEL2 Procedure".

If you do not use ADP *and* do not specify transformations in the `xform_list` parameter to `CREATE_MODEL`, you must implement your own transformations separately in the build, test, and

scoring data. You must take special care to implement the exact same transformations in each data set.

If you do not use ADP, but you *do* specify transformations in the `xform_list` parameter to `CREATE_MODEL`, OML4SQL embeds the transformation definitions in the model and prepares the test and scoring data to match the build data.

The values for the `PREP_*` setting are described in the following table.

Table 57-6 PREP_* Setting

Setting Name	Setting Value	Description
`PREP_AUTO`	• `PREP_AUTO_ON` • `PREP_AUTO_OFF`	This setting enables fully automated data preparation. The default is `PREP_AUTO_ON`.
`PREP_SCALE_2DNUM`	• `PREP_SCALE_STDDEV` • `PREP_SCALE_RANGE`	This setting enables scaling data preparation for two-dimensional numeric columns. `PREP_AUTO` must be `OFF` for this setting to take effect. The following are the possible values: • `PREP_SCALE_STDDEV`: A request to divide the column values by the standard deviation of the column and is often provided together with `PREP_SHIFT_MEAN` to yield z-score normalization. • `PREP_SCALE_RANGE`: A request to divide the column values by the range of values and is often provided together with `PREP_SHIFT_MIN` to yield a range of [0,1].
`PREP_SCALE_NNUM`	`PREP_SCALE_MAXABS`	This setting enables scaling data preparation for nested numeric columns. `PREP_AUTO` must be `OFF` for this setting to take effect. If specified, then the valid value for this setting is `PREP_SCALE_MAXABS`, which yields data in the range of [-1,1].
`PREP_SHIFT_2DNUM`	• `PREP_SHIFT_MEAN` • `PREP_SHIFT_MIN`	This setting enables centering data preparation for two-dimensional numeric columns. `PREP_AUTO` must be `OFF` for this setting to take effect. The following are the possible values: • `PREP_SHIFT_MEAN`: Results in subtracting the average of the column from each value. • `PREP_SHIFT_MIN`: Results in subtracting the minimum of the column from each value. .

See Also:

Oracle® Machine Learning for SQL for information about data transformations

57.4.3 DBMS_DATA_MINING — Machine Learning Function Settings

The settings described in this table apply to a machine learning function.

Table 57-7 Machine Learning Function Settings

Machine Learning Function	Setting Name	Setting Value	Description
Association	ASSO_MAX_RULE_LENGTH	TO_CHAR(2<= *numeric_expr* <=20)	Maximum rule length for association rules. Default is 4.
Association	ASSO_MIN_CONFIDENCE	TO_CHAR(0<= *numeric_expr* <=1)	Minimum confidence for association rules. Default is 0.1.
Association	ASSO_MIN_SUPPORT	TO_CHAR(0<= *numeric_expr* <=1)	Minimum support for association rules Default is 0.1.
Association	ASSO_MIN_SUPPORT_INT	a positive integer	Minimum absolute support that each rule must satisfy. The value must be an integer. The default is 1.
Association	ASSO_MIN_REV_CONFIDENCE	TO_CHAR(0<= *numeric_expr* <=1)	Sets the Minimum Reverse Confidence that each rule should satisfy. The Reverse Confidence of a rule is defined as the number of transactions in which the rule occurs divided by the number of transactions in which the consequent occurs. The value is real number between 0 and 1. The default is 0.
Association	ASSO_IN_RULES	NULL	Sets Including Rules applied for each association rule: it specifies the list of items that at least one of them must appear in each reported association rule, either as antecedent or as consequent. It is a comma separated string containing the list of including items. If not set, the default behavior is, the filtering is not applied. For example, `INSERT INTO sett_tab (setting_name, setting_value) VALUES (dbms_data_mining.asso_in_rules, '''a'','b''');`

Table 57-7 (Cont.) Machine Learning Function Settings

Machine Learning Function	Setting Name	Setting Value	Description
Association	ASSO_EX_RULES	NULL	Sets Excluding Rules applied for each association rule: it specifies the list of items that none of them can appear in each reported association rules. It is a comma separated string containing the list of excluded items. No rule can contain any item in the list. The default is NULL. For example, `INSERT INTO sett_tab (setting_name, setting_value) VALUES (dbms_data_mining.asso_ex_rules, '''a'',''b''');`
Association	ASSO_ANT_IN_RULES	NULL	Sets Including Rules for the antecedent: it specifies the list of items that at least one of them must appear in the antecedent part of each reported association rule. It is a comma separated string containing the list of including items. The antecedent part of each rule must contain at least one item in the list. The default is NULL. For example, `INSERT INTO sett_tab (setting_name, setting_value) VALUES (dbms_data_mining.asso_ant_in_rules, '''a'',''b''');`
Association	ASSO_ANT_EX_RULES	NULL	Sets Excluding Rules for the antecedent: it specifies the list of items that none of them can appear in the antecedent part of each reported association rule. It is a comma separated string containing the list of excluded items. No rule can contain any item in the list in its antecedent part. The default is NULL. For example, `INSERT INTO sett_tab (setting_name, setting_value) VALUES (dbms_data_mining.asso_ant_ex_rules, '''a'',''b''');`

Table 57-7 (Cont.) Machine Learning Function Settings

Machine Learning Function	Setting Name	Setting Value	Description

Association	ASSO_CONS_IN_RULES	NULL	Sets Including Rules for the consequent: it specifies the list of items that at least one of them must appear in the consequent part of each reported association rule. It is a comma separated string containing the list of including items. The consequent of each rule must be an item in the list.
			The default is NULL.
			For example,
			`INSERT INTO sett_tab (setting_name, setting_value) VALUES`
			`(dbms_data_mining.asso_cons_in_rules, '''a'',''b''');`

Association	ASSO_CONS_EX_RULES	NULL	Sets Excluding Rules for the consequent: it specifies the list of items that none of them can appear in the consequent part of each reported association rule. It is a comma separated string containing the list of excluded items. No rule can have any item in the list as its consequent.
			The excluding rule can be used to reduce the data that must be stored, but the user may be required to build an extra model for executing different including or Excluding Rules.
			The default is NULL.
			For example,
			`INSERT INTO sett_tab (setting_name, setting_value) VALUES`
			`(dbms_data_mining.asso_cons_ex_rules, '''a'',''b''');`

Table 57-7 (Cont.) Machine Learning Function Settings

Machine Learning Function	Setting Name	Setting Value	Description
Association	ASSO_AGGREGATES	NULL	Specifies the columns to be aggregated. It is a comma separated string containing the names of the columns for aggregation. The number of columns in the list must be <= 10. You can set ASSO_AGGREGATES if ODMS_ITEM_ID_COLUMN_NAME is set indicating transactional input data. See DBMS_DATA_MINING - Global Settings. The data table must have valid column names such as ITEM_ID and CASE_ID which are derived from ODMS_ITEM_ID_COLUMN_NAME and case_id_column_name respectively. Numeric values are supported. ITEM_VALUE is not a mandatory value. The default is NULL. For each item, the user may supply several columns to aggregate. It requires more memory to buffer the extra data. Also, the performance impact can be seen because of the larger input data set and more operation.
Association	ASSO_ABS_ERROR	0<ASSO_ABS_ERRORMAX(ASSO_MIN_SUPPORT, ASSO_MIN_CONFIDENCE).	Specifies the absolute error for the association rules sampling. A smaller value of ASSO_ABS_ERROR obtains a larger sample size which gives accurate results but takes longer computational time. Set a reasonable value for ASSO_ABS_ERROR, such as its default value, to avoid large sample size. The default value is 0.5 * MAX(ASSO_MIN_SUPPORT, ASSO_MIN_CONFIDENCE).
Association	ASSO_CONF_LEVEL	0 ASSO_CONF_LEVEL 1	Specifies the confidence level for an association rules sample. A larger value of ASSO_CONF_LEVEL obtains a larger sample size. Any value between 0.9 and 1 is suitable. The default value is 0.95.
Classification	CLAS_COST_TABLE_NAME	*table_name*	(Decision tree only) Name of a table that stores a cost matrix to be used by the algorithm in building the model. The cost matrix specifies the costs associated with misclassifications. Only decision tree models can use a cost matrix at build time. All classification algorithms can use a cost matrix at apply time. The cost matrix table is user-created. See "ADD_COST_MATRIX Procedure" for the column requirements. See *Oracle Machine Learning for SQL Concepts* for information about costs.

Table 57-7 (Cont.) Machine Learning Function Settings

Machine Learning Function	Setting Name	Setting Value	Description
Classification	CLAS_PRIORS_TABLE_NAME	*table_name*	(Naive Bayes) Name of a table that stores prior probabilities to offset differences in distribution between the build data and the scoring data. The priors table is user-created. See *Oracle Machine Learning for SQL User's Guide* for the column requirements. See *Oracle Machine Learning for SQL Concepts* for additional information about priors.
Classification	CLAS_WEIGHTS_TABLE_NAME	*table_name*	(GLM and SVM only) Name of a table that stores weighting information for individual target values in SVM classification and GLM logistic regression models. The weights are used by the algorithm to bias the model in favor of higher weighted classes. The class weights table is user-created. See *Oracle Machine Learning for SQL User's Guide* for the column requirements. See *Oracle Machine Learning for SQL Concepts* for additional information about class weights.
Classification	CLAS_WEIGHTS_BALANCED	ON OFF	This setting indicates that the algorithm must create a model that balances the target distribution. This setting is most relevant in the presence of rare targets, as balancing the distribution may enable better average accuracy (average of per-class accuracy) instead of overall accuracy (which favors the dominant class). The default value is OFF.
Classification	CLAS_MAX_SUP_BINS	For Decision Tree: 2 <= *a number* <=2147483647 For Random Forest: 2 <= *a number* <=254	This parameter specifies the maximum number of bins for each attribute. The default value is 32. See, DBMS_DATA_MINING — Automatic Data Preparation

Table 57-7 (Cont.) Machine Learning Function Settings

Machine Learning Function	Setting Name	Setting Value	Description
Clustering	CLUS_NUM_CLUSTERS	*TO_CHAR(numeric_expr >=1)*	The maximum number of leaf clusters generated by a clustering algorithm. The algorithm may return fewer clusters, depending on the data. Enhanced *k*-Means usually produces the exact number of clusters specified by CLUS_NUM_CLUSTERS, unless there are fewer distinct data points. When Expectation maximization (EM) is used for clustering, it may return fewer clusters than the number specified by CLUS_NUM_CLUSTERS depending on the data. The number of clusters returned by EM cannot be greater than the number of components, which is governed by algorithm-specific settings. (See *Expectation Maximization Settings for Learning* table) Depending on these settings, there may be fewer clusters than components. If component clustering is disabled, the number of clusters equals the number of components. The setting can be used only for EM Clustering algorithm. For EM Clustering algorithm, the default value of CLUS_NUM_CLUSTERS is system-determined. For *k*-Means and O-Cluster, the default is 10.
Feature extraction	FEAT_NUM_FEATURES	TO_CHAR(*numeric_expr* >=1)	The number of features to be extracted by a feature extraction model. The default is estimated from the data by the algorithm. If the matrix rank is smaller than this number, fewer features will be returned. For CUR Matrix Decomposition, the FEAT_NUM_FEATURES value is the same as the CURS_SVD_RANK value.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about machine learning functions

57.4.4 DBMS_DATA_MINING — Global Settings

The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

Table 57-8 Global Settings

Setting Name	Setting Value	Description
ODMS_BOXCOX	ODMS_BOXCOX_ENABLE ODMS_BOXCOX_DISABLE	This setting enables the Box-Cox variance-stabilization transformation. It is useful when the variance increases as the target value increases. It reduces variance and transforms a multiplicative relationship with the target, with a simpler additive relationship. This setting is applicable only to the Exponential Smoothing algorithm. When a value for EXSM_MODEL setting is not specified, the default value is ODMS_BOXCOX_ENABLE and when a value for the EXSM_MODEL setting is provided, the default value is ODMS_BOXCOX_DISABLE.
ODMS_EXPLOSION_MIN_SUPP	A positive integer	It is the minimum required support for categorical values that must be included in the explosion mapping. It removes categorical values with insufficient row instances to have a statistically significant effect on the model, however, they could potentially degrade performance. The default is system determined depending on the number of rows in the dataset. A value of 1 results into mapping all categorical values.
ODMS_ITEM_ID_COLUMN_NAME	*column_name*	(Association rules only) Name of a column that contains the items in a transaction. When this setting is specified, the algorithm expects the data to be presented in a native transactional format, consisting of two columns: • Case ID, either categorical or numeric • Item ID, either categorical or numeric **Note:** Oracle Machine Learning does not support BOOLEAN values for this setting. A typical example of transactional data is market basket data, wherein a case represents a basket that may contain many items. Each item is stored in a separate row, and many rows may be needed to represent a case. The case ID values do not uniquely identify each row. Transactional data is also called multi-record case data. Association rules function is normally used with transactional data, but it can also be applied to single-record case data (similar to other algorithms). For more information about single-record and multi-record case data, see *Oracle SQL Developer Data Modeler User's Guide*.

Table 57-8 (Cont.) Global Settings

Setting Name	Setting Value	Description
ODMS_ITEM_VALUE_COLUMN_NAME	*column_name*	(Association rules only) Name of a column that contains a value associated with each item in a transaction. This setting is only used when a value has been specified for ODMS_ITEM_ID_COLUMN_NAME indicating that the data is presented in native transactional format. If ASSO_AGGREGATES is used, then the build data must include the following three columns and the columns specified in the AGGREGATES setting. • Case ID, either categorical or numeric • Item ID, either categorical or numeric, specified by ODMS_ITEM_ID_COLUMN_NAME • Item value, either categorical or numeric, specified by ODMS_ITEM_VALUE_COLUMN_NAME **Note:** Oracle Machine Learning does not support BOOLEAN values for this setting. If ASSO_AGGREGATES, Case ID, and Item ID column are present, then the Item Value column may or may not appear. The Item Value column may specify information such as the number of items (for example, three apples) or the type of the item (for example, macintosh apples). For details on ASSO_AGGREGATES, see DBMS_DATA_MINING - Mining Function Settings.
ODMS_MISSING_VALUE_TREATMENT	ODMS_MISSING_VALUE_MEAN_MODE ODMS_MISSING_VALUE_DELETE_ROW ODMS_MISSING_VALUE_AUTO	Indicates how to treat missing values in the training data. This setting does not affect the scoring data. The default value is ODMS_MISSING_VALUE_AUTO. ODMS_MISSING_VALUE_MEAN_MODE replaces missing values with the mean (numeric attributes) or the mode (categorical attributes) both at build time and apply time where appropriate. ODMS_MISSING_VALUE_AUTO performs different strategies for different algorithms. When ODMS_MISSING_VALUE_TREATMENT is set to ODMS_MISSING_VALUE_DELETE_ROW, the rows in the training data that contain missing values are deleted. However, if you want to replicate this missing value treatment in the scoring data, then you must perform the transformation explicitly. The value ODMS_MISSING_VALUE_DELETE_ROW applies to all algorithms.

Table 57-8 (Cont.) Global Settings

Setting Name	Setting Value	Description
ODMS_ROW_WEIGHT_COLUMN_NAME	*column_name*	(GLM only) Name of a column in the training data that contains a weighting factor for the rows. The column data type must be numeric. Oracle Machine Learning does not support BOOLEAN values for this setting. Row weights can be used as a compact representation of repeated rows, as in the design of experiments where a specific configuration is repeated several times. Row weights can also be used to emphasize certain rows during model construction. For example, to bias the model towards rows that are more recent and away from potentially obsolete data.
ODMS_TEXT_POLICY_NAME	The name of an Oracle Text POLICY created using CTX_DDL.CREATE_POLICY.	Affects how individual tokens are extracted from unstructured text. For details about CTX_DDL.CREATE_POLICY, see *Oracle Text Reference*.
ODMS_TEXT_MAX_FEATURES	1 <= *value*	The maximum number of distinct features, across all text attributes, to use from a document set passed to CREATE_MODEL. The default is 3000. ESA has the default value of 300000.
ODMS_TEXT_MIN_DOCUMENTS	Non-negative value	This is a text processing setting the controls how in how many documents a token needs to appear to be used as a feature. The default is 1. ESA has a default of 3.
ODMS_PARTITION_COLUMNS	Comma separated list of machine learning attributes	This setting indicates a request to build a partitioned model. The setting value is a comma-separated list of the machine learning attributes used to determine the in-list partition key values. Oracle Machine Learning supports numeric and categorical values including BOOLEAN for this setting. These machine learning attributes are taken from the input columns unless an XFORM_LIST parameter is passed to CREATE_MODEL or CREATE_MODEL2. If the XFORM_LIST parameter is passed to during model building, then the machine learning attributes are taken from the attributes produced by these transformations.
ODMS_MAX_PARTITIONS	1< value <= 1000000	This setting indicates the maximum number of partitions allowed for the model. The default is 1000.
ODMS_SAMPLING	ODMS_SAMPLING_ENABLE ODMS_SAMPLING_DISABLE	This setting allows the user to request a sampling of the build data. The default is ODMS_SAMPLING_DISABLE.
ODMS_SAMPLE_SIZE	0 < Value	This setting determines how many rows will be sampled (approximately). It can be set only if ODMS_SAMPLING is enabled. The default value is the system determined.

Table 57-8 (Cont.) Global Settings

Setting Name	Setting Value	Description
ODMS_PARTITION_BUILD_TYPE	ODMS_PARTITION_BUILD_INTRA ODMS_PARTITION_BUILD_INTER ODMS_PARTITION_BUILD_HYBRID	This setting controls the parallel build of partitioned models. ODMS_PARTITION_BUILD_INTRA — Each partition is built in parallel using all replicas. ODMS_PARTITION_BUILD_INTER — Each partition is built entirely in a single slave, but multiple partitions may be built at the same time since multiple replicas are active. ODMS_PARTITION_BUILD_HYBRID — It is a combination of the other two types and is recommended for most situations to adapt to dynamic environments. The default mode is ODMS_PARTITION_BUILD_HYBRID
ODMS_TABLESPACE_NAME	*tablespace_name*	This setting controls the storage specifications. If you explicitly sets this to the name of a tablespace (for which you have sufficient quota), then the specified tablespace storage creates the resulting model content. If you do not provide this setting, then the default tablespace of the user creates the resulting model content.
ODMS_RANDOM_SEED	The value must be a non-negative integer	The hash function with a random number seed generates a random number with uniform distribution. Users can control the random number seed by this setting. The default is 0. This setting is used by Random Forest, Neural Network, and CUR Matrix Decomposition.
ODMS_DETAILS	• ODMS_ENABLE • ODMS_DISABLE	This setting reduces the space that is used while creating a model, especially a partitioned model. The default value is ODMS_ENABLE. When the setting is ODMS_ENABLE, it creates model tables and views when the model is created. You can query the model with SQL. When the setting is ODMS_DISABLE, model views are not created and tables relevant to model details are not created either. The reduction in space depends on the model. Reduction on the order of 10x can be achieved.

See Also:

Oracle Machine Learning for SQL Concepts for information about GLM

Oracle Machine Learning for SQL Concepts for information about association rules

Oracle Machine Learning for SQL User's Guide for information about machine learning unstructured text

57.5 DBMS_DATA_MINING — Algorithm Specific Model Settings

Oracle Machine Learning for SQL uses algorithm specific settings to define the characteristics of a model.

All settings have default values. If you want to override one or more of the settings for a model, then you must specify those settings.

The information you provide in the settings table is used by the model at build time. The name of the settings table is an optional argument to the CREATE_MODEL Procedure. You can also provide these settings through the CREATE_MODEL2 Procedure.

The settings used by a model can be found by querying the data dictionary view `ALL_MINING_MODEL_SETTINGS`. This view displays the model settings used by the machine learning models to which you have access. All of the default and user-specified setting values are included in the view.

See Also:

- `ALL_MINING_MODEL_SETTINGS` in *Oracle Database Reference*
- *Oracle Machine Learning for SQL User's Guide* for information about specifying model settings

57.5.1 DBMS_DATA_MINING — Algorithm Settings: ALGO_EXTENSIBLE_LANG

The settings listed in the following table configure the behavior of the machine learning model with an extensible algorithm. The model is built in the R language.

The `RALG_*_FUNCTION` specifies the R script that is used to build, score, and view an R model and must be registered in the Oracle Machine Learning for R script repository. The R scripts are registered through Oracle Machine Learning for R with special privileges. When `ALGO_EXTENSIBLE_LANG` is set to R in the `MINING_MODEL_SETTING` table, the machine learning model is built in the R language. After the R model is built, the names of the R scripts are recorded in the `MINING_MODEL_SETTING` table in the `SYS` schema. The scripts must exist in the script repository for the R model to function. The amount of R memory used to build, score, and view the R model through these R scripts can be controlled by Oracle Machine Learning for R.

All algorithm-independent `DBMS_DATA_MINING` subprograms can operate on an R model for machine learning functions such as association, attribute importance, classification, clustering, feature extraction, and regression.

The supported `DBMS_DATA_MINING` subprograms include, but are not limited, to the following:

- ADD_COST_MATRIX Procedure
- COMPUTE_CONFUSION_MATRIX Procedure
- COMPUTE_LIFT Procedure
- COMPUTE_ROC Procedure

- CREATE_MODEL Procedure
- DROP_MODEL Procedure
- EXPORT_MODEL Procedure
- GET_MODEL_COST_MATRIX Function
- IMPORT_MODEL Procedure
- REMOVE_COST_MATRIX Procedure
- RENAME_MODEL Procedure

Table 57-9 ALGO_EXTENSIBLE_LANG Settings

Setting Name	Setting Value	Description
RALG_BUILD_FUNCTION	R_BUILD_FUNCTION_SCRIPT_NAME	Specifies the name of an existing registered R script for the R algorithm machine learning model build function. The R script defines an R function for the first input argument for training data and returns an R model object. For clustering and feature extraction machine learning function model build, the R attributes dm$nclus and dm$nfeat must be set on the R model to indicate the number of clusters and features respectively. The RALG_BUILD_FUNCTION must be set along with ALGO_EXTENSIBLE_LANG in the model_setting_table.
RALG_BUILD_PARAMETER	SELECT *value* param_name, ...FROM DUAL	Specifies a list of numeric and string scalar for optional input parameters of the model build function.
RALG_SCORE_FUNCTION	R_SCORE_FUNCTION_SCRIPT_NAME	Specifies the name of an existing registered R script to score data. The script returns a data.frame containing the corresponding prediction results. The setting is used to score data for machine learning functions such as regression, classification, clustering, and feature extraction. This setting does not apply to the association and the attribute importance functions.
RALG_WEIGHT_FUNCTION	R_WEIGHT_FUNCTION_SCRIPT_NAME	Specifies the name of an existing registered R script for the R algorithm that computes the weight (contribution) for each attribute in scoring. The script returns a data.frame containing the contributing weight for each attribute in a row. This function setting is needed for the PREDICTION_DETAILS SQL function.
RALG_DETAILS_FUNCTION	R_DETAILS_FUNCTION_SCRIPT_NAME	Specifies the name of an existing registered R script for the R algorithm that produces the model information. This setting is required to generate a model view.
RALG_DETAILS_FORMAT	SELECT *type_value column_name*, ... FROM DUAL	Specifies the SELECT query for the list of numeric and string scalars for the output column type and the column name of the generated model view. This setting is required to generate a model view.

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide*

57.5.2 DBMS_DATA_MINING — Algorithm Settings: CUR Matrix Decomposition

The following settings affects the behavior of the CUR Matrix Decomposition algorithm.

Table 57-10 CUR Matrix Decomposition Settings

Setting Name	Setting Value	Description
CURS_APPROX_ATTR_NUM	The value must be a positive integer	Defines the approximate number of attributes to be selected. The default value is the number of attributes.
CURS_ROW_IMPORTANCE	CURS_ROW_IMP_ENABLE CURS_ROW_IMP_DISABLE	Defines the flag indicating whether or not to perform row selection. The default value is CURS_ROW_IMP_DISABLE.
CURS_APPROX_ROW_NUM	The value must be a positive integer	Defines the approximate number of rows to be selected. This parameter is only used when users decide to perform row selection (CURS_ROW_IMP_ENABLE). The default value is the total number of rows.
CURS_SVD_RANK	The value must be a positive integer	Defines the rank parameter used in the column/row leverage score calculation. If users do not provide an input value, the value is determined by the system.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts*

57.5.3 DBMS_DATA_MINING — Algorithm Settings: Decision Tree

These settings configure the behavior of the Decision Tree algorithm. Note that the Decision Tree settings are also used to configure the behavior of Random Forest as it constructs each individual decision tree.

Table 57-11 Decision Tree Settings

Setting Name	Setting Value	Description
TREE_IMPURITY_METRIC	TREE_IMPURITY_ENTROPY TREE_IMPURITY_GINI	Tree impurity metric for Decision Tree. Tree algorithms seek the best test question for splitting data at each node. The best splitter and split values are those that result in the largest increase in target value homogeneity (purity) for the entities in the node. Purity is by a metric. Decision trees can use either Gini (TREE_IMPURITY_GINI) or entropy (TREE_IMPURITY_ENTROPY) as the purity metric. By default, the algorithm uses TREE_IMPURITY_GINI.
TREE_TERM_MAX_DEPTH	For Decision Tree: 2<= *a number* <=20 For Random Forest: 2<= *a number* <=100	Criteria for splits: maximum tree depth (the maximum number of nodes between the root and any leaf node, including the leaf node). For Decision Tree, the default is 7. For Random Forest, the default is 16.
TREE_TERM_MINPCT_NODE	0<= *a number*<=10	The minimum number of training rows in a node expressed as a percentage of the rows in the training data. Default is 0.05, indicating 0.05%.
TREE_TERM_MINPCT_SPLIT	0 < *a number* <=20	The minimum number of rows required to consider splitting a node expressed as a percentage of the training rows. Default is 0.1, indicating 0.1%.
TREE_TERM_MINREC_NODE	*a number*>=0	The minimum number of rows in a node. Default is 10.
TREE_TERM_MINREC_SPLIT	a number > 1	Criteria for splits: minimum number of records in a parent node expressed as a value. No split is attempted if the number of records is below this value. Default is 20.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about Decision Tree

57.5.4 DBMS_DATA_MINING — Algorithm Settings: Expectation Maximization

These algorithm settings configure the behavior of the Expectation Maximization algorithm.

> **See Also:**
>
> *Oracle Data Mining Concepts* for information about Expectation Maximization

Table 57-12 Expectation Maximization Settings for Data Preparation and Analysis

Setting Name	Setting Value	Description
EMCS_ATTRIBUTE_FILTER	EMCS_ATTR_FILTER_ENABLE EMCS_ATTR_FILTER_DISABLE	Whether or not to include uncorrelated attributes in the model. When EMCS_ATTRIBUTE_FILTER is enabled, uncorrelated attributes are not included. **Note:** This setting applies only to attributes that are not nested. For Clustering, the default is system-determined. For anomaly detection, the default is EMCS_ATTR_FILTER_DISABLE.
EMCS_MAX_NUM_ATTR_2D	TO_CHAR(*numeric_expr* >=1)	Maximum number of correlated attributes to include in the model. Note: This setting applies only to attributes that are not nested (2D). Default is 50.
EMCS_NUM_DISTRIBUTION	EMCS_NUM_DISTR_BERNOULLI EMCS_NUM_DISTR_GAUSSIAN EMCS_NUM_DISTR_SYSTEM	The distribution for modeling numeric attributes. Applies to the input table or view as a whole and does not allow per-attribute specifications. The options include Bernoulli, Gaussian, or system-determined distribution. When Bernoulli or Gaussian distribution is chosen, all numeric attributes are modeled using the same type of distribution. When the distribution is system-determined, individual attributes may use different distributions (either Bernoulli or Gaussian), depending on the data. Default is EMCS_NUM_DISTR_SYSTEM.
EMCS_NUM_EQUIWIDTH_BINS	TO_CHAR(1 <*numeric_expr* <=255)	Number of equi-width bins that will be used for gathering cluster statistics for numeric columns. Default is 11.

Table 57-12 (Cont.) Expectation Maximization Settings for Data Preparation and Analysis

Setting Name	Setting Value	Description
EMCS_NUM_PROJECTIONS	TO_CHAR(*numeric_expr* >=1)	Specifies the number of projections that will be used for each nested column. If a column has fewer distinct attributes than the specified number of projections, the data will not be projected. The setting applies to all nested columns. Default is 50.
EMCS_NUM_QUANTILE_BINS	TO_CHAR(1 <*numeric_expr* <=255)	Specifies the number of quantile bins that will be used for modeling numeric columns with multivalued Bernoulli distributions. Default is system-determined.
EMCS_NUM_TOPN_BINS	TO_CHAR(1 <*numeric_expr* <=255)	Specifies the number of top-N bins that will be used for modeling categorical columns with multivalued Bernoulli distributions. Default is system-determined.

Table 57-13 Expectation Maximization Settings for Learning

Setting Name	Setting Value	Description
EMCS_CONVERGENCE_CRITERION	EMCS_CONV_CRIT_HELDASIDE EMCS_CONV_CRIT_BIC	The convergence criterion for EM. The convergence criterion may be based on a held-aside data set, or it may be Bayesian Information Criterion. Default is system determined.
EMCS_LOGLIKE_IMPROVEMENT	TO_CHAR(0 < *numeric_expr* < 1)	When the convergence criterion is based on a held-aside data set (EMCS_CONVERGENCE_CRITERION = EMCS_CONV_CRIT_HELDASIDE), this setting specifies the percentage improvement in the value of the log likelihood function that is required for adding a new component to the model. Default value is 0.001.
EMCS_NUM_COMPONENTS	TO_CHAR(*numeric_expr* >=1)	Maximum number of components in the model. If model search is enabled, the algorithm automatically determines the number of components based on improvements in the likelihood function or based on regularization, up to the specified maximum. For EM Clustering, the number of components must be greater than or equal to the number of clusters. Default is 20 for both EM Clustering and EM Anomaly.
EMCS_NUM_ITERATIONS	TO_CHAR(*numeric_expr* >=1)	Specifies the maximum number of iterations in the EM algorithm. Default is 100.
EMCS_MODEL_SEARCH	EMCS_MODEL_SEARCH_ENABLE EMCS_MODEL_SEARCH_DISABLE (default).	This setting enables model search in EM where different model sizes are explored and a best size is selected. The default is EMCS_MODEL_SEARCH_DISABLE.

Table 57-13 (Cont.) Expectation Maximization Settings for Learning

Setting Name	Setting Value	Description
EMCS_REMOVE_COMPONENTS	EMCS_REMOVE_COMPS_ENABLE (default) EMCS_REMOVE_COMPS_DISABLE	This setting allows the EM algorithm to remove a small component from the solution. The default is EMCS_REMOVE_COMPS_ENABLE.
EMCS_RANDOM_SEED	Non-negative integer	This setting controls the seed of the random generator used in EM. The default is 0.

Table 57-14 Expectation Maximization Settings for Component Clustering

Setting Name	Setting Value	Description
EMCS_CLUSTER_COMPONENTS	EMCS_CLUSTER_COMP_ENABLE EMCS_CLUSTER_COMP_DISABLE	Enables or disables the grouping of EM components into high-level clusters. When disabled, the components themselves are treated as clusters. The setting can be used only for EM Clustering. When component clustering is enabled, model scoring through the SQL CLUSTER function will produce assignments to the higher level clusters. When clustering is disabled, the CLUSTER function will produce assignments to the original components. Default is EMCS_CLUSTER_COMP_ENABLE.
EMCS_CLUSTER_THRESH	TO_CHAR(*numeric_expr* >=1)	Dissimilarity threshold that controls the clustering of EM components. When the dissimilarity measure is less than the threshold, the components are combined into a single cluster. The setting can be used only for EM Clustering. A lower threshold may produce more clusters that are more compact. A higher threshold may produce fewer clusters that are more spread out. Default is 2.
EMCS_LINKAGE_FUNCTION	EMCS_LINKAGE_SINGLE EMCS_LINKAGE_AVERAGE EMCS_LINKAGE_COMPLETE	Allows the specification of a linkage function for the agglomerative clustering step. The setting can be used only for EM Clustering. EMCS_LINKAGE_SINGLE uses the nearest distance within the branch. The clusters tend to be larger and have arbitrary shapes. EMCS_LINKAGE_AVERAGE uses the average distance within the branch. There is less chaining effect and the clusters are more compact. EMCS_LINKAGE_COMPLETE uses the maximum distance within the branch. The clusters are smaller and require strong component overlap. Default is EMCS_LINKAGE_SINGLE.

Table 57-15 Expectation Maximization Settings for Cluster Statistics

Setting Name	Setting Value	Description
EMCS_CLUSTER_STATISTICS	EMCS_CLUS_STATS_ENABLE EMCS_CLUS_STATS_DISABLE	Enables or disables the gathering of descriptive statistics for clusters (centroids, histograms, and rules). When statistics are disabled, model size is reduced, and GET_MODEL_DETAILS_EM only returns taxonomy (hierarchy) and cluster counts. The setting can be used only for EM Clustering. Default is EMCS_CLUS_STATS_ENABLE.
EMCS_MIN_PCT_ATTR_SUPPORT	TO_CHAR(0 < *numeric_expr* < 1)	Minimum support required for including an attribute in the cluster rule. The support is the percentage of the data rows assigned to a cluster that must have non-null values for the attribute. The setting can be used only for EM Clustering. Default is 0.1.

Table 57-16 Expectation Maximization Settings for Anomaly Detection

Setting Name	Setting Value	Description
EMCS_OUTLIER_RATE	TO_CHAR(0 < numeric_expr < 1)	The desired rate of outliers in the training data. The setting can be used only for EM Anomaly. Default is 0.05.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

57.5.5 DBMS_DATA_MINING — Algorithm Settings: Explicit Semantic Analysis

Explicit Semantic Analysis (ESA) is a useful technique for extracting meaningful and interpretable features.

The settings listed in the following table configure the ESA values.

Table 57-17 Explicit Semantic Analysis Settings

Setting Name	Setting Value	Description
ESAS_EMBEDDINGS	ESAS_EMBEDDINGS_ENABLE ESAS_EMBEDDINGS_DISABLE	This setting applies to feature extraction models. The default value is ESAS_EMBEDDINGS_DISABLE. When you set ESAS_EMBEDDINGS_ENABLE: • ESA generates embeddings during scoring • The FEATURE_ID of the generated embeddings is of the datatype NUMBER • The CASE_ID_COLUMN_NAME argument of the DBMS_DATA_MINING.CREATE_MODEL and DBMS_DATA_MINING.CREATE_MODEL2 function is optional. .
ESAS_EMBEDDING_SIZE	A positive integer less than or equal to 4096	This setting applies to feature extraction models. This setting specifies the size of the vectors representing embeddings. You can set this parameter only if you have enabled ESAS_EMBEDDINGS. The default size is 1024. If this value is less than the number of distinct features in the training set, then the actual number of explicit features is used as the size of embedding vectors instead.
ESAS_MIN_ITEMS	Text input 100 Non-text input is 0	This setting determines the minimum number of non-zero entries that need to be present in an input row. The default is 100 for text input and 0 for non-text input.
ESAS_TOPN_FEATURES	A positive integer	This setting controls the maximum number of features per attribute. The default is 1000.
ESAS_VALUE_THRESHOLD	Non-negative number	This setting thresholds a small value for attribute weights in the transformed build data. The default is 1e-8.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

See Also:

Oracle Machine Learning for SQL Concepts for information about ESA.

57.5.6 DBMS_DATA_MINING — Algorithm Settings: Exponential Smoothing

These settings configure the behavior of the Exponential Smoothing (ESM) algorithm.

The settings listed in the following table specify the setting names and possible values for Exponential Smoothing. The Constant Value column specifies constants using the prefix `DBMS_DATA_MINING`.

For Global settings, see DBMS_DATA_MINING — Global Settings.

Table 57-18 Exponential Smoothing Settings

Setting Name	Setting Value	Description
EXSM_MODEL	EXSM_SIMPLE EXSM_SIMPLE_MULT_ERR EXSM_HOLT EXSM_HOLT_DAMPED EXSM_MULT_TREND EXSM_MULT_TREND_DAMPED EXSM_SEASON_ADD EXSM_SEASON_MUL EXSM_WINTERS EXSM_WINTERS_DAMPED EXSM_ADDWINTERS EXSM_ADDWINTERS_DAMPED EXSM_WINTERS_MUL_TREND EXSM_WINTERS_MUL_TREND_DMP	This setting specifies the model. EXSM_SIMPLE: Simple exponential smoothing model is applied. EXSM_SIMPLE_MULT_ERR: Simple exponential smoothing model with multiplicative error is applied. EXSM_HOLT: Holt linear exponential smoothing model is applied. EXSM_HOLT_DAMPED: Holt linear exponential smoothing model with damped trend is applied. EXSM_MULT_TREND: Exponential smoothing model with multiplicative trend is applied. EXSM_MULT_TREND_DAMPED: Exponential smoothing model with multiplicative damped trend is applied. EXSM_SEASON_ADD: Exponential smoothing with additive seasonality, but no trend, is applied. EXSM_SEASON_MUL: Exponential smoothing with multiplicative seasonality, but no trend, is applied. EXSM_WINTERS: Holt-Winters triple exponential smoothing model, additive trend, multiplicative seasonality is applied. EXSM_WINTERS_DAMPED: Holt-Winters multiplicative exponential smoothing model with damped trend, additive trend, multiplicative seasonality is applied. EXSM_ADDWINTERS: Holt-Winters additive exponential smoothing model, additive trend, additive seasonality is applied. EXSM_ADDWINTERS_DAMPED: Holt-Winters additive exponential smoothing model with damped trend, additive trend, additive seasonality is applied. EXSM_WINTERS_MUL_TREND: Holt-Winters multiplicative exponential smoothing model with multiplicative trend, multiplicative trend, multiplicative seasonality is applied. EXSM_WINTERS_MUL_TREND_DMP: Holt-Winters multiplicative exponential smoothing model with damped multiplicative trend, multiplicative trend, multiplicative seasonality is applied. The default value is EXSM_SIMPLE.

Table 57-18 (Cont.) Exponential Smoothing Settings

Setting Name	Setting Value	Description
EXSM_SEASONALITY	positive integer > 1	This setting specifies a positive integer value as the length of seasonal cycle. The value it takes must be larger than 1. For example, setting value 4 means that every group of four observations forms a seasonal cycle. This setting is only applicable and must be provided for models with seasonality, otherwise the model throws an error. When EXSM_INTERVAL is not set, this setting applies to the original input time series. When EXSM_INTERVAL is set, this setting applies to the accumulated time series.
EXSM_INTERVAL	EXSM_INTERVAL_YEAR EXSM_INTERVAL_QTR EXSM_INTERVAL_MONTH EXSM_INTERVAL_WEEK EXSM_INTERVAL_DAY EXSM_INTERVAL_HOUR EXSM_INTERVAL_MINUTE EXSM_INTERVAL_SECOND	This setting only applies and must be provided when the time column (case_id column) has datetime type. It specifies the spacing interval of the accumulated equally spaced time series. The model throws an error if the time column of input table is of datetime type and setting EXSM_INTERVAL is not provided. The model throws an error if the time column of input table is of oracle number type and setting EXSM_INTERVAL is provided.
EXSM_INITVL_OPTIMIZE	EXSM_INITVL_OPTIMIZE_ENABLE EXSM_INITVL_OPTIMIZE_DISABLE	The setting EXSM_INITVL_OPTIMIZE determines whether initial values are optimized during model build. The default value is EXSM_INITVL_OPTIMIZE_ENABLE.

Note:

EXSM_INITVL_OPTIMIZE can only be set to EXSM_INITVL_OPTIMIZE_DISABLE if the user has set EXSM_MODEL to EXSM_HW or EXSM_HW_ADDSEA. If EXSM_MODEL is set to another model type or is not specified, error 40213 (conflicting settings) is thrown and the model is not built.

Table 57-18 (Cont.) Exponential Smoothing Settings

Setting Name	Setting Value	Description
EXSM_ACCUMULATE	EXSM_ACCU_TOTAL EXSM_ACCU_STD EXSM_ACCU_MAX EXSM_ACCU_MIN EXSM_ACCU_AVG EXSM_ACCU_MEDIAN EXSM_ACCU_COUNT	This setting only applies and must be provided when the time column has datetime type. It specifies how to generate the value of the accumulated time series from the input time series.
EXSM_SETMISSING	Specify an option: EXSM_MISS_MIN EXSM_MISS_MAX EXSM_MISS_AVG EXSM_MISS_MEDIAN EXSM_MISS_LAST EXSM_MISS_FIRST EXSM_MISS_PREV EXSM_MISS_NEXT EXSM_MISS_AUTO	This setting specifies how to handle missing values, which may come from input data and/or the accumulation process of time series. You can specify either a number or an option. If a number is specified, all the missing values are set to that number. EXSM_MISS_MIN: Replaces missing value with minimum of the accumulated time series. EXSM_MISS_MAX: Replaces missing value with maximum of the accumulated time series. EXSM_MISS_AVG: Replaces missing value with average of the accumulated time series. EXSM_MISS_MEDIAN: Replaces missing value with median of the accumulated time series. EXSM_MISS_LAST: Replaces missing value with last non-missing value of the accumulated time series. EXSM_MISS_FIRST: Replaces missing value with first non-missing value of the accumulated time series. EXSM_MISS_PREV: Replaces missing value with the previous non-missing value of the accumulated time series. EXSM_MISS_NEXT: Replaces missing value with the next non-missing value of the accumulated time series. EXSM_MISS_AUTO: EXSM model treats the input data as an irregular (non-uniformly spaced) time series. If this setting is not provided, EXSM_MISS_AUTO is the default value. In such a case, the model treats the input time series as irregular time series, viewing missing values as gaps.
EXSM_PREDICTION_STEP	It must be set to a number between 1-30.	This setting specifies how many steps ahead the predictions are to be made. If it is not set, the default value is 1: the model gives one-step-ahead prediction. A value greater than 30 results in an error.

Table 57-18 (Cont.) Exponential Smoothing Settings

Setting Name	Setting Value	Description
EXSM_CONFIDENCE_LEVEL	It must be a number between 0 and 1, exclusive.	This setting specifies the desired confidence level for prediction. The lower and upper bounds of the specified confidence interval is reported. If this setting is not specified, the default confidence level is 95%.
EXSM_OPT_CRITERION	EXSM_OPT_CRIT_LIK EXSM_OPT_CRIT_MSE EXSM_OPT_CRIT_AMSE EXSM_OPT_CRIT_SIG EXSM_OPT_CRIT_MAE	This setting specifies the desired optimization criterion. The optimization criterion is useful as a diagnostic for comparing models' fit to the same data. EXSM_OPT_CRIT_LIK: Minus twice the log-likelihood of a model. EXSM_OPT_CRIT_MSE: Mean square error of a model. EXSM_OPT_CRIT_AMSE: Average mean square error over user-specified time window. EXSM_OPT_CRIT_SIG: Model's standard deviation of residuals. EXSM_OPT_CRIT_MAE: Mean absolute error of a model. The default value is EXSM_OPT_CRIT_LIK.
EXSM_NMSE	positive integer	This setting specifies the length of the window used in computing the error metric average mean square error (AMSE).
EXSM_SERIES_LIST	Comma delimited list of time series columns	This setting allows you to forecast up to twenty predictor series in addition to the target series. The column names in EXSM_SERIES_LIST are enclosed in single quotes. **Note:** The list is enclosed in single quotes, not the individual column names. For example: `INSERT INTO <settings_table_name VALUES(dbms_data_mining.exsm_series_list, '<column1>,<column2>,<column3>,<column4>');` For the prefix DM$ must be added to the build and scoring data sets and column names must be less than 125 characters long. See Model Detail Views for Exponential Smoothing.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about ESM.

57.5.7 DBMS_DATA_MINING — Algorithm Settings: Generalized Linear Model

The settings listed in the following table configure the behavior of the Generalized Linear Model algorithm.

Table 57-19 DBMS_DATA_MINING GLM Settings

Setting Name	Setting Value	Description
GLMS_CONF_LEVEL	TO_CHAR(0< *numeric_expr* <1)	The confidence level for coefficient confidence intervals. The default confidence level is 0.95.
GLMS_FTR_GEN_METHOD	GLMS_FTR_GEN_QUADRATIC GLMS_FTR_GEN_CUBIC	Whether feature generation is quadratic or cubic. When feature generation is enabled, the algorithm automatically chooses the most appropriate feature generation method based on the data.
GLMS_FTR_GENERATION	GLMS_FTR_GENERATION_ENABLE GLMS_FTR_GENERATION_DISABLE	Whether or not feature generation is enabled for GLM. By default, feature generation is not enabled. **Note:** Feature generation can only be enabled when feature selection is also enabled.
GLMS_FTR_SEL_CRIT	GLMS_FTR_SEL_AIC GLMS_FTR_SEL_SBIC GLMS_FTR_SEL_RIC GLMS_FTR_SEL_ALPHA_INV	Feature selection penalty criterion for adding a feature to the model. When feature selection is enabled, the algorithm automatically chooses the penalty criterion based on the data.
GLMS_FTR_SELECTION	GLMS_FTR_SELECTION_ENABLE GLMS_FTR_SELECTION_DISABLE	Whether or not feature selection is enabled for GLM. By default, feature selection is not enabled.
GLMS_MAX_FEATURES	TO_CHAR(0 < *numeric_expr* <= 2000)	When feature selection is enabled, this setting specifies the maximum number of features that can be selected for the final model. By default, the algorithm limits the number of features to ensure sufficient memory.

Table 57-19 (Cont.) DBMS_DATA_MINING GLM Settings

Setting Name	Setting Value	Description
GLMS_PRUNE_MODEL	GLMS_PRUNE_MODEL_ENABLE GLMS_PRUNE_MODEL_DISABLE	Prune enable or disable for features in the final model. Pruning is based on T-Test statistics for linear regression, or Wald Test statistics for logistic regression. Features are pruned in a loop until all features are statistically significant with respect to the full data. When feature selection is enabled, the algorithm automatically prunes based on the data.
GLMS_REFERENCE_CLASS_NAME	*target_value*	The target value used as the reference class in a binary logistic regression model. Probabilities are produced for the other class. By default, the algorithm chooses the value with the highest prevalence (the most cases) for the reference class.
GLMS_RIDGE_REGRESSION	GLMS_RIDGE_REG_ENABLE GLMS_RIDGE_REG_DISABLE	Enable or disable ridge regression. Ridge applies to both regression and classification machine learning functions. When ridge is enabled, prediction bounds are not produced by the PREDICTION_BOUNDS SQL function. **Note**: Ridge may only be enabled when feature selection is not specified, or has been explicitly disabled. If ridge regression and feature selection are both explicitly enabled, then an exception is raised.
GLMS_RIDGE_VALUE	TO_CHAR (*numeric_expr* > 0)	The value of the ridge parameter. This setting is only used when the algorithm is configured to use ridge regression. If ridge regression is enabled internally by the algorithm, then the ridge parameter is determined by the algorithm.
GLMS_ROW_DIAGNOSTICS	GLMS_ROW_DIAG_ENABLE GLMS_ROW_DIAG_DISABLE (default).	Enable or disable row diagnostics.
GLMS_CONV_TOLERANCE	The range is (0, 1) non-inclusive.	Convergence Tolerance setting of the GLM algorithm The default value is system-determined.
GLMS_NUM_ITERATIONS	Positive integer	Maximum number of iterations for the GLM algorithm. The default value is system-determined.
GLMS_BATCH_ROWS	0 or Positive integer	Number of rows in a batch used by the SGD solver. The value of this parameter sets the size of the batch for the SGD solver. An input of 0 triggers a data driven batch size estimate. The default is 2000

Table 57-19 (Cont.) DBMS_DATA_MINING GLM Settings

Setting Name	Setting Value	Description
GLMS_SOLVER	GLMS_SOLVER_SGD (StochasticGradient Descent) GLMS_SOLVER_CHOL (Cholesky) GLMS_SOLVER_QR GLMS_SOLVER_LBFGS_ADMM	This setting allows the user to choose the GLM solver. The solver cannot be selected if GLMS_FTR_SELECTION setting is enabled. The following are the options: • GLMS_SOLVER_SGD: Optimizes generalized linear models by iteratively updating parameters using a subset of the data to minimize errors. : • GLMS_SOLVER_CHOL: Solves generalized linear models using the Cholesky decomposition method, which provides a stable and efficient solution by transforming the model into a simpler form. • GLMS_SOLVER_QR: Utilizes the QR decomposition technique to solve generalized linear models, ensuring numerical stability and accuracy by decomposing the problem into orthogonal and triangular components. • GLMS_SOLVER_LBFGS_ADMM: Combines L-BFGS, an approximation of the Broyden-Fletcher-Goldfarb-Shanno optimization algorithm, with ADMM for solving large-scale generalized linear model problems efficiently. The default value is system determined.
GLMS_SPARSE_SOLVER	GLMS_SPARSE_SOLVER_ENABLE GLMS_SPARSE_SOLVER_DISABLE (default).	This setting allows the user to use sparse solver if it is available. The default value is GLMS_SPARSE_SOLVER_DISABLE.
GLMS_LINK_FUNCTION	GLMS_IDENTITY_LINK GLMS_LOGIT_LINK GLMS_PROBIT_LINK GLMS_CLOGLOG_LINK GLMS_CAUCHIT_LINK	This setting allows the user to specify the link function for building a GLM model. The link functions are specific to the mining function. For classification, the following are applicable: • GLMS_LOGIT_LINK (default) • GLMS_PROBIT_LINK • GLMS_CLOGLOG_LINK • GLMS_CAUCHIT_LINK For regression, the following is applicable: • GLMS_IDENTITY_LINK (default)

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.
- DBMS_DATA_MINING — Algorithm Settings: Neural Network
  The settings listed in the following table configure the behavior of the Neural Network algorithm.

- DBMS_DATA_MINING — Solver Settings: LBFGS
  The settings listed in the following table configure the behavior of L-BFGS. Neural Network and Generalized Linear Model (GLM) use these settings.
- DBMS_DATA_MINING — Solver Settings: ADMM
  The settings listed in the following table configure the behavior of Alternating Direction Method of Multipliers (ADMM). The Generalized Linear Model (GLM) algorithm uses these settings.
- *Oracle Machine Learning for SQL Concepts*

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about GLM.

57.5.8 DBMS_DATA_MINING — Algorithm Settings: *k*-Means

The settings listed in the following table configure the behavior of the *k*-Means algorithm.

Table 57-20 k-Means Settings

Setting Name	Setting Value	Description
KMNS_CONV_TOLERANCE	TO_CHAR(0<*numeric_expr*<1)	Minimum Convergence Tolerance for *k*-Means. The algorithm iterates until the minimum Convergence Tolerance is satisfied or until the maximum number of iterations, specified in KMNS_ITERATIONS, is reached. Decreasing the Convergence Tolerance produces a more accurate solution but may result in longer run times. The default Convergence Tolerance is 0.001.
KMNS_DISTANCE	KMNS_COSINE KMNS_EUCLIDEAN	Distance function for *k*-Means. The default distance function is KMNS_EUCLIDEAN.
KMNS_ITERATIONS	TO_CHAR(*positive_numeric_expr*)	Maximum number of iterations for *k*-Means. The algorithm iterates until either the maximum number of iterations is reached or the minimum Convergence Tolerance, specified in KMNS_CONV_TOLERANCE, is satisfied. The default number of iterations is 20.
KMNS_MIN_PCT_ATTR_SUPPORT	TO_CHAR(0<=*numeric_expr*<=1)	Minimum percentage of attribute values that must be non-null in order for the attribute to be included in the rule description for the cluster. If the data is sparse or includes many missing values, a minimum support that is too high can cause very short rules or even empty rules. The default minimum support is 0.1.
KMNS_NUM_BINS	TO_CHAR(*numeric_expr*>0)	Number of bins in the attribute histogram produced by *k*-means. The bin boundaries for each attribute are computed globally on the entire training data set. The binning method is equi-width. All attributes have the same number of bins with the exception of attributes with a single value that have only one bin. The default number of histogram bins is 11.

Table 57-20 (Cont.) k-Means Settings

Setting Name	Setting Value	Description
KMNS_SPLIT_CRITERION	KMNS_SIZE KMNS_VARIANCE	Split criterion for *k*-means. The split criterion controls the initialization of new *k*-Means clusters. The algorithm builds a binary tree and adds one new cluster at a time. When the split criterion is based on size, the new cluster is placed in the area where the largest current cluster is located. When the split criterion is based on the variance, the new cluster is placed in the area of the most spread-out cluster. The default split criterion is the KMNS_VARIANCE.
KMNS_RANDOM_SEED	Non-negative integer	This setting controls the seed of the random generator used during the *k*-Means initialization. It must be a non-negative integer value. The default is 0.
KMNS_DETAILS	KMNS_DETAILS_NONE KMNS_DETAILS_HIERARCHY KMNS_DETAILS_ALL	This setting determines the level of cluster detail that are computed during the build. KMNS_DETAILS_NONE: No cluster details are computed. Only the scoring information is persisted. KMNS_DETAILS_HIERARCHY: Cluster hierarchy and cluster record counts are computed. This is the default value. KMNS_DETAILS_ALL: Cluster hierarchy, record counts, descriptive statistics (means, variances, modes, histograms, and rules) are computed.
KMNS_WINSORIZE	KMNS_WINSORIZE_ENABLE KMNS_WINSORIZE_DISABLE	To winorize data, enable or disable this parameter. Data is restricted in a window size of six standard deviations around the mean value when winsorize is enabled. This functionality can be used with AUTO_DATA_PREP turned ON and OFF. The values outside the range are replaced with the ends of the interval. Winsorize is not enabled by default. **Note:** Winsorize is only available when the KMNS_EUCLIDEAN distance function is used. An exception is raised if Winsorize is enabled and other distance functions are set.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.

- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> - For generic machine learning function settings related to Clustering, see DBMS_DATA_MINING — Machine Learning Functions.
> - *Oracle Machine Learning for SQL Concepts* for information about *k*-Means

57.5.9 DBMS_DATA_MINING - Algorithm Settings: Multivariate State Estimation Technique - Sequential Probability Ratio Test

Settings that configure the training calibration behavior of the Multivariate State Estimation Technique - Sequential Probability Ratio Test algorithm.

Table 57-21 MSET-SPRT Settings

Setting Name	Setting Value	Description
MSET_ADB_HEIGHT	A positive double	Estimates the band within which signal values normally oscillate. The default value is 0.05.
MSET_ALERT_COUNT	A positive integer	The number of the last *n* signals (the alert window) that should have passed the threshold to raise an alert. The alert count should be lower or equal to the alert window. The default value is 5.
MSET_ALERT_WINDOW	A positive integer greater than or equal to MSET_ALERT_COUNT	The number of signals to consider in the SPRT hypothesis consolidation logic. The default value is 5.
MSET_ALPHA_PROB	A positive double between 0 and 1	False Alarm Probability FAP (false positive). The default is 0.01.
MSET_BETA_PROB	A positive double between 0 and 1	Missed Alarm Probability MAP (false negative). The default is 0.10.
MSET_HELDASIDE	A positive integer	The approximate number of data rows used for MSET model calibration. You can use ODMS_RANDOM_SEED to change the held-aside sample. The default value is 10000.
MSET_MEMORY_VECTORS	A positive integer	The default value is data driven.

Table 57-21 (Cont.) MSET-SPRT Settings

Setting Name	Setting Value	Description
MSET_PROJECTION_THRESHOLD	A positive integer >0, <=10000	Specifies whether to use random projections. When the number of sensors exceeds the setting value, random projections are used. To turn off random projections, set the threshold to a value that is equal to or greater than the number of sensors. The default value is 500.
MSET_STD_TOLERANCE	A positive integer	The tolerance in standard deviations used in the SPRT calculation. The default value is 3.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

57.5.10 DBMS_DATA_MINING — Algorithm Settings: Naive Bayes

The settings listed in the following table configure the behavior of the Naive Bayes algorithm.

Table 57-22 Naive Bayes Settings

Setting Name	Setting Value	Description
NABS_PAIRWISE_THRESHOLD	TO_CHAR(0<= *numeric_expr* <=1)	Value of pairwise threshold for NB algorithm Default is 0.
NABS_SINGLETON_THRESHOLD	TO_CHAR(0<= *numeric_expr* <=1)	Value of singleton threshold for NB algorithm Default value is 0.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about Naive Bayes

57.5.11 DBMS_DATA_MINING — Algorithm Settings: Neural Network

The settings listed in the following table configure the behavior of the Neural Network algorithm.

Table 57-23 DBMS_DATA_MINING Neural Network Settings

Setting Name	Setting Value	Description
NNET_SOLVER	One of the following strings: • NNET_SOLVER_ADAM • NNET_SOLVER_LBFGS	Specifies the method of optimization. The default value is system determined.

NNET_ACTIVATIONS	One or more of the following strings:
	• NNET_ACTIVATIONS_ARCTAN
	• NNET_ACTIVATIONS_BIPOLAR_SIG
	• NNET_ACTIVATIONS_LINEAR
	• NNET_ACTIVATIONS_LOG_SIG
	• NNET_ACTIVATIONS_RELU
	• NNET_ACTIVATIONS_TANH
		Specifies the activation functions for the hidden layers. You can specify a single activation function, which is then applied to each hidden layer, or you can specify an activation function for each layer individually. Different layers can have different activation functions.
		To apply a different activation function to one or more of the layers, you must specify an activation function for each layer. The number of activation functions you specify must be consistent with the NNET_HIDDEN_LAYERS and NNET_NODES_PER_LAYER values.
		For example, if you have three hidden layers, you could specify the use of the same activation function for all three layers with the following settings value:
		`('NNET_ACTIVATIONS', 'NNET_ACTIVATIONS_TANH')`
		The following settings value specifies a different activation function for each layer:
		`('NNET_ACTIVATIONS', '''NNET_ACTIVATIONS_TANH'', ''NNET_ACTIVATIONS_LOG_SIG'', ''NNET_ACTIVATIONS_ARCTAN''')`
		Note:
		You specify the different activation functions as strings within a single string. All quotes are single and two single quotes are used to escape a single quote in SQL statements and PL/SQL blocks.
		The default value is NNET_ACTIVATIONS_LOG_SIG.

Table 57-23 (Cont.) DBMS_DATA_MINING Neural Network Settings

Setting Name	Setting Value	Description
NNET_HELDASIDE_MAX_FAIL	A positive integer	With NNET_REGULARIZER_HELDASIDE, the training process is stopped early if the network performance on the validation data fails to improve or remains the same for NNET_HELDASIDE_MAX_FAIL epochs in a row. The default value is 6.
NNET_HELDASIDE_RATIO	0 <= *numeric_expr* <=1	Define the held ratio for the held-aside method. The default value is 0.25.
NNET_HIDDEN_LAYERS	A positive integer	Defines the topology by the number of hidden layers. The default value is 1.
NNET_ITERATIONS	A positive integer	Specifies the maximum number of iterations in the Neural Network algorithm. For the DMSSET_NN_SOLVER_LBFGS solver, the default value is 200. For the DMSSET_NN_SOLVER_ADAM solver, the default value is 10000.
NNET_NODES_PER_LAYER	A positive integer or a list of positive integers	Defines the topology by the number of nodes per layer. Different layers can have different numbers of nodes. To specify the same number of nodes for each layer, you can provide a single value, which is then applied to each layer. To specify a different number of nodes for one or more layers, provide a list of comma-separated positive integers, one for each layer. For example, '10, 20, 5' for three layers. The setting values must be consistent with the NNET_HIDDEN_LAYERS value. The default number of nodes per layer is the number of attributes or 50 (if the number of attributes > 50).
NNET_REG_LAMBDA	TO_CHAR(*numeric_expr* >=0)	Defines the L2 regularization parameter lambda. This can not be set together with NNET_REGULARIZER_HELDASIDE. The default value is 1.
NNET_REGULARIZER	One of the following strings: • NNET_REGULARIZER_HELDASIDE • NNET_REGULARIZER_L2 • NNET_REGULARIZER_NONE	Regularization setting for Neural Network algorithm. If the total number of training rows is greater than 50000, the default is NNET_REGULARIZER_HELDASIDE. If the total number of training rows is less than or equal to 50000, the default is NNET_REGULARIZER_NONE.
NNET_TOLERANCE	TO_CHAR(0< *numeric_expr* <1)	Defines the convergence tolerance setting of the Neural Network algorithm. The default value is 0.000001.

Table 57-23 (Cont.) DBMS_DATA_MINING Neural Network Settings

Setting Name	Setting Value	Description
NNET_WEIGHT_LOWER_BOUND	A real number	The setting specifies the lower bound of the region where weights are randomly initialized. NNET_WEIGHT_LOWER_BOUND and NNET_WEIGHT_UPPER_BOUND must be set together. Setting one and not setting the other raises an error. NNET_WEIGHT_LOWER_BOUND must not be greater than NNET_WEIGHT_UPPER_BOUND. The default value is -sqrt(6/(l_nodes+r_nodes)). The value of l_nodes for: • input layer dense attributes is (1+number of dense attributes) • input layer sparse attributes is number of sparse attributes • each hidden layer is (1+number of nodes in that hidden layer) The value of r_nodes is the number of nodes in the layer that the weight is connecting to.
NNET_WEIGHT_UPPER_BOUND	A real number	This setting specifies the upper bound of the region where weights are initialized. It should be set in pairs with NNET_WEIGHT_LOWER_BOUND and its value must not be smaller than the value of NNET_WEIGHT_LOWER_BOUND. If not specified, the values of NNET_WEIGHT_LOWER_BOUND and NNET_WEIGHT_UPPER_BOUND are system determined. The default value is sqrt(6/(l_nodes+r_nodes)). See NNET_WEIGHT_LOWER_BOUND.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.
- DBMS_DATA_MINING — Solver Settings: LBFGS
  The settings listed in the following table configure the behavior of L-BFGS. Neural Network and Generalized Linear Model (GLM) use these settings.

See Also:

Oracle Machine Learning for SQL Concepts for information about Neural Network.

57.5.12 DBMS_DATA_MINING — Algorithm Settings: Non-Negative Matrix Factorization

The settings listed in the following table configure the behavior of the Non-negative Matrix Factorization algorithm.

You can query the data dictionary view *_MINING_MODEL_SETTINGS (using the ALL, USER, or DBA prefix) to find the setting values for a model. See *Oracle Database Reference* for information about *_MINING_MODEL_SETTINGS.

Table 57-24 NMF Settings

Setting Name	Setting Value	Description
NMFS_CONV_TOLERANCE	TO_CHAR(0< *numeric_expr* <=0.5)	Convergence tolerance for NMF algorithm Default is 0.05
NMFS_NONNEGATIVE_SCORING	NMFS_NONNEG_SCORING_ENABLE NMFS_NONNEG_SCORING_DISABLE	Whether negative numbers should be allowed in scoring results. When set to NMFS_NONNEG_SCORING_ENABLE, negative feature values will be replaced with zeros. When set to NMFS_NONNEG_SCORING_DISABLE, negative feature values will be allowed. Default is NMFS_NONNEG_SCORING_ENABLE
NMFS_NUM_ITERATIONS	TO_CHAR(1 <= *numeric_expr* <=500)	Number of iterations for NMF algorithm Default is 50
NMFS_RANDOM_SEED	TO_CHAR(*numeric_expr*)	Random seed for NMF algorithm. Default is -1.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about NMF

57.5.13 DBMS_DATA_MINING — Algorithm Settings: O-Cluster

The settings in the table configure the behavior of the O-Cluster algorithm.

Table 57-25 O-CLuster Settings

Setting Name	Setting Value	Description
OCLT_SENSITIVITY	TO_CHAR(0 <=*numeric_expr* <=1)	A fraction that specifies the peak density required for separating a new cluster. The fraction is related to the global uniform density. Default is 0.5.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about O-Cluster

57.5.14 DBMS_DATA_MINING — Algorithm Settings: Random Forest

These settings configure the behavior of the Random Forest algorithm. Random Forest makes use of the Decision Tree settings to configure the construction of individual trees.

Table 57-26 Random Forest Settings

Setting Name	Setting Value	Description
RFOR_MTRY	a number >= 0	Size of the random subset of columns to be considered when choosing a split at a node. For each node, the size of the pool remains the same, but the specific candidate columns change. The default is half of the columns in the model signature. The special value 0 indicates that the candidate pool includes all columns.
RFOR_NUM_TREES	1<= *a number* <=65535	Number of trees in the forest Default is 20.
RFOR_SAMPLING_RATIO	0< *a fraction*<=1	Fraction of the training data to be randomly sampled for use in the construction of an individual tree. The default is half of the number of rows in the training data.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

- DBMS_DATA_MINING — Algorithm Settings: Decision Tree
  These settings configure the behavior of the Decision Tree algorithm. Note that the Decision Tree settings are also used to configure the behavior of Random Forest as it constructs each individual decision tree.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about Random Forest

57.5.15 DBMS_DATA_MINING — Algorithm Constants and Settings: Singular Value Decomposition

The following settings configure the behavior of the Singular Value Decomposition algorithm.

Table 57-27 Singular Value Decomposition Settings

Setting Name	Setting Value	Description
SVDS_U_MATRIX_OUTPUT	SVDS_U_MATRIX_ENABLE SVDS_U_MATRIX_DISABLE	Indicates whether or not to persist the **U** Matrix produced by SVD. The U matrix in SVD has as many rows as the number of rows in the build data. To avoid creating a large model, the U matrix is persisted only when SVDS_U_MATRIX_OUTPUT is enabled. When SVDS_U_MATRIX_OUTPUT is enabled, the build data must include a case ID. If no case ID is present and the U matrix is requested, then an exception is raised. Default is SVDS_U_MATRIX_DISABLE.
SVDS_SCORING_MODE	SVDS_SCORING_SVD SVDS_SCORING_PCA	Whether to use SVD or PCA scoring for the model. When the build data is scored with SVD, the projections will be the same as the U matrix. When the build data is scored with PCA, the projections will be the product of the U and S matrices. Default is SVDS_SCORING_SVD.
SVDS_SOLVER	SVDS_SOLVER_TSSVD SVDS_SOLVER_TSEIGEN SVDS_SOLVER_SSVD SVDS_SOLVER_STEIGEN	This setting indicates the solver to be used for computing SVD of the data. In the case of PCA, the solver setting indicates the type of SVD solver used to compute the PCA for the data. When this setting is not specified the solver type selection is data driven. If the number of attributes is greater than 3240, then the default wide solver is used. Otherwise, the default narrow solver is selected. The following are the group of solvers: • Narrow data solvers: for matrices with up to 11500 attributes (TSEIGEN) or up to 8100 attributes (TSSVD). • Wide data solvers: for matrices up to 1 million attributes. For narrow data solvers: • Tall-Skinny SVD uses QR computation TSVD (SVDS_SOLVER_TSSVD) • Tall-Skinny SVD uses eigenvalue computation, TSEIGEN (SVDS_SOLVER_TSEIGEN), is the default solver for narrow data. For wide data solvers: • Stochastic SVD uses QR computation SSVD (SVDS_SOLVER_SSVD), is the default solver for wide data solvers. • Stochastic SVD uses eigenvalue computations, STEIGEN (SVDS_SOLVER_STEIGEN).

Table 57-27 (Cont.) Singular Value Decomposition Settings

Setting Name	Setting Value	Description
SVDS_TOLERANCE	Range [0, 1]	This setting is used to prune features. Define the minimum value the eigenvalue of a feature as a share of the first eigenvalue to not to prune. Default value is data driven.
SVDS_RANDOM_SEED	Range [0 - 4,294,967,296]	The random seed value is used for initializing the sampling matrix used by the Stochastic SVD solver. The default is 0. The SVD Solver must be set to SSVD or STEIGEN.
SVDS_OVER_SAMPLING	Range [1, 5000].	This setting is configures the number of columns in the sampling matrix used by the Stochastic SVD solver. The number of columns in this matrix is equal to the requested number of features plus the oversampling setting. The SVD Solver must be set to SSVD or STEIGEN.
SVDS_POWER_ITERATIONS	Range [0, 20].	The power iteration setting improves the accuracy of the SSVD solver. The default is 2. The SVD Solver must be set to SSVD or STEIGEN.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts*

57.5.16 DBMS_DATA_MINING — Algorithm Settings: Support Vector Machine

The settings listed in the following table configure the behavior of the Support Vector Machine algorithm.

Table 57-28 SVM Settings

Setting Name	Setting Value	Description
SVMS_COMPLEXITY_FACTOR	TO_CHAR(*numeric_expr* >0)	Regularization setting that balances the complexity of the model against model robustness to achieve good generalization on new data. SVM uses a data-driven approach to finding the complexity factor. Value of complexity factor for SVM algorithm (both classification and regression). Default value estimated from the data by the algorithm.

Table 57-28 (Cont.) SVM Settings

Setting Name	Setting Value	Description
SVMS_CONV_TOLERANCE	TO_CHAR(*numeric_expr* >0)	Convergence tolerance for SVM algorithm. Default is 0.0001.
SVMS_EPSILON	TO_CHAR(*numeric_expr* >0)	Regularization setting for regression, similar to complexity factor. Epsilon specifies the allowable residuals, or noise, in the data. Value of epsilon factor for SVM regression. Default is 0.1.
SVMS_KERNEL_FUNCTION	SVMS_GAUSSIAN SVMS_LINEAR	Kernel for Support Vector Machine. Linear or Gaussian. The default value is SVMS_LINEAR.
SVMS_OUTLIER_RATE	TO_CHAR(*0< numeric_expr* <1)	The desired rate of outliers in the training data. Valid for One-Class SVM models only (anomaly detection). Default is 0.01.
SVMS_STD_DEV	TO_CHAR(*numeric_expr* >0)	Controls the spread of the Gaussian kernel function. SVM uses a data-driven approach to find a standard deviation value that is on the same scale as distances between typical cases. Value of standard deviation for SVM algorithm. This is applicable only for Gaussian kernel. Default value estimated from the data by the algorithm.
SVMS_NUM_ITERATIONS	Positive integer	This setting sets an upper limit on the number of SVM iterations. The default is system determined because it depends on the SVM solver.
SVMS_NUM_PIVOTS	Range [1; 10000]	This setting sets an upper limit on the number of pivots used in the Incomplete Cholesky decomposition. It can be set only for non-linear kernels. The default value is 200.
SVMS_BATCH_ROWS	Positive integer	This setting applies to SVM models with linear kernel. This setting sets the size of the batch for the SGD solver. An input of 0 triggers a data driven batch size estimate. The default is 20000.
SVMS_REGULARIZER	SVMS_REGULARIZER_L1 SVMS_REGULARIZER_L2	This setting controls the type of regularization that the SGD SVM solver uses. The setting can be used only for linear SVM models. The default is system determined because it depends on the potential model size.
SVMS_SOLVER	SVMS_SOLVER_SGD (Sub-Gradient Descend) SVMS_SOLVER_IPM (Interior Point Method)	This setting allows the user to choose the SVM solver. The SGD solver cannot be selected if the kernel is non-linear. The default value is system determined.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about SVM

57.5.17 DBMS_DATA_MINING — Algorithm Settings: XGBoost

Settings that configure the behavior of the XGBoost gradient boosting algorithm.

The XGBoost settings are case sensitive. Enter the settings as they appear in the settings table. These settings match the XGBoost settings available in open source. OML4SQL XGBoost is based on the 1.7.4 version of XGBoost.
For Global settings, see DBMS_DATA_MINING — Global Settings.

For generic machine learning technique settings, see DBMS_DATA_MINING — Machine Learning Functions.

Table 57-29 General Settings

Setting Name	Setting Value	Description

`xgboost_booster`	A string that is one of the following:
	`dart`
	`gblinear`
	`gbtree`
		The booster to use:
		• `dart`
		• `gblinear`
		• `gbtree`
		The `dart` and `gbtree` boosters use tree-based models whereas `gblinear` uses linear functions.
		The default value is `gbtree`.

`xgboost_num_round`	A non-negative integer.	The number of rounds for boosting. The default value is `10`.

Table 57-30 Settings for Tree Boosting

Setting Name	Setting Value	Description
`xgboost_alpha`	A non-negative number	L1 regularization term on weights. Increasing this value makes the model more conservative. The default value is `0`.
`xgboost_colsample_bylevel`	A number in the range (0, 1]	Subsample ratio of columns for each split, in each level. Subsampling occurs each time a new split is made. This parameter has no effect when `tree_method` is set to `hist`. The default value is `1`.
`xgboost_colsample_bynode`	A number in the range (0, 1]	The subsample ratio of columns for each node (split). Subsampling occurs once every time a new split is evaluated. Columns are subsampled from the set of columns chosen for the current level. The default value is `1`.

Table 57-30 (Cont.) Settings for Tree Boosting

Setting Name	Setting Value	Description
xgboost_colsample_bytree	A number in the range (0, 1]	Subsample ratio of columns when constructing each tree. Subsampling occurs once in every boosting iteration. The default value is 1.
xgboost_eta	A number in the range [0, 1]	Step-size shrinkage used in the update step to prevent overfitting. After each boosting step, eta shrinks the feature weights to make the boosting process more conservative. The default value is 0.3.
xgboost_gamma	A number in the range [0, ∞]	Minimum loss reduction required to make a further partition on a leaf node of the tree. The larger gamma value is, the more conservative the algorithm is. The default value is 0.
xgboost_grow_policy	A string; one of the following: • depthwise • lossguide	Controls the way new nodes are added to the tree: • depthwise splits at nodes closest to the root • lossguide splits at nodes with the highest loss change Valid only if tree_method is set to hist. The default value is depthwise.
xgboost_interaction_constraints	[[x0,x1,x2],[x0,x4],[x5,x6]] where xn are feature names or columns	This setting specifies permitted interactions in the model. Specify the constrains in the form of a nested list where each inner list is a group of features (column names) that are allowed to interact with each other. If a single column is passed in the interactions then, the input is ignored. Here, features x0, x1, and x2 are allowed to interact with each other but with no other feature. Similarly, x0 and x4 are allowed to interact with each other but with no other feature and so on. This setting is applicable to 2-Dimensional features. An error occurs if you pass columns of non-supported type and non-existing feature names.
xgboost_lambda	A non-negative number	L2 regularization term on weights. The default value is 1.
xgboost_max_bin	A non-negative integer	Maximum number of discrete bins to bucket continuous features. Increasing this number improves the optimality of splits at the cost of higher computation time. This parameter is valid only when tree_method is set to hist. The default value is 256.

Table 57-30 (Cont.) Settings for Tree Boosting

Setting Name	Setting Value	Description
xgboost_max_delta_step	A number in the range [0, ∞]	Maximum delta step allowed for each leaf output. Setting this to a positive value can help make the update step more conservative. Usually this parameter is not needed, but it might help in logistic regression when the class is extremely imbalanced. Setting it to value from 1 to 10 might help control the update. The default value is 0, which means there is no constraint.
xgboost_max_depth	An integer in the range [0, ∞]	Maximum depth of a tree. Increasing this value makes the model more complex and more likely to overfit. Setting this to 0 indicates no limit. **Note:** You must set a max_depth limit when the grow_policy setting is depthwise. The default value is 6.
xgboost_max_leaves	A non-negative number	Maximum number of nodes to add. Use this setting only when grow_policy is set to lossguide. The default value is 0.
xgboost_min_child_weight	A number in the range [0, ∞]	Minimum sum of instance weight (hessian) needed in a child. If the tree partition step results in a leaf node with a sum of instance weight less than min_child_weight, then the building process stops partitioning. In a linear regression task, this corresponds to the minimum number of instances needed in each node. The larger min_child_weight is, the more conservative the algorithm is. The default value is 1.
xgboost_decrease_constraints	[x4,x5]	This setting specifies the features (column names) that must obey decreasing constraint. The feature names are separated by a comma. For example, setting value 'x4,x5' sets decreasing constraint on features x4 and x5. This setting applies to numeric columns and 2-Dimensional features. An error occurs if you pass columns of non-supported type and non-existing feature names.

Table 57-30 (Cont.) Settings for Tree Boosting

Setting Name	Setting Value	Description
xgboost_increase_constraints	[x0,x3]	This setting specifies the features (column names) that must obey increasing constraint. The feature names are separated by a comma. For example, setting value 'x0,x3' sets increasing constraint on features x0 and x3. This setting is applicable to 2-Dimensional features. An error occurs if you pass columns of non-supported type and non-existing feature names.
xgboost_num_parallel_tree	A non-negative integer	Number of parallel trees constructed during each iteration. Use this option to support a boosted random forest. The default value is 1.
xgboost_scale_pos_weight	A non-negative number	Controls the balance of positive and negative weights, which is useful for unbalanced classes. A typical value to consider: sum(negative cases) / sum(positive cases). The default value is 1.
xgboost_sketch_eps	A number in the range (0, 1)	Increases enumeration accuracy. Valid only for the approximate greedy tree method. Compared to directly selecting the number of bins, this setting comes with a theoretical guarantee with sketch accuracy. You usually do not need to change this setting, but you might consider setting a lower number for more accurate enumeration. The default value is 0.03.
xgboost_subsample	A number in the range (0, 1]	Subsample ratio of the training instances. A setting of 0.5 means that XGBoost randomly samples half of the training data prior to growing trees, which prevents overfitting. Subsampling occurs once in every boosting iteration. The default value is 1.
xgboost_tree_method	A string that is one of the following: • approx • auto • exact • hist	Tree construction algorithm used in XGBoost: • approx: Approximate greedy algorithm using sketching and histogram. • auto: Use a heuristic to choose the faster algorithm: – For a small to medium sized data set, uses the exact greedy algorithm. – For a very large data set, uses the approximate greedy algorithm. • exact: Exact greedy algorithm. • hist: Fast histogram optimized approximate greedy algorithm; uses some performance improvements such as bins caching. The default value is auto.

Table 57-30 (Cont.) Settings for Tree Boosting

Setting Name	Setting Value	Description
xgboost_updater	A comma-separated string; one or more of the following: • grow_colmaker • grow_histmaker • grow_skmaker • grow_quantile_histmaker • prune • sync	Defines the sequence of tree updaters to run, which provides a modular way to construct and to modify the trees. This is an advanced parameter that is usually set automatically, depending on some other parameters. However, you can also explicitly specify a settting. The setting values are: • grow_colmaker: Non-distributed column-based construction of trees. • grow_histmaker: Distributed tree construction with row-based data splitting based on a global proposal of histogram counting. • grow_skmaker: Uses the approximate sketching algorithm. • grow_quantile_histmaker: Grow tree using quantized histogram. • prune: Prunes the splits where loss < min_split_loss (or gamma). • sync: Synchronizes trees in all distributed nodes.

Table 57-31 Settings for the Dart Booster

Setting Name	Setting Value	Description
xgboost_one_drop	A number that is 0 or 1	When set to 1, at least one tree is always dropped during the dropout. When set to 0, at least one tree is not always dropped during the dropout. The default value is 0.
xgboost_normalize_type	A string; either: • forest • tree	Type of normalization algorithm: • forest: New trees have the same weight as the sum of the dropped trees (forest): – The weight of new trees is 1 / (1 + learning_rate) – Dropped trees are scaled by a factor of 1 / (1 + learning_rate) • tree: New trees have the same weight as dropped trees: – The weight of new trees is 1 / (k + learning_rate) – Dropped trees are scaled by a factor of k / (k + learning_rate) The default value is tree.
xgboost_rate_drop	A number in the range [0.0, 1.0]	Dropout rate (a fraction of the previous trees to drop during the dropout). The default value is 0.0.
xgboost_sample_type	A string; either: • uniform • weighted	Type of sampling algorithm: • uniform: Dropped trees are selected uniformly • weighted: Dropped trees are selected in proportion to weight The default value is uniform.

Table 57-31 (Cont.) Settings for the Dart Booster

Setting Name	Setting Value	Description
xgboost_skip_drop	A number in the range [0.0, 1.0]	Probability of skipping the dropout procedure during a boosting iteration. If a dropout is skipped, new trees are added in the same manner as gbtree. A non-zero skip_drop has higher priority than rate_drop or one_drop. The default value is 0.0.

Table 57-32 Settings for the Linear Booster

Setting Name	Setting Value	Description
xgboost_alpha	A non-negative number	L1 regularization term on weights, normalized to the number of training examples. Increasing this value makes the model more conservative. The default value is 0.
xgboost_feature_selector	A string that is one of the following: • cyclic • greedy • random • shuffle • thrifty	Feature selection and ordering method: • cyclic: Deterministic selection by cycling through the features one at a time. • greedy: Selects the coordinate with the greatest gradient magnitude. This method: – Has O(num_feature^2) complexity – Is fully deterministic – Allows restricting the selection to the top_k features per group with the largest magnitude of univariate weight change, by setting the top_k parameter; doing so reduces the complexity to O(num_feature*top_k). • random: A random (with replacement) coordinate selector. • shuffle: Similar to cyclic but with random feature shuffling prior to each update. • thrifty: Thrifty, approximately-greedy feature selector. Prior to cyclic updates, reorders features in descending magnitude of their univariate weight changes. This operation is multithreaded and is a linear complexity approximation of the quadratic greedy selection. Restricts the selection per group to the top_k features with the largest magnitude of univariate weight change. The default value is cyclic.
xgboost_lambda	A non-negative number	L2 regularization term on weights, normalized to the number of training examples. Increasing this value makes the model more conservative. The default value is 0.
xgboost_top_k	A non-negative integer	Number of top features to select for the greedy or thrifty feature selector. The value of 0 uses all of the features. The default value is 0.

Table 57-32 (Cont.) Settings for the Linear Booster

Setting Name	Setting Value	Description
xgboost_updater	A string that is one of the following: • coord_descent • shotgun	Algorithm to fit the linear model: • coord_descent: Ordinary coordinate descent algorithm; multithreaded but still produces a deterministic solution. • shotgun: Parallel coordinate descent algorithm based on the shotgun algorithm; uses "hogwild" parallelism and therefore produces a nondeterministic solution on each run. The default value is shotgun.

Table 57-33 Settings for Tweedie Regression

Setting Name	Setting Value	Description
xgboost_tweedie_variance_power	A number in the range (1, 2)	Controls the variance of the Tweedie distribution var(y) ~ E(y)^tweedie_variance_power. A setting closer to 1 shifts towards a Poisson distribution. A setting closer to 2 shifts towards a gamma distribution. The default value is 1.5.

Some XGBoost objectives apply only to classification function models and other objectives apply only to regression function models. If you specify an incompatible objective value, an error is raised. In the DBMS_DATA_MINING.CREATE_MODEL procedure, if you specify DBMS_DATA_MINING.CLASSIFICATION as the function, then the only objective values that you can use are the binary and multi values. The one exception is binary: logitraw, which produces a continuous value and applies only to a regression model. If you specify DBMS_DATA_MINING.REGRESSION as the function, then you can specify binary: logitraw or any of the count, rank, reg, and survival values as the objective.

Table 57-34 Settings for Learning Tasks

Setting Name	Setting Value	Description
xgboost_objective	For a classification model, a string that is one of the following: • binary:hinge • binary:logistic • multi:softmax • multi:softprob For a regression model, a string that is one of the following: • binary:logitraw • count:poisson • rank:map • rank:ndcg • rank:pairwise • reg:gamma • reg:logistic • reg:tweedie • survival:aft • survival:cox • reg:squarederror • reg:squaredlogerror	**Settings for a Classification model:** • binary:hinge: Hinge loss for binary classification. This setting makes predictions of 0 or 1, rather than producing probabilities. • binary:logistic: Logistic regression for binary classification. The output is the probability. • multi:softmax: Performs multiclass classification using the softmax objective; you must also set num_class(*number_of_classes*). • multi:softprob: : Same as softmax, except the output is a vector of ndata * nclass, which can be further reshaped to an ndata * nclass matrix. The result contains the predicted probability of each data point belonging to each class. The default objective value for classification is multi:softprob. **Settings for a Regression model:** • binary:logitraw: Logistic regression for binary classification; the output is the score before logistic transformation. • count:poisson: Poisson regression for count data; the output is the mean of the Poisson distribution. The max_delta_step value is set to 0.7 by default in Poisson regression to safeguard optimization. • rank:map: Using LambdaMART, performs list-wise ranking in which the Mean Average Precision (MAP) is maximized. • rank:ndcg: Using LambdaMART, performs list-wise ranking in which the Normalized Discounted Cumulative Gain (NDCG) is maximized. • rank:pairwise: Performs ranking by minimizing the pairwise loss. • reg:gamma: Gamma regression with log-link; the output is the mean of the gamma distribution. This setting might be useful for any outcome that might be gamma-distributed, such as modeling insurance claims severity. • reg:logistic: Logistic regression. • reg:tweedie: Tweedie regression with log-link. This setting might be useful for any outcome that might be Tweedie-distributed, such as modeling total loss in insurance. • survival:aft: Applies the Accelerated Failure Time (AFT) model for censored survival time data. When you select this option, eval_metric uses aft-nloglik as the default value. • survival:cox: Cox regression for right-censored survival time data (negative values are considered right-censored). Predictions are returned on the hazard ratio scale (that is, as HR = exp(marginal_prediction) in the proportional hazard function h(t) = h0(t) * HR).

Table 57-34 (Cont.) Settings for Learning Tasks

Setting Name	Setting Value	Description
		• `reg:squarederror`: Regression with squared loss. • `reg:squaredlogerror`: Regression with squared log loss. All input labels must be greater than -1. The default `objective` value for regression is `reg:squarederror`.
`xgboost_aft_loss_distribution`	[normal, logistic, extreme]	Specifies the distribution of the Z term in the AFT model. It specifies the Probabilty Density Function used by `survival:aft` objective and `aft-nloglik` evaluation metric. The default value is `normal`.
`xgboost_aft_loss_distribution_scale`	A positive number	Specifies the scaling factor σ, which scales the size of Z term in the AFT model. The default value is `1`.
`xgboost_aft_right_bound_column_name`	*column_name*	Specifies the column containing the right bounds of the labels for an AFT model. You cannot select this parameter for a non-AFT model. **Note:** Oracle Machine Learning does not support `BOOLEAN` values for this setting.
`xgboost_base_score`	A number	Initial prediction score of all instances, global bias. For a sufficient number of iterations, changing this value does not have much effect. The default value is `0.5`.

Table 57-34 (Cont.) Settings for Learning Tasks

Setting Name	Setting Value	Description
xgboost_eval_metric	A comma-separated string; one or more of the following: • aft-nloglik • auc • aucpr • cox-nloglik • error • error@*t* • gamma-deviance • gamma-nloglik • logloss • mae • map • map@*n* • merror • mlogloss • ndcg • ndcg@*n* • poisson-nloglik • rmse • tweedie-nloglik@*rho* • ndcg- • map- • rmsle	Evaluation metrics for validation data. You can specify one or more of these evaluation metrics: • aft-nloglik: Sets the eval_metric to negative log likelihood of AFT model. • auc: Area under the curve. • aucpr: Area under the PR curve. • cox-nloglik: Negative partial log-likelihood for Cox proportional hazards regression. • error: Binary classification error rate, calculated as the number of wrong cases divided by the number of all cases. For the predictions, the evaluation regards the instances with a prediction value larger than 0.5 as positive instances, and the others as negative instances. • error@*t*: You can specify a binary classification threshold value other than 0.5 by specifying a numerical value *t*; for example, error@0.8. • gamma-deviance: Residual deviance for gamma regression. • gamma-nloglik: Negative log-likelihood for gamma regression. • logloss: Negative log-likelihood. • mae: Mean absolute error. • map: Mean average precision. • map@*n*: Assigns the integer *n* as the cut-off value for the top positions in the lists for evaluation. • merror: Multiclass classification error rate calculated as the number of wrong cases divided by the number of all cases; the objective must be multi:softprob or multi:softmax. • mlogloss: Multiclass logloss; the objective must be multi:softprob or multi:softmax. • ndcg: Normalized Discounted Cumulative Gain. • ndcg@*n*: Assigns the integer *n* as the cut-off value for the top positions in the lists for evaluation. • poisson-nloglik: Negative log-likelihood for Poisson regression • rmse: Root Mean Square Error. • tweedie-nloglik@*rho*: Negative log-likelihood for Tweedie regression (at a specified value rho of the tweedie_variance_power parameter); rho must be a number in the range (1, 2); for example, tweedie-nloglik@1.8. • ndcg- and map-: In XGBoost, NDCG and MAP will evaluate the score of a list without any positive samples as 1. By adding "-" in the evaluation metric XGBoost will evaluate these score as 0 to be consistent under some conditions. • rmsle: It is root mean square log error. This is the default metric of reg:squaredlogerror objective. This metric reduces errors generated by outliers in

Table 57-34 (Cont.) Settings for Learning Tasks

Setting Name	Setting Value	Description
		dataset. But because log function is employed, `rmsle` might output nan when prediction value is less than -1. A default metric is assigned according to the objective: • `error` for classification • `mean average precision` for ranking • `rmse` for regression
`xgboost_seed`	A non-negative integer	Random number seed. The default value is `0`.

Related Topics

- DBMS_DATA_MINING — Machine Learning Functions
  A machine learning **function** refers to the methods for solving a given class of machine learning problems.
- DBMS_DATA_MINING — Global Settings
  The configuration settings in this table are applicable to any type of model, but are currently only implemented for specific algorithms.

> **See Also:**
>
> https://github.com/oracle/oracle-db-examples/tree/master/machine-learning/sql/, select the release, and browse for an example of XGBoost.

57.6 DBMS_DATA_MINING — Solver Settings

Oracle Machine Learning for SQL algorithms can use different solvers. Solver settings can be provided at build time in the settings table.

Related Topics

- DBMS_DATA_MINING - Solver Settings: Adam
  These settings configure the behavior of the Adaptive Moment Estimation (Adam) solver.
- DBMS_DATA_MINING — Solver Settings: ADMM
  The settings listed in the following table configure the behavior of Alternating Direction Method of Multipliers (ADMM). The Generalized Linear Model (GLM) algorithm uses these settings.
- DBMS_DATA_MINING — Solver Settings: LBFGS
  The settings listed in the following table configure the behavior of L-BFGS. Neural Network and Generalized Linear Model (GLM) use these settings.

57.6.1 DBMS_DATA_MINING - Solver Settings: Adam

These settings configure the behavior of the Adaptive Moment Estimation (Adam) solver.

Neural Network models use these settings.

Table 57-35 DBMS_DATA_MINING Adam Settings

Setting Name	Setting Value	Description
ADAM_ALPHA	A non-negative double precision floating point number in the interval (0; 1]	The learning rate for Adam. The default value is 0.001.
ADAM_BATCH_ROWS	A positive integer	The number of rows per batch. The default value is 10000.
ADAM_BETA1	A positive double precision floating point number in the interval [0; 1)	The exponential decay rate for the 1st moment estimates. The default value is 0.9.
ADAM_BETA2	A positive double precision floating point number in the interval [0; 1)	The exponential decay rate for the 2nd moment estimates. The default value is 0.99.
ADAM_GRADIENT_TOLERANCE	A positive double precision floating point number	The gradient infinity norm tolerance for Adam. The default value is 1E-9.

Related Topics

- DBMS_DATA_MINING — Algorithm Settings: Neural Network
  The settings listed in the following table configure the behavior of the Neural Network algorithm.

57.6.2 DBMS_DATA_MINING — Solver Settings: ADMM

The settings listed in the following table configure the behavior of Alternating Direction Method of Multipliers (ADMM). The Generalized Linear Model (GLM) algorithm uses these settings.

Table 57-36 DBMS_DATA_MINING ADMM Settings

Settings Name	Setting Value	Description
ADMM_CONSENSUS	A positive integer	It is a ADMM's consensus parameter. The value must be a positive number. The default value is 0.1.
ADMM_ITERATIONS	A positive integer	The number of ADMM iterations. The value must be a positive integer. The default value is 50.
ADMM_TOLERANCE	A positive integer	It is a tolerance parameter. The value must be a positive number. The default value is 0.0001

Related Topics

- DBMS_DATA_MINING — Algorithm Settings: Generalized Linear Model
  The settings listed in the following table configure the behavior of the Generalized Linear Model algorithm.
- *Oracle Machine Learning for SQL Concepts*

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about neural network

57.6.3 DBMS_DATA_MINING — Solver Settings: LBFGS

The settings listed in the following table configure the behavior of L-BFGS. Neural Network and Generalized Linear Model (GLM) use these settings.

Table 57-37 DBMS_DATA_MINING L-BFGS Settings

Setting Name	Setting Value	Description
LBFGS_GRADIENT_TOLERANCE	TO_CHAR (*numeric_expr* >0)	Defines gradient infinity norm tolerance for L-BFGS. Default value is 1E-9.
LBFGS_HISTORY_DEPTH	The value must be a positive integer.	Defines the number of historical copies kept in L-BFGS solver. The default value is 20.
LBFGS_SCALE_HESSIAN	LBFGS_SCALE_HESSIAN_ENABLE LBFGS_SCALE_HESSIAN_DISABLE	Defines whether to scale Hessian in L-BFGS or not. Default value is LBFGS_SCALE_HESSIAN_ENABLE.

Related Topics

- DBMS_DATA_MINING — Algorithm Settings: Neural Network
  The settings listed in the following table configure the behavior of the Neural Network algorithm.
- DBMS_DATA_MINING — Algorithm Settings: Generalized Linear Model
  The settings listed in the following table configure the behavior of the Generalized Linear Model algorithm.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about neural network

57.7 DBMS_DATA_MINING Datatypes

The DBMS_DATA_MINING package defines object data types for processing transactional data. The package also defines a type for user-specified transformations. These types are called DM_NESTED_*n*, where *n* identifies the Oracle data type of the nested attributes.

The Oracle Machine Learning for SQL object data types are described in the following table:

Table 57-38 DBMS_DATA_MINING Summary of Data Types

Datatype	Description
DM_NESTED_BINARY_DOUBLE	The name and value of a numerical attribute of type BINARY_DOUBLE.
DM_NESTED_BINARY_DOUBLES	A collection of DM_NESTED_BINARY_DOUBLE.
DM_NESTED_BINARY_FLOAT	The name and value of a numerical attribute of type BINARY_FLOAT.
DM_NESTED_BINARY_FLOATS	A collection of DM_NESTED_BINARY_FLOAT.
DM_NESTED_CATEGORICAL	The name and value of a categorical attribute of type CHAR, VARCHAR, or VARCHAR2.
DM_NESTED_CATEGORICALS	A collection of DM_NESTED_CATEGORICAL.
DM_NESTED_NUMERICAL	The name and value of a numerical attribute of type NUMBER or FLOAT.
DM_NESTED_NUMERICALS	A collection of DM_NESTED_NUMERICAL.
ORA_MINING_VARCHAR2_NT	A table of VARCHAR2(4000).
TRANSFORM_LIST	A list of user-specified transformations for a model. Accepted as a parameter by the CREATE_MODEL Procedure. This collection type is defined in the DBMS_DATA_MINING_TRANSFORM package.

For more information about processing nested data, see *Oracle Machine Learning for SQL User's Guide*.

> **Note:**
>
> Starting from Oracle Database 12*c* Release 2, *GET_MODEL_DETAILS are deprecated and are replaced with *Model Detail Views*. See *Oracle Machine Learning for SQL User's Guide*.

57.7.1 Deprecated Types

This topic contains tables listing deprecated types.

The DBMS_DATA_MINING package defines object datatypes for storing information about model attributes. Most of these types are returned by the table functions GET_*n*, where *n* identifies the type of information to return. These functions take a model name as input and return the requested information as a collection of rows.

For a list of the GET functions, see "Summary of DBMS_DATA_MINING Subprograms".

All the table functions use pipelining, which causes each row of output to be materialized as it is read from model storage, without waiting for the generation of the complete table object. For more information on pipelined, parallel table functions, consult the *Oracle Database PL/SQL Language Reference.*

Table 57-39 DBMS_DATA_MINING Summary of Deprecated Datatypes

Datatype	Description
DM_CENTROID	The centroid of a cluster.
DM_CENTROIDS	A collection of DM_CENTROID. A member of DM_CLUSTER.
DM_CHILD	A child node of a cluster.
DM_CHILDREN	A collection of DM_CHILD. A member of DM_CLUSTER.
DM_CLUSTER	A cluster. A cluster includes DM_PREDICATES, DM_CHILDREN, DM_CENTROIDS, and DM_HISTOGRAMS. It also includes a DM_RULE. See also, DM_CLUSTER Fields.
DM_CLUSTERS	A collection of DM_CLUSTER. Returned by GET_MODEL_DETAILS_KM Function, GET_MODEL_DETAILS_OC Function, and GET_MODEL_DETAILS_EM Function. See also, DM_CLUSTER Fields.
DM_CONDITIONAL	The conditional probability of an attribute in a Naive Bayes model.
DM_CONDITIONALS	A collection of DM_CONDITIONAL. Returned by GET_MODEL_DETAILS_NB Function.
DM_COST_ELEMENT	The actual and predicted values in a cost matrix.
DM_COST_MATRIX	A collection of DM_COST_ELEMENT. Returned by GET_MODEL_COST_MATRIX Function.
DM_EM_COMPONENT	A component of an Expectation Maximization model.
DM_EM_COMPONENT_SET	A collection of DM_EM_COMPONENT. Returned by GET_MODEL_DETAILS_EM_COMP Function.
DM_EM_PROJECTION	A projection of an Expectation Maximization model.
DM_EM_PROJECTION_SET	A collection of DM_EM_PROJECTION. Returned by GET_MODEL_DETAILS_EM_PROJ Function.
DM_GLM_COEFF	The coefficient and associated statistics of an attribute in a Generalized Linear Model.
DM_GLM_COEFF_SET	A collection of DM_GLM_COEFF. Returned by GET_MODEL_DETAILS_GLM Function.
DM_HISTOGRAM_BIN	A histogram associated with a cluster.
DM_HISTOGRAMS	A collection of DM_HISTOGRAM_BIN. A member of DM_CLUSTER. See also, DM_CLUSTER Fields.
DM_ITEM	An item in an association rule.
DM_ITEMS	A collection of DM_ITEM.
DM_ITEMSET	A collection of DM_ITEMS.
DM_ITEMSETS	A collection of DM_ITEMSET. Returned by GET_FREQUENT_ITEMSETS Function.
DM_MODEL_GLOBAL_DETAIL	High-level statistics about a model.
DM_MODEL_GLOBAL_DETAILS	A collection of DM_MODEL_GLOBAL_DETAIL. Returned by GET_MODEL_DETAILS_GLOBAL Function.
DM_NB_DETAIL	Information about an attribute in a Naive Bayes model.

Table 57-39 (Cont.) DBMS_DATA_MINING Summary of Deprecated Datatypes

Datatype	Description
DM_NB_DETAILS	A collection of DM_DB_DETAIL. Returned by GET_MODEL_DETAILS_NB Function.
DM_NMF_ATTRIBUTE	An attribute in a feature of a Non-Negative Matrix Factorization model.
DM_NMF_ATTRIBUTE_SET	A collection of DM_NMF_ATTRIBUTE. A member of DM_NMF_FEATURE.
DM_NMF_FEATURE	A feature in a Non-Negative Matrix Factorization model.
DM_NMF_FEATURE_SET	A collection of DM_NMF_FEATURE. Returned by GET_MODEL_DETAILS_NMF Function.
DM_PREDICATE	Antecedent and consequent in a rule.
DM_PREDICATES	A collection of DM_PREDICATE. A member of DM_RULE and DM_CLUSTER. Predicates are returned by GET_ASSOCIATION_RULES Function, GET_MODEL_DETAILS_EM Function, GET_MODEL_DETAILS_KM Function, and GET_MODEL_DETAILS_OC Function. See also, DM_CLUSTER Fields.
DM_RANKED_ATTRIBUTE	An attribute ranked by its importance in an Attribute Importance model.
DM_RANKED_ATTRIBUTES	A collection of DM_RANKED_ATTRIBUTE. Returned by GET_MODEL_DETAILS_AI Function.
DM_RULE	A rule that defines a conditional relationship. The rule can be one of the association rules returned by GET_ASSOCIATION_RULES Function, or it can be a rule associated with a cluster in the collection of clusters returned by GET_MODEL_DETAILS_KM Function and GET_MODEL_DETAILS_OC Function. See also, DM_CLUSTER Fields.
DM_RULES	A collection of DM_RULE. Returned by GET_ASSOCIATION_RULES Function. See also, DM_CLUSTER Fields.
DM_SVD_MATRIX	A factorized matrix S, V, or U returned by a Singular Value Decomposition model.
DM_SVD_MATRIX_SET	A collection of DM_SVD_MATRIX. Returned by GET_MODEL_DETAILS_SVD Function.
DM_SVM_ATTRIBUTE	The name, value, and coefficient of an attribute in a Support Vector Machine model.
DM_SVM_ATTRIBUTE_SET	A collection of DM_SVM_ATTRIBUTE. Returned by GET_MODEL_DETAILS_SVM Function. Also a member of DM_SVM_LINEAR_COEFF.
DM_SVM_LINEAR_COEFF	The linear coefficient of each attribute in a Support Vector Machine model.
DM_SVM_LINEAR_COEFF_SET	A collection of DM_SVM_LINEAR_COEFF. Returned by GET_MODEL_DETAILS_SVM Function for an SVM model built using the linear kernel.

Table 57-39 (Cont.) DBMS_DATA_MINING Summary of Deprecated Datatypes

Datatype	Description
DM_TRANSFORM	The transformation and reverse transformation expressions for an attribute.
DM_TRANSFORMS	A collection of DM_TRANSFORM. Returned by GET_MODEL_TRANSFORMATIONS Function.

Return Values for Clustering Algorithms

The table contains description of DM_CLUSTER return value columns, nested table columns, and rows.

Table 57-40 DM_CLUSTER Return Values for Clustering Algorithms

Return Value	Description
DM_CLUSTERS	A set of rows of type DM_CLUSTER. The rows have the following columns: (id NUMBER, cluster_id VARCHAR2(4000), record_count NUMBER, parent NUMBER, tree_level NUMBER, dispersion NUMBER, split_predicate DM_PREDICATES, child DM_CHILDREN, centroid DM_CENTROIDS, histogram DM_HISTOGRAMS, rule DM_RULE)
DM_PREDICATE	The antecedent and consequent columns each return nested tables of type DM_PREDICATES. The rows, of type DM_PREDICATE, have the following columns: (attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), conditional_operator CHAR(2)/*=,<>,<,>,<=,>=*/, attribute_num_value NUMBER, attribute_str_value VARCHAR2(4000), attribute_support NUMBER, attribute_confidence NUMBER)

DM_CLUSTER Fields

The following table describes DM_CLUSTER fields.

Table 57-41 DM_CLUSTER Fields

Column Name	Description
id	Cluster identifier
cluster_id	The ID of a cluster in the model
record_count	Specifies the number of records
parent	Parent ID

Table 57-41 (Cont.) DM_CLUSTER Fields

Column Name	Description
tree_level	Specifies the number of splits from the root
dispersion	A measure used to quantify whether a set of observed occurrences are dispersed compared to a standard statistical model.
split_predicate	The split_predicate column of DM_CLUSTER returns a nested table of type DM_PREDICATES. Each row, of type DM_PREDICATE, has the following columns: (attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), conditional_operator CHAR(2) / *=,<>,<,>,<=,>=*/, attribute_num_value NUMBER, attribute_str_value VARCHAR2(4000), attribute_support NUMBER, attribute_confidence NUMBER) Note: The Expectation Maximization algorithm uses all the fields except dispersion and split_predicate.
child	The child column of DM_CLUSTER returns a nested table of type DM_CHILDREN. The rows, of type DM_CHILD, have a single column of type NUMBER, which contains the identifiers of each child.
centroid	The centroid column of DM_CLUSTER returns a nested table of type DM_CENTROIDS. The rows, of type DM_CENTROID, have the following columns: (attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), mean NUMBER, mode_value VARCHAR2(4000), variance NUMBER)
histogram	The histogram column of DM_CLUSTER returns a nested table of type DM_HISTOGRAMS. The rows, of type DM_HISTOGRAM_BIN, have the following columns: (attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), bin_id NUMBER, lower_bound NUMBER, upper_bound NUMBER, label VARCHAR2(4000), count NUMBER)

Table 57-41 (Cont.) DM_CLUSTER Fields

Column Name	Description
rule	The rule column of DM_CLUSTER returns a single row of type DM_RULE. The columns are: (rule_id INTEGER, antecedent DM_PREDICATES, consequent DM_PREDICATES, rule_support NUMBER, rule_confidence NUMBER, rule_lift NUMBER, antecedent_support NUMBER, consequent_support NUMBER, number_of_items INTEGER)

Usage Notes

- The table function pipes out rows of type DM_CLUSTER. For information on Oracle Machine Learning for SQL data types and piped output from table functions, see "Data Types".
- For descriptions of predicates (DM_PREDICATE) and rules (DM_RULE), see GET_ASSOCIATION_RULES Function.

57.8 Summary of DBMS_DATA_MINING Subprograms

This table summarizes the subprograms included in the DBMS_DATA_MINING package.

The GET_* interfaces are replaced by model views. Oracle recommends that users leverage model detail views instead. For more information, refer to Model Detail Views in *Oracle Machine Learning for SQL User's Guide* and Static Data Dictionary Views: ALL_ALL_TABLES to ALL_OUTLINES in *Oracle Database Reference*.

Table 57-42 DBMS_DATA_MINING Package Subprograms

Subprogram	Purpose
ADD_COST_MATRIX Procedure	Adds a cost matrix to a classification model
ADD_PARTITION Procedure	Adds single or multiple partitions in an existing partition model
ALTER_REVERSE_EXPRESSION Procedure	Changes the reverse transformation expression to an expression that you specify
APPLY Procedure	Applies a model to a data set (scores the data)
COMPUTE_CONFUSION_MATRIX Procedure	Computes the confusion matrix for a classification model
COMPUTE_CONFUSION_MATRIX_PART Procedure	Computes the evaluation matrix for partitioned models
COMPUTE_LIFT Procedure	Computes lift for a classification model
COMPUTE_LIFT_PART Procedure	Computers lift for partitioned models
COMPUTE_ROC Procedure	Computes Receiver Operating Characteristic (ROC) for a classification model
COMPUTE_ROC_PART Procedure	Computes Receiver Operating Characteristic (ROC) for a partitioned model

Table 57-42 (Cont.) DBMS_DATA_MINING Package Subprograms

Subprogram	Purpose
CREATE_MODEL Procedure	Creates a model
CREATE_MODEL2 Procedure	Creates a model without extra persistent stages
Create Model Using Registration Information	Fetches setting information from JSON object
DROP_ALGORITHM Procedure	Drops the registered algorithm information.
DROP_PARTITION Procedure	Drops a single partition
DROP_MODEL Procedure	Drops a model
EXPORT_MODEL Procedure	Exports a model to a dump file
EXPORT_SERMODEL Procedure	Exports a model in a serialized format
FETCH_JSON_SCHEMA Procedure	Fetches and reads JSON schema from `all_mining_algorithms` view
GET_MODEL_COST_MATRIX Function	Returns the cost matrix for a model
IMPORT_MODEL Procedure	Imports a model into a user schema
IMPORT_ONNX_MODEL Procedure	Imports an ONNX model into the Database
IMPORT_SERMODEL Procedure	Imports a serialized model back into the database
JSON Schema for R Extensible Algorithm	Displays flexibility in creating JSON schema for R Extensible
REGISTER_ALGORITHM Procedure	Registers a new algorithm
RANK_APPLY Procedure	Ranks the predictions from the `APPLY` results for a classification model
REMOVE_COST_MATRIX Procedure	Removes a cost matrix from a model
RENAME_MODEL Procedure	Renames a model

Deprecated GET_MODEL_DETAILS

Starting from Oracle Database 12*c* Release 2, the following `GET_MODEL_DETAILS` are deprecated:

Table 57-43 Deprecated `GET_MODEL_DETAILS` Functions

Subprogram	Purpose
GET_ASSOCIATION_RULES Function	Returns the rules from an association model
GET_FREQUENT_ITEMSETS Function	Returns the frequent itemsets for an association model
GET_MODEL_DETAILS_AI Function	Returns details about an attribute importance model
GET_MODEL_DETAILS_EM Function	Returns details about an Expectation Maximization model
GET_MODEL_DETAILS_EM_COMP Function	Returns details about the parameters of an Expectation Maximization model
GET_MODEL_DETAILS_EM_PROJ Function	Returns details about the projects of an Expectation Maximization model
GET_MODEL_DETAILS_GLM Function	Returns details about a Generalized Linear Model model

Table 57-43 (Cont.) Deprecated `GET_MODEL_DETAILS` Functions

Subprogram	Purpose
GET_MODEL_DETAILS_GLOBAL Function	Returns high-level statistics about a model
GET_MODEL_DETAILS_KM Function	Returns details about a *k*-Means model
GET_MODEL_DETAILS_NB Function	Returns details about a Naive Bayes model
GET_MODEL_DETAILS_NMF Function	Returns details about a Non-Negative Matrix Factorization model
GET_MODEL_DETAILS_OC Function	Returns details about an O-Cluster model
GET_MODEL_SETTINGS Function	Returns the settings used to build the given model This function is replaced with `USER/ALL/DBA_MINING_MODEL_SETTINGS`
GET_MODEL_SIGNATURE Function	Returns the list of columns from the build input table This function is replaced with `USER/ALL/DBA_MINING_MODEL_ATTRIBUTES`
GET_MODEL_DETAILS_SVD Function	Returns details about a Singular Value Decomposition model
GET_MODEL_DETAILS_SVM Function	Returns details about a Support Vector Machine model with a linear kernel
GET_MODEL_TRANSFORMATIONS Function	Returns the transformations embedded in a model This function is replaced with `USER/ALL/DBA_MINING_MODEL_XFORMS`
GET_MODEL_DETAILS_XML Function	Returns details about a Decision Tree model
GET_TRANSFORM_LIST Procedure	Converts between two different transformation specification formats

Related Topics

- *Oracle Machine Learning for SQL User's Guide*
- *Oracle Database Reference*

57.8.1 ADD_COST_MATRIX Procedure

The `ADD_COST_MATRIX` procedure associates a cost matrix table with a classification model. The cost matrix biases the model by assigning costs or benefits to specific model outcomes.

The cost matrix is stored with the model and taken into account when the model is scored.

You can also specify a cost matrix inline when you invoke an Oracle Machine Learning for SQL function for scoring. To view the scoring matrix for a model, query the `DM$VC` prefixed model view. Refer to Model Detail View for Classification Algorithm.

To obtain the default scoring matrix for a model, query the `DM$VC` prefixed model view. To remove the default scoring matrix from a model, use the `REMOVE_COST_MATRIX` procedure. See REMOVE_COST_MATRIX Procedure.

> **See Also:**
>
> - "Biasing a Classification Model" in *Oracle Machine Learning for SQL Concepts* for more information about costs
> - *Oracle Database SQL Language Reference* for syntax of inline cost matrix
> - Specifying Costs in *Oracle Machine Learning for SQL User's Guide*

Syntax

```
DBMS_DATA_MINING.ADD_COST_MATRIX (
       model_name                IN VARCHAR2,
       cost_matrix_table_name    IN VARCHAR2,
       cost_matrix_schema_name   IN VARCHAR2 DEFAULT NULL);
       partition_name            IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-44 ADD_COST_MATRIX Procedure Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is assumed.
`cost_matrix_table_name`	Name of the cost matrix table (described in Table 57-45).
`cost_matrix_schema_name`	Schema of the cost matrix table. If no schema is specified, then the current schema is used.
`partition_name`	Name of the partition in a partitioned model

Usage Notes

1. If the model is not in your schema, then `ADD_COST_MATRIX` requires the `ALTER ANY MINING MODEL` system privilege or the `ALTER` object privilege for the machine learning model.
2. The cost matrix table must have the columns shown in Table 57-45.

Table 57-45 Required Columns in a Cost Matrix Table

Column Name	Data Type
`ACTUAL_TARGET_VALUE`	Valid target data type
`PREDICTED_TARGET_VALUE`	Valid target data type
`COST`	`NUMBER`, `FLOAT`, `BINARY_DOUBLE`, or `BINARY_FLOAT`

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for valid target data types

3. The types of the actual and predicted target values must be the same as the type of the model target. For example, if the target of the model is `BINARY_DOUBLE`, then the actual and

predicted values must be BINARY_DOUBLE. If the actual and predicted values are CHAR or VARCHAR, then ADD_COST_MATRIX treats them as VARCHAR2 internally.

If the types do not match, or if the actual or predicted value is not a valid target value, then the ADD_COST_MATRIX procedure raises an error.

> **Note:**
>
> If a reverse transformation is associated with the target, then the actual and predicted values must be consistent with the target after the reverse transformation has been applied.
>
> See "Reverse Transformations and Model Transparency" under the "About Transformation Lists" section in DBMS_DATA_MINING_TRANSFORM Operational Notes for more information.

4. Since a benefit can be viewed as a negative cost, you can specify a benefit for a given outcome by providing a negative number in the costs column of the cost matrix table.
5. All classification algorithms can use a cost matrix for scoring. The Decision Tree algorithm can also use a cost matrix at build time. If you want to build a Decision Tree model with a cost matrix, specify the cost matrix table name in the CLAS_COST_TABLE_NAME setting in the settings table for the model. See Table 57-7.

   The cost matrix used to create a Decision Tree model becomes the default scoring matrix for the model. If you want to specify different costs for scoring, use the REMOVE_COST_MATRIX procedure to remove the cost matrix and the ADD_COST_MATRIX procedure to add a new one.
6. Scoring on a partitioned model is partition-specific. Scoring cost matrices can be added to or removed from an individual partition in a partitioned model. If PARTITION_NAME is NOT NULL, then the model must be a partitioned model. The COST_MATRIX is added to that partition of the partitioned model.

   If the PARTITION_NAME is NULL, but the model is a partitioned model, then the COST_MATRIX table is added to every partition in the model.

Example

This example creates a cost matrix table called COSTS_NB and adds it to a Naive Bayes model called NB_SH_CLAS_SAMPLE. The model has a binary target: 1 means that the customer responds to a promotion; 0 means that the customer does not respond. The cost matrix assigns a cost of .25 to misclassifications of customers who do not respond and a cost of .75 to misclassifications of customers who do respond. This means that it is three times more costly to misclassify responders than it is to misclassify non-responders.

```
CREATE TABLE costs_nb (
  actual_target_value           NUMBER,
  predicted_target_value        NUMBER,
  cost                          NUMBER);
INSERT INTO costs_nb values (0, 0, 0);
INSERT INTO costs_nb values (0, 1, .25);
INSERT INTO costs_nb values (1, 0, .75);
INSERT INTO costs_nb values (1, 1, 0);
COMMIT;

EXEC dbms_data_mining.add_cost_matrix('nb_sh_clas_sample', 'costs_nb');

SELECT cust_gender, COUNT(*) AS cnt, ROUND(AVG(age)) AS avg_age
```

```
   FROM mining_data_apply_v
   WHERE PREDICTION(nb_sh_clas_sample COST MODEL
      USING cust_marital_status, education, household_size) = 1
   GROUP BY cust_gender
   ORDER BY cust_gender;

C        CNT    AVG_AGE
- ---------- ----------
F         72         39
M        555         44
```

57.8.2 ADD_PARTITION Procedure

ADD_PARTITION procedure supports a single or multiple partition addition to an existing partitioned model.

The ADD_PARTITION procedure derives build settings and user-defined expressions from the existing model. The target column must exist in the input data query when adding partitions to a supervised model.

Syntax

```
DBMS_DATA_MINING.ADD_PARTITION (
      model_name                IN VARCHAR2,
      data_query                IN CLOB,
      add_options               IN VARCHAR2 DEFAULT ERROR);
```

Parameters

Table 57-46 ADD_PARTITION Procedure Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
data_query	An arbitrary SQL statement that provides data to the model build. The user must have privilege to evaluate this query.
add_options	Allows users to control the conditional behavior of ADD for cases where rows in the input dataset conflict with existing partitions in the model. The following are the possible values: • REPLACE: Replaces the existing partition for which the conflicting keys are found. • ERROR: Terminates the ADD operation without adding any partitions. • IGNORE: Eliminates the rows having the conflicting keys.

> **Note:**
>
> For better performance, Oracle recommends using DROP_PARTITION followed by the ADD_PARTITION instead of using the REPLACE option.

57.8.3 ALTER_REVERSE_EXPRESSION Procedure

This procedure replaces a reverse transformation expression with an expression that you specify. If the attribute does not have a reverse expression, the procedure creates one from the specified expression.

You can also use this procedure to customize the output of clustering, feature extraction, and anomaly detection models.

Syntax

```
DBMS_DATA_MINING.ALTER_REVERSE_EXPRESSION (
        model_name             VARCHAR2,
        expression             CLOB,
        attribute_name         VARCHAR2 DEFAULT NULL,
        attribute_subname      VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-47 ALTER_REVERSE_EXPRESSION Procedure Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, your own schema is used.
`expression`	An expression to replace the reverse transformation associated with the attribute.
`attribute_name`	Name of the attribute. Specify `NULL` if you wish to apply *`expression`* to a cluster, feature, or One-Class SVM prediction.
`attribute_subname`	Name of the nested attribute if *`attribute_name`* is a nested column, otherwise `NULL`.

Usage Notes

1. For purposes of model transparency, Oracle Machine Learning for SQL provides reverse transformations for transformations that are embedded in a model. Reverse transformations are applied to the attributes returned in model detail views and to the scored target of predictive models.

   > **See Also:**
   >
   > - "About Transformation Lists" under DBMS_DATA_MINING_TRANSFORM Operational Notes
   > - Model Detail Views in *Oracle Machine Learning for SQL User's Guide*

2. If you alter the reverse transformation for the target of a model that has a cost matrix, you must specify a transformation expression that has the same type as the actual and predicted values in the cost matrix. Also, the reverse transformation that you specify must result in values that are present in the cost matrix.

> **See Also:**
>
> "ADD_COST_MATRIX Procedure" and *Oracle Machine Learning for SQL Concepts* for information about cost matrixes.

3. To prevent reverse transformation of an attribute, you can specify `NULL` for *`expression`*.
4. The reverse transformation expression can contain a reference to a PL/SQL function that returns a valid Oracle data type. For example, you could define a function like the following for a categorical attribute named `blood_pressure` that has values 'Low', 'Medium' and 'High'.

```
CREATE OR REPLACE FUNCTION numx(c char) RETURN NUMBER IS
  BEGIN
    CASE c WHEN ''Low'' THEN RETURN 1;
           WHEN ''Medium'' THEN RETURN 2;
           WHEN ''High'' THEN RETURN 3;
           ELSE RETURN null;
    END CASE;
  END numx;
```

   Then you could invoke `ALTER_REVERSE_EXPRESION` for `blood_pressure` as follows.

```
EXEC dbms_data_mining.alter_reverse_expression(
             '<model_name>', 'NUMX(blood_pressure)', 'blood_pressure');
```

5. You can use `ALTER_REVERSE_EXPRESSION` to label clusters produced by clustering models and features produced by feature extraction.

   You can use `ALTER_REVERSE_EXPRESSION` to replace the zeros and ones returned by anomaly-detection models. By default, anomaly-detection models label anomalous records with 0 and all other records with 1.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for information about anomaly detection

Examples

1. In this example, the target (`affinity_card`) of the model `CLASS_MODEL` is manipulated internally as `yes` or `no` instead of `1` or `0` but returned as `1`s and `0`s when scored. The `ALTER_REVERSE_EXPRESSION` procedure causes the target values to be returned as `TRUE` or `FALSE`.

```
DECLARE
        v_xlst dbms_data_mining_transform.TRANSFORM_LIST;
  BEGIN
    dbms_data_mining_transform.SET_TRANSFORM(v_xlst,
          'affinity_card', NULL,
          'decode(affinity_card, 1, ''yes'', ''no'')',
          'decode(affinity_card, ''yes'', 1, 0)');
    dbms_data_mining.CREATE_MODEL(
      model_name             => 'CLASS_MODEL',
      mining_function        => dbms_data_mining.classification,
      data_table_name        => 'mining_data_build',
      case_id_column_name    => 'cust_id',
```

```
      target_column_name     => 'affinity_card',
      settings_table_name    => NULL,
      data_schema_name       => 'oml_user',
      settings_schema_name   => NULL,
      xform_list             => v_xlst );
  END;
/
SELECT cust_income_level, occupation,
          PREDICTION(CLASS_MODEL USING *) predict_response
      FROM mining_data_test WHERE age = 60 AND cust_gender IN 'M'
      ORDER BY cust_income_level;

CUST_INCOME_LEVEL              OCCUPATION                PREDICT_RESPONSE
------------------------------ -------------------- --------------------
A: Below 30,000                Transp.                                 1
E: 90,000 - 109,999            Transp.                                 1
E: 90,000 - 109,999            Sales                                   1
G: 130,000 - 149,999           Handler                                 0
G: 130,000 - 149,999           Crafts                                  0
H: 150,000 - 169,999           Prof.                                   1
J: 190,000 - 249,999           Prof.                                   1
J: 190,000 - 249,999           Sales                                   1

BEGIN
  dbms_data_mining.ALTER_REVERSE_EXPRESSION (
     model_name      => 'CLASS_MODEL',
     expression      => 'decode(affinity_card, ''yes'', ''TRUE'', ''FALSE'')',
     attribute_name  => 'affinity_card');
END;
/
column predict_response on
column predict_response format a20
SELECT cust_income_level, occupation,
             PREDICTION(CLASS_MODEL USING *) predict_response
      FROM mining_data_test WHERE age = 60 AND cust_gender IN 'M'
      ORDER BY cust_income_level;

CUST_INCOME_LEVEL              OCCUPATION           PREDICT_RESPONSE
------------------------------ -------------------- --------------------
A: Below 30,000                Transp.              TRUE
E: 90,000 - 109,999            Transp.              TRUE
E: 90,000 - 109,999            Sales                TRUE
G: 130,000 - 149,999           Handler              FALSE
G: 130,000 - 149,999           Crafts               FALSE
H: 150,000 - 169,999           Prof.                TRUE
J: 190,000 - 249,999           Prof.                TRUE
J: 190,000 - 249,999           Sales                TRUE
```

2. This example specifies labels for the clusters that result from the `sh_clus` model. The labels consist of the word "Cluster" and the internal numeric identifier for the cluster.

```
BEGIN
  dbms_data_mining.ALTER_REVERSE_EXPRESSION( 'sh_clus', '''Cluster ''||value');
END;
/

SELECT cust_id, cluster_id(sh_clus using *) cluster_id
   FROM sh_aprep_num
       WHERE cust_id < 100011
       ORDER by cust_id;

CUST_ID CLUSTER_ID
------- --------------------------------------------------
```

```
100001 Cluster 18
100002 Cluster 14
100003 Cluster 14
100004 Cluster 18
100005 Cluster 19
100006 Cluster 7
100007 Cluster 18
100008 Cluster 14
100009 Cluster 8
100010 Cluster 8
```

57.8.4 APPLY Procedure

The APPLY procedure applies a machine learning model to the data of interest, and generates the results in a table. The APPLY procedure is also referred to as **scoring**.

For predictive machine learning functions, the APPLY procedure generates predictions in a target column. For descriptive machine learning functions such as Clustering, the APPLY process assigns each case to a cluster with a probability.

In Oracle Machine Learning for SQL, the APPLY procedure is not applicable to Association models and Attribute Importance models.

> **Note:**
>
> Scoring can also be performed directly in SQL using the OML4SQL functions. See
>
> - Oracle Machine Learning for SQL Functions in *Oracle Database SQL Language Reference*
> - Scoring and Deployment in *Oracle Machine Learning for SQL User's Guide*

Syntax

```
DBMS_DATA_MINING.APPLY (
      model_name           IN VARCHAR2,
      data_table_name      IN VARCHAR2,
      case_id_column_name  IN VARCHAR2,
      result_table_name    IN VARCHAR2,
      data_schema_name     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-48 *APPLY Procedure Parameters*

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
data_table_name	Name of table or view containing the data to be scored
case_id_column_name	Name of the case identifier column
result_table_name	Name of the table in which to store apply results
data_schema_name	Name of the schema containing the data to be scored

Usage Notes

1. The data provided for `APPLY` must undergo the same preprocessing as the data used to create and test the model. When you use Automatic Data Preparation, the preprocessing required by the algorithm is handled for you by the model: both at build time and apply time. (See "Automatic Data Preparation".)
2. `APPLY` creates a table in the user's schema to hold the results. The columns are algorithm-specific.

   The columns in the results table are listed in Table 57-49 through Table 57-53. The case ID column name in the results table will match the case ID column name provided by you. The type of the incoming case ID column is also preserved in `APPLY` output.

   > **Note:**
   >
   > Make sure that the case ID column does not have the same name as one of the columns that will be created by `APPLY`. For example, when applying a Classification model, the case ID in the scoring data must not be `PREDICTION` or `PROBABILITY` (See Table 57-49).

3. The data type for the `PREDICTION`, `CLUSTER_ID`, and `FEATURE_ID` output columns is influenced by any reverse expression that is embedded in the model by the user. If the user does not provide a reverse expression that alters the scored value type, then the types will conform to the descriptions in the following tables. See "ALTER_REVERSE_EXPRESSION Procedure".
4. If the model is partitioned, the `result_table_name` can contain results from different partitions depending on the data from the input data table. An additional column called `PARTITION_NAME` is added to the result table indicating the partition name that is associated with each row.

   For a non-partitioned model, the behavior does not change.

Classification

The results table for Classification has the columns described in Table 57-49. If the target of the model is categorical, the `PREDICTION` column will have a `VARCHAR2` data type. If the target has a binary type, the `PREDICTION` column will have the binary type of the target.

Table 57-49 APPLY Results Table for Classification

Column Name	Data type
Case ID column name	Type of the case ID
PREDICTION	Type of the target
PROBABILITY	BINARY_DOUBLE

Anomaly Detection

The results table for Anomaly Detection has the columns described in Table 57-50.

Table 57-50 APPLY Results Table for Anomaly Detection

Column Name	Data Type
Case ID column name	Type of the case ID
PREDICTION	NUMBER
PROBABILITY	BINARY_DOUBLE

Regression

The results table for Regression has the columns described in APPLY Procedure.

Table 57-51 APPLY Results Table for Regression

Column Name	Data Type
Case ID column name	Type of the case ID
PREDICTION	Type of the target

Clustering

Clustering is an unsupervised machine learning function, and hence there are no targets. The results of an APPLY procedure contain simply the cluster identifier corresponding to a case, and the associated probability. The results table has the columns described in Table 57-52.

Table 57-52 APPLY Results Table for Clustering

Column Name	Data Type
Case ID column name	Type of the case ID
CLUSTER_ID	NUMBER
PROBABILITY	BINARY_DOUBLE

Feature Extraction

Feature Extraction is also an unsupervised machine learning function, hence there are no targets. The results of an APPLY procedure will contain simply the feature identifier corresponding to a case, and the associated match quality. The results table has the columns described in Table 57-53.

Table 57-53 APPLY Results Table for Feature Extraction

Column Name	Data Type
Case ID column name	Type of the case ID
FEATURE_ID	NUMBER
MATCH_QUALITY	BINARY_DOUBLE

Examples

This example applies the GLM Regression model `GLMR_SH_REGR_SAMPLE` to the data in the `MINING_DATA_APPLY_V` view. The `APPLY` results are output of the table `REGRESSION_APPLY_RESULT`.

```
SQL> BEGIN
       DBMS_DATA_MINING.APPLY (
       model_name     => 'glmr_sh_regr_sample',
       data_table_name     => 'mining_data_apply_v',
       case_id_column_name => 'cust_id',
       result_table_name   => 'regression_apply_result');
    END;
    /

SQL> SELECT * FROM regression_apply_result WHERE cust_id >  101485;

   CUST_ID PREDICTION
---------- ----------
    101486 22.8048824
    101487 25.0261101
    101488 48.6146619
    101489   51.82595
    101490 22.6220714
    101491 61.3856816
    101492 24.1400748
    101493  58.034631
    101494 45.7253149
    101495 26.9763318
    101496 48.1433425
    101497 32.0573434
    101498 49.8965531
    101499  56.270656
    101500 21.1153047
```

57.8.5 COMPUTE_CONFUSION_MATRIX Procedure

This procedure computes a confusion matrix, stores it in a table in the user's schema, and returns the model accuracy.

A confusion matrix is a test metric for classification models. It compares the predictions generated by the model with the actual target values in a set of test data. The confusion matrix lists the number of times each class was correctly predicted and the number of times it was predicted to be one of the other classes.

`COMPUTE_CONFUSION_MATRIX` accepts three input streams:

- The predictions generated on the test data. The information is passed in three columns:
  - Case ID column
  - Prediction column
  - Scoring criterion column containing either probabilities or costs
- The known target values in the test data. The information is passed in two columns:
  - Case ID column
  - Target column containing the known target values

- (Optional) A cost matrix table with predefined columns. See the Usage Notes for the column requirements.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more details about confusion matrixes and other test metrics for classification
>
> "COMPUTE_LIFT Procedure"
>
> "COMPUTE_ROC Procedure"

Syntax

```
DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX (
      accuracy                    OUT NUMBER,
      apply_result_table_name      IN  VARCHAR2,
      target_table_name            IN  VARCHAR2,
      case_id_column_name          IN  VARCHAR2,
      target_column_name           IN  VARCHAR2,
      confusion_matrix_table_name  IN  VARCHAR2,
      score_column_name            IN  VARCHAR2 DEFAULT 'PREDICTION',
      score_criterion_column_name  IN  VARCHAR2 DEFAULT 'PROBABILITY',
      cost_matrix_table_name       IN  VARCHAR2 DEFAULT NULL,
      apply_result_schema_name     IN  VARCHAR2 DEFAULT NULL,
      target_schema_name           IN  VARCHAR2 DEFAULT NULL,
      cost_matrix_schema_name      IN  VARCHAR2 DEFAULT NULL,
      score_criterion_type         IN  VARCHAR2 DEFAULT 'PROBABILITY');
```

Parameters

Table 57-54 COMPUTE_CONFUSION_MATRIX Procedure Parameters

Parameter	Description
accuracy	Output parameter containing the overall percentage accuracy of the predictions.
apply_result_table_name	Table containing the predictions.
target_table_name	Table containing the known target values from the test data.
case_id_column_name	Case ID column in the apply results table. Must match the case identifier in the targets table.
target_column_name	Target column in the targets table. Contains the known target values from the test data.
confusion_matrix_table_name	Table containing the confusion matrix. The table will be created by the procedure in the user's schema. The columns in the confusion matrix table are described in the Usage Notes.
score_column_name	Column containing the predictions in the apply results table. The default column name is PREDICTION, which is the default name created by the APPLY procedure (See "APPLY Procedure").

Table 57-54 (Cont.) COMPUTE_CONFUSION_MATRIX Procedure Parameters

Parameter	Description
score_criterion_column_name	Column containing the scoring criterion in the apply results table. Contains either the probabilities or the costs that determine the predictions. By default, scoring is based on probability; the class with the highest probability is predicted for each case. If scoring is based on cost, the class with the lowest cost is predicted. The score_criterion_type parameter indicates whether probabilities or costs will be used for scoring. The default column name is 'PROBABILITY', which is the default name created by the APPLY procedure (See "APPLY Procedure"). See the Usage Notes for additional information.
cost_matrix_table_name	(Optional) Table that defines the costs associated with misclassifications. If a cost matrix table is provided and the score_criterion_type parameter is set to 'COSTS', the costs in this table will be used as the scoring criteria. The columns in a cost matrix table are described in the Usage Notes.
apply_result_schema_name	Schema of the apply results table. If null, the user's schema is assumed.
target_schema_name	Schema of the table containing the known targets. If null, the user's schema is assumed.
cost_matrix_schema_name	Schema of the cost matrix table, if one is provided. If null, the user's schema is assumed.
score_criterion_type	Whether to use probabilities or costs as the scoring criterion. Probabilities or costs are passed in the column identified in the score_criterion_column_name parameter. The default value of score_criterion_type is 'PROBABILITY'. To use costs as the scoring criterion, specify 'COST'. If score_criterion_type is set to 'COST' but no cost matrix is provided and if there is a scoring cost matrix associated with the model, then the associated costs are used for scoring. See the Usage Notes and the Examples.

Usage Notes

- The predictive information you pass to COMPUTE_CONFUSION_MATRIX may be generated using SQL PREDICTION functions, the DBMS_DATA_MINING.APPLY procedure, or some other mechanism. As long as you pass the appropriate data, the procedure can compute the confusion matrix.
- Instead of passing a cost matrix to COMPUTE_CONFUSION_MATRIX, you can use a scoring cost matrix associated with the model. A scoring cost matrix can be embedded in the model or it can be defined dynamically when the model is applied. To use a scoring cost matrix, invoke the SQL PREDICTION_COST function to populate the score criterion column.
- The predictions that you pass to COMPUTE_CONFUSION_MATRIX are in a table or view specified in apply_result_table_name.

```
CREATE TABLE apply_result_table_name AS (
            case_id_column_name           VARCHAR2,
            score_column_name             VARCHAR2,
            score_criterion_column_name    VARCHAR2);
```

- A cost matrix must have the columns described in Table 57-55.

Table 57-55 Columns in a Cost Matrix

Column Name	Data Type
actual_target_value	Type of the target column in the build data
predicted_target_value	Type of the predicted target in the test data. The type of the predicted target must be the same as the type of the actual target unless the predicted target has an associated reverse transformation.
cost	BINARY_DOUBLE

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for valid target data types
>
> *Oracle Machine Learning for SQL Concepts* for more information about cost matrixes

- The confusion matrix created by COMPUTE_CONFUSION_MATRIX has the columns described in Table 57-56.

Table 57-56 Columns in a Confusion Matrix

Column Name	Data Type
actual_target_value	Type of the target column in the build data
predicted_target_value	Type of the predicted target in the test data. The type of the predicted target is the same as the type of the actual target unless the predicted target has an associated reverse transformation.
value	BINARY_DOUBLE

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more information about confusion matrixes

Examples

These examples use the Naive Bayes model nb_sh_clas_sample.

Compute a Confusion Matrix Based on Probabilities

The following statement applies the model to the test data and stores the predictions and probabilities in a table.

```
CREATE TABLE nb_apply_results AS
       SELECT cust_id,
              PREDICTION(nb_sh_clas_sample USING *) prediction,
              PREDICTION_PROBABILITY(nb_sh_clas_sample USING *) probability
       FROM mining_data_test_v;
```

Using probabilities as the scoring criterion, you can compute the confusion matrix as follows.

```
DECLARE
   v_accuracy    NUMBER;
      BEGIN
        DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX (
                  accuracy                     => v_accuracy,
                  apply_result_table_name      => 'nb_apply_results',
                  target_table_name            => 'mining_data_test_v',
                  case_id_column_name          => 'cust_id',
                  target_column_name           => 'affinity_card',
                  confusion_matrix_table_name  => 'nb_confusion_matrix',
                  score_column_name            => 'PREDICTION',
                  score_criterion_column_name  => 'PROBABILITY'
                  cost_matrix_table_name       =>  null,
                  apply_result_schema_name     =>  null,
                  target_schema_name           =>  null,
                  cost_matrix_schema_name      =>  null,
                  score_criterion_type         => 'PROBABILITY');
        DBMS_OUTPUT.PUT_LINE('**** MODEL ACCURACY ****: ' || ROUND(v_accuracy,4));
      END;
      /
```

The confusion matrix and model accuracy are shown as follows.

```
 **** MODEL ACCURACY ****: .7847

SQL>SELECT * from nb_confusion_matrix;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE      VALUE
------------------- ---------------------- ----------
                  1                      0         60
                  0                      0        891
                  1                      1        286
                  0                      1        263
```

Compute a Confusion Matrix Based on a Cost Matrix Table

The confusion matrix in the previous example shows a high rate of false positives. For 263 cases, the model predicted 1 when the actual value was 0. You could use a cost matrix to minimize this type of error.

The cost matrix table `nb_cost_matrix` specifies that a false positive is 3 times more costly than a false negative.

```
SQL> SELECT * from nb_cost_matrix;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE       COST
------------------- ---------------------- ----------
                  0                      0          0
                  0                      1        .75
                  1                      0        .25
                  1                      1          0
```

This statement shows how to generate the predictions using `APPLY`.

```
BEGIN
    DBMS_DATA_MINING.APPLY(
          model_name          => 'nb_sh_clas_sample',
```

```
          data_table_name     => 'mining_data_test_v',
          case_id_column_name => 'cust_id',
          result_table_name   => 'nb_apply_results');
 END;
/
```

This statement computes the confusion matrix using the cost matrix table. The score criterion column is named 'PROBABILITY', which is the name generated by APPLY.

```
DECLARE
  v_accuracy    NUMBER;
     BEGIN
       DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX (
                accuracy                     => v_accuracy,
                apply_result_table_name      => 'nb_apply_results',
                target_table_name            => 'mining_data_test_v',
                case_id_column_name          => 'cust_id',
                target_column_name           => 'affinity_card',
                confusion_matrix_table_name  => 'nb_confusion_matrix',
                score_column_name            => 'PREDICTION',
                score_criterion_column_name  => 'PROBABILITY',
                cost_matrix_table_name       => 'nb_cost_matrix',
                apply_result_schema_name     => null,
                target_schema_name           => null,
                cost_matrix_schema_name      => null,
                score_criterion_type         => 'COST');
       DBMS_OUTPUT.PUT_LINE('**** MODEL ACCURACY ****: ' || ROUND(v_accuracy,4));
    END;
    /
```

The resulting confusion matrix shows a decrease in false positives (212 instead of 263).

```
**** MODEL ACCURACY ****: .798

SQL> SELECT * FROM nb_confusion_matrix;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE      VALUE
------------------- ---------------------- ----------
                  1                      0         91
                  0                      0        942
                  1                      1        255
                  0                      1        212
```

Compute a Confusion Matrix Based on Embedded Costs

You can use the ADD_COST_MATRIX procedure to embed a cost matrix in a model. The embedded costs can be used instead of probabilities for scoring. This statement adds the previously-defined cost matrix to the model.

```
BEGIN    DBMS_DATA_MINING.ADD_COST_MATRIX ('nb_sh_clas_sample', 'nb_cost_matrix');END;/
```

The following statement applies the model to the test data using the embedded costs and stores the results in a table.

```
CREATE TABLE nb_apply_results AS
         SELECT cust_id,
              PREDICTION(nb_sh_clas_sample COST MODEL USING *) prediction,
              PREDICTION_COST(nb_sh_clas_sample COST MODEL USING *) cost
          FROM mining_data_test_v;
```

You can compute the confusion matrix using the embedded costs.

```
DECLARE
   v_accuracy         NUMBER;
```

```
BEGIN
    DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX (
         accuracy                      => v_accuracy,
         apply_result_table_name       => 'nb_apply_results',
         target_table_name             => 'mining_data_test_v',
         case_id_column_name           => 'cust_id',
         target_column_name            => 'affinity_card',
         confusion_matrix_table_name   => 'nb_confusion_matrix',
         score_column_name             => 'PREDICTION',
         score_criterion_column_name   => 'COST',
         cost_matrix_table_name        => null,
         apply_result_schema_name      => null,
         target_schema_name            => null,
         cost_matrix_schema_name       => null,
         score_criterion_type          => 'COST');
END;
/
```

The results are:

```
**** MODEL ACCURACY ****: .798

SQL> SELECT * FROM nb_confusion_matrix;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE      VALUE
------------------- ---------------------- ----------
                  1                      0         91
                  0                      0        942
                  1                      1        255
                  0                      1        212
```

57.8.6 COMPUTE_CONFUSION_MATRIX_PART Procedure

The `COMPUTE_CONFUSION_MATRIX_PART` procedure computes a confusion matrix, stores it in a table in the user's schema, and returns the model accuracy.

`COMPUTE_CONFUSION_MATRIX_PART` provides support to computation of evaluation metrics per-partition for partitioned models. For non-partitioned models, refer to COMPUTE_CONFUSION_MATRIX Procedure.

A confusion matrix is a test metric for classification models. It compares the predictions generated by the model with the actual target values in a set of test data. The confusion matrix lists the number of times each class was correctly predicted and the number of times it was predicted to be one of the other classes.

`COMPUTE_CONFUSION_MATRIX_PART` accepts three input streams:

- The predictions generated on the test data. The information is passed in three columns:
  - Case ID column
  - Prediction column
  - Scoring criterion column containing either probabilities or costs
- The known target values in the test data. The information is passed in two columns:
  - Case ID column
  - Target column containing the known target values
- (Optional) A cost matrix table with predefined columns. See the Usage Notes for the column requirements.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more details about confusion matrixes and other test metrics for classification
>
> "COMPUTE_LIFT_PART Procedure"
>
> "COMPUTE_ROC_PART Procedure"

Syntax

```
DBMS_DATA_MINING.compute_confusion_matrix_part(
      accuracy                    OUT DM_NESTED_NUMERICALS,
      apply_result_table_name     IN  VARCHAR2,
      target_table_name           IN  VARCHAR2,
      case_id_column_name         IN  VARCHAR2,
      target_column_name          IN  VARCHAR2,
      confusion_matrix_table_name IN  VARCHAR2,
      score_column_name           IN  VARCHAR2 DEFAULT 'PREDICTION',
      score_criterion_column_name IN  VARCHAR2 DEFAULT 'PROBABILITY',
      score_partition_column_name IN  VARCHAR2 DEFAULT 'PARTITION_NAME',
      cost_matrix_table_name      IN  VARCHAR2 DEFAULT NULL,
      apply_result_schema_name    IN  VARCHAR2 DEFAULT NULL,
      target_schema_name          IN  VARCHAR2 DEFAULT NULL,
      cost_matrix_schema_name     IN  VARCHAR2 DEFAULT NULL,
      score_criterion_type        IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-57 COMPUTE_CONFUSION_MATRIX_PART Procedure Parameters

Parameter	Description
accuracy	Output parameter containing the overall percentage accuracy of the predictions The output argument is changed from NUMBER to DM_NESTED_NUMERICALS
apply_result_table_name	Table containing the predictions
target_table_name	Table containing the known target values from the test data
case_id_column_name	Case ID column in the apply results table. Must match the case identifier in the targets table.
target_column_name	Target column in the targets table. Contains the known target values from the test data.
confusion_matrix_table_name	Table containing the confusion matrix. The table will be created by the procedure in the user's schema. The columns in the confusion matrix table are described in the Usage Notes.
score_column_name	Column containing the predictions in the apply results table. The default column name is PREDICTION, which is the default name created by the APPLY procedure (See "APPLY Procedure").

Table 57-57 (Cont.) COMPUTE_CONFUSION_MATRIX_PART Procedure Parameters

Parameter	Description
score_criterion_column_name	Column containing the scoring criterion in the apply results table. Contains either the probabilities or the costs that determine the predictions. By default, scoring is based on probability; the class with the highest probability is predicted for each case. If scoring is based on cost, then the class with the lowest cost is predicted. The score_criterion_type parameter indicates whether probabilities or costs will be used for scoring. The default column name is PROBABILITY, which is the default name created by the APPLY procedure (See "APPLY Procedure"). See the Usage Notes for additional information.
score_partition_column_name	(Optional) Parameter indicating the column which contains the name of the partition. This column slices the input test results such that each partition has independent evaluation matrices computed.
cost_matrix_table_name	(Optional) Table that defines the costs associated with misclassifications. If a cost matrix table is provided and the score_criterion_type parameter is set to COSTS, the costs in this table will be used as the scoring criteria. The columns in a cost matrix table are described in the Usage Notes.
apply_result_schema_name	Schema of the apply results table. If null, then the user's schema is assumed.
target_schema_name	Schema of the table containing the known targets. If null, then the user's schema is assumed.
cost_matrix_schema_name	Schema of the cost matrix table, if one is provided. If null, then the user's schema is assumed.
score_criterion_type	Whether to use probabilities or costs as the scoring criterion. Probabilities or costs are passed in the column identified in the score_criterion_column_name parameter. The default value of score_criterion_type is PROBABILITY. To use costs as the scoring criterion, specify COST. If score_criterion_type is set to COST but no cost matrix is provided and if there is a scoring cost matrix associated with the model, then the associated costs are used for scoring. See the Usage Notes and the Examples.

Usage Notes

- The predictive information you pass to COMPUTE_CONFUSION_MATRIX_PART may be generated using SQL PREDICTION functions, the DBMS_DATA_MINING.APPLY procedure, or some other mechanism. As long as you pass the appropriate data, the procedure can compute the confusion matrix.

- Instead of passing a cost matrix to `COMPUTE_CONFUSION_MATRIX_PART`, you can use a scoring cost matrix associated with the model. A scoring cost matrix can be embedded in the model or it can be defined dynamically when the model is applied. To use a scoring cost matrix, invoke the SQL `PREDICTION_COST` function to populate the score criterion column.
- The predictions that you pass to `COMPUTE_CONFUSION_MATRIX_PART` are in a table or view specified in `apply_result_table_name`.

```
CREATE TABLE apply_result_table_name AS (
            case_id_column_name           VARCHAR2,
            score_column_name             VARCHAR2,
            score_criterion_column_name    VARCHAR2);
```

- A cost matrix must have the columns described in Table 57-55.

Table 57-58 Columns in a Cost Matrix

Column Name	Data Type
`actual_target_value`	Type of the target column in the test data
`predicted_target_value`	Type of the predicted target in the test data. The type of the predicted target must be the same as the type of the actual target unless the predicted target has an associated reverse transformation.
`cost`	`BINARY_DOUBLE`

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for valid target data types
>
> *Oracle Machine Learning for SQL Concepts* for more information about cost matrixes

- The confusion matrix created by `COMPUTE_CONFUSION_MATRIX_PART` has the columns described in Table 57-56.

Table 57-59 Columns in a Confusion Matrix Part

Column Name	Data Type
`actual_target_value`	Type of the target column in the test data
`predicted_target_value`	Type of the predicted target in the test data. The type of the predicted target is the same as the type of the actual target unless the predicted target has an associated reverse transformation.
`value`	`BINARY_DOUBLE`

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more information about confusion matrixes

Examples

These examples use the Naive Bayes model `nb_sh_clas_sample`.

Compute a Confusion Matrix Based on Probabilities

The following statement applies the model to the test data and stores the predictions and probabilities in a table.

```
CREATE TABLE nb_apply_results AS
       SELECT cust_id,
              PREDICTION(nb_sh_clas_sample USING *) prediction,
              PREDICTION_PROBABILITY(nb_sh_clas_sample USING *) probability
       FROM mining_data_test_v;
```

Using probabilities as the scoring criterion, you can compute the confusion matrix as follows.

```
DECLARE
   v_accuracy    NUMBER;
      BEGIN
        DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX_PART (
                   accuracy                     => v_accuracy,
                   apply_result_table_name      => 'nb_apply_results',
                   target_table_name            => 'mining_data_test_v',
                   case_id_column_name          => 'cust_id',
                   target_column_name           => 'affinity_card',
                   confusion_matrix_table_name  => 'nb_confusion_matrix',
                   score_column_name            => 'PREDICTION',
                   score_criterion_column_name  => 'PROBABILITY'
                   score_partition_column_name  => 'PARTITION_NAME'
                   cost_matrix_table_name       =>  null,
                   apply_result_schema_name     =>  null,
                   target_schema_name           =>  null,
                   cost_matrix_schema_name      =>  null,
                   score_criterion_type         => 'PROBABILITY');
        DBMS_OUTPUT.PUT_LINE('**** MODEL ACCURACY ****: ' || ROUND(v_accuracy,4));
      END;
      /
```

The confusion matrix and model accuracy are shown as follows.

```
 **** MODEL ACCURACY ****: .7847

SELECT * FROM NB_CONFUSION_MATRIX;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE      VALUE
------------------- ---------------------- ----------
                  1                      0         60
                  0                      0        891
                  1                      1        286
                  0                      1        263
```

Compute a Confusion Matrix Based on a Cost Matrix Table

The confusion matrix in the previous example shows a high rate of false positives. For 263 cases, the model predicted 1 when the actual value was 0. You could use a cost matrix to minimize this type of error.

The cost matrix table `nb_cost_matrix` specifies that a false positive is 3 times more costly than a false negative.

```
 SELECT * from NB_COST_MATRIX;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE       COST
```

```
------------------ --------------------- ----------
                 0                     0          0
                 0                     1        .75
                 1                     0        .25
                 1                     1          0
```

This statement shows how to generate the predictions using `APPLY`.

```
BEGIN
    DBMS_DATA_MINING.APPLY(
          model_name          => 'nb_sh_clas_sample',
          data_table_name     => 'mining_data_test_v',
          case_id_column_name => 'cust_id',
          result_table_name   => 'nb_apply_results');
 END;
/
```

This statement computes the confusion matrix using the cost matrix table. The score criterion column is named '`PROBABILITY`', which is the name generated by `APPLY`.

```
DECLARE
  v_accuracy    NUMBER;
     BEGIN
       DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX_PART (
                accuracy                     => v_accuracy,
                apply_result_table_name      => 'nb_apply_results',
                target_table_name            => 'mining_data_test_v',
                case_id_column_name          => 'cust_id',
                target_column_name           => 'affinity_card',
                confusion_matrix_table_name  => 'nb_confusion_matrix',
                score_column_name            => 'PREDICTION',
                score_criterion_column_name  => 'PROBABILITY',
                score_partition_column_name  => 'PARTITION_NAME'
                cost_matrix_table_name       => 'nb_cost_matrix',
                apply_result_schema_name     => null,
                target_schema_name           => null,
                cost_matrix_schema_name      => null,
                score_criterion_type         => 'COST');
       DBMS_OUTPUT.PUT_LINE('**** MODEL ACCURACY ****: ' || ROUND(v_accuracy,4));
    END;
    /
```

The resulting confusion matrix shows a decrease in false positives (212 instead of 263).

```
**** MODEL ACCURACY ****: .798

 SELECT * FROM NB_CONFUSION_MATRIX;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE      VALUE
------------------- ---------------------- ----------
                  1                      0         91
                  0                      0        942
                  1                      1        255
                  0                      1        212
```

Compute a Confusion Matrix Based on Embedded Costs

You can use the `ADD_COST_MATRIX` procedure to embed a cost matrix in a model. The embedded costs can be used instead of probabilities for scoring. This statement adds the previously-defined cost matrix to the model.

```
BEGIN
DBMS_DATA_MINING.ADD_COST_MATRIX ('nb_sh_clas_sample', 'nb_cost_matrix');
END;/
```

The following statement applies the model to the test data using the embedded costs and stores the results in a table.

```
CREATE TABLE nb_apply_results AS
         SELECT cust_id,
              PREDICTION(nb_sh_clas_sample COST MODEL USING *) prediction,
              PREDICTION_COST(nb_sh_clas_sample COST MODEL USING *) cost
          FROM mining_data_test_v;
```

You can compute the confusion matrix using the embedded costs.

```
DECLARE
   v_accuracy         NUMBER;
   BEGIN
       DBMS_DATA_MINING.COMPUTE_CONFUSION_MATRIX_PART (
            accuracy                      => v_accuracy,
            apply_result_table_name       => 'nb_apply_results',
            target_table_name             => 'mining_data_test_v',
            case_id_column_name           => 'cust_id',
            target_column_name            => 'affinity_card',
            confusion_matrix_table_name   => 'nb_confusion_matrix',
            score_column_name             => 'PREDICTION',
            score_criterion_column_name   => 'COST',
            score_partition_column_name   => 'PARTITION_NAME'
            cost_matrix_table_name        => null,
            apply_result_schema_name      => null,
            target_schema_name            => null,
            cost_matrix_schema_name       => null,
            score_criterion_type          => 'COST');
   END;
   /
```

The results are:

```
**** MODEL ACCURACY ****: .798

 SELECT * FROM NB_CONFUSION_MATRIX;
ACTUAL_TARGET_VALUE PREDICTED_TARGET_VALUE      VALUE
------------------- ---------------------- ----------
                  1                      0         91
                  0                      0        942
                  1                      1        255
                  0                      1        212
```

57.8.7 COMPUTE_LIFT Procedure

This procedure computes lift and stores the results in a table in the user's schema.

Lift is a test metric for binary classification models. To compute lift, one of the target values must be designated as the positive class. `COMPUTE_LIFT` compares the predictions generated by the model with the actual target values in a set of test data. Lift measures the degree to which the model's predictions of the positive class are an improvement over random chance.

Lift is computed on scoring results that have been ranked by probability (or cost) and divided into quantiles. Each quantile includes the scores for the same number of cases.

`COMPUTE_LIFT` calculates quantile-based and cumulative statistics. The number of quantiles and the positive class are user-specified. Additionally, `COMPUTE_LIFT` accepts three input streams:

- The predictions generated on the test data. The information is passed in three columns:

  - Case ID column
  - Prediction column
  - Scoring criterion column containing either probabilities or costs associated with the predictions
- The known target values in the test data. The information is passed in two columns:
  - Case ID column
  - Target column containing the known target values
- (Optional) A cost matrix table with predefined columns. See the Usage Notes for the column requirements.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more details about lift and test metrics for classification
>
> "COMPUTE_CONFUSION_MATRIX Procedure"
>
> "COMPUTE_ROC Procedure"

Syntax

```
DBMS_DATA_MINING.COMPUTE_LIFT (
      apply_result_table_name      IN VARCHAR2,
      target_table_name            IN VARCHAR2,
      case_id_column_name          IN VARCHAR2,
      target_column_name           IN VARCHAR2,
      lift_table_name              IN VARCHAR2,
      positive_target_value        IN VARCHAR2,
      score_column_name            IN VARCHAR2 DEFAULT 'PREDICTION',
      score_criterion_column_name  IN VARCHAR2 DEFAULT 'PROBABILITY',
      num_quantiles                IN NUMBER DEFAULT 10,
      cost_matrix_table_name       IN VARCHAR2 DEFAULT NULL,
      apply_result_schema_name     IN VARCHAR2 DEFAULT NULL,
      target_schema_name           IN VARCHAR2 DEFAULT NULL,
      cost_matrix_schema_name      IN VARCHAR2 DEFAULT NULL
      score_criterion_type         IN VARCHAR2 DEFAULT 'PROBABILITY');
```

Parameters

Table 57-60 COMPUTE_LIFT Procedure Parameters

Parameter	Description
apply_result_table_name	Table containing the predictions.
target_table_name	Table containing the known target values from the test data.
case_id_column_name	Case ID column in the apply results table. Must match the case identifier in the targets table.
target_column_name	Target column in the targets table. Contains the known target values from the test data.

Table 57-60 (Cont.) COMPUTE_LIFT Procedure Parameters

Parameter	Description
lift_table_name	Table containing the lift statistics. The table will be created by the procedure in the user's schema. The columns in the lift table are described in the Usage Notes.
positive_target_value	The positive class. This should be the class of interest, for which you want to calculate lift. If the target column is a NUMBER, you can use the TO_CHAR() operator to provide the value as a string.
score_column_name	Column containing the predictions in the apply results table. The default column name is 'PREDICTION', which is the default name created by the APPLY procedure (See "APPLY Procedure").
score_criterion_column_name	Column containing the scoring criterion in the apply results table. Contains either the probabilities or the costs that determine the predictions. By default, scoring is based on probability; the class with the highest probability is predicted for each case. If scoring is based on cost, the class with the lowest cost is predicted. The score_criterion_type parameter indicates whether probabilities or costs will be used for scoring. The default column name is 'PROBABILITY', which is the default name created by the APPLY procedure (See "APPLY Procedure"). See the Usage Notes for additional information.
num_quantiles	Number of quantiles to be used in calculating lift. The default is 10.
cost_matrix_table_name	(Optional) Table that defines the costs associated with misclassifications. If a cost matrix table is provided and the score_criterion_type parameter is set to 'COST', the costs will be used as the scoring criteria. The columns in a cost matrix table are described in the Usage Notes.
apply_result_schema_name	Schema of the apply results table. If null, the user's schema is assumed.
target_schema_name	Schema of the table containing the known targets. If null, the user's schema is assumed.
cost_matrix_schema_name	Schema of the cost matrix table, if one is provided. If null, the user's schema is assumed.

Table 57-60 (Cont.) COMPUTE_LIFT Procedure Parameters

Parameter	Description
score_criterion_type	Whether to use probabilities or costs as the scoring criterion. Probabilities or costs are passed in the column identified in the score_criterion_column_name parameter. The default value of score_criterion_type is 'PROBABILITY'. To use costs as the scoring criterion, specify 'COST'. If score_criterion_type is set to 'COST' but no cost matrix is provided and if there is a scoring cost matrix associated with the model, then the associated costs are used for scoring. See the Usage Notes and the Examples.

Usage Notes

- The predictive information you pass to COMPUTE_LIFT may be generated using SQL PREDICTION functions, the DBMS_DATA_MINING.APPLY procedure, or some other mechanism. As long as you pass the appropriate data, the procedure can compute the lift.
- Instead of passing a cost matrix to COMPUTE_LIFT, you can use a scoring cost matrix associated with the model. A scoring cost matrix can be embedded in the model or it can be defined dynamically when the model is applied. To use a scoring cost matrix, invoke the SQL PREDICTION_COST function to populate the score criterion column.
- The predictions that you pass to COMPUTE_LIFT are in a table or view specified in apply_results_table_name.

```
CREATE TABLE apply_result_table_name AS (
            case_id_column_name           VARCHAR2,
            score_column_name             VARCHAR2,
            score_criterion_column_name    VARCHAR2);
```

- A cost matrix must have the columns described in Table 57-61.

Table 57-61 Columns in a Cost Matrix

Column Name	Data Type
actual_target_value	Type of the target column in the build data
predicted_target_value	Type of the predicted target in the test data. The type of the predicted target must be the same as the type of the actual target unless the predicted target has an associated reverse transformation.
cost	NUMBER

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more information about cost matrixes

- The table created by COMPUTE_LIFT has the columns described in Table 57-62

Table 57-62 Columns in a Lift Table

Column Name	Data Type
quantile_number	NUMBER
probability_threshold	NUMBER
gain_cumulative	NUMBER
quantile_total_count	NUMBER
quantile_target_count	NUMBER
percent_records_cumulative	NUMBER
lift_cumulative	NUMBER
target_density_cumulative	NUMBER
targets_cumulative	NUMBER
non_targets_cumulative	NUMBER
lift_quantile	NUMBER
target_density	NUMBER

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for details about the information in the lift table

- When a cost matrix is passed to `COMPUTE_LIFT`, the cost threshold is returned in the `probability_threshold` column of the lift table.

Examples

This example uses the Naive Bayes model `nb_sh_clas_sample`.

The example illustrates lift based on probabilities. For examples that show computation based on costs, see "COMPUTE_CONFUSION_MATRIX Procedure".

The following statement applies the model to the test data and stores the predictions and probabilities in a table.

```
CREATE TABLE nb_apply_results AS
    SELECT cust_id, t.prediction, t.probability
    FROM mining_data_test_v, TABLE(PREDICTION_SET(nb_sh_clas_sample USING *)) t;
```

Using probabilities as the scoring criterion, you can compute lift as follows.

```
BEGIN
       DBMS_DATA_MINING.COMPUTE_LIFT (
            apply_result_table_name           => 'nb_apply_results',
            target_table_name                 => 'mining_data_test_v',
            case_id_column_name               => 'cust_id',
            target_column_name                => 'affinity_card',
            lift_table_name                       => 'nb_lift',
            positive_target_value             =>  to_char(1),
            score_column_name                 => 'PREDICTION',
            score_criterion_column_name    => 'PROBABILITY',
            num_quantiles                         =>  10,
            cost_matrix_table_name            =>  null,
```

```
            apply_result_schema_name           =>  null,
            target_schema_name                 =>  null,
            cost_matrix_schema_name            =>  null,
            score_criterion_type               =>  'PROBABILITY');
      END;
      /
```

This query displays some of the statistics from the resulting lift table.

```
SQL>SELECT quantile_number, probability_threshold, gain_cumulative,
           quantile_total_count
           FROM nb_lift;

QUANTILE_NUMBER PROBABILITY_THRESHOLD GAIN_CUMULATIVE QUANTILE_TOTAL_COUNT
--------------- --------------------- --------------- --------------------
              1            .989335775       .15034965                   55
              2            .980534911       .26048951                   55
              3            .968506098      .374125874                   55
              4            .958975196      .493006993                   55
              5            .946705997      .587412587                   55
              6            .927454174       .66958042                   55
              7            .904403627      .748251748                   55
              8            .836482525      .839160839                   55
             10            .500184953               1                   54
```

57.8.8 COMPUTE_LIFT_PART Procedure

The `COMPUTE_LIFT_PART` procedure computes lift and stores the results in a table in the user's schema. This procedure provides support to the computation of evaluation metrics per-partition for partitioned models.

Lift is a test metric for binary classification models. To compute lift, one of the target values must be designated as the positive class. `COMPUTE_LIFT_PART` compares the predictions generated by the model with the actual target values in a set of test data. Lift measures the degree to which the model's predictions of the positive class are an improvement over random chance.

Lift is computed on scoring results that have been ranked by probability (or cost) and divided into quantiles. Each quantile includes the scores for the same number of cases.

`COMPUTE_LIFT_PART` calculates quantile-based and cumulative statistics. The number of quantiles and the positive class are user-specified. Additionally, `COMPUTE_LIFT_PART` accepts three input streams:

- The predictions generated on the test data. The information is passed in three columns:
  - Case ID column
  - Prediction column
  - Scoring criterion column containing either probabilities or costs associated with the predictions
- The known target values in the test data. The information is passed in two columns:
  - Case ID column
  - Target column containing the known target values
- (Optional) A cost matrix table with predefined columns. See the Usage Notes for the column requirements.

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more details about Lift and test metrics for classification
>
> "COMPUTE_LIFT Procedure"
>
> "COMPUTE_CONFUSION_MATRIX Procedure"
>
> "COMPUTE_CONFUSION_MATRIX_PART Procedure"
>
> "COMPUTE_ROC Procedure"
>
> "COMPUTE_ROC_PART Procedure"

Syntax

```
DBMS_DATA_MINING.COMPUTE_LIFT_PART (
      apply_result_table_name     IN VARCHAR2,
      target_table_name           IN VARCHAR2,
      case_id_column_name         IN VARCHAR2,
      target_column_name          IN VARCHAR2,
      lift_table_name             IN VARCHAR2,
      positive_target_value       IN VARCHAR2,
      score_column_name           IN VARCHAR2 DEFAULT 'PREDICTION',
      score_criterion_column_name IN VARCHAR2 DEFAULT 'PROBABILITY',
      score_partition_column_name IN VARCHAR2 DEFAULT 'PARTITION_NAME',
      num_quantiles               IN NUMBER   DEFAULT 10,
      cost_matrix_table_name      IN VARCHAR2 DEFAULT NULL,
      apply_result_schema_name    IN VARCHAR2 DEFAULT NULL,
      target_schema_name          IN VARCHAR2 DEFAULT NULL,
      cost_matrix_schema_name     IN VARCHAR2 DEFAULT NULL,
      score_criterion_type        IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-63 COMPUTE_LIFT_PART Procedure Parameters

Parameter	Description
apply_result_table_name	Table containing the predictions
target_table_name	Table containing the known target values from the test data
case_id_column_name	Case ID column in the apply results table. Must match the case identifier in the targets table.
target_column_name	Target column in the targets table. Contains the known target values from the test data.
lift_table_name	Table containing the Lift statistics. The table will be created by the procedure in the user's schema. The columns in the Lift table are described in the Usage Notes.
positive_target_value	The positive class. This should be the class of interest, for which you want to calculate Lift. If the target column is a NUMBER, then you can use the TO_CHAR() operator to provide the value as a string.

Table 57-63 (Cont.) COMPUTE_LIFT_PART Procedure Parameters

Parameter	Description
score_column_name	Column containing the predictions in the apply results table. The default column name is PREDICTION, which is the default name created by the APPLY procedure (See "APPLY Procedure").
score_criterion_column_name	Column containing the scoring criterion in the apply results table. Contains either the probabilities or the costs that determine the predictions. By default, scoring is based on probability; the class with the highest probability is predicted for each case. If scoring is based on cost, then the class with the lowest cost is predicted. The score_criterion_type parameter indicates whether probabilities or costs will be used for scoring. The default column name is PROBABILITY, which is the default name created by the APPLY procedure (See "APPLY Procedure"). See the Usage Notes for additional information.
score_partition_column_name	Optional parameter indicating the column containing the name of the partition. This column slices the input test results such that each partition has independent evaluation matrices computed.
num_quantiles	Number of quantiles to be used in calculating Lift. The default is 10.
cost_matrix_table_name	(Optional) Table that defines the costs associated with misclassifications. If a cost matrix table is provided and the score_criterion_type parameter is set to COST, then the costs will be used as the scoring criteria. The columns in a cost matrix table are described in the Usage Notes.
apply_result_schema_name	Schema of the apply results table If null, then the user's schema is assumed.
target_schema_name	Schema of the table containing the known targets If null, then the user's schema is assumed.
cost_matrix_schema_name	Schema of the cost matrix table, if one is provided If null, then the user's schema is assumed.
score_criterion_type	Whether to use probabilities or costs as the scoring criterion. Probabilities or costs are passed in the column identified in the score_criterion_column_name parameter. The default value of score_criterion_type is PROBABILITY. To use costs as the scoring criterion, specify COST. If score_criterion_type is set to COST but no cost matrix is provided and if there is a scoring cost matrix associated with the model, then the associated costs are used for scoring. See the Usage Notes and the Examples.

Usage Notes

- The predictive information you pass to `COMPUTE_LIFT_PART` may be generated using SQL `PREDICTION` functions, the `DBMS_DATA_MINING.APPLY` procedure, or some other mechanism. As long as you pass the appropriate data, the procedure can compute the Lift.
- Instead of passing a cost matrix to `COMPUTE_LIFT_PART`, you can use a scoring cost matrix associated with the model. A scoring cost matrix can be embedded in the model or it can be defined dynamically when the model is applied. To use a scoring cost matrix, invoke the SQL `PREDICTION_COST` function to populate the score criterion column.
- The predictions that you pass to `COMPUTE_LIFT_PART` are in a table or view specified in `apply_results_table_name`.

```
CREATE TABLE apply_result_table_name AS (
           case_id_column_name           VARCHAR2,
           score_column_name            VARCHAR2,
           score_criterion_column_name    VARCHAR2);
```

- A cost matrix must have the columns described in Table 57-61.

Table 57-64 Columns in a Cost Matrix

Column Name	Data Type
actual_target_value	Type of the target column in the test data
predicted_target_value	Type of the predicted target in the test data. The type of the predicted target must be the same as the type of the actual target unless the predicted target has an associated reverse transformation.
cost	NUMBER

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more information about cost matrixes

- The table created by `COMPUTE_LIFT_PART` has the columns described in Table 57-62

Table 57-65 Columns in a COMPUTE_LIFT_PART Table

Column Name	Data Type
quantile_number	NUMBER
probability_threshold	NUMBER
gain_cumulative	NUMBER
quantile_total_count	NUMBER
quantile_target_count	NUMBER
percent_records_cumulative	NUMBER
lift_cumulative	NUMBER
target_density_cumulative	NUMBER
targets_cumulative	NUMBER

Table 57-65 (Cont.) Columns in a COMPUTE_LIFT_PART Table

Column Name	Data Type
non_targets_cumulative	NUMBER
lift_quantile	NUMBER
target_density	NUMBER

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for details about the information in the Lift table

- When a cost matrix is passed to `COMPUTE_LIFT_PART`, the cost threshold is returned in the `probability_threshold` column of the Lift table.

Examples

This example uses the Naive Bayes model `nb_sh_clas_sample`.

The example illustrates Lift based on probabilities. For examples that show computation based on costs, see "COMPUTE_CONFUSION_MATRIX Procedure".

For a partitioned model example, see "COMPUTE_CONFUSION_MATRIX_PART Procedure".

The following statement applies the model to the test data and stores the predictions and probabilities in a table.

```
CREATE TABLE nb_apply_results AS
    SELECT cust_id, t.prediction, t.probability
    FROM mining_data_test_v, TABLE(PREDICTION_SET(nb_sh_clas_sample USING *)) t;
```

Using probabilities as the scoring criterion, you can compute Lift as follows.

```
BEGIN
      DBMS_DATA_MINING.COMPUTE_LIFT_PART (
             apply_result_table_name     => 'nb_apply_results',
             target_table_name           => 'mining_data_test_v',
             case_id_column_name         => 'cust_id',
             target_column_name          => 'affinity_card',
             lift_table_name             => 'nb_lift',
             positive_target_value       =>  to_char(1),
             score_column_name           => 'PREDICTION',
             score_criterion_column_name => 'PROBABILITY',
             score_partition_column_name => 'PARTITITON_NAME',
             num_quantiles               =>  10,
             cost_matrix_table_name      =>  null,
             apply_result_schema_name    =>  null,
             target_schema_name          =>  null,
             cost_matrix_schema_name     =>  null,
             score_criterion_type        =>  'PROBABILITY');
END;
/
```

This query displays some of the statistics from the resulting Lift table.

```
SELECT quantile_number, probability_threshold, gain_cumulative,
           quantile_total_count
           FROM nb_lift;

QUANTILE_NUMBER PROBABILITY_THRESHOLD GAIN_CUMULATIVE QUANTILE_TOTAL_COUNT
--------------- -------------------- --------------- --------------------
              1            .989335775       .15034965                   55
              2            .980534911       .26048951                   55
              3            .968506098      .374125874                   55
              4            .958975196      .493006993                   55
              5            .946705997      .587412587                   55
              6            .927454174       .66958042                   55
              7            .904403627      .748251748                   55
              8            .836482525      .839160839                   55
             10            .500184953               1                   54
```

57.8.9 COMPUTE_ROC Procedure

This procedure computes the receiver operating characteristic (ROC), stores the results in a table in the user's schema, and returns a measure of the model accuracy.

ROC is a test metric for binary classification models. To compute ROC, one of the target values must be designated as the positive class. `COMPUTE_ROC` compares the predictions generated by the model with the actual target values in a set of test data.

ROC measures the impact of changes in the probability threshold. The probability threshold is the decision point used by the model for predictions. In binary classification, the default probability threshold is 0.5. The value predicted for each case is the one with a probability greater than 50%.

ROC can be plotted as a curve on an X-Y axis. The false positive rate is placed on the X axis. The true positive rate is placed on the Y axis. A false positive is a positive prediction for a case that is negative in the test data. A true positive is a positive prediction for a case that is positive in the test data.

`COMPUTE_ROC` accepts two input streams:

- The predictions generated on the test data. The information is passed in three columns:
  - Case ID column
  - Prediction column
  - Scoring criterion column containing probabilities
- The known target values in the test data. The information is passed in two columns:
  - Case ID column
  - Target column containing the known target values

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more details about ROC and test metrics for classification
>
> "COMPUTE_CONFUSION_MATRIX Procedure"
>
> "COMPUTE_LIFT Procedure"

Syntax

```
DBMS_DATA_MINING.COMPUTE_ROC (
      roc_area_under_curve         OUT NUMBER,
      apply_result_table_name      IN  VARCHAR2,
      target_table_name            IN  VARCHAR2,
      case_id_column_name          IN  VARCHAR2,
      target_column_name           IN  VARCHAR2,
      roc_table_name               IN  VARCHAR2,
      positive_target_value        IN  VARCHAR2,
      score_column_name            IN  VARCHAR2 DEFAULT 'PREDICTION',
      score_criterion_column_name  IN  VARCHAR2 DEFAULT 'PROBABILITY',
      apply_result_schema_name     IN  VARCHAR2 DEFAULT NULL,
      target_schema_name           IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-66 COMPUTE_ROC Procedure Parameters

Parameter	Description
roc_area_under_the_curve	Output parameter containing the area under the ROC curve (AUC). The AUC measures the likelihood that an actual positive will be predicted as positive. The greater the AUC, the greater the flexibility of the model in accommodating trade-offs between positive and negative class predictions. AUC can be especially important when one target class is rarer or more important to identify than another.
apply_result_table_name	Table containing the predictions.
target_table_name	Table containing the known target values from the test data.
case_id_column_name	Case ID column in the apply results table. Must match the case identifier in the targets table.
target_column_name	Target column in the targets table. Contains the known target values from the test data.
roc_table_name	Table containing the ROC output. The table will be created by the procedure in the user's schema. The columns in the ROC table are described in the Usage Notes.
positive_target_value	The positive class. This should be the class of interest, for which you want to calculate ROC. If the target column is a NUMBER, you can use the TO_CHAR() operator to provide the value as a string.

Table 57-66 (Cont.) COMPUTE_ROC Procedure Parameters

Parameter	Description
score_column_name	Column containing the predictions in the apply results table. The default column name is 'PREDICTION', which is the default name created by the APPLY procedure (See "APPLY Procedure").
score_criterion_column_name	Column containing the scoring criterion in the apply results table. Contains the probabilities that determine the predictions. The default column name is 'PROBABILITY', which is the default name created by the APPLY procedure (See "APPLY Procedure").
apply_result_schema_name	Schema of the apply results table. If null, the user's schema is assumed.
target_schema_name	Schema of the table containing the known targets. If null, the user's schema is assumed.

Usage Notes

- The predictive information you pass to COMPUTE_ROC may be generated using SQL PREDICTION functions, the DBMS_DATA_MINING.APPLY procedure, or some other mechanism. As long as you pass the appropriate data, the procedure can compute the receiver operating characteristic.
- The predictions that you pass to COMPUTE_ROC are in a table or view specified in apply_results_table_name.

```
CREATE TABLE apply_result_table_name AS (
            case_id_column_name           VARCHAR2,
            score_column_name             VARCHAR2,
            score_criterion_column_name    VARCHAR2);
```

- The table created by COMPUTE_ROC has the columns shown in Table 57-67.

Table 57-67 COMPUTE_ROC Output

Column	Datatype
probability	BINARY_DOUBLE
true_positives	NUMBER
false_negatives	NUMBER
false_positives	NUMBER
true_negatives	NUMBER
true_positive_fraction	NUMBER
false_positive_fraction	NUMBER

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for details about the output of `COMPUTE_ROC`

- ROC is typically used to determine the most desirable probability threshold. This can be done by examining the true positive fraction and the false positive fraction. The true positive fraction is the percentage of all positive cases in the test data that were correctly predicted as positive. The false positive fraction is the percentage of all negative cases in the test data that were incorrectly predicted as positive.

  Given a probability threshold, the following statement returns the positive predictions in an apply result table ordered by probability.

```
SELECT case_id_column_name
       FROM apply_result_table_name
       WHERE probability > probability_threshold
       ORDER BY probability DESC;
```

- There are two approaches to identifying the most desirable probability threshold. Which approach you use depends on whether or not you know the relative cost of positive versus negative class prediction errors.

  If the costs are known, you can apply the relative costs to the ROC table to compute the minimum cost probability threshold. Suppose the relative cost ratio is: Positive Class Error Cost / Negative Class Error Cost = 20. Then execute a query like this.

```
WITH cost AS (
  SELECT probability_threshold, 20 * false_negatives + false_positives cost
    FROM ROC_table
  GROUP BY probability_threshold),
    minCost AS (
      SELECT min(cost) minCost
        FROM cost)
      SELECT max(probability_threshold)probability_threshold
        FROM cost, minCost
    WHERE cost = minCost;
```

  If relative costs are not well known, you can simply scan the values in the ROC table (in sorted order) and make a determination about which of the displayed trade-offs (misclassified positives versus misclassified negatives) is most desirable.

```
SELECT * FROM ROC_table
         ORDER BY probability_threshold;
```

Examples

This example uses the Naive Bayes model `nb_sh_clas_sample`.

The following statement applies the model to the test data and stores the predictions and probabilities in a table.

```
CREATE TABLE nb_apply_results AS
    SELECT cust_id, t.prediction, t.probability
    FROM mining_data_test_v, TABLE(PREDICTION_SET(nb_sh_clas_sample USING *)) t;
```

Using the predictions and the target values from the test data, you can compute ROC as follows.

```
DECLARE
     v_area_under_curve NUMBER;
BEGIN
     DBMS_DATA_MINING.COMPUTE_ROC (
         roc_area_under_curve        => v_area_under_curve,
         apply_result_table_name     => 'nb_apply_results',
         target_table_name           => 'mining_data_test_v',
         case_id_column_name         => 'cust_id',
         target_column_name          => 'mining_data_test_v',
         roc_table_name              => 'nb_roc',
         positive_target_value       => '1',
         score_column_name           => 'PREDICTION',
         score_criterion_column_name => 'PROBABILITY');
     DBMS_OUTPUT.PUT_LINE('**** AREA UNDER ROC CURVE ****: ' ||
     ROUND(v_area_under_curve,4));
END;
/
```

The resulting AUC and a selection of columns from the ROC table are shown as follows.

```
**** AREA UNDER ROC CURVE ****: .8212

 SELECT PROBABILITY, TRUE_POSITIVE_FRACTION, FALSE_POSITIVE_FRACTION
            FROM NB_ROC;

PROBABILITY  TRUE_POSITIVE_FRACTION  FALSE_POSITIVE_FRACTION
-----------  ----------------------  -----------------------
     .00000                       1                        1
     .50018             .826589595              .227902946
     .53851             .823699422              .221837088
     .54991             .820809249              .217504333
     .55628             .815028902              .215771231
     .55628             .817919075              .215771231
     .57563             .800578035              .214904679
     .57563             .812138728              .214904679
      .                        .                       .
      .                        .                       .
      .                        .                       .
```

57.8.10 COMPUTE_ROC_PART Procedure

The `COMPUTE_ROC_PART` procedure computes Receiver Operating Characteristic (ROC), stores the results in a table in the user's schema, and returns a measure of the model accuracy. This procedure provides support to computation of evaluation metrics per-partition for partitioned models.

ROC is a test metric for binary classification models. To compute ROC, one of the target values must be designated as the positive class. `COMPUTE_ROC_PART` compares the predictions generated by the model with the actual target values in a set of test data.

ROC measures the impact of changes in the probability threshold. The probability threshold is the decision point used by the model for predictions. In binary classification, the default probability threshold is `0.5`. The value predicted for each case is the one with a probability greater than 50%.

ROC can be plotted as a curve on an x-y axis. The false positive rate is placed on the x-axis. The true positive rate is placed on the y-axis. A false positive is a positive prediction for a case that is negative in the test data. A true positive is a positive prediction for a case that is positive in the test data.

`COMPUTE_ROC_PART` accepts two input streams:

- The predictions generated on the test data. The information is passed in three columns:
  - Case ID column
  - Prediction column
  - Scoring criterion column containing probabilities
- The known target values in the test data. The information is passed in two columns:
  - Case ID column
  - Target column containing the known target values

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for more details about ROC and test metrics for Classification
>
> "COMPUTE_ROC Procedure"
>
> "COMPUTE_CONFUSION_MATRIX Procedure"
>
> "COMPUTE_LIFT_PART Procedure"
>
> "COMPUTE_LIFT Procedure"

Syntax

```
DBMS_DATA_MINING.compute_roc_part(
      roc_area_under_curve        OUT DM_NESTED_NUMERICALS,
      apply_result_table_name     IN  VARCHAR2,
      target_table_name           IN  VARCHAR2,
      case_id_column_name         IN  VARCHAR2,
      target_column_name          IN  VARCHAR2,
      roc_table_name              IN  VARCHAR2,
      positive_target_value       IN  VARCHAR2,
      score_column_name           IN  VARCHAR2 DEFAULT 'PREDICTION',
      score_criterion_column_name IN  VARCHAR2 DEFAULT 'PROBABILITY',
      score_partition_column_name IN  VARCHAR2 DEFAULT 'PARTITION_NAME',
      apply_result_schema_name    IN  VARCHAR2 DEFAULT NULL,
      target_schema_name          IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-68 COMPUTE_ROC_PART Procedure Parameters

Parameter	Description
roc_area_under_the_curve	Output parameter containing the area under the ROC curve (AUC). The AUC measures the likelihood that an actual positive will be predicted as positive. The greater the AUC, the greater the flexibility of the model in accommodating trade-offs between positive and negative class predictions. AUC can be especially important when one target class is rarer or more important to identify than another. The output argument is changed from NUMBER to DM_NESTED_NUMERICALS.
apply_result_table_name	Table containing the predictions.
target_table_name	Table containing the known target values from the test data.
case_id_column_name	Case ID column in the apply results table. Must match the case identifier in the targets table.
target_column_name	Target column in the targets table. Contains the known target values from the test data.
roc_table_name	Table containing the ROC output. The table will be created by the procedure in the user's schema. The columns in the ROC table are described in the Usage Notes.
positive_target_value	The positive class. This should be the class of interest, for which you want to calculate ROC. If the target column is a NUMBER, then you can use the TO_CHAR() operator to provide the value as a string.
score_column_name	Column containing the predictions in the apply results table. The default column name is PREDICTION, which is the default name created by the APPLY procedure (See "APPLY Procedure").
score_criterion_column_name	Column containing the scoring criterion in the apply results table. Contains the probabilities that determine the predictions. The default column name is PROBABILITY, which is the default name created by the APPLY procedure (See "APPLY Procedure").
score_partition_column_name	Optional parameter indicating the column which contains the name of the partition. This column slices the input test results such that each partition has independent evaluation matrices computed.
apply_result_schema_name	Schema of the apply results table. If null, then the user's schema is assumed.
target_schema_name	Schema of the table containing the known targets. If null, then the user's schema is assumed.

Usage Notes

- The predictive information you pass to `COMPUTE_ROC_PART` may be generated using SQL `PREDICTION` functions, the `DBMS_DATA_MINING.APPLY` procedure, or some other mechanism. As long as you pass the appropriate data, the procedure can compute the receiver operating characteristic.
- The predictions that you pass to `COMPUTE_ROC_PART` are in a table or view specified in `apply_results_table_name`.

```
CREATE TABLE apply_result_table_name AS (
            case_id_column_name           VARCHAR2,
            score_column_name            VARCHAR2,
            score_criterion_column_name    VARCHAR2);
```

- The `COMPUTE_ROC_PART` table has the following columns:

Table 57-69 COMPUTE_ROC_PART Output

Column	Data Type
probability	BINARY_DOUBLE
true_positives	NUMBER
false_negatives	NUMBER
false_positives	NUMBER
true_negatives	NUMBER
true_positive_fraction	NUMBER
false_positive_fraction	NUMBER

> **See Also:**
>
> *Oracle Machine Learning for SQL Concepts* for details about the output of `COMPUTE_ROC_PART`

- ROC is typically used to determine the most desirable probability threshold. This can be done by examining the true positive fraction and the false positive fraction. The true positive fraction is the percentage of all positive cases in the test data that were correctly predicted as positive. The false positive fraction is the percentage of all negative cases in the test data that were incorrectly predicted as positive.

  Given a probability threshold, the following statement returns the positive predictions in an apply result table ordered by probability.

```
SELECT case_id_column_name
       FROM apply_result_table_name
       WHERE probability > probability_threshold
       ORDER BY probability DESC;
```

- There are two approaches to identify the most desirable probability threshold. The approach you use depends on whether you know the relative cost of positive versus negative class prediction errors.

If the costs are known, then you can apply the relative costs to the ROC table to compute the minimum cost probability threshold. Suppose the relative cost ratio is: Positive Class Error Cost / Negative Class Error Cost = 20. Then execute a query as follows:

```
WITH cost AS (
  SELECT probability_threshold, 20 * false_negatives + false_positives cost
    FROM ROC_table
  GROUP BY probability_threshold),
    minCost AS (
      SELECT min(cost) minCost
        FROM cost)
      SELECT max(probability_threshold)probability_threshold
        FROM cost, minCost
    WHERE cost = minCost;
```

If relative costs are not well known, then you can simply scan the values in the ROC table (in sorted order) and make a determination about which of the displayed trade-offs (misclassified positives versus misclassified negatives) is most desirable.

```
SELECT * FROM ROC_table
         ORDER BY probability_threshold;
```

Examples

This example uses the Naive Bayes model `nb_sh_clas_sample`.

The following statement applies the model to the test data and stores the predictions and probabilities in a table.

```
CREATE TABLE nb_apply_results AS
    SELECT cust_id, t.prediction, t.probability
    FROM mining_data_test_v, TABLE(PREDICTION_SET(nb_sh_clas_sample USING *)) t;
```

Using the predictions and the target values from the test data, you can compute ROC as follows.

```
DECLARE
     v_area_under_curve NUMBER;
BEGIN
     DBMS_DATA_MINING.COMPUTE_ROC_PART (
         roc_area_under_curve        => v_area_under_curve,
         apply_result_table_name     => 'nb_apply_results',
         target_table_name           => 'mining_data_test_v',
         case_id_column_name         => 'cust_id',
         target_column_name          => 'affinity_card',
         roc_table_name              => 'nb_roc',
         positive_target_value       => '1',
         score_column_name           => 'PREDICTION',
         score_criterion_column_name => 'PROBABILITY');
         score_partition_column_name => 'PARTITION_NAME'
     DBMS_OUTPUT.PUT_LINE('**** AREA UNDER ROC CURVE ****: ' ||
     ROUND(v_area_under_curve,4));
END;
/
```

The resulting AUC and a selection of columns from the ROC table are shown as follows.

```
**** AREA UNDER ROC CURVE ****: .8212

 SELECT PROBABILITY, TRUE_POSITIVE_FRACTION, FALSE_POSITIVE_FRACTION
           FROM NB_ROC;

PROBABILITY  TRUE_POSITIVE_FRACTION  FALSE_POSITIVE_FRACTION
-----------  ----------------------  -----------------------
     .00000                       1                        1
     .50018               .826589595               .227902946
     .53851               .823699422               .221837088
     .54991               .820809249               .217504333
     .55628               .815028902               .215771231
     .55628               .817919075               .215771231
     .57563               .800578035               .214904679
     .57563               .812138728               .214904679
      .                    .                        .
      .                    .                        .
      .                    .                        .
```

57.8.11 CREATE_MODEL Procedure

This procedure creates an Oracle Machine Learning for SQL model with a given machine learning function.

Syntax

```
DBMS_DATA_MINING.CREATE_MODEL (
      model_name            IN VARCHAR2,
      mining_function       IN VARCHAR2,
      data_table_name       IN VARCHAR2,
      case_id_column_name   IN VARCHAR2,
      target_column_name    IN VARCHAR2 DEFAULT NULL,
      settings_table_name   IN VARCHAR2 DEFAULT NULL,
      data_schema_name      IN VARCHAR2 DEFAULT NULL,
      settings_schema_name  IN VARCHAR2 DEFAULT NULL,
      xform_list            IN TRANSFORM_LIST DEFAULT NULL);
```

Parameters

Table 57-70 CREATE_MODEL Procedure Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used. See the Usage Notes for model naming restrictions.
`mining_function`	The machine learning function. Values are listed in Table 57-3.
`data_table_name`	Table or view containing the build data
`case_id_column_name`	Case identifier column in the build data.
`target_column_name`	For supervised models, the target column in the build data. `NULL` for unsupervised models.
`settings_table_name`	Table containing build settings for the model. `NULL` if there is no settings table (only default settings are used).
`data_schema_name`	Schema hosting the build data. If `NULL`, then the user's schema is assumed.

Table 57-70 (Cont.) CREATE_MODEL Procedure Parameters

Parameter	Description
settings_schema_name	Schema hosting the settings table. If NULLthen the user's schema is assumed.

xform_list	A list of transformations to be used in addition to or instead of automatic transformations, depending on the value of the PREP_AUTO setting. (See "Automatic Data Preparation".)
	The datatype of xform_list is TRANSFORM_LIST, which consists of records of type TRANSFORM_REC. Each TRANSFORM_REC specifies the transformation information for a single attribute.
	`TYPE`
	`  TRANFORM_REC     IS RECORD (`
	`    attribute_name       VARCHAR2(4000),`
	`    attribute_subname    VARCHAR2(4000),`
	`    expression           EXPRESSION_REC,`
	`    reverse_expression   EXPRESSION_REC,`
	`    attribute_spec       VARCHAR2(4000));`
	The expression field stores a SQL expression for transforming the attribute. The reverse_expression field stores a SQL expression for reversing the transformation in model details and, if the attribute is a target, in the results of scoring. The SQL expressions are manipulated by routines in the DBMS_DATA_MINING_TRANSFORM package:
	• SET_EXPRESSION Procedure
	• GET_EXPRESSION Function
	• SET_TRANSFORM Procedure
	The attribute_spec field identifies individualized treatment for the attribute. See the Usage Notes for details.
	See Table 58-1for details about the TRANSFORM_REC type.

Usage Notes

1. You can use the attribute_spec field of the xform_list argument to identify an attribute as unstructured text or to disable Automatic Data Preparation for the attribute. The attribute_spec can have the following values:

   - TEXT: Indicates that the attribute contains unstructured text. The TEXT value may optionally be followed by POLICY_NAME, TOKEN_TYPE, MAX_FEATURES, and MIN_DOCUMENTS parameters.

     TOKEN_TYPE has the following possible values: NORMAL, STEM, THEME, SYNONYM, BIGRAM, STEM_BIGRAM. SYNONYM may be optionally followed by a thesaurus name in square brackets.

     MAX_FEATURES specifies the maximum number of tokens extracted from the text.

     MIN_DOCUMENTS specifies the minimal number of documents in which every selected token shall occur. (For information about creating a text policy, see CTX_DDL.CREATE_POLICY in *Oracle Text Reference*).

     Oracle Machine Learning for SQL can process columns of VARCHAR2/CHAR, CLOB, BLOB, and BFILE as text. If the column is VARCHAR2 or CHAR and you do not specify TEXT, then OML4SQL processes the column as categorical data. If the column is CLOB, then OML4SQL processes it as text by default (You do not need to specify it as TEXT. However, you do need to provide an Oracle Text Policy in the settings). If the column is

BLOB or BFILE, then you must specify it as TEXT, otherwise CREATE_MODEL returns an error.

If you specify TEXT for a nested column or for an attribute in a nested column, then CREATE_MODEL returns an error.

- NOPREP: Disables ADP for the attribute. When ADP is OFF, the NOPREP value is ignored.

  You can specify NOPREP for a nested column, but not for an attribute in a nested column. If you specify NOPREP for an attribute in a nested column when ADP is on, then CREATE_MODEL will return an error.

2. You can obtain information about a model by querying the Data Dictionary views.

```
ALL/USER/DBA_MINING_MODELS
ALL/USER/DBA_MINING_MODEL_ATTRIBUTES
ALL/USER/DBA_MINING_MODEL_SETTINGS
ALL/USER/DBA_MINING_MODEL_VIEWS
ALL/USER/DBA_MINING_MODEL_PARTITIONS
ALL/USER/DBA_MINING_MODEL_XFORMS
```

   You can obtain information about model attributes by querying the model details through model views. Refer to *Oracle Machine Learning for SQL User's Guide*.

3. The naming rules for models are more restrictive than the naming rules for most database schema objects. A model name must satisfy the following additional requirements:

   - It must be 123 or fewer characters long.
   - It must be a nonquoted identifier. Oracle requires that nonquoted identifiers contain only alphanumeric characters, the underscore (_), dollar sign ($), and pound sign (#); the initial character must be alphabetic. Oracle strongly discourages the use of the dollar sign and pound sign in nonquoted literals.

   Naming requirements for schema objects are fully documented in *Oracle Database SQL Language Reference*.

4. To build a partitioned model, you must provide additional settings.

   The setting for partitioning columns are as follows:

```
INSERT INTO settings_table VALUES ('ODMS_PARTITION_COLUMNS', 'GENDER,
AGE');
```

   To set user-defined partition number for a model, the setting is as follows:

```
INSERT INTO settings_table VALUES ('ODMS_MAX_PARTITIONS', '10');
```

   The default value for maximum number of partitions is 1000.

5. By passing an xform_list to CREATE_MODEL, you can specify a list of transformations to be performed on the input data. If the PREP_AUTO setting is ON, the transformations are used in addition to the automatic transformations. If the PREP_AUTO setting is OFF, the specified transformations are the only ones implemented by the model. In both cases, transformation definitions are embedded in the model and run automatically whenever the model is applied. See "Automatic Data Preparation". Other transforms that can be specified with xform_list include FORCE_IN. Refer to *Oracle Machine Learning for SQL User's Guide*.

Examples

The first example builds a classification model using the Support Vector Machine algorithm.

```
-- Create the settings table
CREATE TABLE svm_model_settings (
  setting_name  VARCHAR2(30),
  setting_value VARCHAR2(30));

-- Populate the settings table
-- Specify SVM. By default, Naive Bayes is used for classification.
-- Specify ADP. By default, ADP is not used.
BEGIN
  INSERT INTO svm_model_settings (setting_name, setting_value) VALUES
     (dbms_data_mining.algo_name, dbms_data_mining.algo_support_vector_machines);
  INSERT INTO svm_model_settings (setting_name, setting_value) VALUES
     (dbms_data_mining.prep_auto,dbms_data_mining.prep_auto_on);
  COMMIT;
END;
/
-- Create the model using the specified settings
BEGIN
  DBMS_DATA_MINING.CREATE_MODEL(
    model_name          => 'svm_model',
    mining_function     => dbms_data_mining.classification,
    data_table_name     => 'mining_data_build_v',
    case_id_column_name => 'cust_id',
    target_column_name  => 'affinity_card',
    settings_table_name => 'svm_model_settings');
END;
/
```

You can display the model settings with the following query:

```
SELECT * FROM user_mining_model_settings
       WHERE model_name IN 'SVM_MODEL';

MODEL_NAME     SETTING_NAME            SETTING_VALUE                 SETTING
-------------  ----------------------  ----------------------------  -------
SVM_MODEL      ALGO_NAME               ALGO_SUPPORT_VECTOR_MACHINES  INPUT

SVM_MODEL      SVMS_STD_DEV            3.004524                      DEFAULT
SVM_MODEL      PREP_AUTO               ON                            INPUT
SVM_MODEL      SVMS_COMPLEXITY_FACTOR  1.887389                      DEFAULT
SVM_MODEL      SVMS_KERNEL_FUNCTION    SVMS_LINEAR                      DEFAULT
SVM_MODEL      SVMS_CONV_TOLERANCE     .001                          DEFAULT
```

The following is an example of querying a model view instead of the older `GEL_MODEL_DETAILS_SVM` routine.

```
SELECT target_value, attribute_name, attribute_value, coefficient   FROM
DM$VLSVM_MODEL;
```

The second example creates an anomaly detection model. Anomaly detection uses SVM classification without a target. This example uses the same settings table created for the SVM classification model in the first example.

```
BEGIN
  DBMS_DATA_MINING.CREATE_MODEL(
    model_name          => 'anomaly_detect_model',
```

```
    mining_function     => dbms_data_mining.classification,
    data_table_name     => 'mining_data_build_v',
    case_id_column_name => 'cust_id',
    target_column_name  => null,
    settings_table_name => 'svm_model_settings');
END;
/
```

This query shows that the models created in these examples are the only ones in your schema.

```
SELECT model_name, mining_function, algorithm FROM user_mining_models;

MODEL_NAME              MINING_FUNCTION      ALGORITHM
----------------------  -------------------- ------------------------------
SVM_MODEL               CLASSIFICATION       SUPPORT_VECTOR_MACHINES
ANOMALY_DETECT_MODEL    CLASSIFICATION       SUPPORT_VECTOR_MACHINES
```

This query shows that only the SVM classification model has a target.

```
SELECT model_name, attribute_name, attribute_type, target
       FROM user_mining_model_attributes
       WHERE target = 'YES';

MODEL_NAME          ATTRIBUTE_NAME   ATTRIBUTE_TYPE     TARGET
------------------  ---------------  -----------------  ------
SVM_MODEL           AFFINITY_CARD    CATEGORICAL         YES
```

57.8.12 CREATE_MODEL2 Procedure

The `CREATE_MODEL2` procedure is an alternate procedure to the `CREATE_MODEL` procedure, which enables creating a model without extra persistence stages. In the `CREATE_MODEL` procedure, the input is a table or a view and if such an object is not already present, the user must create it. By using the `CREATE_MODEL2` procedure, the user does not need to create such transient database objects.

Syntax

```
DBMS_DATA_MINING.CREATE_MODEL2 (
     model_name            IN VARCHAR2,
     mining_function       IN VARCHAR2,
     data_query            IN CLOB,
     set_list              IN SETTING_LIST,
     case_id_column_name   IN VARCHAR2 DEFAULT NULL,
     target_column_name    IN VARCHAR2 DEFAULT NULL,
     xform_list            IN TRANSFORM_LIST DEFAULT NULL);
```

Parameters

Table 57-71 CREATE_MODEL2 Procedure Parameters

Parameter	Description
`model_name`	Name of the model in the form [`schema_name`.]`model_name`. If you do not specify a schema, then the current schema is used. See the Usage Notes, CREATE_MODEL Procedure for model naming restrictions.
`mining_function`	The machine learning function. Values are listed in DBMS_DATA_MINING — Machine Learning Function Settings.

Table 57-71 (Cont.) CREATE_MODEL2 Procedure Parameters

Parameter	Description
data_query	A query which provides training data for building the model.
set_list	Specifies the SETTING_LIST SETTING_LIST is a table of CLOB index by VARCHAR2(30); Where the index is the setting name and the CLOB is the setting value for that name.
case_id_column_name	Case identifier column in the build data.
target_column_name	For supervised models, the target column in the build data. NULL for unsupervised models.
xform_list	Refer to CREATE_MODEL Procedure.

Usage Notes

Refer to CREATE_MODEL Procedure for Usage Notes.

Examples

The following example uses the Support Vector Machine algorithm.

```
declare
 v_setlst DBMS_DATA_MINING.SETTING_LIST;

BEGIN
  v_setlst(dbms_data_mining.algo_name) :=
dbms_data_mining.algo_support_vector_machines;
  v_setlst(dbms_data_mining.prep_auto) := dbms_data_mining.prep_auto_on;

DBMS_DATA_MINING.CREATE_MODEL2(
    model_name          => 'svm_model',
    mining_function     => dbms_data_mining.classification,
    data_query          => 'select * from mining_data_build_v',
    data_table_name     => 'mining_data_build_v',
    case_id_column_name=> 'cust_id',
    target_column_name => 'affinity_card',
    set_list            => v_setlst,
    case_id_column_name=> 'cust_id',
    target_column_name => 'affinity_card');
END;
/
```

57.8.13 Create Model Using Registration Information

Create model function fetches the setting information from JSON object.

Usage Notes

If an algorithm is registered, user can create model using the registered algorithm name. Since all R scripts and default setting values are already registered, providing the value through the setting table is not necessary. This makes the use of this algorithm easier.

Examples

The first example builds a Classification model using the GLM algorithm.

```
CREATE TABLE GLM_RDEMO_SETTINGS_CL (

   setting_name  VARCHAR2(30),
   setting_value VARCHAR2(4000));
   BEGIN
        INSERT INTO GLM_RDEMO_SETTINGS_CL VALUES
         ('ALGO_EXTENSIBLE_LANG', 'R');
        INSERT INTO GLM_RDEMO_SETTINGS_CL VALUES
         (dbms_data_mining.ralg_registration_algo_name, 't1');
        INSERT INTO GLM_RDEMO_SETTINGS_CL VALUES
        (dbms_data_mining.odms_formula,
        'AGE + EDUCATION + HOUSEHOLD_SIZE + OCCUPATION');
        INSERT INTO GLM_RDEMO_SETTINGS_CL VALUES
         ('RALG_PARAMETER_FAMILY',   'binomial(logit)' );
   END;
   /
     BEGIN
          DBMS_DATA_MINING.CREATE_MODEL(
          model_name                    =>    'GLM_RDEMO_CLASSIFICATION',
          mining_function               =>     dbms_data_mining.classification,
          data_table_name               =>    'mining_data_build_v',
          case_id_column_name           =>    'CUST_ID',
          target_column_name            =>    'AFFINITY_CARD',
          settings_table_name           =>    'GLM_RDEMO_SETTINGS_CL');
      END;
      /
```

57.8.14 DROP_ALGORITHM Procedure

This function is used to drop the registered algorithm information.

Syntax

```
DBMS_DATA_MINING.DROP_ALGORITHM (algorithm_name  IN  VARCHAR2(30),
                                 cascade         IN  BOOLEAN default FALSE)
```

Parameters

Table 57-72 DROP_ALGORITHM Procedure Parameters

Parameter	Description
`algorithm_name`	Name of the algorithm.
`cascade`	If the cascade option is `TRUE`, all the models with this algorithms are forced to drop. There after, the algorithm is dropped. The default value is `FALSE`.

Usage Note

- To drop a machine learning model, you must be the owner or you must have the `RQADMIN` privilege. See *Oracle Machine Learning for SQL User's Guide* for information about privileges for machine learning.
- Make sure a model is not built on the algorithm, then drop the algorithm from the system table.

- If you try to drop an algorithm with a model built on it, then an error is displayed.

57.8.15 DROP_PARTITION Procedure

Syntax

```
DBMS_DATA_MINING.DROP_PARTITION (
      model_name                IN VARCHAR2,
      partition_name            IN VARCHAR2);
```

Parameters

Table 57-73 DROP_PARTITION Procedure Parameters

Parameters	Description
`model_name`	Name of the machine learning model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Name of the partition that must be dropped.

57.8.16 DROP_MODEL Procedure

This procedure deletes the specified machine learning model.

Syntax

```
DBMS_DATA_MINING.DROP_MODEL (model_name IN VARCHAR2,
                             force      IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 57-74 DROP_MODEL Procedure Parameters

Parameter	Description
`model_name`	Name of the machine learning model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`force`	Forces the machine learning model to be dropped even if it is invalid. A machine learning model may be invalid if a serious system error interrupted the model build process.

Usage Note

To drop a machine learning model, you must be the owner or you must have the `DROP ANY MINING MODEL` privilege. See *Oracle Data Mining User's Guide* for information about privileges for Oracle Machine Learning for SQL.

Example

You can use the following command to delete a valid machine learning model named `nb_sh_clas_sample` that exists in your schema.

```
BEGIN
  DBMS_DATA_MINING.DROP_MODEL(model_name => 'nb_sh_clas_sample');
```

```
END;
/
```

57.8.17 EXPORT_MODEL Procedure

This procedure exports the specified machine learning models to a dump file set.

To import the models from the dump file set, use the IMPORT_MODEL Procedure. `EXPORT_MODEL` and `IMPORT_MODEL` use Oracle Data Pump technology.

When Oracle Data Pump is used to export/import an entire schema or database, the machine learning models in the schema or database are included. However, `EXPORT_MODEL` and `IMPORT_MODEL` are the only utilities that support the export/import of individual models.

> **See Also:**
>
> *Oracle Database Utilities* for information about Oracle Data Pump
>
> *Oracle Machine Learning for SQL User's Guide* for more information about exporting and importing machine learning models

Syntax

```
DBMS_DATA_MINING.EXPORT_MODEL (
      filename          IN VARCHAR2,
      directory         IN VARCHAR2,
      model_filter      IN VARCHAR2 DEFAULT NULL,
      filesize          IN VARCHAR2 DEFAULT NULL,
      operation         IN VARCHAR2 DEFAULT NULL,
      remote_link       IN VARCHAR2 DEFAULT NULL,
      jobname           IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-75 EXPORT_MODEL Procedure Parameters

Parameter	Description
`filename`	Name of the dump file set to which the models should be exported. The name must be unique within the schema. The dump file set can contain one or more files. The number of files in a dump file set is determined by the size of the models being exported (both metadata and data) and a specified or estimated maximum file size. You can specify the file size in the `filesize` parameter, or you can use the `operation` parameter to cause Oracle Data Pump to estimate the file size. If the size of the models to export is greater than the maximum file size, one or more additional files are created. When the export operation completes successfully, the name of the dump file set is automatically expanded to *`filename01.dmp`*, even if there is only one file in the dump set. If there are additional files, they are named sequentially as *`filename02.dmp`*, *`filename03.dmp`*, and so forth.

Table 57-75 (Cont.) EXPORT_MODEL Procedure Parameters

Parameter	Description
directory	Name of a pre-defined directory object that specifies where the dump file set should be created. The exporting user must have read/write privileges on the directory object and on the file system directory that it identifies. See *Oracle Database SQL Language Reference* for information about directory objects.
model_filter	Optional parameter that specifies which model or models to export. If you do not specify a value for model_filter, all models in the schema are exported. You can also specify NULL (the default) or 'ALL' to export all models. You can export individual models by name and groups of models based on machine learning function or algorithm. For instance, you could export all regression models or all Naive Bayes models. Examples are provided in Table 57-76.
filesize	Optional parameter that specifies the maximum size of a file in the dump file set. The size may be specified in bytes, kilobytes (K), megabytes (M), or gigabytes (G). The default size is 50 MB. If the size of the models to export is larger than filesize, one or more additional files are created within the dump set. See the description of the filename parameter for more information.
operation	Optional parameter that specifies whether or not to estimate the size of the files in the dump set. By default the size is not estimated and the value of the filesize parameter determines the size of the files. You can specify either of the following values for operation: • 'EXPORT' — Export all or the specified models. (Default) • 'ESTIMATE' — Estimate the size of the exporting models.
remote_link	Optional parameter that specifies the name of a database link to a remote system. The default value is NULL. A database link is a schema object in a local database that enables access to objects in a remote database. When you specify a value for remote_link, you can export the models in the remote database. The EXP_FULL_DATABASE role is required for exporting the remote models. The EXP_FULL_DATABASE privilege, the CREATE DATABASE LINK privilege, and other privileges may also be required.
jobname	Optional parameter that specifies the name of the export job. By default, the name has the form *username*_exp_*nnnn*, where *nnnn* is a number. For example, a job name in the SCOTT schema might be SCOTT_exp_134. If you specify a job name, it must be unique within the schema. The maximum length of the job name is 30 characters. A log file for the export job, named *jobname*.log, is created in the same directory as the dump file set.

Usage Notes

The model_filter parameter specifies which models to export. You can list the models by name, or you can specify all models that have the same machine learning function or algorithm. You can query the USER_MINING_MODELS view to list the models in your schema.

```
SQL> describe user_mining_models
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 MODEL_NAME                                NOT NULL VARCHAR2(30)
 MINING_FUNCTION                                    VARCHAR2(30)
```

```
ALGORITHM                                          VARCHAR2(30)
CREATION_DATE                             NOT NULL DATE
BUILD_DURATION                                     NUMBER
MODEL_SIZE                                         NUMBER
COMMENTS                                           VARCHAR2(4000)
```

Examples of model filters are provided in Table 57-76.

Table 57-76 Sample Values for the Model Filter Parameter

Sample Value	Meaning
`'mymodel'`	Export the model named `mymodel`
`'name= ''mymodel'''`	Export the model named `mymodel`
`'name IN (''mymodel2'',''mymodel3'')'`	Export the models named `mymodel2` and `mymodel3`
`'ALGORITHM_NAME = ''NAIVE_BAYES'''`	Export all Naive Bayes models. See Table 57-5 for a list of algorithm names.
`'FUNCTION_NAME =''CLASSIFICATION'''`	Export all classification models. See Table 57-3 for a list of machine learning functions.

Examples

1. The following statement exports all the models in the `oml_user3` schema to a dump file set called `models_out` in the directory `$ORACLE_HOME/rdbms/log`. This directory is mapped to a directory object called `DATA_PUMP_DIR`. The `oml_user3` user has read/write access to the directory and to the directory object.

   ```
   SQL>execute dbms_data_mining.export_model ('models_out', 'DATA_PUMP_DIR');
   ```

   You can exit SQL*Plus and list the resulting dump file and log file.

   ```
   SQL>EXIT
   >cd $ORACLE_HOME/rdbms/log
   >ls
   >oml_user3_exp_1027.log  models_out01.dmp
   ```

2. The following example uses the same directory object and is run by the same user. This example exports the models called `NMF_SH_SAMPLE` and `SVMR_SH_REGR_SAMPLE` to a different dump file set in the same directory.

   ```
   SQL>EXECUTE DBMS_DATA_MINING.EXPORT_MODEL ( 'models2_out', 'DATA_PUMP_DIR',
              'name in (''NMF_SH_SAMPLE'', ''SVMR_SH_REGR_SAMPLE'')');
   SQL>EXIT
   >cd $ORACLE_HOME/rdbms/log
   >ls
   >oml_user3_exp_1027.log  models_out01.dmp
    oml_user3_exp_924.log  models2_out01.dmp
   ```

3. The following examples show how to export models with specific algorithm and machine learning function names.

   ```
   SQL>EXECUTE DBMS_DATA_MINING.EXPORT_MODEL('algo.dmp','DM_DUMP',
           'ALGORITHM_NAME IN (''O_CLUSTER'',''GENERALIZED_LINEAR_MODEL'',
           ''SUPPORT_VECTOR_MACHINES'',''NAIVE_BAYES'')');

   SQL>EXECUTE DBMS_DATA_MINING.EXPORT_MODEL('func.dmp', 'DM_DUMP',
           'FUNCTION_NAME IN (CLASSIFICATION,CLUSTERING,FEATURE_EXTRACTION)');
   ```

57.8.18 EXPORT_SERMODEL Procedure

This procedure exports the model in a serialized format so that they can be moved to another platform for scoring.

When exporting a model in serialized format, the user must pass in an empty `BLOB` locator and specify the model name to be exported. If the model is partitioned, the user can optionally select an individual partition to export, otherwise all partitions are exported. The returned `BLOB` contains the content that can be deployed.

Syntax

```
DBMS_DATA_MINING.EXPORT_SERMODEL (
      model_data     IN OUT NOCOPY BLOB,
      model_name     IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-77 EXPORT_SERMODEL Procedure Parameters

Parameter	Description
`model_data`	Provides serialized model data.
`model_name`	Name of the machine learning model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Name of the partition that must be exported.

Examples

The following statement exports all of the models in a serialized format.

```
DECLARE
 v_blob blob;
BEGIN
 dbms_lob.createtemporary(v_blob, FALSE);
 dbms_data_mining.export_sermodel(v_blob, 'MY_MODEL');
-- save v_blob somewhere (e.g., bfile, etc.)
 dbms_lob.freetemporary(v_blob);
END;
/
```

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for more information about exporting and importing machine learning models

57.8.19 FETCH_JSON_SCHEMA Procedure

User can fetch and read JSON schema from the `ALL_MINING_ALGORITHMS` view. This function returns the pre-registered JSON schema for R extensible algorithms.

Syntax

```
DBMS_DATA_MINING.FETCH_JSON_SCHEMA RETURN CLOB;
```

Parameters

Table 57-78 FETCH_JSON_SCHEMA Procedure Parameters

Parameter	Description
`RETURN`	This function returns the pre-registered JSON schema for R extensibility. The default value is `CLOB`.

Usage Note

If a user wants to register a new algorithm using the algorithm registration function, they must fetch and follow the pre-registered JSON schema using this function, when they create the required JSON object metadata, and then pass it to the registration function.

57.8.20 GET_ASSOCIATION_RULES Function

The `GET_ASSOCIATION_RULES` function returns the rules produced by an association model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

You can specify filtering criteria to `GET_ASSOCIATION_RULES` to return a subset of the rules. Filtering criteria can improve the performance of the table function. If the number of rules is large, then the greatest performance improvement will result from specifying the `topn` parameter.

Syntax

```
DBMS_DATA_MINING.get_association_rules(
      model_name       IN VARCHAR2,
      topn             IN NUMBER DEFAULT NULL,
      rule_id          IN INTEGER DEFAULT NULL,
      min_confidence   IN NUMBER DEFAULT NULL,
      min_support      IN NUMBER DEFAULT NULL,
      max_rule_length  IN INTEGER DEFAULT NULL,
      min_rule_length  IN INTEGER DEFAULT NULL,
      sort_order       IN ORA_MINING_VARCHAR2_NT DEFAULT NULL,
      antecedent_items IN DM_ITEMS DEFAULT NULL,
      consequent_items IN DM_ITEMS DEFAULT NULL,
      min_lift         IN NUMBER DEFAULT NULL,
      partition_name   IN VARCHAR2 DEFAULT NULL)
  RETURN DM_Rules PIPELINED;
```

Parameters

Table 57-79 GET_ASSOCIATION_RULES Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used. This is the only required parameter of GET_ASSOCIATION_RULES. All other parameters specify optional filters on the rules to return.
topn	Returns the *n* top rules ordered by confidence and then support, both descending. If you specify a sort order, then the top *n* rules are derived after the sort is performed. If topn is specified and no maximum or minimum rule length is specified, then the only columns allowed in the sort order are RULE_CONFIDENCE and RULE_SUPPORT. If topn is specified and a maximum or minimum rule length *is* specified, then RULE_CONFIDENCE, RULE_SUPPORT, and NUMBER_OF_ITEMS are allowed in the sort order.
rule_id	Identifier of the rule to return. If you specify a value for rule_id, do not specify values for the other filtering parameters.
min_confidence	Returns the rules with confidence greater than or equal to this number.
min_support	Returns the rules with support greater than or equal to this number.
max_rule_length	Returns the rules with a length less than or equal to this number. Rule length refers to the number of items in the rule (See NUMBER_OF_ITEMS in Table 57-80). For example, in the rule A=>B (if A, then B), the number of items is 2. If max_rule_length is specified, then the NUMBER_OF_ITEMS column is permitted in the sort order.
min_rule_length	Returns the rules with a length greater than or equal to this number. See max_rule_length for a description of rule length. If min_rule_length is specified, then the NUMBER_OF_ITEMS column is permitted in the sort order.
sort_order	Sorts the rules by the values in one or more of the returned columns. Specify one or more column names, each followed by ASC for ascending order or DESC for descending order. (See Table 57-80 for the column names.) For example, to sort the result set in descending order first by the NUMBER_OF_ITEMS column, then by the RULE_CONFIDENCE column, you must specify: `ORA_MINING_VARCHAR2_NT('NUMBER_OF_ITEMS DESC', 'RULE_CONFIDENCE DESC')` If you specify topn, the results will vary depending on the sort order. By default, the results are sorted by Confidence in descending order, then by Support in descending order.
antecedent_items	Returns the rules with these items in the antecedent.
consequent_items	Returns the rules with this item in the consequent.
min_lift	Returns the rules with lift greater than or equal to this number.
partition_name	Specifies a partition in a partitioned model.

Return Values

The object type returned by `GET_ASSOCIATION_RULES` is described in Table 57-80. For descriptions of each field, see the Usage Notes.

Table 57-80 GET_ASSOCIATION RULES Function Return Values

Return Value	Description

`DM_RULES`	A set of rows of type `DM_RULE`. The rows have the following columns:
	`(rule_id INTEGER,`
	`antecedent DM_PREDICATES,`
	`consequent DM_PREDICATES,`
	`rule_support NUMBER,`
	`rule_confidence NUMBER,`
	`rule_lift NUMBER,`
	`antecedent_support NUMBER,`
	`consequent_support NUMBER,`
	`number_of_items INTEGER )`

`DM_PREDICATES`	The `antecedent` and `consequent` columns each return nested tables of type `DM_PREDICATES`. The rows, of type `DM_PREDICATE`, have the following columns:
	`(attribute_name VARCHAR2(4000),`
	`attribute_subname VARCHAR2(4000),`
	`conditional_operator CHAR(2)/*=,<>,<,>,<=,>=*/,`
	`attribute_num_value NUMBER,`
	`attribute_str_value VARCHAR2(4000),`
	`attribute_support NUMBER,`
	`attribute_confidence NUMBER)`

Usage Notes

1. This table function pipes out rows of type `DM_RULES`. For information on machine learning data types and piped output from table functions, see "Datatypes".
2. The columns returned by `GET_ASSOCIATION_RULES` are described as follows:

Column in DM_RULES	Description
`rule_id`	Unique identifier of the rule

Column in DM_RULES	Description
antecedent	The independent condition in the rule. When this condition exists, the dependent condition in the consequent also exists. The condition is a combination of attribute values called a predicate (DM_PREDICATE). The predicate specifies a condition for each attribute. The condition may specify equality (=), inequality (<>), greater than (>), less than (<), greater than or equal to (>=), or less than or equal to (<=) a given value. Support and Confidence for each attribute condition in the antecedent is returned in the predicate. Support is the number of transactions that satisfy the antecedent. Confidence is the likelihood that a transaction will satisfy the antecedent. **Note:** The occurrence of the attribute as a DM_PREDICATE indicates the presence of the item in the transaction. The actual value for attribute_num_value or attribute_str_value is meaningless. For example, the following predicate indicates that 'Mouse Pad' is present in the transaction *even though* the attribute value is NULL. `DM_PREDICATE('PROD_NAME', 'Mouse Pad', '= ', NULL, NULL, NULL, NULL))`
consequent	The dependent condition in the rule. This condition exists when the antecedent exists. The consequent, like the antecedent, is a predicate (DM_PREDICATE). Support and confidence for each attribute condition in the consequent is returned in the predicate. Support is the number of transactions that satisfy the consequent. Confidence is the likelihood that a transaction will satisfy the consequent.
rule_support	The number of transactions that satisfy the rule.
rule_confidence	The likelihood of a transaction satisfying the rule.
rule_lift	The degree of improvement in the prediction over random chance when the rule is satisfied.
antecedent_support	The ratio of the number of transactions that satisfy the antecedent to the total number of transactions.
consequent_support	The ratio of the number of transactions that satisfy the consequent to the total number of transactions.
number_of_items	The total number of attributes referenced in the antecedent and consequent of the rule.

Examples

The following example demonstrates an association model build followed by several invocations of the GET_ASSOCIATION_RULES table function:

```
-- prepare a settings table to override default settings
CREATE TABLE market_settings AS
SELECT *
  FROM TABLE(DBMS_DATA_MINING.GET_DEFAULT_SETTINGS)
 WHERE setting_name LIKE 'ASSO_%';
BEGIN
-- update the value of the minimum confidence
UPDATE market_settings
   SET setting_value = TO_CHAR(0.081)
 WHERE setting_name = DBMS_DATA_MINING.asso_min_confidence;

-- build an AR model
```

```
DBMS_DATA_MINING.CREATE_MODEL(
  model_name => 'market_model',
  function => DBMS_DATA_MINING.ASSOCIATION,
  data_table_name => 'market_build',
  case_id_column_name => 'item_id',
  target_column_name => NULL,
  settings_table_name => 'market_settings');
END;
/
-- View the (unformatted) rules
SELECT rule_id, antecedent, consequent, rule_support,
       rule_confidence
  FROM TABLE(DBMS_DATA_MINING.GET_ASSOCIATION_RULES('market_model'));
```

In the previous example, you view all rules. To view just the top 20 rules, use the following statement.

```
-- View the top 20 (unformatted) rules
SELECT rule_id, antecedent, consequent, rule_support,
       rule_confidence
  FROM TABLE(DBMS_DATA_MINING.GET_ASSOCIATION_RULES('market_model', 20));
```

The following query uses the association model `AR_SH_SAMPLE`.

```
SELECT * FROM TABLE (
   DBMS_DATA_MINING.GET_ASSOCIATION_RULES (
      'AR_SH_SAMPLE', 10, NULL, 0.5, 0.01, 2, 1,
         ORA_MINING_VARCHAR2_NT (
         'NUMBER_OF_ITEMS DESC', 'RULE_CONFIDENCE DESC', 'RULE_SUPPORT DESC'),
         DM_ITEMS(DM_ITEM('CUSTPRODS', 'Mouse Pad', 1, NULL),
                  DM_ITEM('CUSTPRODS', 'Standard Mouse', 1, NULL)),
         DM_ITEMS(DM_ITEM('CUSTPRODS', 'Extension Cable', 1, NULL))));
```

The query returns three rules, shown as follows:

```
13  DM_PREDICATES(
       DM_PREDICATE('CUSTPRODS', 'Mouse Pad', '= ', 1, NULL, NULL, NULL),
       DM_PREDICATE('CUSTPRODS', 'Standard Mouse', '= ', 1, NULL, NULL, NULL))
    DM_PREDICATES(
       DM_PREDICATE('CUSTPRODS', 'Extension Cable', '= ', 1, NULL, NULL, NULL))
    .15532      .84393   2.7075     .18404     .3117   2

11  DM_PREDICATES(
       DM_PREDICATE('CUSTPRODS', 'Standard Mouse', '= ', 1, NULL, NULL, NULL))
    DM_PREDICATES(
       DM_PREDICATE('CUSTPRODS', 'Extension Cable', '= ', 1, NULL, NULL, NULL))
    .18085      .56291   1.8059     .32128     .3117   1

9   DM_PREDICATES(
       DM_PREDICATE('CUSTPRODS', 'Mouse Pad', '= ', 1, NULL, NULL, NULL))
    DM_PREDICATES(
       DM_PREDICATE('CUSTPRODS', 'Extension Cable', '= ', 1, NULL, NULL, NULL))
      .17766    .55116   1.7682     .32234     .3117   1
```

> **See Also:**
>
> Table 57-80 for the DM_RULE column data types.

57.8.21 GET_FREQUENT_ITEMSETS Function

The GET_FREQUENT_ITEMSETS function returns a set of rows that represent the frequent itemsets from an association model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead..

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

For a detailed description of frequent itemsets, consult *Oracle Machine Learning for SQL Concepts*.

Syntax

```
DBMS_DATA_MINING.get_frequent_itemsets(
      model_name IN VARCHAR2,
      topn IN NUMBER DEFAULT NULL,
      max_itemset_length IN NUMBER DEFAULT NULL,
      partition_name     IN VARCHAR2 DEFAULT NULL)
  RETURN DM_ItemSets PIPELINED;
```

Parameters

Table 57-81 GET_FREQUENT_ITEMSETS Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
topn	When not NULL, return the top *n* rows ordered by support in descending order
max_itemset_length	Maximum length of an item set.
partition_name	Specifies a partition in a partitioned model.

> **Note:**
>
> The partition_name columns applies only when the model is partitioned.

Return Values

Table 57-82 GET_FREQUENT_ITEMSETS Function Return Values

Return Value	Description
DM_ITEMSETS	A set of rows of type DM_ITEMSET. The rows have the following columns: (partition_name VARCHAR2(128) itemsets_id NUMBER, items DM_ITEMS, support NUMBER, number_of_items NUMBER) **Note:** The partition_name columns applies only when the model is partitioned. The items column returns a nested table of type DM_ITEMS. The rows have type DM_ITEM: (attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), attribute_num_value NUMBER, attribute_str_value VARCHAR2(4000))

Usage Notes

This table function pipes out rows of type DM_ITEMSETS. For information on machine learning data types and piped output from table functions, see "Data Types".

Examples

The following example demonstrates an association model build followed by an invocation of GET_FREQUENT_ITEMSETS table function from Oracle SQL.

```
-- prepare a settings table to override default settings
CREATE TABLE market_settings AS

    SELECT *

  FROM TABLE(DBMS_DATA_MINING.GET_DEFAULT_SETTINGS)
 WHERE setting_name LIKE 'ASSO_%';
BEGIN
-- update the value of the minimum confidence
UPDATE market_settings
   SET setting_value = TO_CHAR(0.081)
 WHERE setting_name = DBMS_DATA_MINING.asso_min_confidence;

/* build a AR model */
DBMS_DATA_MINING.CREATE_MODEL(
  model_name           => 'market_model',
  function             => DBMS_DATA_MINING.ASSOCIATION,
  data_table_name      => 'market_build',
  case_id_column_name  => 'item_id',
  target_column_name   => NULL,
```

```
  settings_table_name  => 'market_settings');
END;
/

-- View the (unformatted) Itemsets from SQL*Plus
SELECT itemset_id, items, support, number_of_items
  FROM TABLE(DBMS_DATA_MINING.GET_FREQUENT_ITEMSETS('market_model'));
```

In the example above, you view all itemsets. To view just the top 20 itemsets, use the following statement:

```
-- View the top 20 (unformatted) Itemsets from SQL*Plus
SELECT itemset_id, items, support, number_of_items
  FROM TABLE(DBMS_DATA_MINING.GET_FREQUENT_ITEMSETS('market_model', 20));
```

57.8.22 GET_MODEL_COST_MATRIX Function

The `GET_*` interfaces are replaced by model views, and Oracle recommends that users leverage the views instead.

The GET_MODEL_COST_MATRIX function is replaced by the `DM$VC` prefixed view, Scoring Cost Matrix. The cost matrix used when building a Decision Tree is made available by the `DM$VM` prefixed view, Decision Tree build cost matrix.

Refer to Model Detail View for Classification Algorithm.

The `GET_MODEL_COST_MATRIX` function returns the rows of a cost matrix associated with the specified model.

By default, this function returns the scoring cost matrix that was added to the model with the `ADD_COST_MATRIX` procedure. If you wish to obtain the cost matrix used to create a model, specify `cost_matrix_type_create` as the `matrix_type`. See Table 57-83.

See also ADD_COST_MATRIX Procedure.

Syntax

```
DBMS_DATA_MINING.GET_MODEL_COST_MATRIX (
      model_name                IN VARCHAR2,
      matrix_type               IN VARCHAR2 DEFAULT cost_matrix_type_score)
      partition_name            IN VARCHAR2 DEFAULT NULL);
RETURN DM_COST_MATRIX PIPELINED;
```

Parameters

Table 57-83 GET_MODEL_COST_MATRIX Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`matrix_type`	The type of cost matrix. `COST_MATRIX_TYPE_SCORE` — cost matrix used for scoring. (Default.) `COST_MATRIX_TYPE_CREATE` — cost matrix used to create the model (Decision Tree only).
`partition_name`	Name of the partition in a partitioned model

Return Values

Table 57-84 GET_MODEL_COST_MATRIX Function Return Values

Return Value	Description
DM_COST_MATRIX	A set of rows of type DM_COST_ELEMENT. The rows have the following columns: actual VARCHAR2(4000), predicted VARCHAR2(4000), cost NUMBER)

Usage Notes

Only Decision Tree models can be built with a cost matrix. If you want to build a Decision Tree model with a cost matrix, specify the cost matrix table name in the CLAS_COST_TABLE_NAME setting in the settings table for the model. See Table 57-7.

The cost matrix used to create a Decision Tree model becomes the default scoring matrix for the model. If you want to specify different costs for scoring, you can use the REMOVE_COST_MATRIX procedure to remove the cost matrix and the ADD_COST_MATRIX procedure to add a new one.

The GET_MODEL_COST_MATRIX may return either the build or scoring cost matrix defined for a model or model partition.

If you do not specify a partitioned model name, then an error is displayed.

Example

This example returns the scoring cost matrix associated with the Naive Bayes model NB_SH_CLAS_SAMPLE.

```
column actual format a10
column predicted format a10
SELECT *
    FROM TABLE(dbms_data_mining.get_model_cost_matrix('nb_sh_clas_sample'))
    ORDER BY predicted, actual;

ACTUAL     PREDICTED   COST
---------- ---------- -----
0          0            .00
1          0            .75
0          1            .25
1          1            .00
```

57.8.23 GET_MODEL_DETAILS_AI Function

The GET_MODEL_DETAILS_AI function returns a set of rows that provide the details of an attribute importance model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

Syntax

```
DBMS_DATA_MINING.get_model_details_ai(
      model_name IN VARCHAR2,
```

```
    partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN dm_ranked_attributes pipelined;
```

Parameters

Table 57-85 GET_MODEL_DETAILS_AI Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
partition_name	Specifies a partition in a partitioned model.

Return Values

Table 57-86 GET_MODEL_DETAILS_AI Function Return Values

Return Value	Description
DM_RANKED_ATTRIBUTES	A set of rows of type DM_RANKED_ATTRIBUTE. The rows have the following columns: (attribute_name VARCHAR2(4000, attribute_subname VARCHAR2(4000), importance_value NUMBER, rank NUMBER(38))

Examples

The following example returns model details for the attribute importance model AI_SH_sample, which was created by the sample program dmaidemo.sql.

```
SELECT attribute_name, importance_value, rank
    FROM TABLE(DBMS_DATA_MINING.GET_MODEL_DETAILS_AI('AI_SH_sample'))
    ORDER BY RANK;

ATTRIBUTE_NAME                          IMPORTANCE_VALUE       RANK
--------------------------------------- ---------------- ----------
HOUSEHOLD_SIZE                               .151685183          1
CUST_MARITAL_STATUS                          .145294546          2
YRS_RESIDENCE                                 .07838928          3
AGE                                          .075027496          4
Y_BOX_GAMES                                  .063039952          5
EDUCATION                                    .059605314          6
HOME_THEATER_PACKAGE                         .056458722          7
OCCUPATION                                   .054652937          8
CUST_GENDER                                  .035264741          9
BOOKKEEPING_APPLICATION                      .019204751         10
PRINTER_SUPPLIES                                      0         11
OS_DOC_SET_KANJI                            -.00050013         12
FLAT_PANEL_MONITOR                          -.00509564         13
BULK_PACK_DISKETTES                         -.00540822         14
COUNTRY_NAME                                -.01201116         15
CUST_INCOME_LEVEL                           -.03951311         16
```

57.8.24 GET_MODEL_DETAILS_EM Function

The `GET_MODEL_DETAILS_EM` function returns a set of rows that provide statistics about the clusters produced by an expectation maximization model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

By default, the EM algorithm groups components into high-level clusters, and `GET_MODEL_DETAILS_EM` returns only the high-level clusters with their hierarchies. Alternatively, you can configure EM model to disable the grouping of components into high-level clusters. In this case, `GET_MODEL_DETAILS_EM` returns the components themselves as clusters with their hierarchies. See Table 57-12.

Syntax

```
DBMS_DATA_MINING.get_model_details_em(
      model_name VARCHAR2,
      cluster_id NUMBER   DEFAULT NULL,
      attribute  VARCHAR2 DEFAULT NULL,
      centroid   NUMBER   DEFAULT 1,
      histogram  NUMBER   DEFAULT 1,
      rules      NUMBER   DEFAULT 2,
      attribute_subname  VARCHAR2 DEFAULT NULL,
      topn_attributes NUMBER DEFAULT NULL,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN dm_clusters PIPELINED;
```

Parameters

Table 57-87 GET_MODEL_DETAILS_EM Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`cluster_id`	The ID of a cluster in the model. When a valid cluster ID is specified, only the details of this cluster are returned. Otherwise, the details for all clusters are returned.
`attribute`	The name of an attribute. When a valid attribute name is specified, only the details of this attribute are returned. Otherwise, the details for all attributes are returned
`centroid`	This parameter accepts the following values: • 1: Details about centroids are returned (default) • 0: Details about centroids are not returned
`histogram`	This parameter accepts the following values: • 1: Details about histograms are returned (default) • 0: Details about histograms are not returned
`rules`	This parameter accepts the following values: • 2: Details about rules are returned (default) • 1: Rule summaries are returned • 0: No information about rules is returned

Table 57-87 (Cont.) GET_MODEL_DETAILS_EM Function Parameters

Parameter	Description
attribute_subname	The name of a nested attribute. The full name of a nested attribute has the form: *attribute_name.attribute_subname* where *attribute_name* is the name of the column and *attribute_subname* is the name of the nested attribute in that column. If the attribute is not nested, then attribute_subname is null.
topn_attributes	Restricts the number of attributes returned in the centroid, histogram, and rules objects. Only the *n* attributes with the highest confidence values in the rules are returned. If the number of attributes included in the rules is less than *topn*, then, up to *n* additional attributes in alphabetical order are returned. If both the attribute and topn_attributes parameters are specified, then topn_attributes is ignored.
partition_name	Specifies a partition in a partitioned model.

Usage Notes

1. For information on Oracle Machine Learning for SQL data types and return values for Clustering algorithms piped output from table functions, see "Data Types".
2. GET_MODEL_DETAILS functions preserve model transparency by automatically reversing the transformations applied during the build process. Thus the attributes returned in the model details are the original attributes (or a close approximation of the original attributes) used to build the model.
3. When cluster statistics are disabled (EMCS_CLUSTER_STATISTICS is set to EMCS_CLUS_STATS_DISABLE), GET_MODEL_DETAILS_EM does not return centroids, histograms, or rules. Only taxonomy (hierarchy) and cluster counts are returned.
4. When the partition_name is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

57.8.25 GET_MODEL_DETAILS_EM_COMP Function

he GET_MODEL_DETAILS_EM_COMP table function returns a set of rows that provide details about the parameters of an expectation maximization model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

Syntax

```
DBMS_DATA_MINING.get_model_details_em_comp(
      model_name IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN DM_EM_COMPONENT_SET PIPELINED;
```

Parameters

Table 57-88 GET_MODEL_DETAILS_EM_COMP Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
partition_name	Specifies a partition in a partitioned model to retrieve details for.

Return Values

Table 57-89 GET_MODEL_DETAILS_EM_COMP Function Return Values

Return Value	Description
DM_EM_COMPONENT_SET	A set of rows of type DM_EM_COMPONENT. The rows have the following columns: (info_type VARCHAR2(30), component_id NUMBER, cluster_id NUMBER, attribute_name VARCHAR2(4000), covariate_name VARCHAR2(4000), attribute_value VARCHAR2(4000), value NUMBER)

Usage Notes

1. This table function pipes out rows of type DM_EM_COMPONENT. For information on Oracle Machine Learning for SQL data types and piped output from table functions, see "Data Types".

   The columns in each row returned by GET_MODEL_DETAILS_EM_COMP are described as follows:

Column in DM_EM_COMPONENT	Description
info_type	The type of information in the row. The following information types are supported: • cluster • prior • mean • covariance • frequency
component_id	Unique identifier of a component
cluster_id	Unique identifier of the high-level leaf cluster for each component
attribute_name	Name of an original attribute or a derived feature ID. The derived feature ID is used in models built on data with nested columns. The derived feature definitions can be obtained from the GET_MODEL_DETAILS_EM_PROJ Function.

Column in DM_EM_COMPONENT	Description
covariate_name	Name of an original attribute or a derived feature ID used in variance/covariance definition
attribute_value	Categorical value or bin interval for binned numerical attributes
value	Encodes different information depending on the value of info_type, as follows: • cluster — The value field is NULL • prior — The value field returns the component prior • mean — The value field returns the mean of the attribute specified in attribute_name • covariance — The value field returns the covariance of the attributes specified in attribute_name and covariate_name. Using the same attribute in attribute_name and covariate_name, returns the variance. • frequency— The value field returns the multivalued Bernoulli frequency parameter for the attribute/value combination specified by attribute_name and attribute_value See Usage Note 2 for details.

2. The following table shows which fields are used for each info_type. The blank cells represent NULLs.

info_type	component_id	cluster_id	attribute_name	covariate_name	attribute_value	value
cluster	X	X				
prior	X	X				X
mean	X	X	X			X
covariance	X	X	X	X		X
frequency	X	X	X		X	X

3. GET_MODEL_DETAILS functions preserve model transparency by automatically reversing the transformations applied during the build process. Thus the attributes returned in the model details are the original attributes (or a close approximation of the original attributes) used to build the model.
4. When the value is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

57.8.26 GET_MODEL_DETAILS_EM_PROJ Function

The GET_MODEL_DETAILS_EM_PROJ function returns a set of rows that provide statistics about the projections produced by an expectation maximization model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

Syntax

```
DBMS_DATA_MINING.get_model_details_em_proj(
      model_name IN VARCHAR2,
```

```
    partition_name IN VARCHAR2 DEFAULT NULL)
RETURN DM_EM_PROJECTION_SET PIPELINED;
```

Parameters

Table 57-90 GET_MODEL_DETAILS_EM_PROJ Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
partition_name	Specifies a partition in a partitioned model

Return Values

Table 57-91 GET_MODEL_DETAILS_EM_PROJ Function Return Values

Return Value	Description
DM_EM_PROJECTION_SET	A set of rows of type DM_EM_PROJECTION. The rows have the following columns: (feature_name VARCHAR2(4000), attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), attribute_value VARCHAR2(4000), coefficient NUMBER) See Usage Notes for details.

Usage Notes

1. This table function pipes out rows of type DM_EM_PROJECTION. For information on machine learning data types and piped output from table functions, see "Datatypes".

   The columns in each row returned by GET_MODEL_DETAILS_EM_PROJ are described as follows:

Column in DM_EM_PROJECTION	Description
feature_name	Name of a derived feature. The feature maps to the attribute_name returned by the GET_MODEL_DETAILS_EM Function.
attribute_name	Name of a column in the build data
attribute_subname	Subname in a nested column
attribute_value	Categorical value
coefficient	Projection coefficient. The representation is sparse; only the non-zero coefficients are returned.

2. GET_MODEL_DETAILS functions preserve model transparency by automatically reversing the transformations applied during the build process. Thus the attributes returned in the model details are the original attributes (or a close approximation of the original attributes) used to build the model.

   The coefficients are related to the transformed, not the original, attributes. When returned directly with the model details, the coefficients may not provide meaningful information.

3. When the value is `NULL` for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.27 GET_MODEL_DETAILS_GLM Function

The `GET_MODEL_DETAILS_GLM` function returns the coefficient statistics for a generalized linear model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

The same set of statistics is returned for both linear and logistic regression, but statistics that do not apply to the machine learning function are returned as `NULL`. For more details, see the Usage Notes.

Syntax

```
DBMS_DATA_MINING.get_model_details_glm(
      model_name IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN DM_GLM_Coeff_Set PIPELINED;
```

Parameters

Table 57-92 GET_MODEL_DETAILS_GLM Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Specifies a partition in a partitioned model

Return Values

Table 57-93 GET_MODEL_DETAILS_GLM Return Values

Return Value	Description
DM_GLM_COEFF_SET	A set of rows of type DM_GLM_COEFF. The rows have the following columns: (class VARCHAR2(4000), attribute_name VARCHAR2(4000), attribute_subname VARCHAR2(4000), attribute_value VARCHAR2(4000), feature_expression VARCHAR2(4000), coefficient NUMBER, std_error NUMBER, test_statistic NUMBER, p_value NUMBER, VIF NUMBER, std_coefficient NUMBER, lower_coeff_limit NUMBER, upper_coeff_limit NUMBER, exp_coefficient BINARY_DOUBLE, exp_lower_coeff_limit BINARY_DOUBLE, exp_upper_coeff_limit BINARY_DOUBLE)

GET_MODEL_DETAILS_GLM returns a row of statistics for each attribute and one extra row for the intercept, which is identified by a null value in the attribute name. Each row has the DM_GLM_COEFF data type. The statistics are described in Table 57-94.

Table 57-94 DM_GLM_COEFF Data Type Description

Column	Description
class	The non-reference target class for logistic regression. The model is built to predict the probability of this class. The other class (the reference class) is specified in the model setting GLMS_REFERENCE_CLASS_NAME. See Table 57-19. For Linear Regression, class is null.
attribute_name	The attribute name when there is no subname, or first part of the attribute name when there is a subname. The value of attribute_name is also the name of the column in the case table that is the source for this attribute. For the intercept, attribute_name is null. Intercepts are equivalent to the bias term in SVM models.
attribute_subname	The name of an attribute in a nested table. The full name of a nested attribute has the form: *attribute_name.attribute_subname* where *attribute_name* is the name of the nested column in the case table that is the source for this attribute. If the attribute is not nested, then attribute_subname is null. If the attribute is an intercept, then both the attribute_name and the attribute_subname are null.

Table 57-94 (Cont.) DM_GLM_COEFF Data Type Description

Column	Description
attribute_value	The value of the attribute (categorical attribute only). For numeric attributes, attribute_value is null.
feature_expression	The feature name constructed by the algorithm when feature generation is enabled and higher-order features are found. If feature selection is not enabled, then the feature name is simply the fully-qualified attribute name (*attribute_name.attribute_subname* if the attribute is in a nested column). For categorical attributes, the algorithm constructs a feature name that has the following form: *fully-qualified_attribute_name.attribute_value* For numeric attributes, the algorithm constructs a name for the higher-order feature by taking the product of the resulting values: (*attrib1*)*(*attrib2*))*..... where *attrib1* and *attrib2* are fully-qualified attribute names.
coefficient	The linear coefficient estimate.
std_error	Standard error of the coefficient estimate.
test_statistic	For linear regression, the t-value of the coefficient estimate. For logistic regression, the Wald chi-square value of the coefficient estimate.
p-value	Probability of the test_statistic. Used to analyze the significance of specific attributes in the model.
VIF	Variance Inflation Factor. The value is zero for the intercept. For logistic regression, VIF is null. VIF is not computed if the solver is Cholesky.
std_coefficient	Standardized estimate of the coefficient.
lower_coeff_limit	Lower confidence bound of the coefficient.
upper_coeff_limit	Upper confidence bound of the coefficient.
exp_coefficient	Exponentiated coefficient for logistic regression. For linear regression, exp_coefficient is null.
exp_lower_coeff_limit	Exponentiated coefficient for lower confidence bound of the coefficient for logistic regression. For linear regression, exp_lower_coeff_limit is null.
exp_upper_coeff_limit	Exponentiated coefficient for upper confidence bound of the coefficient for logistic regression. For linear regression, exp_lower_coeff_limit is null.

Usage Notes

Not all statistics are necessarily returned for each coefficient. Statistics will be null if:

- They do not apply to the machine learning function. For example, exp_coefficient does not apply to linear regression.
- They cannot be computed from a theoretical standpoint. For information on ridge regression, see Table 57-19.
- They cannot be computed because of limitations in system resources.
- Their values would be infinity.

- When the value is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Examples

The following example returns some of the model details for the GLM regression model GLMR_SH_Regr_sample.

```
SET line 120
SET pages 99
column attribute_name format a30
column attribute_subname format a20
column attribute_value format a20
col coefficient format 990.9999
col std_error format 990.9999
SQL> SELECT * FROM
(SELECT attribute_name, attribute_value, coefficient, std_error
  FROM DM$VDGLMR_SH_REGR_SAMPLE order by 1,2)
WHERE rownum < 11;

ATTRIBUTE_NAME                 ATTRIBUTE_VALUE     COEFFICIENT     STD_ERROR
------------------------------ -------------------- ----------- ---------
AFFINITY_CARD                                         -0.5797         0.5283
BOOKKEEPING_APPLICATION                              -0.4689         3.8872
BULK_PACK_DISKETTES                                  -0.9819         2.5430
COUNTRY_NAME                    Argentina             -1.2020         1.1876
COUNTRY_NAME                    Australia             -0.0071         5.1146
COUNTRY_NAME                    Brazil                 5.2931         1.9233
COUNTRY_NAME                    Canada                  4.0191         2.4108
COUNTRY_NAME                    China                  0.8706         3.5889
COUNTRY_NAME                    Denmark              -2.9822         3.1803
COUNTRY_NAME                    France                -1.1044         7.1811
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.28 GET_MODEL_DETAILS_GLOBAL Function

The GET_MODEL_DETAILS_GLOBAL function returns statistics about the model as a whole. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

Global details are available for Generalized Linear Models, Association Rules, Singular Value Decomposition, and Expectation Maximization. There are new Global model views which show global information for all algorithms. Oracle recommends that users leverage the views instead. Refer to Model Details View Global.

Syntax

```
DBMS_DATA_MINING.get_model_details_global(
      model_name IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN DM_model_global_details PIPELINED;
```

Parameters

Table 57-95 GET_MODEL_DETAILS_GLOBAL Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
partition_name	Specifies a partition in a partitioned model.

Return Values

Table 57-96 GET_MODEL_DETAILS_GLOBAL Function Return Values

Return Value	Description
DM_MODEL_GLOBAL_DETAILS	A collection of rows of type DM_MODEL_GLOBAL_DETAIL. The rows have the following columns: (global_detail_name VARCHAR2(30), global_detail_value NUMBER)

Examples

The following example returns the global model details for the GLM regression model GLMR_SH_Regr_sample.

```
SELECT *
  FROM TABLE(dbms_data_mining.get_model_details_global(
              'GLMR_SH_Regr_sample'))
ORDER BY global_detail_name;
GLOBAL_DETAIL_NAME             GLOBAL_DETAIL_VALUE
------------------------------ -------------------
ADJUSTED_R_SQUARE                      .731412557
AIC                                      5931.814
COEFF_VAR                              18.1711243
CORRECTED_TOTAL_DF                           1499
CORRECTED_TOT_SS                       278740.504
DEPENDENT_MEAN                             38.892
ERROR_DF                                     1433
ERROR_MEAN_SQUARE                      49.9440956
ERROR_SUM_SQUARES                      71569.8891
F_VALUE                                62.8492452
GMSEP                                   52.280819
HOCKING_SP                             .034877162
J_P                                    52.1749319
MODEL_CONVERGED                                1
MODEL_DF                                      66
MODEL_F_P_VALUE                                0
MODEL_MEAN_SQUARE                      3138.94871
MODEL_SUM_SQUARES                      207170.615
NUM_PARAMS                                    67
NUM_ROWS                                    1500
ROOT_MEAN_SQ                           7.06711367
R_SQ                                   .743238288
SBIC                                   6287.79977
VALID_COVARIANCE_MATRIX                        1
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.29 GET_MODEL_DETAILS_KM Function

The `GET_MODEL_DETAILS_KM` function returns a set of rows that provide the details of a *k*-means clustering model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

You can provide input to `GET_MODEL_DETAILS_KM` to request specific information about the model, thus improving the performance of the query. If you do not specify filtering parameters, then `GET_MODEL_DETAILS_KM` returns all the information about the model.

Syntax

```
DBMS_DATA_MINING.get_model_details_km(
      model_name VARCHAR2,
      cluster_id NUMBER   DEFAULT NULL,
      attribute  VARCHAR2 DEFAULT NULL,
      centroid   NUMBER   DEFAULT 1,
      histogram  NUMBER   DEFAULT 1,
      rules      NUMBER   DEFAULT 2,
      attribute_subname  VARCHAR2 DEFAULT NULL,
      topn_attributes NUMBER DEFAULT NULL,
      partition_name VARCHAR2 DEFAULT NULL)
  RETURN dm_clusters PIPELINED;
```

Parameters

Table 57-97 GET_MODEL_DETAILS_KM Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`cluster_id`	The ID of a cluster in the model. When a valid cluster ID is specified, only the details of this cluster are returned. Otherwise the details for all clusters are returned.
`attribute`	The name of an attribute. When a valid attribute name is specified, only the details of this attribute are returned. Otherwise, the details for all attributes are returned
`centroid`	This parameter accepts the following values: • 1: Details about centroids are returned (default) • 0: Details about centroids are not returned
`histogram`	This parameter accepts the following values: • 1: Details about histograms are returned (default) • 0: Details about histograms are not returned
`rules`	This parameter accepts the following values: • 2: Details about rules are returned (default) • 1: Rule summaries are returned • 0: No information about rules is returned

Table 57-97 (Cont.) GET_MODEL_DETAILS_KM Function Parameters

Parameter	Description
attribute_subname	The name of a nested attribute. The full name of a nested attribute has the form: *attribute_name.attribute_subname* where *attribute_name* is the name of the column and *attribute_subname* is the name of the nested attribute in that column. If the attribute is not nested, attribute_subname is null.
topn_attributes	Restricts the number of attributes returned in the centroid, histogram, and rules objects. Only the *n* attributes with the highest confidence values in the rules are returned. If the number of attributes included in the rules is less than *topn*, then up to *n* additional attributes in alphabetical order are returned. If both the attribute and topn_attributes parameters are specified, then topn_attributes is ignored.
partition_name	Specifies a partition in a partitioned model.

Usage Notes

1. The table function pipes out rows of type DM_CLUSTERS. For information on machine learning data types and Return Value for Clustering Algorithms piped output from table functions, see "Data Types".
2. When the value is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Examples

The following example returns model details for the *k*-means clustering model KM_SH_Clus_sample.

```
SELECT T.id           clu_id,
       T.record_count rec_cnt,
       T.parent       parent,
       T.tree_level   tree_level,
       T.dispersion   dispersion
  FROM (SELECT *
          FROM TABLE(DBMS_DATA_MINING.GET_MODEL_DETAILS_KM(
                     'KM_SH_Clus_sample'))
        ORDER BY id) T
 WHERE ROWNUM < 6;

    CLU_ID    REC_CNT     PARENT TREE_LEVEL DISPERSION
---------- ---------- ---------- ---------- ----------
         1       1500                     1  5.9152211
         2        638          1          2 3.98458982
         3        862          1          2 5.83732097
         4        376          3          3 5.05192137
         5        486          3          3 5.42901522
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.30 GET_MODEL_DETAILS_NB Function

The `GET_MODEL_DETAILS_NB` function returns a set of rows that provide the details of a naive Bayes model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

Syntax

```
DBMS_DATA_MINING.get_model_details_nb(
      model_name IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN DM_NB_Details PIPELINED;
```

Parameters

Table 57-98 GET_MODEL_DETAILS_NB Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Specifies a partition in a partitioned model

Return Values

Table 57-99 GET_MODEL_DETAILS_NB Function Return Values

Return Value	Description

`DM_NB_DETAILS`	A set of rows of type `DM_NB_DETAIL`. The rows have the following columns:
	`(target_attribute_name VARCHAR2(30),`
	`target_attribute_str_value VARCHAR2(4000),`
	`target_attribute_num_value NUMBER,`
	`prior_probability NUMBER,`
	`conditionals DM_CONDITIONALS)`
	The `conditionals` column of `DM_NB_DETAIL` returns a nested table of type `DM_CONDITIONALS`. The rows, of type `DM_CONDITIONAL`, have the following columns:
	`(attribute_name VARCHAR2(4000),`
	`attribute_subname VARCHAR2(4000),`
	`attribute_str_value VARCHAR2(4000),`
	`attribute_num_value NUMBER,`
	`conditional_probability NUMBER)`

Usage Notes

- The table function pipes out rows of type `DM_NB_DETAILS`. For information on machine learning data types and piped output from table functions, see "Data Types".
- When the value is `NULL` for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Examples

The following query is from the sample program `dmnbdemo.sql`. It returns model details about the model `NB_SH_Clas_sample`. For information about the sample programs, see *Oracle Machine Learning for SQL User's Guide*.

The query creates labels from the bin boundary tables that were used to bin the training data. It replaces the attribute values with the labels. For numeric bins, the labels are `(lower_boundary,upper_boundary]`; for categorical bins, the label matches the value it represents. (This method of categorical label representation will only work for cases where one value corresponds to one bin.) The target was not binned.

```
WITH
    bin_label_view AS (
    SELECT col, bin, (DECODE(bin,'1','[','(') || lv || ',' || val || ']') label
      FROM (SELECT col,
                   bin,
                   LAST_VALUE(val) OVER (
                   PARTITION BY col ORDER BY val
                   ROWS BETWEEN UNBOUNDED PRECEDING AND 1 PRECEDING) lv,
                   val
              FROM nb_sh_sample_num)
   UNION ALL
   SELECT col, bin, val label
     FROM nb_sh_sample_cat
   ),
   model_details AS (
   SELECT T.target_attribute_name                                          tname,
          NVL(TO_CHAR(T.target_attribute_num_value,T.target_attribute_str_value)) tval,
          C.attribute_name                                                 pname,
          NVL(L.label, NVL(C.attribute_str_value, C.attribute_num_value)) pval,
          T.prior_probability                                             priorp,
          C.conditional_probability                                       condp
     FROM TABLE(DBMS_DATA_MINING.GET_MODEL_DETAILS_NB('NB_SH_Clas_sample')) T,
          TABLE(T.conditionals) C,
          bin_label_view L
    WHERE C.attribute_name = L.col (+) AND
          (NVL(C.attribute_str_value,C.attribute_num_value) = L.bin(+))
   ORDER BY 1,2,3,4,5,6
   )
   SELECT tname, tval, pname, pval, priorp, condp
     FROM model_details
    WHERE ROWNUM < 11;

TNAME          TVAL PNAME                     PVAL           PRIORP   CONDP
-------------- ---- ------------------------- ------------- ------- -------
AFFINITY_CARD  0    AGE                       (24,30]         .6500   .1714
AFFINITY_CARD  0    AGE                       (30,35]         .6500   .1509
AFFINITY_CARD  0    AGE                       (35,40]         .6500   .1125
AFFINITY_CARD  0    AGE                       (40,46]         .6500   .1134
AFFINITY_CARD  0    AGE                       (46,53]         .6500   .1071
AFFINITY_CARD  0    AGE                       (53,90]         .6500   .1312
AFFINITY_CARD  0    AGE                       [17,24]         .6500   .2134
AFFINITY_CARD  0    BOOKKEEPING_APPLICATION   0               .6500   .1500
AFFINITY_CARD  0    BOOKKEEPING_APPLICATION   1               .6500   .8500
AFFINITY_CARD  0    BULK_PACK_DISKETTES       0               .6500   .3670
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.31 GET_MODEL_DETAILS_NMF Function

The `GET_MODEL_DETAILS_NMF` function returns a set of rows that provide the details of a non-negative matrix factorization model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

Syntax

```
DBMS_DATA_MINING.get_model_details_nmf(
      model_name IN VARCHAR2,
      partition_name VARCHAR2 DEFAULT NULL)
   RETURN DM_NMF_Feature_Set PIPELINED;
```

Parameters

Table 57-100 GET_MODEL_DETAILS_NMF Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Specifies a partition in a partitioned model

Return Values

Table 57-101 GET_MODEL_DETAILS_NMF Function Return Values

Return Value	Description

`DM_NMF_FEATURE_SET`	A set of rows of `DM_NMF_FEATURE`. The rows have the following columns:
	`(feature_id NUMBER,`
	`mapped_feature_id VARCHAR2(4000),`
	`attribute_set DM_NMF_ATTRIBUTE_SET)`
	The `attribute_set` column of `DM_NMF_FEATURE` returns a nested table of type `DM_NMF_ATTRIBUTE_SET`. The rows, of type `DM_NMF_ATTRIBUTE`, have the following columns:
	`(attribute_name VARCHAR2(4000),`
	`attribute_subname VARCHAR2(4000),`
	`attribute_value VARCHAR2(4000),`
	`coefficient NUMBER)`

Usage Notes

- The table function pipes out rows of type `DM_NMF_FEATURE_SET`. For information on machine learning data types and piped output from table functions, see "Data Types".
- When the value is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Examples

The following example returns model details for the feature extraction model `NMF_SH_Sample`.

```
SELECT * FROM (
SELECT F.feature_id,
       A.attribute_name,
       A.attribute_value,
       A.coefficient
  FROM TABLE(DBMS_DATA_MINING.GET_MODEL_DETAILS_NMF('NMF_SH_Sample')) F,
       TABLE(F.attribute_set) A
ORDER BY feature_id,attribute_name,attribute_value
) WHERE ROWNUM < 11;

FEATURE_ID ATTRIBUTE_NAME          ATTRIBUTE_VALUE          COEFFICIENT
---------- ----------------------  ---------------- -------------------
         1 AFFINITY_CARD                              .051208078859308
         1 AGE                                       .0390513260041573
         1 BOOKKEEPING_APPLICATION                   .0512734004239326
         1 BULK_PACK_DISKETTES                        .232471260895683
         1 COUNTRY_NAME            Argentina         .00766817464479959
         1 COUNTRY_NAME            Australia        .000157637881096675
         1 COUNTRY_NAME            Brazil             .0031409632415604
         1 COUNTRY_NAME            Canada            .00144213099311427
         1 COUNTRY_NAME            China             .000102279310968754
         1 COUNTRY_NAME            Denmark           .000242424084307513
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.32 GET_MODEL_DETAILS_OC Function

The `GET_MODEL_DETAILS_OC` function returns a set of rows that provide the details of an O-cluster clustering model. The rows are an enumeration of the clustering patterns generated during the creation of the model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

You can provide input to `GET_MODEL_DETAILS_OC` to request specific information about the model, thus improving the performance of the query. If you do not specify filtering parameters, then `GET_MODEL_DETAILS_OC` returns all the information about the model.

Syntax

```
DBMS_DATA_MINING.get_model_details_oc(
      model_name VARCHAR2,
      cluster_id NUMBER   DEFAULT NULL,
      attribute  VARCHAR2 DEFAULT NULL,
      centroid   NUMBER   DEFAULT 1,
      histogram  NUMBER   DEFAULT 1,
      rules      NUMBER   DEFAULT 2,
      topn_attributes NUMBER DEFAULT NULL,
      partition_name VARCHAR2 DEFAULT NULL)
  RETURN dm_clusters PIPELINED;
```

Parameters

Table 57-102 GET_MODEL_DETAILS_OC Function Parameters

Parameter	Description
model_name	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
cluster_id	The ID of a cluster in the model. When a valid cluster ID is specified, only the details of this cluster are returned. Otherwise the details for all clusters are returned.
attribute	The name of an attribute. When a valid attribute name is specified, only the details of this attribute are returned. Otherwise, the details for all attributes are returned
centroid	This parameter accepts the following values: • 1: Details about centroids are returned (default) • 0: Details about centroids are not returned
histogram	This parameter accepts the following values: • 1: Details about histograms are returned (default) • 0: Details about histograms are not returned
rules	This parameter accepts the following values: • 2: Details about rules are returned (default) • 1: Rule summaries are returned • 0: No information about rules is returned
topn_attributes	Restricts the number of attributes returned in the centroid, histogram, and rules objects. Only the *n* attributes with the highest confidence values in the rules are returned. If the number of attributes included in the rules is less than *topn*, then up to *n* additional attributes in alphabetical order are returned. If both the attribute and topn_attributes parameters are specified, then topn_attributes is ignored.
partition_name	Specifies a partition in a partitioned model.

Usage Notes

1. For information about machine learning data types and return values for clustering algorithms piped output from table functions, see "Data Types".
2. When the value is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Examples

The following example returns model details for the clustering model OC_SH_Clus_sample.

For each cluster in this example, the split predicate indicates the attribute and the condition used to assign records to the cluster's children during model build. It provides an important piece of information on how the population within a cluster can be divided up into two smaller clusters.

```
SELECT clu_id, attribute_name, op, s_value
     FROM (SELECT a.id clu_id, sp.attribute_name, sp.conditional_operator op,
                  sp.attribute_str_value s_value
             FROM TABLE(DBMS_DATA_MINING.GET_MODEL_DETAILS_OC(
                    'OC_SH_Clus_sample')) a,
```

```
                TABLE(a.split_predicate) sp
          ORDER BY a.id, op, s_value)
    WHERE ROWNUM < 11;

    CLU_ID ATTRIBUTE_NAME       OP S_VALUE
---------- -------------------- --------------------------------
         1 OCCUPATION           IN ?
         1 OCCUPATION           IN Armed-F
         1 OCCUPATION           IN Cleric.
         1 OCCUPATION           IN Crafts
         2 OCCUPATION           IN ?
         2 OCCUPATION           IN Armed-F
         2 OCCUPATION           IN Cleric.
         3 OCCUPATION           IN Exec.
         3 OCCUPATION           IN Farming
         3 OCCUPATION           IN Handler
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.33 GET_MODEL_SETTINGS Function

The `GET_MODEL_SETTINGS` function returns the settings used to build the given model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. See "Static Data Dictionary Views: `ALL_ALL_TABLES` to `ALL_OUTLINES`" in *Oracle Database Reference*.

Syntax

```
FUNCTION get_model_settings(model_name IN VARCHAR2)
  RETURN DM_Model_Settings PIPELINED;
```

Parameters

Table 57-103 GET_MODEL_SETTINGS Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.

Return Values

Table 57-104 GET_MODEL_SETTINGS Function Return Values

Return Value	Description

`DM_MODEL_SETTINGS`	A set of rows of type `DM_MODEL_SETTINGS`. The rows have the following columns:
	`DM_MODEL_SETTINGS TABLE OF SYS.DM_MODEL_SETTING`
	`Name Type`
	`SETTING_NAME VARCHAR2(30)`
	`SETTING_VALUE VARCHAR2(4000)`

Usage Notes

1. This table function pipes out rows of type `DM_MODEL_SETTINGS`. For information on machine learning data types and piped output from table functions, see "DBMS_DATA_MINING Datatypes".
2. The setting names/values include both those specified by the user and any defaults assigned by the build process.

Examples

The following example returns model model settings for an example naive Bayes model.

```
SETTING_NAME                   SETTING_VALUE
------------------------------ ------------------------------
ALGO_NAME                       ALGO_NAIVE_BAYES
PREP_AUTO                       ON
ODMS_MAX_PARTITIONS             1000
NABS_SINGLETON_THRESHOLD       0
CLAS_WEIGHTS_BALANCED          OFF
NABS_PAIRWISE_THRESHOLD        0
ODMS_PARTITION_COLUMNS         GENDER,Y_BOX_GAMES
ODMS_MISSING_VALUE_TREATMENT   ODMS_MISSING_VALUE_AUTO
ODMS_SAMPLING                  ODMS_SAMPLING_DISABLE

9 rows selected.
```

Related Topics

- *Oracle Database Reference*

57.8.34 GET_MODEL_SIGNATURE Function

The `GET_MODEL_SIGNATURE` function returns the list of columns from the build input table that were used by the build process to train the model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. See "Static Data Dictionary Views: `ALL_ALL_TABLES` to `ALL_OUTLINES`" in *Oracle Database Reference*.

Syntax

```
FUNCTION get_model_signature (model_name IN VARCHAR2)
RETURN DM_Model_Signature PIPELINED;
```

Parameters

Table 57-105 GET_MODEL_SIGNATURE Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.

Return Values

Table 57-106 GET_MODEL_SIGNATURE Function Return Values

Return Value	Description
DM_MODEL_SIGNATURE	A set of rows of type DM_MODEL_SIGNATURE. The rows have the following columns: DM_MODEL_SIGNATURE TABLE OF SYS.DM_MODEL_SIGNATURE_ATTRIBUTE Name — Type ATTRIBUTE_NAME — VARCHAR2(130) ATTRIBUTE_TYPE — VARCHAR2(106)

Usage Notes

1. This table function pipes out rows of type DM_MODEL_SIGNATURE. For information on machine learning data types and piped output from table functions, see "DBMS_DATA_MINING Datatypes".
2. The signature names or types include only those attributes used by the build process.

Examples

The following example returns model settings for an example naive Bayes model.

```
ATTRIBUTE_NAME                 ATTRIBUTE_TYPE
------------------------------ ------------------
AGE                            NUMBER
ANNUAL_INCOME                  NUMBER
AVERAGE___ITEMS_PURCHASED      NUMBER
BOOKKEEPING_APPLICATION        NUMBER
BULK_PACK_DISKETTES            NUMBER
BULK_PURCH_AVE_AMT             NUMBER
DISABLE_COOKIES                NUMBER
EDUCATION                      VARCHAR2
FLAT_PANEL_MONITOR             NUMBER
GENDER                         VARCHAR2
HOME_THEATER_PACKAGE           NUMBER
HOUSEHOLD_SIZE                 VARCHAR2
MAILING_LIST                   NUMBER
MARITAL_STATUS                 VARCHAR2
NO_DIFFERENT_KIND_ITEMS        NUMBER
OCCUPATION                     VARCHAR2
OS_DOC_SET_KANJI               NUMBER
PETS                           NUMBER
PRINTER_SUPPLIES               NUMBER
PROMO_RESPOND                  NUMBER
SHIPPING_ADDRESS_COUNTRY       VARCHAR2
SR_CITIZEN                     NUMBER
TOP_REASON_FOR_SHOPPING        VARCHAR2
WKS_SINCE_LAST_PURCH           NUMBER
WORKCLASS                      VARCHAR2
YRS_RESIDENCE                  NUMBER
Y_BOX_GAMES                    NUMBER

27 rows selected.
```

Related Topics

- *Oracle Database Reference*

57.8.35 GET_MODEL_DETAILS_SVD Function

The `GET_MODEL_DETAILS_SVD` function returns a set of rows that provide the details of a singular value decomposition model. Oracle recommends to use model details view settings. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

Refer to Model Details View for Singular Value Decomposition.

Syntax

```
DBMS_DATA_MINING.get_model_details_svd(
      model_name IN VARCHAR2,
      matrix_type IN VARCHAR2 DEFAULT NULL,
      partition_name VARCHAR2 DEFAULT NULL)
   RETURN DM_SVD_MATRIX_Set PIPELINED;
```

Parameters

Table 57-107 GET_MODEL_DETAILS_SVD Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`matrix_type`	Specifies which of the three SVD matrix types to return. Values are: `U`, `S`, `V`, and `NULL`. When `matrix_type` is null (default), all matrices are returned. The U matrix is only computed when the `SVDS_U_MATRIX_OUTPUT` setting is enabled. It is not computed by default. If the model does not contain U matrices and you set `matrix_type` to `U`, an empty set of rows is returned. See Table 57-27.
`partition_name`	A partition in a partitioned model.

Return Values

Table 57-108 GET_MODEL_DETAILS_SVD Function Return Values

Return Value	Description

`DM_SVD_MATRIX_SET`	A set of rows of type `DM_SVD_MATRIX`. The rows have the following columns:
	`(matrix_type CHAR(1),`
	`feature_id NUMBER,`
	`mapped_feature_id VARCHAR2(4000),`
	`attribute_name VARCHAR2(4000),`
	`attribute_subname VARCHAR2(4000),`
	`case_id VARCHAR2(4000),`
	`value NUMBER,`
	`variance NUMBER,`
	`pct_cum_variance NUMBER)`
	See Usage Notes for details.

Usage Notes

1. This table function pipes out rows of type `DM_SVD_MATRIX`. For information on machine learning data types and piped output from table functions, see "Data Types".

   The columns in each row returned by `GET_MODEL_DETAILS_SVD` are described as follows:

Column in DM_SVD_MATRIX_SET	Description
`matrix_type`	The type of matrix. Possible values are **S**, **V**, and **U**. This field is never null.
`feature_id`	The feature that the matrix entry refers to.
`mapped_feature_id`	A descriptive name for the feature.
`attribute_name`	Column name in the **V** matrix component bases. This field is null for the **S** and **U** matrices.
`attribute_subname`	Subname in the **V** matrix component bases. This is relevant only in the case of a nested column. This field is null for the **S** and **U** matrices.
`case_id`	Unique identifier of the row in the build data described by the **U** matrix projection. This field is null for the **S** and **V** matrices.
`value`	The matrix entry value.
`variance`	The variance explained by a component. It is non-null only for **S** matrix entries. This column is non-null only for **S** matrix entries and for SVD models with setting `dbms_data_mining.svds_scoring_mode` set to `dbms_data_mining.svds_scoring_pca` and the build data is centered, either manually or because the setting `dbms_data_mining.prep_auto` is set to `dbms_data_mining.prep_auto_on`.
`pct_cum_variance`	The percent cumulative variance explained by the components thus far. The components are ranked by the explained variance in descending order. This column is non-null only for **S** matrix entries and for SVD models with setting `dbms_data_mining.svds_scoring_mode` set to `dbms_data_mining.svds_scoring_pca` and the build data is centered, either manually or because the setting `dbms_data_mining.prep_auto` is set to `dbms_data_mining.prep_auto_on`.

2. The output of `GET_MODEL_DETAILS` is in sparse format. Zero values are not returned. Only the diagonal elements of the **S** matrix, the non-zero coefficients in the **V** matrix bases, and the non-zero **U** matrix projections are returned.

   There is one exception: If the data row does not produce non-zero **U** Matrix projections, the case ID for that row is returned with `NULL` for the `feature_id` and `value`. This is done to avoid losing any records from the original data.

3. `GET_MODEL_DETAILS` functions preserve model transparency by automatically reversing the transformations applied during the build process. Thus the attributes returned in the model details are the original attributes (or a close approximation of the original attributes) used to build the model.

4. When the value is `NULL` for a partitioned model, an exception is thrown. When the value is not null, it must contain the preferred partition name.

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.36 GET_MODEL_DETAILS_SVM Function

The `GET_MODEL_DETAILS_SVM` function returns a set of rows that provide the details of a linear support vector machines (SVM) model. If invoked for nonlinear SVM, it returns `ORA-40215`. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views in *Oracle Machine Learning for SQL User's Guide*.

In linear SVM models, only nonzero coefficients are stored. This reduces storage and speeds up model loading. As a result, if an attribute is missing in the coefficient list returned by `GET_MODEL_DETAILS_SVM`, then the coefficient of this attribute should be interpreted as zero.

Syntax

```
DBMS_DATA_MINING.get_model_details_svm(
      model_name   VARCHAR2,
      reverse_coef NUMBER DEFAULT 0,
      partition_name VARCHAR2 DEFAULT NULL)
  RETURN DM_SVM_Linear_Coeff_Set PIPELINED;
```

Parameters

Table 57-109 GET_MODEL_DETAILS_SVM Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`reverse_coef`	Whether or not `GET_MODEL_DETAILS_SVM` should transform the attribute coefficients using the original attribute transformations. When `reverse_coef` is set to 0 (default), `GET_MODEL_DETAILS_SVM` returns the coefficients directly from the model without applying transformations. When `reverse_coef` is set to 1, `GET_MODEL_DETAILS_SVM` transforms the coefficients and bias by applying the normalization shifts and scales that were generated using automatic data preparation. See Usage Note 4.
`partition_name`	Specifies a partition in a partitioned model.

Return Values

Table 57-110 GET_MODEL_DETAILS_SVM Function Return Values

Return Value	Description

DM_SVM_LINEAR_COEFF_SET	A set of rows of type DM_SVM_LINEAR_COEFF. The rows have the following columns:
	`(class VARCHAR2(4000),`
	`attribute_set DM_SVM_ATTRIBUTE_SET)`
	The attribute_set column returns a nested table of type DM_SVM_ATTRIBUTE_SET. The rows, of type DM_SVM_ATTRIBUTE, have the following columns:
	`(attribute_name VARCHAR2(4000),`
	`attribute_subname VARCHAR2(4000),`
	`attribute_value VARCHAR2(4000),`
	`coefficient NUMBER)`
	See Usage Notes.

Usage Notes

1. This table function pipes out rows of type DM_SVM_LINEAR_COEFF. For information on machine learning data types and piped output from table functions, see "Data Types".
2. The class column of DM_SVM_LINEAR_COEFF contains classification target values. For SVM Regression models, class is null. For each classification target value, a set of coefficients is returned. For binary classification, one-class classification, and regression models, only a single set of coefficients is returned.
3. The attribute_value column in DM_SVM_ATTRIBUTE_SET is used for categorical attributes.
4. GET_MODEL_DETAILS functions preserve model transparency by automatically reversing the transformations applied during the build process. Thus the attributes returned in the model details are the original attributes (or a close approximation of the original attributes) used to build the model.

   The coefficients are related to the transformed, not the original, attributes. When returned directly with the model details, the coefficients may not provide meaningful information. If you want GET_MODEL_DETAILS_SVM to transform the coefficients such that they relate to the original attributes, set the reverse_coef parameter to 1.
5. When the value is NULL for a partitioned model, an exception is thrown. When the value is not null, it must contain the desired partition name.

Examples

The following example returns model details for the SVM classification model SVMC_SH_Clas_sample, which was created by the sample program dmsvcdem.sql. For information about the sample programs, see *Oracle Machine Learning for SQL User's Guide*.

```
WITH
  mod_dtls AS (
  SELECT *
    FROM TABLE(DBMS_DATA_MINING.GET_MODEL_DETAILS_SVM('SVMC_SH_Clas_sample'))
  ),
  model_details AS (
```

```
  SELECT D.class, A.attribute_name, A.attribute_value, A.coefficient
    FROM mod_dtls D,
          TABLE(D.attribute_set) A
    ORDER BY D.class, ABS(A.coefficient) DESC
  )
  SELECT class, attribute_name aname, attribute_value aval, coefficient coeff
    FROM model_details
    WHERE ROWNUM < 11;

CLASS      ANAME                     AVAL                      COEFF
---------- ------------------------- ------------------------- -----
1                                                             -2.85
1          BOOKKEEPING_APPLICATION                             1.11
1          OCCUPATION                Other                     -.94
1          HOUSEHOLD_SIZE            4-5                        .88
1          CUST_MARITAL_STATUS       Married                    .82
1          YRS_RESIDENCE                                        .76
1          HOUSEHOLD_SIZE            6-8                       -.74
1          OCCUPATION                Exec.                      .71
1          EDUCATION                 11th                      -.71
1          EDUCATION                 Masters                    .63
```

Related Topics

- *Oracle Machine Learning for SQL User's Guide*

57.8.37 GET_MODEL_DETAILS_XML Function

This function returns an XML object that provides the details of a decision tree model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. Use model detail views instead.

See Model Detail Views for Decision Tree in *Oracle Machine Learning for SQL User's Guide*.

Syntax

```
DBMS_DATA_MINING.get_model_details_xml(
      model_name IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN XMLType;
```

Parameters

Table 57-111 GET_MODEL_DETAILS_XML Function Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Specifies a partition in a partitioned model.

Return Values

Table 57-112 GET_MODEL_DETAILS_XML Function Return Value

Return Value	Description
XMLTYPE	The XML definition for the decision tree model. See "XMLTYPE" for details. The XML definition conforms to the Data Mining Group Predictive Model Markup Language (PMML) version 2.1 specification. The specification is available at https://dmg.org. If a nested attribute is used as a splitter, the attribute will appear in the XML document as field="'<column_name>'.<subname>", as opposed to the non-nested attributes which appear in the document as field="<column_name>". **Note:** The column names are surrounded by single quotes and a period separates the column_name from the subname. The rest of the document style remains unchanged.

Usage Notes

Special characters that cannot be displayed by Oracle XML are converted to '#'.

Examples

The following statements in SQL*Plus return the details of the decision tree model dt_sh_clas_sample.

Note: The "&quot" characters you will see in the XML output are a result of SQL*Plus behavior. To display the XML in proper format, cut and past it into a file and open the file in a browser.

```
column dt_details format a320
SELECT
 dbms_data_mining.get_model_details_xml('dt_sh_clas_sample')
 AS DT_DETAILS
FROM dual;


DT_DETAILS
--------------------------------------------------------------------------------
<PMML version="2.1">
  <Header copyright="Copyright (c) 2004, Oracle Corporation. All rights
      reserved."/>
  <DataDictionary numberOfFields="9">
    <DataField name="AFFINITY_CARD" optype="categorical"/>
    <DataField name="AGE" optype="continuous"/>
    <DataField name="BOOKKEEPING_APPLICATION" optype="continuous"/>
    <DataField name="CUST_MARITAL_STATUS" optype="categorical"/>
    <DataField name="EDUCATION" optype="categorical"/>
    <DataField name="HOUSEHOLD_SIZE" optype="categorical"/>
    <DataField name="OCCUPATION" optype="categorical"/>
    <DataField name="YRS_RESIDENCE" optype="continuous"/>
    <DataField name="Y_BOX_GAMES" optype="continuous"/>
  </DataDictionary>
```

```
<TreeModel modelName="DT_SH_CLAS_SAMPLE" functionName="classification"
    splitCharacteristic="binarySplit">
  <Extension name="buildSettings">
    <Setting name="TREE_IMPURITY_METRIC" value="TREE_IMPURITY_GINI"/>
    <Setting name="TREE_TERM_MAX_DEPTH" value="7"/>
    <Setting name="TREE_TERM_MINPCT_NODE" value=".05"/>
    <Setting name="TREE_TERM_MINPCT_SPLIT" value=".1"/>
    <Setting name="TREE_TERM_MINREC_NODE" value="10"/>
    <Setting name="TREE_TERM_MINREC_SPLIT" value="20"/>
    <costMatrix>
      <costElement>
        <actualValue>0</actualValue>
        <predictedValue>0</predictedValue>
        <cost>0</cost>
      </costElement>
      <costElement>
        <actualValue>0</actualValue>
        <predictedValue>1</predictedValue>
        <cost>1</cost>
      </costElement>
      <costElement>
        <actualValue>1</actualValue>
        <predictedValue>0</predictedValue>
        <cost>8</cost>
      </costElement>
      <costElement>
        <actualValue>1</actualValue>
        <predictedValue>1</predictedValue>
        <cost>0</cost>
      </costElement>
    </costMatrix>
  </Extension>
  <MiningSchema>
    .
    .
    .
    .
    .
    .
    </Node>
  </Node>
 </TreeModel>
</PMML>
```

57.8.38 GET_MODEL_TRANSFORMATIONS Function

This function returns the transformation expressions embedded in the specified model. Starting from Oracle Database 12*c* Release 2, this function is deprecated. See "Static Data Dictionary Views: `ALL_ALL_TABLES` to `ALL_OUTLINES`" in *Oracle Database Reference*.

All `GET_*` interfaces are replaced by model views, and Oracle recommends that users reference the model views to retrieve the relevant information. The `GET_MODEL_TRANSFORMATIONS` function is replaced by the following:

- USER(/DBA/ALL)_MINING_MODEL_XFORMS: provides the user-embedded transformations
- DM$VX prefixed model view: provides text feature extraction information
- D$VN prefixed mode view: provides normalization and missing value information
- DM$VB: provides binning information

See Also:

"About Transformation Lists" in DBMS_DATA_MINING_TRANSFORM Operational Notes

"GET_TRANSFORM_LIST Procedure"

"CREATE_MODEL Procedure"

"ALL_MINING_MODEL_XFORMS" in *Oracle Database Reference*

"DBA_MINING_MODEL_XFORMS" in *Oracle Database Reference*

"USER_MINING_MODEL_XFORMS" in *Oracle Database Reference*

Model Details View for Binning

Normalization and Missing Value Handling

Data Preparation for Text Features

Syntax

```
DBMS_DATA_MINING.get_model_transformations(
      model_name IN VARCHAR2,
      partition_name IN VARCHAR2 DEFAULT NULL)
  RETURN DM_Transforms PIPELINED;
```

Parameters

Table 57-113 GET_MODEL_TRANSFORMATIONS Function Parameters

Parameter	Description
`model_name`	Indicates the name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`partition_name`	Specifies a partition in a partitioned model

Return Values

Table 57-114 GET_MODEL_TRANSFORMATIONS Function Return Value

Return Value	Description

`DM_TRANSFORMS`	The transformation expressions embedded in *`model_name`*.
	The `DM_TRANSFORMS` type is a table of `DM_TRANSFORM` objects. Each `DM_TRANSFORM` has these fields:
	`attribute_name VARCHAR2(4000)`
	`attribute_subname VARCHAR2(4000)`
	`expression CLOB`
	`reverse_expression CLOB`

Usage Notes

When Automatic Data Preparation (ADP) is enabled, both automatic and user-defined transformations may be associated with an attribute. In this case, the user-defined transformations are evaluated before the automatic transformations.

When invoked for a partitioned model, the `partition_name` parameter must be specified.

Examples

In this example, several columns in the `SH.CUSTOMERS` table are used to create a naive Bayes model. A transformation expression is specified for one of the columns. The model does not use ADP.

```
CREATE OR REPLACE VIEW mining_data AS
   SELECT cust_id, cust_year_of_birth, cust_income_level,cust_credit_limit
   FROM sh.customers;

describe mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                        NOT NULL NUMBER(4)
 CUST_INCOME_LEVEL                                  VARCHAR2(30)
 CUST_CREDIT_LIMIT                                  NUMBER

CREATE TABLE settings_nb(
      setting_name  VARCHAR2(30),
      setting_value VARCHAR2(30));
BEGIN
     INSERT INTO settings_nb (setting_name, setting_value) VALUES
           (dbms_data_mining.algo_name, dbms_data_mining.algo_naive_bayes);
     INSERT INTO settings_nb (setting_name, setting_value) VALUES
           (dbms_data_mining.prep_auto, dbms_data_mining.prep_auto_off);
     COMMIT;
END;
/
DECLARE
    mining_data_xforms   dbms_data_mining_transform.TRANSFORM_LIST;
  BEGIN
    dbms_data_mining_transform.SET_TRANSFORM (
         xform_list           =>  mining_data_xforms,
         attribute_name       => 'cust_year_of_birth',
         attribute_subname    =>  null,
         expression           => 'cust_year_of_birth + 10',
         reverse_expression   => 'cust_year_of_birth - 10');
    dbms_data_mining.CREATE_MODEL (
        model_name           =>  'new_model',
        mining_function      =>   dbms_data_mining.classification,
        data_table_name      =>  'mining_data',
        case_id_column_name  =>  'cust_id',
        target_column_name   =>  'cust_income_level',
        settings_table_name  =>  'settings_nb',
        data_schema_name     =>   nulL,
        settings_schema_name =>   null,
        xform_list           =>   mining_data_xforms );
  END;
 /
SELECT attribute_name, TO_CHAR(expression), TO_CHAR(reverse_expression)
      FROM TABLE (dbms_data_mining.GET_MODEL_TRANSFORMATIONS('new_model'));
```

```
ATTRIBUTE_NAME      TO_CHAR(EXPRESSION)       TO_CHAR(REVERSE_EXPRESSION)
------------------  ------------------------  ----------------------------
CUST_YEAR_OF_BIRTH  cust_year_of_birth + 10   cust_year_of_birth - 10
```

Related Topics

- *Oracle Database Reference*

57.8.39 GET_TRANSFORM_LIST Procedure

This procedure converts transformation expressions specified as `DM_TRANSFORMS` to a transformation list (`TRANSFORM_LIST`) that can be used in creating a model. `DM_TRANSFORMS` is returned by the `GET_MODEL_TRANSFORMATIONS` function.

You can also use routines in the `DBMS_DATA_MINING_TRANSFORM` package to construct a transformation list.

> **See Also:**
>
> "About Transformation Lists" in DBMS_DATA_MINING_TRANSFORM
>
> "GET_MODEL_TRANSFORMATIONS Function"
>
> "CREATE_MODEL Procedure"

Syntax

```
DBMS_DATA_MINING.GET_TRANSFORM_LIST (
      xform_list           OUT NOCOPY TRANSFORM_LIST,
      model_xforms         IN  DM_TRANSFORMS);
```

Parameters

Table 57-115 GET_TRANSFORM_LIST Procedure Parameters

Parameter	Description

`xform_list`	A list of transformation specifications that can be embedded in a model. Accepted as a parameter to the CREATE_MODEL Procedure.
	The `TRANSFORM_LIST` type is a table of `TRANSFORM_REC` objects. Each `TRANSFORM_REC` has these fields:
	`attribute_name VARCHAR2(30)`
	`attribute_subname VARCHAR2(4000)`
	`expression EXPRESSION_REC`
	`reverse_expression EXPRESSION_REC`
	`attribute_spec VARCHAR2(4000)`
	For details about the `TRANSFORM_LIST` collection type, see Table 58-1.

Table 57-115 (Cont.) GET_TRANSFORM_LIST Procedure Parameters

Parameter	Description
model_xforms	A list of embedded transformation expressions returned by the GET_MODEL_TRANSFORMATIONS Function for a specific model. The DM_TRANSFORMS type is a table of DM_TRANSFORM objects. Each DM_TRANSFORM has these fields: attribute_name VARCHAR2(4000) attribute_subname VARCHAR2(4000) expression CLOB reverse_expression CLOB

Examples

In this example, a model mod1 is trained using several columns in the SH.CUSTOMERS table. The model uses ADP, which automatically bins one of the columns.

A second model mod2 is trained on the same data without ADP, but it uses a transformation list that was obtained from mod1. As a result, both mod1 and mod2 have the same embedded transformation expression.

```
CREATE OR REPLACE VIEW mining_data AS
     SELECT cust_id, cust_year_of_birth, cust_income_level, cust_credit_limit
     FROM sh.customers;

describe mining_data
 Name                                    Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                 NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                      NOT NULL NUMBER(4)
 CUST_INCOME_LEVEL                                VARCHAR2(30)
 CUST_CREDIT_LIMIT                                NUMBER

CREATE TABLE setmod1(setting_name  VARCHAR2(30),setting_value VARCHAR2(30));
BEGIN
   INSERT INTO setmod1 VALUES (dbms_data_mining.algo_name, dbms_data_mining.algo_naive_bayes);
   INSERT INTO setmod1 VALUES (dbms_data_mining.prep_auto,dbms_data_mining.prep_auto_on);
   dbms_data_mining.CREATE_MODEL (
               model_name            => 'mod1',
               mining_function       => dbms_data_mining.classification,
               data_table_name       => 'mining_data',
               case_id_column_name   => 'cust_id',
               target_column_name    => 'cust_income_level',
               settings_table_name   => 'setmod1');
    COMMIT;
END;
/
CREATE TABLE setmod2(setting_name  VARCHAR2(30),setting_value VARCHAR2(30));
BEGIN
  INSERT INTO setmod2
      VALUES (dbms_data_mining.algo_name, dbms_data_mining.algo_naive_bayes);
  COMMIT;
END;
/
DECLARE
  v_xform_list       dbms_data_mining_transform.TRANSFORM_LIST;
  dmxf               DM_TRANSFORMS;
```

```
BEGIN
   EXECUTE IMMEDIATE
    'SELECT dm_transform(attribute_name, attribute_subname,expression, reverse_expression)
     FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS (''mod1''))'
     BULK COLLECT INTO dmxf;
   dbms_data_mining.GET_TRANSFORM_LIST (
        xform_list             =>  v_xform_list,
        model_xforms           =>  dmxf);
   dbms_data_mining.CREATE_MODEL(
         model_name            => 'mod2',
         mining_function       =>  dbms_data_mining.classification,
         data_table_name       => 'mining_data',
         case_id_column_name   => 'cust_id',
         target_column_name    => 'cust_income_level',
         settings_table_name   => 'setmod2',
         xform_list            =>  v_xform_list);
END;
/

-- Transformation expression embedded in mod1
SELECT TO_CHAR(expression) FROM TABLE (dbms_data_mining.GET_MODEL_TRANSFORMATIONS('mod1'));

TO_CHAR(EXPRESSION)
--------------------------------------------------------------------------------
CASE WHEN "CUST_YEAR_OF_BIRTH"<1915 THEN 0 WHEN "CUST_YEAR_OF_BIRTH"<=1915 THEN 0
WHEN "CUST_YEAR_OF_BIRTH"<=1920.5 THEN 1 WHEN "CUST_YEAR_OF_BIRTH"<=1924.5 THEN 2
.
.
.
.5 THEN 29 WHEN "CUST_YEAR_OF_BIRTH" IS NOT NULL THEN 30 END

-- Transformation expression embedded in mod2
SELECT TO_CHAR(expression) FROM TABLE (dbms_data_mining.GET_MODEL_TRANSFORMATIONS('mod2'));

TO_CHAR(EXPRESSION)
--------------------------------------------------------------------------------
CASE WHEN "CUST_YEAR_OF_BIRTH"<1915 THEN 0 WHEN "CUST_YEAR_OF_BIRTH"<=1915 THEN 0
WHEN "CUST_YEAR_OF_BIRTH"<=1920.5 THEN 1 WHEN "CUST_YEAR_OF_BIRTH"<=1924.5 THEN 2
.
.
.
.5 THEN 29 WHEN "CUST_YEAR_OF_BIRTH" IS NOT NULL THEN 30 END

-- Reverse transformation expression embedded in mod1
SELECT TO_CHAR(reverse_expression)FROM TABLE (dbms_data_mining.GET_MODEL_TRANSFORMATIONS('mod1'));

TO_CHAR(REVERSE_EXPRESSION)
--------------------------------------------------------------------------------
DECODE("CUST_YEAR_OF_BIRTH",0,'( ; 1915), [1915; 1915]',1,'(1915; 1920.5]',2,'(1
920.5; 1924.5]',3,'(1924.5; 1928.5]',4,'(1928.5; 1932.5]',5,'(1932.5; 1936.5]',6
.
.
.
8,'(1987.5; 1988.5]',29,'(1988.5; 1989.5]',30,'(1989.5;  )',NULL,'NULL')

-- Reverse transformation expression embedded in mod2
SELECT TO_CHAR(reverse_expression) FROM TABLE (dbms_data_mining.GET_MODEL_TRANSFORMATIONS('mod2'));

TO_CHAR(REVERSE_EXPRESSION)
--------------------------------------------------------------------------------
DECODE("CUST_YEAR_OF_BIRTH",0,'( ; 1915), [1915; 1915]',1,'(1915; 1920.5]',2,'(1
920.5; 1924.5]',3,'(1924.5; 1928.5]',4,'(1928.5; 1932.5]',5,'(1932.5; 1936.5]',6
```

```
.
.
.
8,'(1987.5; 1988.5]',29,'(1988.5; 1989.5]',30,'(1989.5;  )',NULL,'NULL')
```

57.8.40 IMPORT_MODEL Procedure

This procedure imports one or more machine learning models. The procedure is overloaded. You can call it to import machine learning models from a dump file set, or you can call it to import a single machine learning model from a PMML document.

Import from a dump file set

You can import machine learning models from a dump file set that was created by the EXPORT_MODEL Procedure. `IMPORT_MODEL` and `EXPORT_MODEL` use Oracle Data Pump technology to export to and import from a dump file set.

When Oracle Data Pump is used directly to export/import an entire schema or database, the machine learning models in the schema or database are included. `EXPORT_MODEL` and `IMPORT_MODEL` export/import machine learning models only.

Import from PMML

You can import a machine learning model represented in Predictive Model Markup Language (PMML). The model must be of type `RegressionModel`, either linear regression or binary logistic regression.

PMML is an XML-based standard specified by the Data Mining Group (`https://dmg.org`). Applications that are PMML-compliant can deploy PMML-compliant models that were created by any vendor. Oracle Machine Learning for SQL supports the core features of PMML 3.1 for regression models.

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for more information about exporting and importing machine learning models
>
> *Oracle Database Utilities* for information about Oracle Data Pump
>
> `https://dmg.org/dmg-faq.html` for more information about PMML

Syntax

Imports a machine learning model from a dump file set:

```
DBMS_DATA_MINING.IMPORT_MODEL (
      filename          IN  VARCHAR2,
      directory         IN  VARCHAR2,
      model_filter      IN  VARCHAR2 DEFAULT NULL,
      operation         IN  VARCHAR2 DEFAULT NULL,
      remote_link       IN  VARCHAR2 DEFAULT NULL,
      jobname           IN  VARCHAR2 DEFAULT NULL,
      schema_remap      IN  VARCHAR2 DEFAULT NULL,
      tablespace_remap  IN  VARCHAR2 DEFAULT NULL);
```

Imports a machine learning model from a PMML document:

```
DBMS_DATA_MINING.IMPORT_MODEL (
      model_name         IN  VARCHAR2,
      pmmldoc            IN  XMLTYPE
      strict_check       IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 57-116 IMPORT_MODEL Procedure Parameters

Parameter	Description
filename	Name of the dump file set from which the models should be imported. The dump file set must have been created by the EXPORT_MODEL procedure or the expdp export utility of Oracle Data Pump. The dump file set can contain one or more files. (Refer to "EXPORT_MODEL Procedure" for details.) If the dump file set contains multiple files, you can specify '*filename*%U' instead of listing them. For example, if your dump file set contains 3 files, archive01.dmp, archive02.dmp, and archive03.dmp, you can import them by specifying 'archive%U'.
directory	Name of a pre-defined directory object that specifies where the dump file set is located. Both the exporting and the importing user must have read/write access to the directory object and to the file system directory that it identifies. Note: The target database must have also have read/write access to the file system directory.

model_filter	Optional parameter that specifies one or more models to import. If you do not specify a value for model_filter, all models in the dump file set are imported. You can also specify NULL (the default) or 'ALL' to import all models.
	The value of model_filter can be one or more model names. The following are valid filters.
	`'mymodel1'`
	`'name IN (''mymodel2'',''mymodel3'')'`
	The first causes IMPORT_MODEL to import a single model named mymodel1. The second causes IMPORT_MODEL to import two models, mymodel2 and mymodel3.

operation	Optional parameter that specifies whether to import the models or the SQL statements that create the models. By default, the models are imported. You can specify either of the following values for operation: • 'IMPORT' — Import the models (Default) • 'SQL_FILE'— Write the SQL DDL for creating the models to a text file. The text file is named *job_name*.sql and is located in the dump set directory.
remote_link	Optional parameter that specifies the name of a database link to a remote system. The default value is NULL. A database link is a schema object in a local database that enables access to objects in a remote database. When you specify a value for remote_link, you can import models into the local database from the remote database. The import is fileless; no dump file is involved. The IMP_FULL_DATABASE role is required for importing the remote models. The EXP_FULL_DATABASE privilege, the CREATE DATABASE LINK privilege, and other privileges may also be required. (See Example 2.)
jobname	Optional parameter that specifies the name of the import job. By default, the name has the form *username*_imp_*nnnn*, where *nnnn* is a number. For example, a job name in the SCOTT schema might be SCOTT_imp_134. If you specify a job name, it must be unique within the schema. The maximum length of the job name is 30 characters. A log file for the import job, named *jobname*.log, is created in the same directory as the dump file set.

Table 57-116 (Cont.) IMPORT_MODEL Procedure Parameters

Parameter	Description
schema_remap	Optional parameter for importing into a different schema. By default, models are exported and imported within the same schema. If the dump file set belongs to a different schema, you must specify a schema mapping in the form *export_user*:*import_user*. For example, you would specify 'SCOTT:MARY' to import a model exported by SCOTT into the MARY schema. Note: In some cases, you may need to have the IMP_FULL_DATABASE privilege or the SYS role to import a model from a different schema.
tablespace_remap	Optional parameter for importing into a different tablespace. By default, models are exported and imported within the same tablespace. If the dump file set belongs to a different tablespace, you must specify a tablespace mapping in the form *export_tablespace*:*import_tablespace*. For example, you would specify 'TBLSPC01:TBLSPC02' to import a model that was exported from tablespace TBLSPC01 into tablespace TBLSPC02. Note: In some cases, you may need to have the IMP_FULL_DATABASE privilege or the SYS role to import a model from a different tablespace.
model_name	Name for the new model that will be created in the database as a result of an import from PMML The name must be unique within the user's schema.
pmmldoc	The PMML document representing the model to be imported. The PMML document has an XMLTYPE object type. See "XMLTYPE" for details.
strict_check	Whether or not an error occurs when the PMML document contains sections that are not part of core PMML (for example, Output or Targets). OML4SQL supports only core PMML; any non-core features may affect the scoring representation. If the PMML does not strictly conform to core PMML and strict_check is set to TRUE, then IMPORT_MODEL returns an error. If strict_check is FALSE (the default), then the error is suppressed. The model may be imported and scored.

Examples

1. This example shows a model being exported and imported within the schema oml_user2. Then the same model is imported into the oml_user3 schema. The oml_user3 user has the IMP_FULL_DATABASE privilege. The oml_user2 user has been assigned the USER2 tablespace; oml_user3 has been assigned the USER3 tablespace.

```
SQL> connect oml_user2
Enter password: oml_user2_password
Connected.
SQL> select model_name from user_mining_models;

MODEL_NAME
------------------------------
NMF_SH_SAMPLE
SVMO_SH_CLAS_SAMPLE
SVMR_SH_REGR_SAMPLE

-- export the model called NMF_SH_SAMPLE to a dump file in same schema
SQL>EXECUTE DBMS_DATA_MINING.EXPORT_MODEL (
            filename =>'NMF_SH_SAMPLE_out',
            directory =>'DATA_PUMP_DIR',
            model_filter => 'name = ''NMF_SH_SAMPLE''');

-- import the model back into the same schema
```

```
SQL>EXECUTE DBMS_DATA_MINING.IMPORT_MODEL (
            filename => 'NMF_SH_SAMPLE_out01.dmp',
            directory => 'DATA_PUMP_DIR',
            model_filter => 'name = ''NMF_SH_SAMPLE''');

-- connect as different user
-- import same model into that schema
SQL> connect oml_user3
Enter password: oml_user3_password
Connected.
SQL>EXECUTE DBMS_DATA_MINING.IMPORT_MODEL (
            filename => 'NMF_SH_SAMPLE_out01.dmp',
            directory => 'DATA_PUMP_DIR',
            model_filter => 'name = ''NMF_SH_SAMPLE''',
            operation =>'IMPORT',
            remote_link => NULL,
            jobname => 'nmf_imp_job',
            schema_remap => 'oml_user2:oml_user3',
            tablespace_remap => 'USER2:USER3');
```

The following example shows user `MARY` importing all models from a dump file, `model_exp_001.dmp`, which was created by user `SCOTT`. User `MARY` has been assigned a tablespace named `USER2`; user `SCOTT` was assigned the tablespace `USERS` when the models were exported into the dump file `model_exp_001.dmp`.The dump file is located in the file system directory mapped to a directory object called `DM_DUMP`. If user `MARY` does not have `IMP_FULL_DATABASE` privileges, `IMPORT_MODEL` will raise an error.

```
-- import all models
DECLARE
  file_name  VARCHAR2(40);
BEGIN
  file_name := 'model_exp_001.dmp';
  DBMS_DATA_MINING.IMPORT_MODEL(
            filename=> 'file_name',
            directory=>'DM_DUMP',
            schema_remap=>'SCOTT:MARY',
            tablespace_remap=>'USERS:USER2');
  DBMS_OUTPUT.PUT_LINE(
           'DBMS_DATA_MINING.IMPORT_MODEL of all models from SCOTT done!');
END;
/
```

2. This example shows how the user `xuser` could import the model `oml_user.r1mod` from a remote database. The SQL*Net connection alias for the remote database is `R1DB`. The user `xuser` is assigned the `SYSAUX` tablespace; the user `oml_user` is assigned the `TBS_1` tablespace.

```
CONNECT / AS SYSDBA;
GRANT CREATE DATABASE LINK TO xuser;
GRANT imp_full_database TO xuser;
CONNECT xuser/xuserpassword
CREATE DATABASE LINK oml_user_link
        CONNECT TO oml_user IDENTIFIED BY oml_userpassword USING 'R1DB';
EXEC dbms_data_mining.import_model (
    NULL,
   'oml_user_DIR',
   'R1MOD',
    remote_link => 'oml_user_LINK', schema_remap => 'oml_user:XUSER',
                   tablespace_remap => 'TBS_1:SYSAUX' );
SELECT name FROM dm_user_models;

NAME
```

```
--------------------------------------------------------------------------------
R1MOD
```

3. This example shows how a PMML document called `SamplePMML1.xml` could be imported from a location referenced by directory object `PMMLDIR` into the schema of the current user. The imported model will be called `PMMLMODEL1`.

```
BEGIN
    dbms_data_mining.import_model ('PMMLMODEL1',
        XMLType (bfilename ('PMMLDIR', 'SamplePMML1.xml'),
          nls_charset_id ('AL32UTF8')
        ));
END;
```

57.8.41 IMPORT_ONNX_MODEL Procedure

This procedure enables you to import an ONNX model into the Database.

Syntax

```
DBMS_DATA_MINING.IMPORT_ONNX_MODEL(
model_name  IN  VARCHAR2,
model_data  IN  BLOB,
metadata    IN  JSON);
```

Parameters

Table 57-117 IMPORT_ONNX_MODEL Procedure Parameters

Parameter	Description
`model_name`	Name of the model in the form `[schema_name.]model_name`. If you do not specify a schema, then your own schema is used.
`model_data`	It is a `BLOB` holding the ONNX representation of the model. The `BLOB` contains the identical byte sequence as the one stored in an ONNX file.
`metadata`	A JSON description of the metadata describing the model. The metadata at minimum must describe the machine learning function supported by the model. The model's metadata parameters are described in JSON Metadata Parameters for ONNX Models .

Example

The following example illustrates a code snippet of using the `DBMS_DATA_MINING.IMPORT_ONNX_MODEL` procedure. The complete step-by-step example is illustrated in Import ONNX Models and Generate Embeddings and Alternate Method to Import ONNX Models.

```
DBMS_DATA_MINING.IMPORT_ONNX_MODEL('my_embedding_model.onnx',
                                                :blob_bind_variable,
                                                 JSON('{"function" :
"embedding",
                                                        "embeddingOutput" :
"embedding" ,
                                                         "input":{"input":
["DATA"]}}'));
```

For a complete example to illustrate how you can define a BLOB variable and use it in the IMPORT_ONNX_MODEL procedure, you can have the following:

```
CREATE OR REPLACE MY_LOAD_EMBEDDING_MODEL(embedding_model_name VARCHAR2,
onnx_blob BLOB) IS
BEGIN
DBMS_DATA_MINING.IMPORT_ONNX_MODEL(embedding_model_name,
                            onnx_blob,
                            JSON('{"function" : "embedding",
                                   "embeddingOutput" : "embedding" ,
                                   "input":{"input": ["DATA"]}}'));
END;
/
```

Usage Notes

The name of the model follows the same restrictions as those used for other machine learning models, namely:

- The schema name, if provided, is limited to 128 characters.
- The model name is limited to 123 characters and must follow the rules of unquoted identifiers: they contain only alphanumeric characters, the underscore (_), dollar sign ($), and pound sign (#). The initial character must be alphabetic.
- The model size is limited to 1 gigabyte.
- The model must not depend on external initializers. To know more about initializers and other ONNX concepts, see https://onnx.ai/onnx/intro/concepts.html.

57.8.42 IMPORT_SERMODEL Procedure

This procedure imports the serialized format of the model back into a database.

The import routine takes the serialized content in the BLOB and the name of the model to be created with the content. This import does not create model views or tables that are needed for querying model details. The import procedure only provides the ability to score the model.

Syntax

```
DBMS_DATA_MINING.IMPORT_SERMODEL (
      model_data     IN BLOB,
      model_name     IN VARCHAR2,);
```

Parameters

Table 57-118 IMPORT_SERMODEL Procedure Parameters

Parameter	Description
model_data	Provides model data in BLOB format.
model_name	Name of the machine learning model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.

Examples

The following statement imports the serialized format of the models.

```
declare
 v_blob blob;
BEGIN
 dbms_lob.createtemporary(v_blob, FALSE);
-- fill in v_blob from somewhere (e.g., bfile, etc.)
 dbms_data_mining.import_sermodel(v_blob, 'MY_MODEL');
 dbms_lob.freetemporary(v_blob);
END;
/
```

Related Topics

- EXPORT_SERMODEL Procedure
  This procedure exports the model in a serialized format so that they can be moved to another platform for scoring.

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for more information about exporting and importing machine learning models

57.8.43 JSON Schema for R Extensible Algorithm

Provides some flexibility when creating a new JSON object following the JSON schema.

Usage Note

Some flexibility when creating a new JSON object is as follows:

- Partial registration is allowed. For example, the detail function can be missing.
- Different orders are allowed. For example, the detail function can be written before the build function or after it.

Example 57-1 JSON Schema

JSON schema 1.1 for R extensible algorithm:

```
{
    "type": "object",
    "properties": {
        "algo_name_display": { "type" : "object",
                                                "properties" : {
                                                "language" : { "type" :
"string",

"enum" : ["English", "Spanish", "French"],

"default" : "English"},
```

```
                                                  "name" : { "type" : "string"}}
                                               },

            "function_language": {"type": "string" },
            "mining_function": {
                     "type" : "array",
                     "items" : [
                         { "type" : "object",
                            "properties" : {
                               "mining_function_name"  : { "type" : "string"},
                               "build_function": {
                                       "type": "object",
                                       "properties": {
                                            "function_body": { "type": "CLOB" }
                                                           }
                                        },

            "detail_function": {
                     "type" : "array",
                      "items" : [
                          {"type": "object",
                            "properties": {
                                 "function_body": { "type": "CLOB" },
                                 "view_columns": { "type" : "array",
                                                                        "items" : {

"type" : "object",

"properties" : {

 "name" : { "type" : "string"},

 "type" : { "type" : "string",

                     "enum" : ["VARCHAR2",

                                       "NUMBER",

                                       "DATE",

                                       "BOOLEAN"]

                    }
                                                                                  }
                                                                 }
                                                }
                                       }
                     ]
            },

           "score_function": {
                     "type": "object",
                     "properties": {
                           "function_body": { "type": "CLOB" }
                            }
                     },
```

```
        "weight_function": {
                        "type": "object",
                        "properties": {
                            "function_body": { "type": "CLOB" },
                        }
                    }
                                }
           }]
        },

       "algo_setting": {
                "type" : "array",
                "items" : [
                    { "type" : "object",
                       "properties" : {
                          "name"              : { "type" : "string"},
                          "name_display": { "type" : "object",
                                                    "properties" : {
                                                    "language" :
{ "type" : "string",

   "enum" : ["English", "Spanish", "French"],

   "default" : "English"},
                                                    "name" : { "type" :
"string"}}
                                                },
                          "type" : { "type" : "string",
                                        "enum" : ["string", "integer",
"number", "boolean"]},

                          "optional": {"type" : "BOOLEAN",
                                        "default" : "FALSE"},

                          "value" : { "type" :  "string"},

                          "min_value" : { "type": "object",
                                            "properties": {
                                                  "min_value":
{"type": "number"},
                                                   "inclusive":
{ "type": "boolean",

     "default" : TRUE},
                                                }
                                            },
                           "max_value" : {"type": "object",
                                            "properties": {
                                                 "max_value":
{"type": "number"},
                                                 "inclusive":
{ "type": "boolean",

   "default" : TRUE},
                                                    }
                                            },
```

```
                          "categorical choices" : { "type": "array",
                                                                  "items": {
                                                                      "type":
"string"
                                                                   }
                                                               },

                          "description_display": { "type" : "object",

"properties" : {

"language" : { "type" : "string",

             "enum" : ["English", "Spanish", "French"],

             "default" : "English"},
                                                                  "name" :
{ "type" : "string"}}
                                                               }
                        }
                     }
                  ]
            }
    }
}
```

Example 57-2 JSON object example

The following is an JSON object example that must be passed to the registration procedure:

```
{  "algo_name_display"   :     {"English", "t1"},
                          "function_language"    :      "R",
                          "mining_function" : {
  "mining_function_name" : "CLASSIFICATION",
                          "build_function" : {"function_body": "function(dat,
formula, family)
{
                                                    set.seed(1234);
                                          mod <- glm(formula = formula,
data=dat,
                                                       family=
eval(parse(text=family))); mod}"},
           "score_function" :  { "function_body": "function(mod, dat) {
                                             res <- predict(mod, newdata =
dat,
type=''response
                                         '');
                                             res2=data.frame(1-res, res);
res2}"}}
                           },
                           "algo_setting" :   [{"name"                :
"dbms_data_mining.odms_m


                                                 issing_value_treatment",
```

```
                              "name_display"   : {"English",
"dbms_data_mining.odms_missing_value
_treatment"},
                              "type"                  : "string",
                              "optional"            :  "TRUE",
                              "value"                :
"dbms_data_mining.odms_missing_value_mean_mode",
                              "categorical choices"   :
[    "dbms_data_mining.odms_missing_value_mean_mode",

"dbms_data_mining.odms_missing_value_auto",

"dbms_data_mining.odms_missing_value_delete_row"],
                              "description"               : {"English",
                                                                  "how to
treat missing values"}
                           },

{"name"                : "RALG_PARAMETER_FAMILY",
                              "name_display"   : {"English",
"RALG_PARAMETER_FAMILY"},
                              "type"                  : "string",
                              "optional"            :  "TRUE",
                              "value"                 :  "",
                              "description"       : {"English", "R family
parameter in build function"}
                           }
],
                         }
```

57.8.44 REGISTER_ALGORITHM Procedure

Use this function to register a new algorithm by providing the algorithm name, machine learning function, and all other algorithm metadata.

Syntax

```
DBMS_DATA_MINING.REGISTER_ALGORITHM (
                  algorithm_name          IN VARCHAR2,
                  algorithm_metadata      IN CLOB,
                  algorithm_description   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-119 ***REGISTER_ALGORITHM Procedure Parameters***

Parameter	Description
algorithm_name	Name of the algorithm.
algorithm_metadata	Metadata of the algorithm.
algorithm_description	Description of the algorithm.

Usage Notes

The registration procedure performs the following:

- Checks whether `algorithm_metadata` has correct JSON syntax.
- Checks whether the input JSON object follows the predefined JSON schema.
- Checks whether current user has `RQADMIN` privilege.
- Checks duplicate algorithms so that the same algorithm is not registered twice.
- Checks for missing entries. For example, algorithm name, algorithm type, metadata, and build function.

Register Algorithms After the JSON Object Is Created

SQL users can register new algorithms by creating a JSON object following the JSON schema and passing it to the `REGISTER_ALGORITHM` procedure.

```
BEGIN
  DBMS_DATA_MINING.register_algorithm(
    algorithm_name                 =>   't1',
    algorithm_metadata            =>
    '{"function_language" : "R",
      "mining_function" :
        { "mining_function_name" : "CLASSIFICATION",
          "build_function" : {"function_body": "function(dat, formula,
family) { set.seed(1234);
                                          mod <- glm(formula = formula,
data=dat,

family=eval(parse(text=family)));
mod}"},
           "score_function" :  {"function_body": "function(mod, dat) {
                                              res <- predict(mod, newdata =
dat, type=''response'');
                                              res2=data.frame(1-res, res);
res2}"}}
    }',
    algorithm_description  => 't1');
END;
/
```

57.8.45 RANK_APPLY Procedure

This procedure ranks the results of an `APPLY` operation based on a top-N specification for predictive and descriptive model results.

For classification models, you can provide a cost matrix as input, and obtain the ranked results with costs applied to the predictions.

Syntax

```
DBMS_DATA_MINING.RANK_APPLY (
      apply_result_table_name        IN VARCHAR2,
      case_id_column_name            IN VARCHAR2,
      score_column_name              IN VARCHAR2,
      score_criterion_column_name    IN VARCHAR2,
      ranked_apply_table_name        IN VARCHAR2,
      top_N                          IN NUMBER (38) DEFAULT 1,
      cost_matrix_table_name         IN VARCHAR2    DEFAULT NULL,
```

```
    apply_result_schema_name       IN VARCHAR2    DEFAULT NULL,
    cost_matrix_schema_name        IN VARCHAR2    DEFAULT NULL);
```

Parameters

Table 57-120 RANK_APPLY Procedure Parameters

Parameter	Description
apply_result_table_name	Name of the table or view containing the results of an APPLY operation on the test data set (see Usage Notes)
case_id_column_name	Name of the case identifier column. This must be the same as the one used for generating APPLY results.
score_column_name	Name of the prediction column in the apply results table
score_criterion_column_name	Name of the probability column in the apply results table
ranked_apply_result_tab_name	Name of the table containing the ranked apply results
top_N	Top N predictions to be considered from the APPLY results for precision recall computation
cost_matrix_table_name	Name of the cost matrix table
apply_result_schema_name	Name of the schema hosting the APPLY results table
cost_matrix_schema_name	Name of the schema hosting the cost matrix table

Usage Notes

You can use RANK_APPLY to generate ranked apply results, based on a top-N filter and also with application of cost for predictions, if the model was built with costs.

The behavior of RANK_APPLY is similar to that of APPLY with respect to other DDL-like operations such as CREATE_MODEL, DROP_MODEL, and RENAME_MODEL. The procedure does not depend on the model; the only input of relevance is the apply results generated in a fixed schema table from APPLY.

The main intended use of RANK_APPLY is for the generation of the final APPLY results against the scoring data in a production setting. You can apply the model against test data using APPLY, compute various test metrics against various cost matrix tables, and use the candidate cost matrix for RANK_APPLY.

The schema for the apply results from each of the supported algorithms is listed in subsequent sections. The case_id column will be the same case identifier column as that of the apply results.

Classification Models — NB and SVM

For numerical targets, the ranked results table will have the definition as shown:

```
(case_id       VARCHAR2/NUMBER,
prediction     NUMBER,
probability    NUMBER,
cost           NUMBER,
rank           INTEGER)
```

For categorical targets, the ranked results table will have the following definition:

```
(case_id       VARCHAR2/NUMBER,
prediction     VARCHAR2,
probability    NUMBER,
cost           NUMBER,
rank           INTEGER)
```

Clustering Using *k*-Means or O-Cluster

Clustering is an unsupervised machine learning function, and hence there are no targets. The results of an `APPLY` operation contains simply the cluster identifier corresponding to a case, and the associated probability. Cost matrix is not considered here. The ranked results table will have the definition as shown, and contains the cluster ids ranked by `top-N`.

```
(case_id       VARCHAR2/NUMBER,
cluster_id     NUMBER,
probability    NUMBER,
rank           INTEGER)
```

Feature Extraction using NMF

Feature extraction is also an unsupervised machine learning function, and hence there are no targets. The results of an `APPLY` operation contains simply the feature identifier corresponding to a case, and the associated match quality. Cost matrix is not considered here. The ranked results table will have the definition as shown, and contains the feature ids ranked by `top-N`.

```
(case_id        VARCHAR2/NUMBER,
feature_id      NUMBER,
match_quality   NUMBER,
rank            INTEGER)
```

Examples

```
BEGIN
/* build a model with name census_model.
 * (See example under CREATE_MODEL)
 */

/* if training data was pre-processed in any manner,
 * perform the same pre-processing steps on apply
 * data also.
 * (See examples in the section on DBMS_DATA_MINING_TRANSFORM)
 */

/* apply the model to data to be scored */
DBMS_DATA_MINING.RANK_APPLY(
  apply_result_table_name       => 'census_apply_result',
  case_id_column_name           => 'person_id',
  score_column_name             => 'prediction',
  score_criterion_column_name   => 'probability
  ranked_apply_result_tab_name  => 'census_ranked_apply_result',
  top_N                         => 3,
  cost_matrix_table_name        => 'census_cost_matrix');
END;
/

-- View Ranked Apply Results
SELECT *
  FROM census_ranked_apply_result;
```

57.8.46 REMOVE_COST_MATRIX Procedure

The `REMOVE_COST_MATRIX` procedure removes the default scoring matrix from a classification model.

> **See Also:**
>
> - "ADD_COST_MATRIX Procedure"
> - "REMOVE_COST_MATRIX Procedure"

Syntax

```
DBMS_DATA_MINING.REMOVE_COST_MATRIX (
      model_name   IN  VARCHAR2);
```

Parameters

Table 57-121 Remove_Cost_Matrix Procedure Parameters

Parameter	Description
`model_name`	Name of the model in the form [*schema_name*.]*model_name*. If you do not specify a schema, your own schema is used.

Usage Notes

If the model is not in your schema, then `REMOVE_COST_MATRIX` requires the `ALTER ANY MINING MODEL` system privilege or the `ALTER` object privilege for the machine learning model.

Example

The naive Bayes model `NB_SH_CLAS_SAMPLE` has an associated cost matrix that can be used for scoring the model.

```
SQL>SELECT *
      FROM TABLE(dbms_data_mining.get_model_cost_matrix('nb_sh_clas_sample'))
      ORDER BY predicted, actual;

ACTUAL     PREDICTED        COST
---------- ---------- ----------
0          0                 0
1          0                .75
0          1                .25
1          1                 0
```

You can remove the cost matrix with `REMOVE_COST_MATRIX`.

```
SQL>EXECUTE dbms_data_mining.remove_cost_matrix('nb_sh_clas_sample');

SQL>SELECT *
      FROM TABLE(dbms_data_mining.get_model_cost_matrix('nb_sh_clas_sample'))
      ORDER BY predicted, actual;

no rows selected
```

57.8.47 RENAME_MODEL Procedure

This procedure changes the name of the machine learning model indicated by *model_name* to the name that you specify as *new_model_name*.

If a model with *new_model_name* already exists, then the procedure optionally renames *new_model_name* to *versioned_model_name* before renaming *model_name* to *new_model_name*.

The model name is in the form [*schema_name*.]*model_name*. If you do not specify a schema, your own schema is used. For machine learning model naming restrictions, see the Usage Notes for "CREATE_MODEL Procedure".

Syntax

```
DBMS_DATA_MINING.RENAME_MODEL (
     model_name            IN VARCHAR2,
     new_model_name        IN VARCHAR2,
     versioned_model_name  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 57-122 RENAME_MODEL Procedure Parameters

Parameter	Description
model_name	Model to be renamed.
new_model_name	New name for the model *model_name*.
versioned_model_name	New name for the model *new_model_name* if it already exists.

Usage Notes

If you attempt to rename a model while it is being applied, then the model will be renamed but the apply operation will return indeterminate results.

Examples

1. This example changes the name of model census_model to census_model_2012.

```
BEGIN
  DBMS_DATA_MINING.RENAME_MODEL(
    model_name      => 'census_model',
    new_model_name  => 'census_model_2012');
END;
/
```

2. In this example, there are two classification models in the user's schema: clas_mod, the working model, and clas_mod_tst, a test model. The RENAME_MODEL procedure preserves clas_mod as clas_mod_old and makes the test model the new working model.

```
SELECT model_name FROM user_mining_models;
MODEL_NAME
------------------------------------------------------------------
CLAS_MOD
CLAS_MOD_TST

BEGIN
  DBMS_DATA_MINING.RENAME_MODEL(
```

```
    model_name           => 'clas_mod_tst',
    new_model_name       => 'clas_mod',
    versioned_model_name => 'clas_mod_old');
END;
/

SELECT model_name FROM user_mining_models;
MODEL_NAME
--------------------------------------------------------------------------------
CLAS_MOD
CLAS_MOD_OLD
```

58
DBMS_DATA_MINING_TRANSFORM

DBMS_DATA_MINING_TRANSFORM implements a set of transformations that are commonly used in machine learning.

This chapter contains the following topics:

- Overview
- Operational Notes
- Security Model
- Datatypes
- Constants
- Summary of DBMS_DATA_MINING_TRANSFORM Subprograms

> **See Also:**
>
> - DBMS_DATA_MINING
> - *Oracle Machine Learning for SQL User's Guide*

58.1 DBMS_DATA_MINING_TRANSFORM Overview

A transformation is a SQL expression that modifies the data in one or more columns.

Data must typically undergo certain transformations before it can be used to build a machine learning model. Many machine learning algorithms have specific transformation requirements.

Data that will be scored must be transformed in the same way as the data that was used to create (train) the model.

External or Embedded Transformations

DBMS_DATA_MINING_TRANSFORM offers two approaches to implementing transformations. For a given model, you can either:

- Create a list of transformation expressions and pass it to the CREATE_MODEL Procedure

  or

- Create a view that implements the transformations and pass the name of the view to the CREATE_MODEL Procedure

If you create a transformation list and pass it to CREATE_MODEL, the transformation expressions are embedded in the model and automatically implemented whenever the model is applied.

If you create a view, the transformation expressions are external to the model. You will need to re-create the transformations whenever you apply the model.

> **Note:**
>
> Embedded transformations significantly enhance the model's usability while simplifying the process of model management.

Automatic Transformations

Oracle Machine Learning for SQL supports an Automatic Data Preparation (ADP) mode. When ADP is enabled, most algorithm-specific transformations are *automatically* embedded. Any additional transformations must be explicitly provided in an embedded transformation list or in a view.

If ADP is enabled and you create a model with a transformation list, both sets of transformations are embedded. The model will execute the user-specified transformations from the transformation list before executing the automatic transformations specified by ADP.

Within a transformation list, you can selectively disable ADP for individual attributes.

> **See Also:**
>
> "Automatic Data Preparation"
>
> *Oracle Machine Learning for SQL User's Guide* for a more information about ADP
>
> "DBMS_DATA_MINING_TRANSFORM-About Transformation Lists"

Transformations in DBMS_DATA_MINING_TRANSFORM

The transformations supported by `DBMS_DATA_MINING_TRANSFORM` are summarized in this section.

Binning

Binning refers to the mapping of continuous or discrete values to discrete values of reduced cardinality.

- Supervised Binning (Categorical and Numerical)

  Binning is based on intrinsic relationships in the data as determined by a decision tree model.

  See "INSERT_BIN_SUPER Procedure".
- Top-N Frequency Categorical Binning

  Binning is based on the number of cases in each category.

  See "INSERT_BIN_CAT_FREQ Procedure"
- Equi-Width Numerical Binning

  Binning is based on equal-range partitions.

  See "INSERT_BIN_NUM_EQWIDTH Procedure".
- Quantile Numerical Binning

  Binning is based on quantiles computed using the SQL `NTILE` function.

See "INSERT_BIN_NUM_QTILE Procedure".

Linear Normalization

Normalization is the process of scaling continuous values down to a specific range, often between zero and one. Normalization transforms each numerical value by subtracting a number (the **shift**) and dividing the result by another number (the **scale**).

```
x_new = (x_old-shift)/scale
```

- Min-Max Normalization

  Normalization is based on the minimum and maximum with the following shift and scale:

  ```
  shift = min
  scale = max-min
  ```

  See "INSERT_NORM_LIN_MINMAX Procedure".

- Scale Normalization

  Normalization is based on the minimum and maximum with the following shift and scale:

  ```
  shift = 0
  scale = max{abs(max), abs(min)}
  ```

  See "INSERT_NORM_LIN_SCALE Procedure".

- Z-Score Normalization

  Normalization is based on the mean and standard deviation with the following shift and scale:

  ```
  shift = mean
  scale = standard_deviation
  ```

  See "INSERT_NORM_LIN_ZSCORE Procedure".

Outlier Treatment

An outlier is a numerical value that is located far from the rest of the data. Outliers can artificially skew the results of machine learning.

- Winsorizing

  Outliers are replaced with the nearest value that is not an outlier.

  See "INSERT_CLIP_WINSOR_TAIL Procedure"

- Trimming

  Outliers are set to `NULL`.

  See "INSERT_CLIP_TRIM_TAIL Procedure".

Missing Value Treatment

Missing data may indicate sparsity or it may indicate that some values are missing at random. `DBMS_DATA_MINING_TRANSFORM` supports the following transformations for minimizing the effects of missing values:

- Missing numerical values are replaced with the mean.

  See "INSERT_MISS_NUM_MEAN Procedure".

- Missing categorical values are replaced with the mode.

  See "INSERT_MISS_CAT_MODE Procedure".

> **Note:**
>
> Oracle Machine Learning for SQL also has default mechanisms for handling missing data. See *Oracle Machine Learning for SQL User's Guide* for details.

58.2 DBMS_DATA_MINING_TRANSFORM Operational Notes

The `DBMS_DATA_MINING_TRANSFORM` package offers a flexible framework for specifying data transformations. If you choose to embed transformations in the model (the preferred method), you create a **transformation list** object and pass it to the CREATE_MODEL Procedure. If you choose to transform the data without embedding, you create a view.

When specified in a transformation list, the transformation expressions are run by the model. When specified in a view, the transformation expressions are run by the view.

Transformation Definitions

Transformation definitions are used to generate the SQL expressions that transform the data. For example, the transformation definitions for normalizing a numeric column are the shift and scale values for that data.

With the `DBMS_DATA_MINING_TRANSFORM` package, you can call procedures to compute the transformation definitions, or you can compute them yourself, or you can do both.

Transformation Definition Tables

`DBMS_DATA_MINING_TRANSFORM` provides **`INSERT`** procedures that compute transformation definitions and insert them in transformation definition tables. You can modify the values in the transformation definition tables or populate them yourself.

`XFORM` routines use populated definition tables to transform data in external views. **`STACK`** routines use populated definition tables to build transformation lists.

To specify transformations based on definition tables, follow these steps:

1. Use **`CREATE`** routines to create transformation definition tables.

   The tables have columns to hold the transformation definitions for a given type of transformation. For example, the CREATE_BIN_NUM Procedure creates a definition table that has a column for storing data values and another column for storing the associated bin identifiers.

2. Use **`INSERT`** routines to compute and insert transformation definitions in the tables.

   Each `INSERT` routine uses a specific technique for computing the transformation definitions. For example, the INSERT_BIN_NUM_EQWIDTH Procedure computes bin boundaries by identifying the minimum and maximum values then setting the bin boundaries at equal intervals.

3. Use **`STACK`** or **`XFORM`** routines to generate transformation expressions based on the information in the definition tables:

   - Use **`STACK`** routines to add the transformation expressions to a transformation list. Pass the transformation list to the CREATE_MODEL Procedure. The transformation expressions will be assembled into one long SQL query and embedded in the model.

- Use **XFORM** routines to execute the transformation expressions within a view. The transformations will be external to the model and will need to be re-created whenever the model is applied to new data.

Transformations Without Definition Tables

STACK routines are not the only method for adding transformation expressions to a transformation list. You can also build a transformation list without using definition tables.

To specify transformations without using definition tables, follow these steps:

1. Write a SQL expression for transforming an attribute.
2. Write a SQL expression for reversing the transformation. (See "Reverse Transformations and Model Transparency" in "DBMS_DATA_MINING_TRANSFORM-About Transformation Lists".)
3. Determine whether or not to disable ADP for the attribute. By default ADP is enabled for the attribute if it is specified for the model. (See "Disabling Automatic Data Preparation" in "DBMS_DATA_MINING_TRANSFORM - About Transformation Lists".)
4. Specify the SQL expressions and ADP instructions in a call to the SET_TRANSFORM Procedure, which adds the information to a transformation list.
5. Repeat steps 1 through 4 for each attribute that you wish to transform.
6. Pass the transformation list to the CREATE_MODEL Procedure. The transformation expressions will be assembled into one long SQL query and embedded in the model.

> **Note:**
>
> SQL expressions that you specify with SET_TRANSFORM must fit within a VARCHAR2. To specify a longer expression, you can use the SET_EXPRESSION Procedure. With SET_EXPRESSION, you can build an expression by appending rows to a VARCHAR2 array.

About Stacking

Transformation lists are built by stacking transformation records. Transformation lists are evaluated from bottom to top. Each transformation expression depends on the result of the transformation expression below it in the stack.

Related Topics

- CREATE_MODEL Procedure
  This procedure creates an Oracle Machine Learning for SQL model with a given machine learning function.
- DBMS_DATA_MINING_TRANSFORM — About Transformation Lists
  The elements of a transformation list are **transformation records**. Each transformation record provides all the information needed by the model for managing the transformation of a single attribute.
- DBMS_DATA_MINING_TRANSFORM — About Stacking and Stack Procedures
  Transformation lists are built by stacking transformation records. Transformation lists are evaluated from bottom to top. Each transformation expression depends on the result of the transformation expression below it in the stack.

- DBMS_DATA_MINING_TRANSFORM — Nested Data Transformations
  The `CREATE` routines create transformation definition tables that include two columns, `col` and `att`, for identifying attributes.

58.2.1 DBMS_DATA_MINING_TRANSFORM — About Transformation Lists

The elements of a transformation list are **transformation records**. Each transformation record provides all the information needed by the model for managing the transformation of a single attribute.

Each transformation record includes the following fields:

- *attribute_name* — Name of the column of data to be transformed
- *attribute_subname* — Name of the nested attribute if *attribute_name* is a nested column, otherwise `NULL`
- *expression* — SQL expression for transforming the attribute
- *reverse_expression* — SQL expression for reversing the transformation
- *attribute_spec* — Identifies special treatment for the attribute during the model build. See Table 58-33 for details.

> **See Also:**
>
> - Table 58-1 for details about the `TRANSFORM_LIST` and `TRANSFORM_REC` object types
> - SET_TRANSFORM Procedure
> - CREATE_MODEL Procedure

Reverse Transformations and Model Transparency

An algorithm manipulates transformed attributes to train and score a model. The transformed attributes, however, may not be meaningful to an end user. For example, if attribute *x* has been transformed into bins 1 — 4, the bin names 1, 2 , 3, and 4 are manipulated by the algorithm, but a user is probably not interested in the model details about bins 1 — 4 or in predicting the numbers 1 — 4.

To return original attribute values in model details and predictions, you can provide a reverse expression in the transformation record for the attribute. For example, if you specify the transformation expression `'log(10, y)'` for attribute *y*, you could specify the reverse transformation expression `'power(10, y)'`.

Reverse transformations enable **model transparency**. They make internal processing transparent to the user.

> **Note:**
>
> `STACK` procedures automatically reverse normalization transformations, but they do not provide a mechanism for reversing binning, clipping, or missing value transformations.
>
> You can use the `DBMS_DATA_MINING.ALTER_REVERSE_EXPRESSION` procedure to specify or update reverse transformations expressions for an existing model.

> **See Also:**
>
> Table 58-1
>
> "ALTER_REVERSE_EXPRESSION Procedure"
>
> "Summary of DBMS_DATA_MINING Subprograms" for links to the model details functions

Disabling Automatic Data Preparation

ADP is controlled by a model-specific setting (`PREP_AUTO`). The `PREP_AUTO` setting affects all model attributes unless you disable it for individual attributes.

If ADP is enabled and you set *`attribute_spec`* to `NOPREP`, only the transformations that you specify for that attribute will be evaluated. If ADP is enabled and you do *not* set *`attribute_spec`* to `NOPREP`, the automatic transformations will be evaluated *after* the transformations that you specify for the attribute.

If ADP is not enabled for the model, the *`attribute_spec`* field of the transformation record is ignored.

> **See Also:**
>
> "Automatic Data Preparation" for information about the `PREP_AUTO` setting

Adding Transformation Records to a Transformation List

A transformation list is a stack of transformation records. When a new transformation record is added, it is appended to the top of the stack. (See "About Stacking" for details.)

When you use `SET_TRANSFORM` to add a transformation record to a transformation list, you can specify values for all the fields in the transformation record.

When you use `STACK` procedures to add transformation records to a transformation list, only the transformation expression field is populated. For normalization transformations, the reverse transformation expression field is also populated.

You can use both `STACK` procedures and `SET_TRANSFORM` to build one transformation list. Each `STACK` procedure call adds transformation records for all the attributes in a specified

transformation definition table. Each `SET_TRANSFORM` call adds a transformation record for a single attribute.

58.2.2 DBMS_DATA_MINING_TRANSFORM — About Stacking and Stack Procedures

Transformation lists are built by stacking transformation records. Transformation lists are evaluated from bottom to top. Each transformation expression depends on the result of the transformation expression below it in the stack.

Stack Procedures

`STACK` procedures create transformation records from the information in transformation definition tables. For example `STACK_BIN_NUM` builds a transformation record for each attribute specified in a definition table for numeric binning. `STACK` procedures stack the transformation records as follows:

- If an attribute is specified in the definition table but not in the transformation list, the `STACK` procedure creates a transformation record, computes the reverse transformation (if possible), inserts the transformation and reverse transformation in the transformation record, and appends the transformation record to the top of the transformation list.
- If an attribute is specified in the transformation list but not in the definition table, the `STACK` procedure takes no action.
- If an attribute is specified in the definition table *and* in the transformation list, the `STACK` procedure stacks the transformation expression from the definition table on top of the transformation expression in the transformation record and updates the reverse transformation. See Table 58-1and Example 58-4.

Example 58-1 Stacking a Clipping Transformation

This example shows how STACK_CLIP Procedure would add transformation records to a transformation list. Note that the clipping transformations are not reversed in `COL1` and `COL2` after stacking (as described in "Reverse Transformations and Model Transparency" in "DBMS_DATA_MINING_TRANSFORM-About Transformation Lists").

Refer to:

- CREATE_CLIP Procedure — Creates the definition table
- INSERT_CLIP_TRIM_TAIL Procedure — Inserts definitions in the table
- INSERT_CLIP_WINSOR_TAIL Procedure — Inserts definitions in the table
- Table 58-1 — Describes the structure of the transformation list (`TRANSFORM_LIST` object)

Assume a clipping definition table populated as follows.

col	att	lcut	lval	rcut	rval
COL1	null	-1.5	-1.5	4.5	4.5
COL2	null	0	0	1	1

Assume the following transformation list before stacking.

```
------------------------
transformation record #1:
------------------------
```

```
     attribute_name       = COL1
     attribute_subname    = null
     expression           = log(10, COL1)
     reverse_expression   = power(10, COL1)
-------------------------
transformation record #2:
-------------------------
     attribute_name       = COL3
     attribute_subname    = null
     expression           = ln(COL3)
     reverse_expression   = exp(COL3)
```

After stacking, the transformation list is as follows.

```
-------------------------
transformation record #1:
-------------------------
     attribute_name       = COL1
     attribute_subname    = null
     expression           = CASE WHEN log(10, COL1) < -1.5 THEN -1.5
                                 WHEN log(10, COL1) > 4.5  THEN 4.5
                                 ELSE log(10, COL1)
                            END;
     reverse_expression   = power(10, COL1)
-------------------------
transformation record #2:
-------------------------
     attribute_name       = COL3
     attribute_subname    = null
     expression           = ln(COL3)
     reverse_expression   = exp(COL3)
-------------------------
transformation record #3:
-------------------------
     attribute_name       = COL2
     attribute_subname    = null
     expression           = CASE WHEN COL2 < 0 THEN 0
                                 WHEN COL2 > 1 THEN 1
                                 ELSE COL2
                            END;
     reverse_expression   = null
```

58.2.3 DBMS_DATA_MINING_TRANSFORM — Nested Data Transformations

The `CREATE` routines create transformation definition tables that include two columns, `col` and `att`, for identifying attributes.

The column `col` holds the name of a column in the data table. If the data column is not nested, then `att` is null, and the name of the attribute is *`col`*. If the data column is nested, then `att` holds the name of the nested attribute, and the name of the attribute is *`col.att`*. The `INSERT` and `XFORM` routines ignore the `att` column in the definition tables. Neither the `INSERT` nor the `XFORM` routines support nested data.

Only the `STACK` procedures and `SET_TRANSFORM` support nested data. Nested data transformations are always embedded in the model.

Nested columns in Oracle Machine Learning for SQL can have the following types:

```
DM_NESTED_NUMERICALS
DM_NESTED_CATEGORICALS
DM_NESTED_BINARY_DOUBLES
DM_NESTED_BINARY_FLOATS
```

> **See Also:**
>
> "Constants"
>
> *Oracle Machine Learning for SQL User's Guide* for details about nested attributes in Oracle Machine Learning for SQL

Specifying Nested Attributes in a Transformation Record

A transformation record (`TRANSFORM_REC`) includes two fields, `attribute_name` and `attribute_subname`, for identifying the attribute. The field `attribute_name` holds the name of a column in the data table. If the data column is not nested, then `attribute_subname` is null, and the name of the attribute is *`attribute_name`*. If the data column is nested, then `attribute_subname` holds the name of the nested attribute, and the name of the attribute is *`attribute_name.attribute_subname`*.

Transforming Individual Nested Attributes

You can specify different transformations for different attributes in a nested column, and you can specify a default transformation for all the remaining attributes in the column. To specify a default nested transformation, specify null in the `attribute_name` field and the name of the nested column in the `attribute_subname` field as shown in Example 58-2. Note that the keyword `VALUE` is used to represent the value of a nested attribute in a transformation expression.

Example 58-2 Transforming a Nested Column

The following statement transforms two of the nested attributes in `COL_N1`. Attribute `ATTR1` is transformed with normalization; Attribute `ATTR2` is set to null, which causes attribute removal transformation (`ATTR2` is not used in training the model). All the remaining attributes in `COL_N1` are divided by 10.

```
DECLARE
  stk dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
  dbms_data_mining_transform.SET_TRANSFORM(
      stk,'COL_N1', 'ATTR1', '(VALUE - (-1.5))/20', 'VALUE *20 + (-1.5)');
  dbms_data_mining_transform.SET_TRANSFORM(
      stk,'COL_N1', 'ATTR2', NULL, NULL);
  dbms_data_mining_transform.SET_TRANSFORM(
      stk, NULL, 'COL_N1', 'VALUE/10', 'VALUE*10');
END;
/
```

The following SQL is generated from this statement.

```
CAST(MULTISET(SELECT DM_NESTED_NUMERICAL(
                                "ATTRIBUTE_NAME",
                                DECODE("ATTRIBUTE_NAME",
                                  'ATTR1', ("VALUE" - (-1.5))/20,
                                  "VALUE"/10))
                        FROM TABLE("COL_N1")
```

```
                    WHERE "ATTRIBUTE_NAME" IS NOT IN ('ATTR2'))
            AS DM_NESTED_NUMERICALS)
```

If transformations are not specified for `COL_N1.ATTR1` and `COL_N1.ATTR2`, then the default transformation is used for all the attributes in `COL_N1`, and the resulting SQL does not include a `DECODE`.

```
    CAST(MULTISET(SELECT DM_NESTED_NUMERICAL(
                            "ATTRIBUTE_NAME",
                            "VALUE"/10)
                    FROM TABLE("COL_N1"))
            AS DM_NESTED_NUMERICALS)
```

Since `DECODE` is limited to 256 arguments, multiple `DECODE` functions are nested to support an arbitrary number of individual nested attribute specifications.

Adding a Nested Column

You can specify a transformation that adds a nested column to the data, as shown in Example 58-3.

Example 58-3 Adding a Nested Column to a Transformation List

```
DECLARE
    v_xlst dbms_data_mining_transform.TRANSFORM_LIST;
  BEGIN
    dbms_data_mining_transform.SET_TRANSFORM(v_xlst,
      'YOB_CREDLIM', NULL,
      'dm_nested_numericals(
           dm_nested_numerical(
                ''CUST_YEAR_OF_BIRTH'', cust_year_of_birth),
           dm_nested_numerical(
                ''CUST_CREDIT_LIMIT'', cust_credit_limit))',
        NULL);
    dbms_data_mining_transform.SET_TRANSFORM(
              v_xlst, 'CUST_YEAR_OF_BIRTH', NULL, NULL, NULL);
    dbms_data_mining_transform.SET_TRANSFORM(
              v_xlst, 'CUST_CREDIT_LIMIT', NULL, NULL, NULL);
    dbms_data_mining_transform.XFORM_STACK(
              v_xlst, 'mining_data', 'mining_data_v');
END;
/

set long 2000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_V';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_POSTAL_CODE",dm_nested_numericals(
        dm_nested_numerical(
           'CUST_YEAR_OF_BIRTH', cust_year_of_birth),
        dm_nested_numerical(
           'CUST_CREDIT_LIMIT', cust_credit_limit)) "YOB_CREDLIM" FROM mining_data

SELECT * FROM mining_data_v WHERE cust_id = 104500;

CUST_ID CUST_POSTAL_CODE YOB_CREDLIM(ATTRIBUTE_NAME, VALUE)
------- ---------------- ------------------------------------------------------
 104500 68524             DM_NESTED_NUMERICALS(DM_NESTED_NUMERICAL(
                         'CUST_YEAR_OF_BIRTH', 1962),
                          DM_NESTED_NUMERICAL('CUST_CREDIT_LIMIT', 15000))
```

Stacking Nested Transformations

Example 58-4 shows how the STACK_NORM_LIN Procedure would add transformation records for nested column `COL_N` to a transformation list.

Refer to:

- CREATE_NORM_LIN Procedure — Creates the definition table
- INSERT_NORM_LIN_MINMAX Procedure — Inserts definitions in the table
- INSERT_NORM_LIN_SCALE Procedure — Inserts definitions in the table
- INSERT_NORM_LIN_ZSCORE Procedure — Inserts definitions in the table
- Table 58-1 — Describes the structure of the transformation list

Example 58-4 Stacking a Nested Normalization Transformation

Assume a linear normalization definition table populated as follows.

col	att	shift	scale
COL_N	ATT2	0	20
null	COL_N	0	10

Assume the following transformation list before stacking.

```
-------------------------
transformation record #1:
-------------------------
     attribute_name       = COL_N
     attribute_subname    = ATT1
     expression           = log(10, VALUE)
     reverse_expression   = power(10, VALUE)
-------------------------
transformation record #2:
-------------------------
     attribute_name       = null
     attribute_subname    = COL_N
     expression           = ln(VALUE)
     reverse_expression   = exp(VALUE)
```

After stacking, the transformation list is as follows.

```
-------------------------
transformation record #1:
-------------------------
     attribute_name       = COL_N
     attribute_subname    = ATT1
     expression           = (log(10, VALUE) - 0)/10
     reverse_expression   = power(10, VALUE*10 + 0)
-------------------------
transformation record #2:
-------------------------
     attribute_name       = NULL
     attribute_subname    = COL_N
     expression           = (ln(VALUE)- 0)/10
     reverse_expression   = exp(VALUE *10 + 0)
-------------------------
transformation record #3:
-------------------------
```

```
attribute_name      = COL_N
attribute_subname   = ATT2
expression          = (ln(VALUE) - 0)/20
reverse_expression  = exp(VALUE * 20 + 0)
```

58.3 DBMS_DATA_MINING_TRANSFORM Security Model

The DBMS_DATA_MINING_TRANSFORM package is owned by user SYS and is installed as part of database installation. Execution privilege on the package is granted to public. The routines in the package are run with invokers' rights (run with the privileges of the current user).

The DBMS_DATA_MINING_TRANSFORM.INSERT_* procedures have a *data_table_name* parameter that enables the user to provide the input data for transformation purposes. The value of *data_table_name* can be the name of a physical table or a view. The *data_table_name* parameter can also accept an inline query.

> **Note:**
>
> Because an inline query can be used to specify the data for transformation, Oracle strongly recommends that the calling routine perform any necessary SQL injection checks on the input string.

> **See Also:**
>
> "Operational Notes" for a description of the DBMS_DATA_MINING_TRANSFORM.INSERT_* procedures

58.4 DBMS_DATA_MINING_TRANSFORM Datatypes

DBMS_DATA_MINING_TRANSFORM defines the datatypes described in the following table.

Table 58-1 Datatypes in DBMS_DATA_MINING_TRANSFORM

List Type	List Elements	Description
COLUMN_LIST	VARRAY(1000) OF varchar2(32)	COLUMN_LIST stores quoted and non-quoted identifiers for column names. COLUMN_LIST is the datatype of the *exclude_list* parameter in the INSERT procedures. See "INSERT_AUTOBIN_NUM_EQWIDTH Procedure" for an example. See *Oracle Database PL/SQL Language Reference* for information about populating VARRAY structures.

Table 58-1 (Cont.) Datatypes in DBMS_DATA_MINING_TRANSFORM

List Type	List Elements	Description
DESCRIBE_LIST	DBMS_SQL.DESC_TAB2 TYPE desc_tab2 IS TABLE OF desc_rec2 INDEX BY BINARY_INTEGER TYPE desc_rec2 IS RECORD (col_type BINARY_INTEGER := 0, col_max_len BINARY_INTEGER := 0, col_name VARCHAR2(32767):= '', col_name_len BINARY_INTEGER := 0, col_schema_name VARCHAR2(32) := '', col_schema_name_len BINARY_INTEGER := 0, col_precision BINARY_INTEGER := 0, col_scale BINARY_INTEGER := 0, col_charsetid BINARY_INTEGER := 0, col_charsetform BINARY_INTEGER := 0, col_null_ok BOOLEAN := TRUE);	DESCRIBE_LIST describes the columns of the data table after the transformation list has been applied. A DESCRIBE_LIST is returned by the DESCRIBE_STACK Procedure. The DESC_TAB2 and DESC_REC2 types are defined in the DBMS_SQL package. See "DESC_REC2 Record Type". The col_type field of DESC_REC2 identifies the datatype of the column. The datatype is expressed as a numeric constant that represents a built-in datatype. For example, a 1 indicates a variable length character string. The codes for Oracle built-in datatypes are listed in *Oracle Database SQL Language Reference*. The codes for the Oracle Machine Learning for SQL nested types are described in "Constants". The col_name field of DESC_REC2 identifies the column name. It may be populated with a column name, an alias, or an expression. If the column name is a SELECT expression, it may be very long. If the expression is longer than 30 bytes, it cannot be used in a view unless it is given an alias.
TRANSFORM_LIST	**TABLE OF transform_rec** **TYPE transform_rec IS RECORD** (attribute_name VARCHAR2(30), attribute_subname VARCHAR2(4000), expression **EXPRESSION_REC**, reverse_expression **EXPRESSION_REC**, attribute_spec VARCHAR2(4000)); **TYPE expression_rec IS RECORD** (lstmt DBMS_SQL.VARCHAR2A, lb BINARY_INTEGER DEFAULT 1, ub BINARY_INTEGER DEFAULT 0); TYPE varchar2a IS TABLE OF VARCHAR2(32767) INDEX BY BINARY_INTEGER;	TRANSFORM_LIST is a list of transformations that can be embedded in a model. A TRANSFORM_LIST is accepted as an argument by the CREATE_MODEL Procedure. Each element in a TRANSFORM_LIST is a TRANSFORM_REC that specifies how to transform a single attribute. The attribute_name is a column name. The attribute_subname is the nested attribute name if the column is nested, otherwise attribute_subname is null. The expression field holds a SQL expression for transforming the attribute. See "About Transformation Lists" for an explanation of reverse expressions. The attribute_spec field can be used to cause the attribute to be handled in a specific way during the model build. See Table 58-33 for details. The expressions in a TRANSFORM_REC have type EXPRESSION_REC. The lstmt field stores a VARCHAR2A, which is a table of VARCHAR2(32767). The VARCHAR2A datatype allows transformation expressions to be very long, as they can be broken up across multiple rows of VARCHAR2. The VARCHAR2A type is defined in the DBMS_SQL package. See "VARCHAR2A Table Type". The ub (upper bound) and lb (lower bound) fields indicate how many rows there are in the VARCHAR2A table. If ub < lb (default) the EXPRESSION_REC is empty; if lb=ub=1 there is one row; if lb=1 and ub=2 there are 2 rows, and so on.

58.5 DBMS_DATA_MINING_TRANSFORM Constants

DBMS_DATA_MINING_TRANSFORM defines the constants described in the following table.

Table 58-2 Constants in DBMS_DATA_MINING_TRANSFORM

Constant	Value	Description
NEST_NUM_COL_TYPE	100001	Indicates that an attribute in the transformation list comes from a row in a column of DM_NESTED_NUMERICALS. Nested numerical attributes are defined as follows: attribute_name VARCHAR2(4000) value NUMBER
NEST_CAT_COL_TYPE	100002	Indicates that an attribute in the transformation list comes from a row in a column of DM_NESTED_CATEGORICALS. Nested categorical attributes are defined as follows: attribute_name VARCHAR2(4000) value VARCHAR2(4000)
NEST_BD_COL_TYPE	100003	Indicates that an attribute in the transformation list comes from a row in a column of DM_NESTED_BINARY_DOUBLES. Nested binary double attributes are defined as follows: attribute_name VARCHAR2(4000) value BINARY_DOUBLE
NEST_BF_COL_TYPE	100004	Indicates that an attribute in the transformation list comes from a row in a column of DM_NESTED_BINARY_FLOATS. attribute_name VARCHAR2(4000) value BINARY_FLOAT

See Also:

Oracle Machine Learning for SQL User's Guide for information about nested data in Oracle Machine Learning for SQL

58.6 Summary of DBMS_DATA_MINING_TRANSFORM Subprograms

This table lists the DBMS_DATA_MINING_TRANSFORM subprograms in alphabetical order and briefly describes them.

Table 58-3 DBMS_DATA_MINING_TRANSFORM Package Subprograms

Subprogram	Purpose
CREATE_BIN_CAT Procedure	Creates a transformation definition table for categorical binning
CREATE_BIN_NUM Procedure	Creates a transformation definition table for numerical binning
CREATE_CLIP Procedure	Creates a transformation definition table for clipping
CREATE_COL_REM Procedure	Creates a transformation definition table for column removal

Table 58-3 (Cont.) DBMS_DATA_MINING_TRANSFORM Package Subprograms

Subprogram	Purpose
CREATE_MISS_CAT Procedure	Creates a transformation definition table for categorical missing value treatment
CREATE_MISS_NUM Procedure	Creates a transformation definition table for numerical missing values treatment
CREATE_NORM_LIN Procedure	Creates a transformation definition table for linear normalization
DESCRIBE_STACK Procedure	Describes the transformation list
GET_EXPRESSION Function	Returns a `VARCHAR2` chunk from a transformation expression
INSERT_AUTOBIN_NUM_EQWIDTH Procedure	Inserts numeric automatic equi-width binning definitions in a transformation definition table
INSERT_BIN_CAT_FREQ Procedure	Inserts categorical frequency-based binning definitions in a transformation definition table
INSERT_BIN_NUM_EQWIDTH Procedure	Inserts numeric equi-width binning definitions in a transformation definition table
INSERT_BIN_NUM_QTILE Procedure	Inserts numeric quantile binning expressions in a transformation definition table
INSERT_BIN_SUPER Procedure	Inserts supervised binning definitions in numerical and categorical transformation definition tables
INSERT_CLIP_TRIM_TAIL Procedure	Inserts numerical trimming definitions in a transformation definition table
INSERT_CLIP_WINSOR_TAIL Procedure	Inserts numerical winsorizing definitions in a transformation definition table
INSERT_MISS_CAT_MODE Procedure	Inserts categorical missing value treatment definitions in a transformation definition table
INSERT_MISS_NUM_MEAN Procedure	Inserts numerical missing value treatment definitions in a transformation definition table
INSERT_NORM_LIN_MINMAX Procedure	Inserts linear min-max normalization definitions in a transformation definition table
INSERT_NORM_LIN_SCALE Procedure	Inserts linear scale normalization definitions in a transformation definition table
INSERT_NORM_LIN_ZSCORE Procedure	Inserts linear zscore normalization definitions in a transformation definition table
SET_EXPRESSION Procedure	Adds a `VARCHAR2` chunk to an expression
SET_TRANSFORM Procedure	Adds a transformation record to a transformation list
STACK_BIN_CAT Procedure	Adds a categorical binning expression to a transformation list
STACK_BIN_NUM Procedure	Adds a numerical binning expression to a transformation list
STACK_CLIP Procedure	Adds a clipping expression to a transformation list
STACK_COL_REM Procedure	Adds a column removal expression to a transformation list
STACK_MISS_CAT Procedure	Adds a categorical missing value treatment expression to a transformation list
STACK_MISS_NUM Procedure	Adds a numerical missing value treatment expression to a transformation list
STACK_NORM_LIN Procedure	Adds a linear normalization expression to a transformation list
XFORM_BIN_CAT Procedure	Creates a view of the data table with categorical binning transformations

Table 58-3 (Cont.) DBMS_DATA_MINING_TRANSFORM Package Subprograms

Subprogram	Purpose
XFORM_BIN_NUM Procedure	Creates a view of the data table with numerical binning transformations
XFORM_CLIP Procedure	Creates a view of the data table with clipping transformations
XFORM_COL_REM Procedure	Creates a view of the data table with column removal transformations
XFORM_EXPR_NUM Procedure	Creates a view of the data table with the specified numeric transformations
XFORM_EXPR_STR Procedure	Creates a view of the data table with the specified categorical transformations
XFORM_MISS_CAT Procedure	Creates a view of the data table with categorical missing value treatment
XFORM_MISS_NUM Procedure	Creates a view of the data table with numerical missing value treatment
XFORM_NORM_LIN Procedure	Creates a view of the data table with linear normalization transformations
XFORM_STACK Procedure	Creates a view of the transformation list

58.6.1 CREATE_BIN_CAT Procedure

This procedure creates a transformation definition table for categorical binning.

The columns are described in the following table.

Table 58-4 Columns in a Transformation Definition Table for Categorical Binning

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of categorical data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is a nested column. If *col* is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.
val	VARCHAR2(4000)	Values of the attribute
bin	VARCHAR2(4000)	Bin assignments for the values

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_BIN_CAT (
     bin_table_name     IN VARCHAR2,
     bin_schema_name    IN VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-5 CREATE_BIN_CAT Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table to be created
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about categorical data.
2. See "Nested Data Transformations" for information about transformation definition tables and nested data.
3. You can use the following procedures to populate the transformation definition table:
   - INSERT_BIN_CAT_FREQ Procedure — frequency-based binning
   - INSERT_BIN_SUPER Procedure — supervised binning

> **See Also:**
>
> "Binning" in DBMS_DATA_MINING_TRANSFORM Overview
>
> "Operational Notes"

Examples

The following statement creates a table called `bin_cat_xtbl` in the current schema. The table has columns that can be populated with bin assignments for categorical attributes.

```
BEGIN
   DBMS_DATA_MINING_TRANSFORM.CREATE_BIN_CAT('bin_cat_xtbl');
END;
/
DESCRIBE bin_cat_xtbl
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                                VARCHAR2(30)
 ATT                                                VARCHAR2(4000)
 VAL                                                VARCHAR2(4000)
 BIN                                                VARCHAR2(4000)
```

58.6.2 CREATE_BIN_NUM Procedure

This procedure creates a transformation definition table for numerical binning.

The columns are described in the following table.

Table 58-6 Columns in a Transformation Definition Table for Numerical Binning

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of numerical data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is a nested column. If *col* is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.
val	NUMBER	Values of the attribute
bin	VARCHAR2(4000)	Bin assignments for the values

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_BIN_NUM (
     bin_table_name    IN VARCHAR2,
     bin_schema_name   IN VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-7 CREATE_BIN_NUM Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table to be created
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. See "Nested Data Transformations" for information about transformation definition tables and nested data.
3. You can use the following procedures to populate the transformation definition table:
   - INSERT_AUTOBIN_NUM_EQWIDTH Procedure — automatic equi-width binning
   - INSERT_BIN_NUM_EQWIDTH Procedure — user-specified equi-width binning
   - INSERT_BIN_NUM_QTILE Procedure — quantile binning
   - INSERT_BIN_SUPER Procedure — supervised binning

> **See Also:**
>
> "Binning" in DBMS_DATA_MINING_TRANSFORM Overview
>
> "Operational Notes"

Examples

The following statement creates a table called `bin_num_xtbl` in the current schema. The table has columns that can be populated with bin assignments for numerical attributes.

```
BEGIN
  DBMS_DATA_MINING_TRANSFORM.CREATE_BIN_NUM('bin_num_xtbl');
END;
/

DESCRIBE bin_num_xtbl
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                                VARCHAR2(30)
 ATT                                                VARCHAR2(4000)
 VAL                                                NUMBER
 BIN                                                VARCHAR2(4000)
```

58.6.3 CREATE_CLIP Procedure

This procedure creates a transformation definition table for clipping or winsorizing to minimize the effect of outliers.

The columns are described in the following table.

Table 58-8 Columns in a Transformation Definition Table for Clipping or Winsorizing

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of numerical data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is a nested column of DM_NESTED_NUMERICALS. If *col* is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.
lcut	NUMBER	The lowest typical value for the attribute. If the attribute values were plotted on an *xy* axis, *lcut* would be the left-most boundary of the range of values considered typical for this attribute. Any values to the left of *lcut* are outliers.
lval	NUMBER	Value assigned to an outlier to the left of *lcut*
rcut	NUMBER	The highest typical value for the attribute If the attribute values were plotted on an *xy* axis, *rcut* would be the right-most boundary of the range of values considered typical for this attribute. Any values to the right of *rcut* are outliers.
rval	NUMBER	Value assigned to an outlier to the right of *rcut*

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_CLIP (
     clip_table_name    IN VARCHAR2,
     clip_schema_name   IN VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-9 CREATE_CLIP Procedure Parameters

Parameter	Description
clip_table_name	Name of the transformation definition table to be created
clip_schema_name	Schema of *clip_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. See "Nested Data Transformations" for information about transformation definition tables and nested data.
3. You can use the following procedures to populate the transformation definition table:
   - INSERT_CLIP_TRIM_TAIL Procedure — replaces outliers with nulls
   - INSERT_CLIP_WINSOR_TAIL Procedure — replaces outliers with an average value

> **See Also:**
>
> "Outlier Treatment" in DBMS_DATA_MINING_TRANSFORM Overview
>
> "Operational Notes"

Examples

The following statement creates a table called clip_xtbl in the current schema. The table has columns that can be populated with clipping instructions for numerical attributes.

```
BEGIN
  DBMS_DATA_MINING_TRANSFORM.CREATE_CLIP('clip_xtbl');
END;
/

DESCRIBE clip_xtbl
 Name                                    Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                               VARCHAR2(30)
 ATT                                               VARCHAR2(4000)
 LCUT                                              NUMBER
 LVAL                                              NUMBER
 RCUT                                              NUMBER
 RVAL                                              NUMBER
```

58.6.4 CREATE_COL_REM Procedure

This procedure creates a transformation definition table for removing columns from the data table.

The columns are described in the following table.

Table 58-10 Columns in a Transformation Definition Table for Column Removal

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is nested (DM_NESTED_NUMERICALS or DM_NESTED_CATEGORICALS). If col is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_COL_REM (
      rem_table_name           VARCHAR2,
      rem_schema_name          VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-11 CREATE_COL_REM Procedure Parameters

Parameter	Description
rem_table_name	Name of the transformation definition table to be created
rem_schema_name	Schema of *rem_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See "Nested Data Transformations" for information about transformation definition tables and nested data.
2. See "Operational Notes".

Examples

The following statement creates a table called rem_att_xtbl in the current schema. The table has columns that can be populated with the names of attributes to exclude from the data to be mined.

```
BEGIN
    DBMS_DATA_MINING_TRANSFORM.CREATE_COL_REM ('rem_att_xtbl');
END;
 /
DESCRIBE rem_att_xtbl
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                                VARCHAR2(30)
 ATT                                                VARCHAR2(4000)
```

58.6.5 CREATE_MISS_CAT Procedure

This procedure creates a transformation definition table for replacing categorical missing values.

The columns are described in the following table.

Table 58-12 Columns in a Transformation Definition Table for Categorical Missing Value Treatment

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of categorical data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is a nested column of DM_NESTED_CATEGORICALS. If *col* is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.
val	VARCHAR2(4000)	Replacement for missing values in the attribute

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_MISS_CAT (
     miss_table_name       IN VARCHAR2,
     miss_schema_name      IN VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-13 CREATE_MISS_CAT Procedure Parameters

Parameter	Description
miss_table_name	Name of the transformation definition table to be created
miss_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about categorical data.
2. See "Nested Data Transformations" for information about transformation definition tables and nested data.
3. You can use the INSERT_MISS_CAT_MODE Procedure to populate the transformation definition table.

> **See Also:**
>
> "Missing Value Treatment" in DBMS_DATA_MINING_TRANSFORM Overview
>
> "Operational Notes"

Examples

The following statement creates a table called `miss_cat_xtbl` in the current schema. The table has columns that can be populated with values for missing data in categorical attributes.

```
BEGIN

  DBMS_DATA_MINING_TRANSFORM.CREATE_MISS_CAT('miss_cat_xtbl');
END;
/

DESCRIBE miss_cat_xtbl
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                                VARCHAR2(30)
 ATT                                                VARCHAR2(4000)
 VAL                                                VARCHAR2(4000)
```

58.6.6 CREATE_MISS_NUM Procedure

This procedure creates a transformation definition table for replacing numerical missing values.

The columns are described in Table 58-14.

Table 58-14 Columns in a Transformation Definition Table for Numerical Missing Value Treatment

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of numerical data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is a nested column of DM_NESTED_NUMERICALS. If *col* is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.
val	NUMBER	Replacement for missing values in the attribute

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_MISS_NUM (
     miss_table_name       IN VARCHAR2,
     miss_schema_name      IN VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-15 CREATE_MISS_NUM Procedure Parameters

Parameter	Description
miss_table_name	Name of the transformation definition table to be created
miss_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. See "Nested Data Transformations" for information about transformation definition tables and nested data.
3. You can use the INSERT_MISS_NUM_MEAN Procedure to populate the transformation definition table.

> **See Also:**
>
> "Missing Value Treatment" in DBMS_DATA_MINING_TRANSFORM Overview
>
> "Operational Notes"

Example

The following statement creates a table called `miss_num_xtbl` in the current schema. The table has columns that can be populated with values for missing data in numerical attributes.

```
BEGIN
    DBMS_DATA_MINING_TRANSFORM.CREATE_MISS_NUM('miss_num_xtbl');
END;
/

DESCRIBE miss_num_xtbl
 Name                                    Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                               VARCHAR2(30)
 ATT                                               VARCHAR2(4000)
 VAL                                               NUMBER
```

58.6.7 CREATE_NORM_LIN Procedure

This procedure creates a transformation definition table for linear normalization.

The columns are described in Table 58-16.

Table 58-16 Columns in a Transformation Definition Table for Linear Normalization

Name	Datatype	Description
col	VARCHAR2(30)	Name of a column of numerical data. If the column is not nested, the column name is also the attribute name. For information about attribute names, see *Oracle Machine Learning for SQL User's Guide*.
att	VARCHAR2(4000)	The attribute subname if *col* is a nested column of DM_NESTED_NUMERICALS. If *col* is nested, the attribute name is *col.att*. If *col* is not nested, *att* is null.
shift	NUMBER	A constant to subtract from the attribute values
scale	NUMBER	A constant by which to divide the shifted values

Syntax

```
DBMS_DATA_MINING_TRANSFORM.CREATE_NORM_LIN (
     norm_table_name       IN VARCHAR2,
     norm_schema_name      IN VARCHAR2 DEFAULT NULL );
```

Parameters

Table 58-17 CREATE_NORM_LIN Procedure Parameters

Parameter	Description
norm_table_name	Name of the transformation definition table to be created
norm_schema_name	Schema of *norm_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. See "Nested Data Transformations" for information about transformation definition tables and nested data.
3. You can use the following procedures to populate the transformation definition table:
   - INSERT_NORM_LIN_MINMAX Procedure — Uses linear min-max normalization
   - INSERT_NORM_LIN_SCALE Procedure — Uses linear scale normalization
   - INSERT_NORM_LIN_ZSCORE Procedure — Uses linear zscore normalization

> **See Also:**
>
> "Linear Normalization" in DBMS_DATA_MINING_TRANSFORM Overview
>
> "Operational Notes"

Examples

The following statement creates a table called norm_xtbl in the current schema. The table has columns that can be populated with shift and scale values for normalizing numerical attributes.

```
BEGIN
     DBMS_DATA_MINING_TRANSFORM.CREATE_NORM_LIN('norm_xtbl');
END;
/

DESCRIBE norm_xtbl
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 COL                                                VARCHAR2(30)
 ATT                                                VARCHAR2(4000)
 SHIFT                                              NUMBER
 SCALE                                              NUMBER
```

58.6.8 DESCRIBE_STACK Procedure

This procedure describes the columns of the data table after a list of transformations has been applied.

Only the columns that are specified in the transformation list are transformed. The remaining columns in the data table are included in the output without changes.

To create a view of the data table after the transformations have been applied, use the XFORM_STACK Procedure.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.DESCRIBE_STACK (
     xform_list           IN  TRANSFORM_LIST,
     data_table_name      IN  VARCHAR2,
     describe_list        OUT DESCRIBE_LIST,
     data_schema_name     IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-18 DESCRIBE_STACK Procedure Parameters

Parameter	Description
xform_list	A list of transformations. See Table 58-1 for a description of the TRANSFORM_LIST object type.
data_table_name	Name of the table containing the data to be transformed
describe_list	Descriptions of the columns in the data table after the transformations specified in *xform_list* have been applied. See Table 58-1 for a description of the DESCRIBE_LIST object type.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes" for information about transformation lists and embedded transformations.

Examples

This example shows the column name and datatype, the column name length, and the column maximum length for the view oml_user.cust_info after the transformation list has been applied. All the transformations are user-specified. The results of DESCRIBE_STACK do not include one of the columns in the original table, because the SET_TRANSFORM procedure sets that column to NULL.

```
CREATE OR REPLACE VIEW cust_info AS
         SELECT a.cust_id, c.country_id, c.cust_year_of_birth,
         CAST(COLLECT(DM_Nested_Numerical(
                  b.prod_name, 1))
                AS DM_Nested_Numericals) custprods
                 FROM sh.sales a, sh.products b, sh.customers c
                  WHERE a.prod_id = b.prod_id AND
                        a.cust_id=c.cust_id and
                        a.cust_id between 100001 AND 105000
         GROUP BY a.cust_id, country_id, cust_year_of_birth;
```

```
describe cust_info
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                   NOT NULL NUMBER
 COUNTRY_ID                                NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                        NOT NULL NUMBER(4)
 CUSTPRODS                                          SYS.DM_NESTED_NUMERICALS

DECLARE
  cust_stack   dbms_data_mining_transform.TRANSFORM_LIST;
  cust_cols    dbms_data_mining_transform.DESCRIBE_LIST;
BEGIN
  dbms_data_mining_transform.SET_TRANSFORM (cust_stack,
     'country_id', NULL, 'country_id/10', 'country_id*10');
  dbms_data_mining_transform.SET_TRANSFORM (cust_stack,
      'cust_year_of_birth', NULL, NULL, NULL);
  dbms_data_mining_transform.SET_TRANSFORM (cust_stack,
      'custprods', 'Mouse Pad', 'value*100', 'value/100');
  dbms_data_mining_transform.DESCRIBE_STACK(
       xform_list => cust_stack,
       data_table_name => 'cust_info',
       describe_list => cust_cols);
  dbms_output.put_line('====');
  for i in 1..cust_cols.COUNT loop
    dbms_output.put_line('COLUMN_NAME:     '||cust_cols(i).col_name);
    dbms_output.put_line('COLUMN_TYPE:     '||cust_cols(i).col_type);
    dbms_output.put_line('COLUMN_NAME_LEN: '||cust_cols(i).col_name_len);
    dbms_output.put_line('COLUMN_MAX_LEN:  '||cust_cols(i).col_max_len);
    dbms_output.put_line('====');
  END loop;
END;
/
====
COLUMN_NAME:     CUST_ID
COLUMN_TYPE:     2
COLUMN_NAME_LEN: 7
COLUMN_MAX_LEN:  22
====
COLUMN_NAME:     COUNTRY_ID
COLUMN_TYPE:     2
COLUMN_NAME_LEN: 10
COLUMN_MAX_LEN:  22
====
COLUMN_NAME:     CUSTPRODS
COLUMN_TYPE:     100001
COLUMN_NAME_LEN: 9
COLUMN_MAX_LEN:  40
====
```

58.6.9 GET_EXPRESSION Function

This function returns a row from a `VARCHAR2` array that stores a transformation expression. The array is built by calls to the SET_EXPRESSION Procedure.

The array can be used for specifying SQL expressions that are too long to be used with the SET_TRANSFORM Procedure.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.GET_EXPRESSION (
     expression           IN EXPRESSION_REC,
     chunk_num            IN PLS_INTEGER DEFAULT NULL);
 RETURN VARCHAR2;
```

Parameters

Table 58-19 GET_EXPRESSION Function Parameters

Parameter	Description
expression	An expression record (EXPRESSION_REC) that specifies a transformation expression or a reverse transformation expression for an attribute. Each expression record includes a VARCHAR2 array and index fields for specifying upper and lower boundaries within the array. There are two EXPRESSION_REC fields within a transformation record (TRANSFORM_REC): one for the transformation expression; the other for the reverse transformation expression. See Table 58-1 for a description of the EXPRESSION_REC type.
chunk	A VARCHAR2 chunk (row) to be appended to *expression*.

Usage Notes

1. Chunk numbering starts with one. For chunks outside of the range, the return value is null. When a chunk number is null the whole expression is returned as a string. If the expression is too big, a VALUE_ERROR is raised.
2. See "About Transformation Lists".
3. See "Operational Notes".

Examples

See the example for the SET_EXPRESSION Procedure.

Related Topics

- SET_EXPRESSION Procedure
  This procedure appends a row to a VARCHAR2 array that stores a SQL expression.
- SET_TRANSFORM Procedure
  This procedure appends the transformation instructions for an attribute to a transformation list.

58.6.10 INSERT_AUTOBIN_NUM_EQWIDTH Procedure

This procedure performs numerical binning and inserts the transformation definitions in a transformation definition table. The procedure identifies the minimum and maximum values and computes the bin boundaries at equal intervals.

INSERT_AUTOBIN_NUM_EQWIDTH computes the number of bins separately for each column. If you want to use equi-width binning with the same number of bins for each column, use the INSERT_BIN_NUM_EQWIDTH Procedure.

INSERT_AUTOBIN_NUM_EQWIDTH bins all the NUMBER and FLOAT columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_AUTOBIN_NUM_EQWIDTH (
     bin_table_name         IN VARCHAR2,
     data_table_name        IN VARCHAR2,
     bin_num                IN PLS_INTEGER DEFAULT 3,
     max_bin_num            IN PLS_INTEGER DEFAULT 100,
     exclude_list           IN COLUMN_LIST DEFAULT NULL,
     round_num              IN PLS_INTEGER DEFAULT 6,
     sample_size            IN PLS_INTEGER DEFAULT 50000,
     bin_schema_name        IN VARCHAR2 DEFAULT NULL,
     data_schema_name       IN VARCHAR2 DEFAULT NULL,
     rem_table_name         IN VARCHAR2 DEFAULT NULL,
     rem_schema_name        IN VARCHAR2 DEFAULT NULL));
```

Parameters

Table 58-20 INSERT_AUTOBIN_NUM_EQWIDTH Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table for numerical binning. You can use the CREATE_BIN_NUM Procedure to create the definition table. The following columns are required: COL VARCHAR2(30) VAL NUMBER BIN VARCHAR2(4000) CREATE_BIN_NUM creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_AUTOBIN_NUM_EQWIDTH.
data_table_name	Name of the table containing the data to be transformed
bin_num	Minimum number of bins. If *bin_num* is 0 or NULL, it is ignored. The default value of *bin_num* is 3.
max_bin_num	Maximum number of bins. If *max_bin_num* is 0 or NULL, it is ignored. The default value of *max_bin_num* is 100.
exclude_list	List of numerical columns to be excluded from the binning process. If you do not specify *exclude_list*, all numerical columns in the data source are binned. The format of *exclude_list* is: dbms_data_mining_transform.COLUMN_LIST('*col1*','*col2*', ...'*coln*')
round_num	Specifies how to round the number in the VAL column of the transformation definition table. When *round_num* is positive, it specifies the most significant digits to retain. When *round_num* is negative, it specifies the least significant digits to remove. In both cases, the result is rounded to the specified number of digits. See the Usage Notes for an example. The default value of *round_num* is 6.

Table 58-20 (Cont.) INSERT_AUTOBIN_NUM_EQWIDTH Procedure Parameters

Parameter	Description
sample_size	Size of the data sample. If *sample_size* is less than the total number of non-NULL values in the column, then *sample_size* is used instead of the SQL COUNT function in computing the number of bins. If *sample_size* is 0 or NULL, it is ignored. See the Usage Notes. The default value of *sample_size* is 50,000.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
rem_table_name	Name of a transformation definition table for column removal. The table must have the columns described in "CREATE_COL_REM Procedure". INSERT_AUTOBIN_NUM_EQWIDTH ignores columns with all nulls or only one unique value. If you specify a value for *rem_table_name*, these columns are removed from the mining data. If you do not specify a value for *rem_table_name*, these unbinned columns remain in the data.
rem_schema_name	Schema of *rem_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. INSERT_AUTOBIN_NUM_EQWIDTH computes the number of bins for a column based on the number of non-null values (COUNT), the maximum (MAX), the minimum (MIN), the standard deviation (STDDEV), and the constant C=3.49/0.9:

   ```
   N=floor(power(COUNT,1/3)*(max-min)/(c*dev))
   ```

   If the *sample_size* parameter is specified, it is used instead of COUNT.

   See *Oracle Machine Learning for SQL User's Guide* for information about the COUNT, MAX, MIN, STDDEV, FLOOR, and POWER functions.
3. INSERT_AUTOBIN_NUM_EQWIDTH uses absolute values to compute the number of bins. The sign of the parameters *bin_num*, *max_bin_num*, and *sample_size* has no effect on the result.
4. In computing the number of bins, INSERT_AUTOBIN_NUM_EQWIDTH evaluates the following criteria in the following order:
   a. The minimum number of bins (*bin_num*)
   b. The maximum number of bins (*max_bin_num*)
   c. The maximum number of bins for integer columns, calculated as the number of distinct values in the range *max-min+1*.
5. The *round_num* parameter controls the rounding of column values in the transformation definition table, as follows:

   ```
   For a value of 308.162:
   when round_num =  1       result is 300
   when round_num =  2       result is 310
   when round_num =  3       result is 308
   when round_num =  0       result is 308.162
   ```

```
    when round_num = -1       result is 308.16
    when round_num = -2       result is 308.2
```

Examples

In this example, `INSERT_AUTOBIN_NUM_EQWIDTH` computes the bin boundaries for the `cust_year_of_birth` column in `sh.customers` and inserts the transformations in a transformation definition table. The STACK_BIN_NUM Procedure creates a transformation list from the contents of the definition table. The CREATE_MODEL Procedure embeds the transformation list in a new model called `nb_model`.

The transformation and reverse transformation expressions embedded in `nb_model` are returned by the GET_MODEL_TRANSFORMATIONS Function.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, cust_year_of_birth, cust_postal_code
       FROM sh.customers;

DESCRIBE mining_data
 Name                           Null?    Type
 ------------------------------ -------- ----------------------------
 CUST_ID                        NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH             NOT NULL NUMBER(4)
 CUST_POSTAL_CODE               NOT NULL VARCHAR2(10)

BEGIN
  dbms_data_mining_transform.CREATE_BIN_NUM(
     bin_table_name   => 'bin_tbl');
  dbms_data_mining_transform.INSERT_AUTOBIN_NUM_EQWIDTH (
     bin_table_name   => 'bin_tbl',
     data_table_name  => 'mining_data',
     bin_num          => 3,
     max_bin_num      => 5,
     exclude_list     => dbms_data_mining_transform.COLUMN_LIST('cust_id'));
END;
/

set numwidth 4
column val off
SELECT col, val, bin FROM bin_tbl
      ORDER BY val ASC;

COL                         VAL BIN
------------------------- ---- -----
CUST_YEAR_OF_BIRTH        1913
CUST_YEAR_OF_BIRTH        1928 1
CUST_YEAR_OF_BIRTH        1944 2
CUST_YEAR_OF_BIRTH        1959 3
CUST_YEAR_OF_BIRTH        1975 4
CUST_YEAR_OF_BIRTH        1990 5

DECLARE
     year_birth_xform   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
     dbms_data_mining_transform.STACK_BIN_NUM (
         bin_table_name           => 'bin_tbl',
         xform_list               =>  year_birth_xform);
     dbms_data_mining.CREATE_MODEL(
         model_name              => 'nb_model',
         mining_function         => dbms_data_mining.classification,
         data_table_name         => 'mining_data',
         case_id_column_name     => 'cust_id',
```

```
         target_column_name      => 'cust_postal_code',
         settings_table_name     => null,
         data_schema_name        => null,
         settings_schema_name    => null,
         xform_list              => year_birth_xform);
END;
/

SELECT attribute_name
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('nb_model'));

ATTRIBUTE_NAME
------------------------
CUST_YEAR_OF_BIRTH

SELECT expression
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('nb_model'));

EXPRESSION
--------------------------------------------------------------------------------
CASE WHEN "CUST_YEAR_OF_BIRTH"<1913 THEN NULL WHEN "CUST_YEAR_OF_BIRTH"<=1928.4
 THEN '1' WHEN "CUST_YEAR_OF_BIRTH"<=1943.8 THEN '2' WHEN "CUST_YEAR_OF_BIRTH"
<=1959.2 THEN '3' WHEN "CUST_YEAR_OF_BIRTH"<=1974.6 THEN '4' WHEN
"CUST_YEAR_OF_BIRTH" <=1990 THEN '5' END

SELECT reverse_expression
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('nb_model'));

REVERSE_EXPRESSION
--------------------------------------------------------------------------------
DECODE("CUST_YEAR_OF_BIRTH",'5','(1974.6; 1990]','1','[1913; 1928.4]','2','(1928
.4; 1943.8]','3','(1943.8; 1959.2]','4','(1959.2; 1974.6]',NULL,'( ; 1913), (199
0;  ), NULL')
```

58.6.11 INSERT_BIN_CAT_FREQ Procedure

This procedure performs categorical binning and inserts the transformation definitions in a transformation definition table. The procedure computes the bin boundaries based on frequency.

`INSERT_BIN_CAT_FREQ` bins all the `CHAR` and `VARCHAR2` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_BIN_CAT_FREQ (
     bin_table_name       IN VARCHAR2,
     data_table_name      IN VARCHAR2,
     bin_num              IN PLS_INTEGER DEFAULT 9,
     exclude_list         IN COLUMN_LIST DEFAULT NULL,
     default_num          IN PLS_INTEGER DEFAULT 2,
     bin_support          IN NUMBER DEFAULT NULL,
     bin_schema_name      IN VARCHAR2 DEFAULT NULL,
     data_schema_name     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-21 INSERT_BIN_CAT_FREQ Procedure Parameters

Parameter	Description

bin_table_name	Name of the transformation definition table for categorical binning. You can use the CREATE_BIN_CAT Procedure to create the definition table.The following columns are required:
	`COL VARCHAR2(30)`
	`VAL VARCHAR2(4000)`
	`BIN VARCHAR2(4000)`
	CREATE_BIN_CAT creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_BIN_CAT_FREQ.

data_table_name	Name of the table containing the data to be transformed
bin_num	The number of bins to fill using frequency-based binning The total number of bins will be *bin_num*+1. The additional bin is the default bin. Classes that are not assigned to a frequency-based bin will be assigned to the default bin. The default binning order is from highest to lowest: the most frequently occurring class is assigned to the first bin, the second most frequently occurring class is assigned to the second bin, and so on.You can reverse the binning order by specifying a negative number for *bin_num*. The negative sign causes the binning order to be from lowest to highest. If the total number of distinct values (classes) in the column is less than *bin_num*, then a separate bin will be created for each value and the default bin will be empty. If you specify NULL or 0 for *bin_num*, no binning is performed. The default value of *bin_num* is 9.
exclude_list	List of categorical columns to be excluded from the binning process. If you do not specify *exclude_list*, all categorical columns in the data source are binned. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
default_num	The number of class occurrences (rows of the same class) required for assignment to the default bin By default, *default_num* is the minimum number of occurrences required for assignment to the default bin. For example, if *default_num* is 3 and a given class occurs only once, it will not be assigned to the default bin. You can change the occurrence requirement from minimum to maximum by specifying a negative number for *default_num*. For example, if *default_num* is -3 and a given class occurs only once, it *will* be assigned to the default bin, but a class that occurs four or more times will not be included. If you specify NULL or 0 for *default_bin*, there are no requirements for assignment to the default bin. The default value of *default_num* is 2.

Table 58-21 (Cont.) INSERT_BIN_CAT_FREQ Procedure Parameters

Parameter	Description
bin_support	The number of class occurrences (rows of the same class) required for assignment to a frequency-based bin. *bin_support* is expressed as a fraction of the total number of rows. By default, *bin_support* is the minimum percentage required for assignment to a frequency-based bin. For example, if there are twenty rows of data and you specify.2 for *bin_support*, then there must be four or more occurrences of a class (.2*20) in order for it to be assigned to a frequency-based bin. You can change *bin_support* from a minimum percentage to a maximum percentage by specifying a negative number for *bin_support*. For example, if there are twenty rows of data and you specify -.2 for *bin_support*, then there must be four or less occurrences of a class in order for it to be assigned to a frequency-based bin. Classes that occur less than a positive *bin_support* or more than a negative *bin_support* will be assigned to the default bin. If you specify NULL or 0 for *bin_support*, then there is no support requirement for frequency-based binning. The default value of *bin_support* is NULL.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about categorical data.
2. If values occur with the same frequency, INSERT_BIN_CAT_FREQ assigns them in descending order when binning is from most to least frequent, or in ascending order when binning is from least to most frequent.

Examples

1. In this example, INSERT_BIN_CAT_FREQ computes the bin boundaries for the cust_postal_code and cust_city columns in sh.customers and inserts the transformations in a transformation definition table. The STACK_BIN_CAT Procedure creates a transformation list from the contents of the definition table, and the CREATE_MODEL Procedure embeds the transformation list in a new model called nb_model.

   The transformation and reverse transformation expressions embedded in nb_model are returned by the GET_MODEL_TRANSFORMATIONS Function.

```
CREATE OR REPLACE VIEW mining_data AS
          SELECT cust_id, cust_year_of_birth, cust_postal_code, cust_city
          FROM sh.customers;

DESCRIBE mining_data
 Name                                Null?    Type
 ----------------------------------- -------- ------------------------------
 CUST_ID                             NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                  NOT NULL NUMBER(4)
 CUST_POSTAL_CODE                    NOT NULL VARCHAR2(10)
 CUST_CITY                           NOT NULL VARCHAR2(30)
```

```
BEGIN
    dbms_data_mining_transform.CREATE_BIN_CAT(
       bin_table_name   => 'bin_tbl_1');
    dbms_data_mining_transform.INSERT_BIN_CAT_FREQ (
       bin_table_name   => 'bin_tbl_1',
       data_table_name  => 'mining_data',
       bin_num          => 4);
END;
/

column col format a18
column val format a15
column bin format a10
SELECT col, val, bin
       FROM bin_tbl_1
       ORDER BY col ASC, bin ASC;

COL                VAL             BIN
------------------ --------------- ----------
CUST_CITY          Los Angeles     1
CUST_CITY          Greenwich       2
CUST_CITY          Killarney       3
CUST_CITY          Montara         4
CUST_CITY                          5
CUST_POSTAL_CODE   38082           1
CUST_POSTAL_CODE   63736           2
CUST_POSTAL_CODE   55787           3
CUST_POSTAL_CODE   78558           4
CUST_POSTAL_CODE                   5

DECLARE
      city_xform   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
      dbms_data_mining_transform.STACK_BIN_CAT (
          bin_table_name          => 'bin_tbl_1',
          xform_list              =>  city_xform);
      dbms_data_mining.CREATE_MODEL(
           model_name             => 'nb_model',
           mining_function        => dbms_data_mining.classification,
           data_table_name        => 'mining_data',
           case_id_column_name    => 'cust_id',
           target_column_name     => 'cust_city',
           settings_table_name    => null,
           data_schema_name       => null,
           settings_schema_name   => null,
           xform_list             => city_xform);
END;
/

SELECT attribute_name
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('nb_model'));

ATTRIBUTE_NAME
--------------------------------------------------------------------------------
CUST_CITY
CUST_POSTAL_CODE

SELECT expression
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('nb_model'));

EXPRESSION
```

```
--------------------------------------------------------------------------------
DECODE("CUST_CITY",'Greenwich','2','Killarney','3','Los Angeles','1',
'Montara','4',NULL,NULL,'5')
DECODE("CUST_POSTAL_CODE",'38082','1','55787','3','63736','2','78558','4',NULL,NULL,'5')

SELECT reverse_expression
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('nb_model'));

REVERSE_EXPRESSION
--------------------------------------------------------------------------------
DECODE("CUST_CITY",'2','''Greenwich''','3','''Killarney''','1',
'''Los Angeles''','4','''Montara''',NULL,'NULL','5','DEFAULT')
DECODE("CUST_POSTAL_CODE",'1','''38082''','3','''55787''','2','''63736''',
'4','''78558''',NULL,'NULL','5','DEFAULT')
```

2. The binning order in example 1 is from most frequent to least frequent. The following example shows reverse order binning (least frequent to most frequent). The binning order is reversed by setting *bin_num* to -4 instead of 4.

```
BEGIN
    dbms_data_mining_transform.CREATE_BIN_CAT(
        bin_table_name   => 'bin_tbl_reverse');
    dbms_data_mining_transform.INSERT_BIN_CAT_FREQ (
        bin_table_name   => 'bin_tbl_reverse',
        data_table_name  => 'mining_data',
        bin_num          => -4);
 END;
 /

column col format a20
SELECT col, val, bin
       FROM bin_tbl_reverse
       ORDER BY col ASC, bin ASC;

COL                  VAL             BIN
-------------------- --------------- ----------
CUST_CITY            Tokyo           1
CUST_CITY            Sliedrecht      2
CUST_CITY            Haarlem         3
CUST_CITY            Diemen          4
CUST_CITY                            5
CUST_POSTAL_CODE     49358           1
CUST_POSTAL_CODE     80563           2
CUST_POSTAL_CODE     74903           3
CUST_POSTAL_CODE     71349           4
CUST_POSTAL_CODE                     5
```

58.6.12 INSERT_BIN_NUM_EQWIDTH Procedure

This procedure performs numerical binning and inserts the transformation definitions in a transformation definition table. The procedure identifies the minimum and maximum values and computes the bin boundaries at equal intervals.

`INSERT_BIN_NUM_EQWIDTH` computes a specified number of bins (*n*) and assigns *(max-min)/n* values to each bin. The number of bins is the same for each column. If you want to use equi-width binning, but you want the number of bins to be calculated on a per-column basis, use the INSERT_AUTOBIN_NUM_EQWIDTH Procedure.

`INSERT_BIN_NUM_EQWIDTH` bins all the `NUMBER` and `FLOAT` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_BIN_NUM_EQWIDTH (
     bin_table_name        IN VARCHAR2,
     data_table_name       IN VARCHAR2,
     bin_num               IN PLS_INTEGER DEFAULT 10,
     exclude_list          IN COLUMN_LIST DEFAULT NULL,
     round_num             IN PLS_INTEGER DEFAULT 6,
     bin_schema_name       IN VARCHAR2 DEFAULT NULL,
     data_schema_name      IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-22 INSERT_BIN_NUM_EQWIDTH Procedure Parameters

Parameter	Description

bin_table_name	Name of the transformation definition table for numerical binning. You can use the CREATE_BIN_NUM Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`VAL NUMBER`
	`BIN VARCHAR2(4000)`
	CREATE_BIN_NUM creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_BIN_NUM_EQWIDTH.

data_table_name	Name of the table containing the data to be transformed
bin_num	Number of bins. No binning occurs if *bin_num* is 0 or NULL. The default number of bins is 10.
exclude_list	List of numerical columns to be excluded from the binning process. If you do not specify *exclude_list*, all numerical columns in the data source are binned. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
round_num	Specifies how to round the number in the VAL column of the transformation definition table. When *round_num* is positive, it specifies the most significant digits to retain. When *round_num* is negative, it specifies the least significant digits to remove. In both cases, the result is rounded to the specified number of digits. See the Usage Notes for an example. The default value of *round_num* is 6.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.

2. The *round_num* parameter controls the rounding of column values in the transformation definition table, as follows:

```
For a value of 308.162:
when round_num =  1       result is 300
when round_num =  2       result is 310
when round_num =  3       result is 308
when round_num =  0       result is 308.162
when round_num = -1       result is 308.16
when round_num = -2       result is 308.2
```

3. INSERT_BIN_NUM_EQWIDTH ignores columns with all NULL values or only one unique value.

Examples

In this example, INSERT_BIN_NUM_EQWIDTH computes the bin boundaries for the affinity_card column in mining_data_build and inserts the transformations in a transformation definition table. The STACK_BIN_NUM Procedure creates a transformation list from the contents of the definition table. The CREATE_MODEL Procedure embeds the transformation list in a new model called glm_model.

The transformation and reverse transformation expressions embedded in glm_model are returned by the GET_MODEL_TRANSFORMATIONS Function.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, cust_income_level, cust_gender, affinity_card
       FROM mining_data_build;

DESCRIBE mining_data
 Name                      Null?    Type
 ------------------------- -------- -----------------
 CUST_ID                   NOT NULL NUMBER
 CUST_INCOME_LEVEL                  VARCHAR2(30)
 CUST_GENDER                        VARCHAR2(1)
 AFFINITY_CARD                      NUMBER(10)

BEGIN
    dbms_data_mining_transform.CREATE_BIN_NUM(
        bin_table_name   => 'bin_tbl');
    dbms_data_mining_transform.INSERT_BIN_NUM_EQWIDTH (
        bin_table_name   => 'bin_tbl',
        data_table_name  => 'mining_data',
        bin_num          => 4,
        exclude_list     => dbms_data_mining_transform.COLUMN_LIST('cust_id'));
END;
/

set numwidth 10
column val off
column col format a20
column bin format a10
SELECT col, val, bin FROM bin_tbl
    ORDER BY val ASC;

COL                          VAL  BIN
-------------------- ----------  ----------
AFFINITY_CARD                 0
AFFINITY_CARD               .25  1
AFFINITY_CARD                .5  2
AFFINITY_CARD               .75  3
AFFINITY_CARD                 1  4
```

```
CREATE TABLE glmsettings(
        setting_name  VARCHAR2(30),
        setting_value VARCHAR2(30));

BEGIN
   INSERT INTO glmsettings (setting_name, setting_value) VALUES
         (dbms_data_mining.algo_name, dbms_data_mining.algo_generalized_linear_model);
   COMMIT;
END;
/

DECLARE
     xforms   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
     dbms_data_mining_transform.STACK_BIN_NUM (
         bin_table_name           => 'bin_tbl',
         xform_list               =>  xforms,
         literal_flag             =>  TRUE);
     dbms_data_mining.CREATE_MODEL(
         model_name              => 'glm_model',
         mining_function         => dbms_data_mining.regression,
         data_table_name         => 'mining_data',
         case_id_column_name     => 'cust_id',
         target_column_name      => 'affinity_card',
         settings_table_name     => 'glmsettings',
         data_schema_name        => null,
         settings_schema_name    => null,
         xform_list              => xforms);
END;
/

SELECT attribute_name
      FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('glm_model'));

ATTRIBUTE_NAME
------------------------
AFFINITY_CARD

SELECT expression
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('glm_model'));

EXPRESSION
--------------------------------------------------------------------------------
CASE WHEN "AFFINITY_CARD"<0 THEN NULL WHEN "AFFINITY_CARD"<=.25 THEN 1 WHEN
"AFFINITY_CARD"<=.5 THEN 2 WHEN "AFFINITY_CARD"<=.75 THEN 3 WHEN
"AFFINITY_CARD"<=1 THEN 4 END

SELECT reverse_expression
       FROM TABLE(dbms_data_mining.GET_MODEL_TRANSFORMATIONS('glm_model'));

REVERSE_EXPRESSION
--------------------------------------------------------------------------------
DECODE("AFFINITY_CARD",4,'(.75; 1]',1,'[0; .25]',2,'(.25; .5]',3,'(.5; .75]',
NULL,'( ; 0), (1;  ), NULL')
```

58.6.13 INSERT_BIN_NUM_QTILE Procedure

This procedure performs numerical binning and inserts the transformation definitions in a transformation definition table. The procedure calls the SQL `NTILE` function to order the data and divide it equally into the specified number of bins (quantiles).

`INSERT_BIN_NUM_QTILE` bins all the `NUMBER` and `FLOAT` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_BIN_NUM_QTILE (
    bin_table_name       IN VARCHAR2,
    data_table_name      IN VARCHAR2,
    bin_num              IN PLS_INTEGER DEFAULT 10,
    exclude_list         IN COLUMN_LIST DEFAULT NULL,
    bin_schema_name      IN VARCHAR2 DEFAULT NULL,
    data_schema_name     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-23 INSERT_BIN_NUM_QTILE Procedure Parameters

Parameter	Description

`bin_table_name`	Name of the transformation definition table for numerical binning. You can use the CREATE_BIN_NUM Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`VAL NUMBER`
	`BIN VARCHAR2(4000)`
	`CREATE_BIN_NUM` creates an additional column, `ATT`, which may be used for specifying nested attributes. This column is not used by `INSERT_BIN_NUM_QTILE`.

`data_table_name`	Name of the table containing the data to be transformed
`bin_num`	Number of bins. No binning occurs if *`bin_num`* is `0` or `NULL`. The default number of bins is 10.

`exclude_list`	List of numerical columns to be excluded from the binning process. If you do not specify *`exclude_list`*, all numerical columns in the data source are binned.
	The format of *`exclude_list`* is:
	`dbms_data_mining_transform.COLUMN_LIST('col1','col2',`
	`...'coln')`

`bin_schema_name`	Schema of *`bin_table_name`*. If no schema is specified, the current schema is used.
`data_schema_name`	Schema of *`data_table_name`*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.

2. After dividing the data into quantiles, the NTILE function distributes any remainder values one for each quantile, starting with the first. See *Oracle Database SQL Language Reference* for details.
3. Columns with all NULL values are ignored by INSERT_BIN_NUM_QTILE.

Examples

In this example, INSERT_BIN_NUM_QTILE computes the bin boundaries for the cust_year_of_birth and cust_credit_limit columns in sh.customers and inserts the transformations in a transformation definition table. The STACK_BIN_NUM Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in STACK_VIEW. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
        SELECT cust_id, cust_year_of_birth, cust_credit_limit, cust_city
        FROM sh.customers;

DESCRIBE mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                        NOT NULL NUMBER(4)
 CUST_CREDIT_LIMIT                         NUMBER
 CUST_CITY                                 NOT NULL VARCHAR2(30)

BEGIN
   dbms_data_mining_transform.CREATE_BIN_NUM(
        bin_table_name   => 'bin_tbl');
   dbms_data_mining_transform.INSERT_BIN_NUM_QTILE (
        bin_table_name   => 'bin_tbl',
        data_table_name  => 'mining_data',
        bin_num          => 3,
        exclude_list     => dbms_data_mining_transform.COLUMN_LIST('cust_id'));
END;
/

set numwidth 8
column val off
column col format a20
column bin format a10
SELECT col, val, bin
      FROM bin_tbl
      ORDER BY col ASC, val ASC;

COL                      VAL BIN
-------------------- -------- ----------
CUST_CREDIT_LIMIT        1500
CUST_CREDIT_LIMIT        3000 1
CUST_CREDIT_LIMIT        9000 2
CUST_CREDIT_LIMIT       15000 3
CUST_YEAR_OF_BIRTH       1913
CUST_YEAR_OF_BIRTH       1949 1
CUST_YEAR_OF_BIRTH       1965 2
CUST_YEAR_OF_BIRTH       1990 3

DECLARE
   xforms   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
   dbms_data_mining_transform.STACK_BIN_NUM (
```

```
        bin_table_name        => 'bin_tbl',
        xform_list            =>  xforms);
   dbms_data_mining_transform.XFORM_STACK (
        xform_list            =>  xforms,
        data_table_name       => 'mining_data',
        xform_view_name       => 'stack_view');
END;
/

set long 3000
SELECT text FROM user_views WHERE view_name in 'STACK_VIEW';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID",CASE WHEN "CUST_YEAR_OF_BIRTH"<1913 THEN NULL WHEN "CUST_YEAR_O
F_BIRTH"<=1949 THEN '1' WHEN "CUST_YEAR_OF_BIRTH"<=1965 THEN '2' WHEN "CUST_YEAR
_OF_BIRTH"<=1990 THEN '3' END "CUST_YEAR_OF_BIRTH",CASE WHEN "CUST_CREDIT_LIMIT"
<1500 THEN NULL WHEN "CUST_CREDIT_LIMIT"<=3000 THEN '1' WHEN "CUST_CREDIT_LIMIT"
<=9000 THEN '2' WHEN "CUST_CREDIT_LIMIT"<=15000 THEN '3' END "CUST_CREDIT_LIMIT"
,"CUST_CITY" FROM mining_data
```

58.6.14 INSERT_BIN_SUPER Procedure

This procedure performs numerical and categorical binning and inserts the transformation definitions in transformation definition tables. The procedure computes bin boundaries based on intrinsic relationships between predictors and a target.

`INSERT_BIN_SUPER` uses an intelligent binning technique known as **supervised binning**. It builds a single-predictor decision tree and derives the bin boundaries from splits within the tree.

`INSERT_BIN_SUPER` bins all the `VARCHAR2`, `CHAR`, `NUMBER`, and `FLOAT` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_BIN_SUPER (
    num_table_name        IN VARCHAR2,
    cat_table_name        IN VARCHAR2,
    data_table_name       IN VARCHAR2,
    target_column_name    IN VARCHAR2,
    max_bin_num           IN PLS_INTEGER  DEFAULT 1000,
    exclude_list          IN COLUMN_LIST  DEFAULT NULL,
    num_schema_name       IN VARCHAR2     DEFAULT NULL,
    cat_schema_name       IN VARCHAR2     DEFAULT NULL,
    data_schema_name      IN VARCHAR2     DEFAULT NULL,
    rem_table_name        IN VARCHAR2     DEFAULT NULL,
    rem_schema_name       IN VARCHAR2     DEFAULT NULL);
```

Parameters

Table 58-24 INSERT_BIN_SUPER Procedure Parameters

Parameter	Description

num_table_name	Name of the transformation definition table for numerical binning. You can use the CREATE_BIN_NUM Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`VAL VNUMBER`
	`BIN VARCHAR2(4000)`
	CREATE_BIN_NUM creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_BIN_SUPER.

cat_table_name	Name of the transformation definition table for categorical binning. You can use the CREATE_BIN_CAT Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`VAL VARCHAR2(4000)`
	`BIN VARCHAR2(4000)`
	CREATE_BIN_CAT creates an additional column, ATT, which is used for specifying nested attributes. This column is not used by INSERT_BIN_SUPER.

data_table_name	Name of the table containing the data to be transformed
target_column_name	Name of a column to be used as the target for the decision tree models
max_bin_num	The maximum number of bins. The default is 1000.
exclude_list	List of columns to be excluded from the binning process. If you do not specify *exclude_list*, all numerical and categorical columns in the data source are binned. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
num_schema_name	Schema of *num_table_name*. If no schema is specified, the current schema is used.
cat_schema_name	Schema of *cat_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
rem_table_name	Name of a column removal definition table. The table must have the columns described in "CREATE_COL_REM Procedure". You can use CREATE_COL_REM to create the table. See Usage Notes.
rem_schema_name	Schema of *rem_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical and categorical data.

2. Columns that have no significant splits are not binned. You can remove the unbinned columns from the mining data by specifying a column removal definition table. If you do not specify a column removal definition table, the unbinned columns remain in the mining data.
3. See *Oracle Machine Learning for SQL Concepts* to learn more about decision trees in Oracle Machine Learning for SQL

Examples

In this example, `INSERT_BIN_SUPER` computes the bin boundaries for predictors of `cust_credit_limit` and inserts the transformations in transformation definition tables. One predictor is numerical, the other is categorical. (`INSERT_BIN_SUPER` determines that the `cust_postal_code` column is not a significant predictor.) `STACK` procedures create transformation lists from the contents of the definition tables.

The SQL expressions that compute the transformations are shown in the views `MINING_DATA_STACK_NUM` and `MINING_DATA_STACK_CAT`. The views are for display purposes only; they cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
     SELECT cust_id, cust_year_of_birth, cust_marital_status,
            cust_postal_code, cust_credit_limit
     FROM sh.customers;

DESCRIBE mining_data
 Name                           Null?    Type
 ------------------------------ -------- -------------------------------------
 CUST_ID                        NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH             NOT NULL NUMBER(4)
 CUST_MARITAL_STATUS                     VARCHAR2(20)
 CUST_POSTAL_CODE               NOT NULL VARCHAR2(10)
 CUST_CREDIT_LIMIT                       NUMBER

BEGIN
    dbms_data_mining_transform.CREATE_BIN_NUM(
         bin_table_name      => 'bin_num_tbl');
    dbms_data_mining_transform.CREATE_BIN_CAT(
         bin_table_name      => 'bin_cat_tbl');
    dbms_data_mining_transform.CREATE_COL_REM(
         rem_table_name      => 'rem_tbl');
END;
/

BEGIN
   COMMIT;
   dbms_data_mining_transform.INSERT_BIN_SUPER (
      num_table_name      => 'bin_num_tbl',
      cat_table_name      => 'bin_cat_tbl',
      data_table_name     => 'mining_data',
      target_column_name  => 'cust_credit_limit',
      max_bin_num         =>  4,
      exclude_list        =>  dbms_data_mining_transform.COLUMN_LIST('cust_id'),
      num_schema_name     => 'oml_user',
      cat_schema_name     => 'oml_user',
      data_schema_name    => 'oml_user',
      rem_table_name      => 'rem_tbl',
      rem_schema_name     => 'oml_user');
   COMMIT;
END;
/

set numwidth 8
```

```
column val off
SELECT col, val, bin FROM bin_num_tbl
      ORDER BY bin ASC;

COL                        VAL BIN
-------------------- -------- ----------
CUST_YEAR_OF_BIRTH     1923.5 1
CUST_YEAR_OF_BIRTH     1923.5 1
CUST_YEAR_OF_BIRTH     1945.5 2
CUST_YEAR_OF_BIRTH     1980.5 3
CUST_YEAR_OF_BIRTH            4

column val on
column val format a20
SELECT col, val, bin FROM bin_cat_tbl
      ORDER BY bin ASC;

COL                  VAL                  BIN
-------------------- -------------------- ----------
CUST_MARITAL_STATUS  married              1
CUST_MARITAL_STATUS  single               2
CUST_MARITAL_STATUS  Mar-AF               3
CUST_MARITAL_STATUS  Mabsent              3
CUST_MARITAL_STATUS  Divorc.              3
CUST_MARITAL_STATUS  Married              3
CUST_MARITAL_STATUS  Widowed              3
CUST_MARITAL_STATUS  NeverM               3
CUST_MARITAL_STATUS  Separ.               3
CUST_MARITAL_STATUS  divorced             4
CUST_MARITAL_STATUS  widow                4

SELECT col from rem_tbl;

COL
--------------------
CUST_POSTAL_CODE

DECLARE
    xforms_num      dbms_data_mining_transform.TRANSFORM_LIST;
    xforms_cat      dbms_data_mining_transform.TRANSFORM_LIST;
    BEGIN
       dbms_data_mining_transform.STACK_BIN_NUM (
            bin_table_name    => 'bin_num_tbl',
            xform_list        => xforms_num);
       dbms_data_mining_transform.XFORM_STACK (
            xform_list         => xforms_num,
            data_table_name    => 'mining_data',
            xform_view_name    => 'mining_data_stack_num');
       dbms_data_mining_transform.STACK_BIN_CAT (
             bin_table_name    => 'bin_cat_tbl',
             xform_list        => xforms_cat);
       dbms_data_mining_transform.XFORM_STACK (
             xform_list         => xforms_cat,
             data_table_name    => 'mining_data',
             xform_view_name    => 'mining_data_stack_cat');
  END;
 /

set long 3000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK_NUM';

TEXT
```

```
--------------------------------------------------------------------------------
SELECT "CUST_ID",CASE WHEN "CUST_YEAR_OF_BIRTH"<1923.5 THEN '1' WHEN "CUST_YEAR_
OF_BIRTH"<=1923.5 THEN '1' WHEN "CUST_YEAR_OF_BIRTH"<=1945.5 THEN '2' WHEN "CUST
_YEAR_OF_BIRTH"<=1980.5 THEN '3' WHEN "CUST_YEAR_OF_BIRTH" IS NOT NULL THEN '4'
END "CUST_YEAR_OF_BIRTH","CUST_MARITAL_STATUS","CUST_POSTAL_CODE","CUST_CREDIT_L
IMIT" FROM mining_data


SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK_CAT';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_YEAR_OF_BIRTH",DECODE("CUST_MARITAL_STATUS",'Divorc.','3'
,'Mabsent','3','Mar-AF','3','Married','3','NeverM','3','Separ.','3','Widowed','3
','divorced','4','married','1','single','2','widow','4') "CUST_MARITAL_STATUS","
CUST_POSTAL_CODE","CUST_CREDIT_LIMIT" FROM mining_data
```

58.6.15 INSERT_CLIP_TRIM_TAIL Procedure

This procedure replaces numeric outliers with nulls and inserts the transformation definitions in a transformation definition table.

`INSERT_CLIP_TRIM_TAIL` computes the boundaries of the data based on a specified percentage. It removes the values that fall outside the boundaries (tail values) from the data. If you wish to replace the tail values instead of removing them, use the INSERT_CLIP_WINSOR_TAIL Procedure.

`INSERT_CLIP_TRIM_TAIL` clips all the `NUMBER` and `FLOAT` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_CLIP_TRIM_TAIL (
    clip_table_name     IN VARCHAR2,
    data_table_name     IN VARCHAR2,
    tail_frac           IN NUMBER DEFAULT 0.025,
    exclude_list        IN COLUMN_LIST DEFAULT NULL,
    clip_schema_name    IN VARCHAR2 DEFAULT NULL,
    data_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-25 INSERT_CLIP_TRIM_TAIL Procedure Parameters

Parameter	Description

`clip_table_name`	Name of the transformation definition table for numerical clipping. You can use the CREATE_CLIP Procedure to create the definition table. The following columns are required:
	
	`COL VARCHAR2(30)`
	`LCUT NUMBER`
	`LVAL NUMBER`
	`RCUT NUMBER`
	`RVAL NUMBER`
	
	`CREATE_CLIP` creates an additional column, `ATT`, which may be used for specifying nested attributes. This column is not used by `INSERT_CLIP_TRIM_TAIL`.

`data_table_name`	Name of the table containing the data to be transformed

Table 58-25 (Cont.) INSERT_CLIP_TRIM_TAIL Procedure Parameters

Parameter	Description
tail_frac	The percentage of non-null values to be designated as outliers at each end of the data. For example, if *tail_frac* is .01, then 1% of the data at the low end and 1% of the data at the high end will be treated as outliers. If *tail_frac* is greater than or equal to .5, no clipping occurs. The default value of *tail_frac* is 0.025.
exclude_list	List of numerical columns to be excluded from the clipping process. If you do not specify *exclude_list*, all numerical columns in the data are clipped. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
clip_schema_name	Schema of *clip_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. The DBMS_DATA_MINING_TRANSFORM package provides two clipping procedures: INSERT_CLIP_TRIM_TAIL and INSERT_CLIP_WINSOR_TAIL. Both procedures compute the boundaries as follows:
   - Count the number of non-null values, ***n***, and sort them in ascending order
   - Calculate the number of outliers, ***t***, as ***n*tail_frac***
   - Define the lower boundary ***lcut*** as the value at position ***1+floor(t)***
   - Define the upper boundary ***rcut*** as the value at position ***n-floor(t)***

     (The SQL FLOOR function returns the largest integer less than or equal to ***t***.)
   - All values that are <= ***lcut*** or => ***rcut*** are designated as outliers.

   INSERT_CLIP_TRIM_TAIL replaces the outliers with nulls, effectively removing them from the data.

   INSERT_CLIP_WINSOR_TAIL assigns ***lcut*** to the low outliers and ***rcut*** to the high outliers.

Examples

In this example, INSERT_CLIP_TRIM_TAIL trims 10% of the data in two columns (5% from the high end and 5% from the low end) and inserts the transformations in a transformation definition table. The STACK_CLIP Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the trimming is shown in the view MINING_DATA_STACK. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, cust_year_of_birth, cust_credit_limit, cust_city
       FROM sh.customers;
```

```
DESCRIBE mining_data
 Name                           Null?    Type
 ------------------------------ -------- -------------------
 CUST_ID                        NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH             NOT NULL NUMBER(4)
 CUST_CREDIT_LIMIT                       NUMBER
 CUST_CITY                      NOT NULL VARCHAR2(30)

BEGIN
   dbms_data_mining_transform.CREATE_CLIP(
      clip_table_name    => 'clip_tbl');
   dbms_data_mining_transform.INSERT_CLIP_TRIM_TAIL(
      clip_table_name    => 'clip_tbl',
      data_table_name    => 'mining_data',
      tail_frac          => 0.05,
      exclude_list       => DBMS_DATA_MINING_TRANSFORM.COLUMN_LIST('cust_id'));
END;
/

SELECT col, lcut, lval, rcut, rval
      FROM clip_tbl
      ORDER BY col ASC;

COL                     LCUT     LVAL     RCUT     RVAL
------------------- -------- -------- -------- --------
CUST_CREDIT_LIMIT       1500             11000
CUST_YEAR_OF_BIRTH      1934              1982

DECLARE
     xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
     dbms_data_mining_transform.STACK_CLIP (
          clip_table_name    => 'clip_tbl',
          xform_list         => xforms);
     dbms_data_mining_transform.XFORM_STACK (
          xform_list         => xforms,
          data_table_name    => 'mining_data',
          xform_view_name    => 'mining_data_stack');
 END;
 /

set long 3000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID",CASE WHEN "CUST_YEAR_OF_BIRTH" < 1934 THEN NULL WHEN "CUST_YEAR
_OF_BIRTH" > 1982 THEN NULL ELSE "CUST_YEAR_OF_BIRTH" END "CUST_YEAR_OF_BIRTH",C
ASE WHEN "CUST_CREDIT_LIMIT" < 1500 THEN NULL WHEN "CUST_CREDIT_LIMIT" > 11000 T
HEN NULL ELSE "CUST_CREDIT_LIMIT" END "CUST_CREDIT_LIMIT","CUST_CITY" FROM minin
g_data
```

58.6.16 INSERT_CLIP_WINSOR_TAIL Procedure

This procedure replaces numeric outliers with the upper or lower boundary values. It inserts the transformation definitions in a transformation definition table.

`INSERT_CLIP_WINSOR_TAIL` computes the boundaries of the data based on a specified percentage. It replaces the values that fall outside the boundaries (tail values) with the related boundary value. If you wish to set tail values to null, use the INSERT_CLIP_TRIM_TAIL Procedure.

INSERT_CLIP_WINSOR_TAIL clips all the NUMBER and FLOAT columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_CLIP_WINSOR_TAIL (
    clip_table_name    IN VARCHAR2,
    data_table_name    IN VARCHAR2,
    tail_frac          IN NUMBER DEFAULT 0.025,
    exclude_list       IN COLUMN_LIST DEFAULT NULL,
    clip_schema_name   IN VARCHAR2 DEFAULT NULL,
    data_schema_name   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-26 INSERT_CLIP_WINSOR_TAIL Procedure Parameters

Parameter	Description
clip_table_name	Name of the transformation definition table for numerical clipping. You can use the CREATE_CLIP Procedure to create the definition table. The following columns are required: COL VARCHAR2(30) LCUT NUMBER LVAL NUMBER RCUT NUMBER RVAL NUMBER CREATE_CLIP creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_CLIP_WINSOR_TAIL.
data_table_name	Name of the table containing the data to be transformed
tail_frac	The percentage of non-null values to be designated as outliers at each end of the data. For example, if *tail_frac* is .01, then 1% of the data at the low end and 1% of the data at the high end will be treated as outliers. If *tail_frac* is greater than or equal to .5, no clipping occurs. The default value of *tail_frac* is 0.025.
exclude_list	List of numerical columns to be excluded from the clipping process. If you do not specify *exclude_list*, all numerical columns in the data are clipped. The format of *exclude_list* is: dbms_data_mining_transform.COLUMN_LIST('col*1*','col*2*', ...'col*n*')
clip_schema_name	Schema of *clip_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. The DBMS_DATA_MINING_TRANSFORM package provides two clipping procedures: INSERT_CLIP_WINSOR_TAIL and INSERT_CLIP_TRIM_TAIL. Both procedures compute the boundaries as follows:

- Count the number of non-null values, ***n***, and sort them in ascending order
- Calculate the number of outliers, ***t***, as ***n*tail_frac***
- Define the lower boundary ***lcut*** as the value at position ***1+floor(t)***
- Define the upper boundary ***rcut*** as the value at position ***n-floor(t)***

  (The SQL `FLOOR` function returns the largest integer less than or equal to ***t***.)
- All values that are <= ***lcut*** or => ***rcut*** are designated as outliers.

`INSERT_CLIP_WINSOR_TAIL` assigns ***lcut*** to the low outliers and ***rcut*** to the high outliers.

`INSERT_CLIP_TRIM_TAIL` replaces the outliers with nulls, effectively removing them from the data.

Examples

In this example, `INSERT_CLIP_WINSOR_TAIL` winsorizes 10% of the data in two columns (5% from the high end, and 5% from the low end) and inserts the transformations in a transformation definition table. The STACK_CLIP Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in the view `MINING_DATA_STACK`. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
        SELECT cust_id, cust_year_of_birth, cust_credit_limit, cust_city
        FROM sh.customers;

describe mining_data
 Name                                    Null?    Type
 ----------------------------------------- -------- -------------
 CUST_ID                                 NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                      NOT NULL NUMBER(4)
 CUST_CREDIT_LIMIT                                NUMBER
 CUST_CITY                               NOT NULL VARCHAR2(30)

BEGIN
  dbms_data_mining_transform.CREATE_CLIP(
      clip_table_name   => 'clip_tbl');
  dbms_data_mining_transform.INSERT_CLIP_WINSOR_TAIL(
      clip_table_name   => 'clip_tbl',
      data_table_name   => 'mining_data',
      tail_frac         => 0.05,
      exclude_list      => DBMS_DATA_MINING_TRANSFORM.COLUMN_LIST('cust_id'));
END;
/

SELECT col, lcut, lval, rcut, rval FROM clip_tbl
   ORDER BY col ASC;
COL                            LCUT     LVAL     RCUT     RVAL
------------------------------ -------- -------- -------- --------
CUST_CREDIT_LIMIT                 1500     1500    11000    11000
CUST_YEAR_OF_BIRTH                1934     1934     1982     1982

DECLARE
   xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
   dbms_data_mining_transform.STACK_CLIP (
   clip_table_name    => 'clip_tbl',
   xform_list        => xforms);
```

```
dbms_data_mining_transform.XFORM_STACK (
   xform_list         => xforms,
   data_table_name    => 'mining_data',
   xform_view_name    => 'mining_data_stack');
END;
/

set long 3000
SQL> SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID",CASE WHEN "CUST_YEAR_OF_BIRTH" < 1934 THEN 1934 WHEN "CUST_YEAR
_OF_BIRTH" > 1982 THEN 1982 ELSE "CUST_YEAR_OF_BIRTH" END "CUST_YEAR_OF_BIRTH",C
ASE WHEN "CUST_CREDIT_LIMIT" < 1500 THEN 1500 WHEN "CUST_CREDIT_LIMIT" > 11000 T
HEN 11000 ELSE "CUST_CREDIT_LIMIT" END "CUST_CREDIT_LIMIT","CUST_CITY" FROM mini
ng_data
```

58.6.17 INSERT_MISS_CAT_MODE Procedure

This procedure replaces missing categorical values with the value that occurs most frequently in the column (the mode). It inserts the transformation definitions in a transformation definition table.

`INSERT_MISS_CAT_MODE` replaces missing values in all `VARCHAR2` and `CHAR` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_MISS_CAT_MODE (
    miss_table_name    IN VARCHAR2,
    data_table_name    IN VARCHAR2,
    exclude_list       IN COLUMN_LIST DEFAULT NULL,
    miss_schema_name   IN VARCHAR2 DEFAULT NULL,
    data_schema_name   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-27 INSERT_MISS_CAT_MODE Procedure Parameters

Parameter	Description

`miss_table_name`	Name of the transformation definition table for categorical missing value treatment. You can use the CREATE_MISS_CAT Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`VAL VARCHAR2(4000)`
	`CREATE_MISS_CAT` creates an additional column, `ATT`, which may be used for specifying nested attributes. This column is not used by `INSERT_MISS_CAT_MODE`.

`data_table_name`	Name of the table containing the data to be transformed

Table 58-27 (Cont.) INSERT_MISS_CAT_MODE Procedure Parameters

Parameter	Description
exclude_list	List of categorical columns to be excluded from missing value treatment. If you do not specify *exclude_list*, all categorical columns are transformed. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
miss_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about categorical data.
2. If you wish to replace categorical missing values with a value other than the mode, you can edit the transformation definition table.

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for information about default missing value treatment in Oracle Machine Learning for SQL

Example

In this example, `INSERT_MISS_CAT_MODE` computes missing value treatment for `cust_city` and inserts the transformation in a transformation definition table. The STACK_MISS_CAT Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in the view `MINING_DATA_STACK`. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
         SELECT cust_id, cust_year_of_birth, cust_city
         FROM sh.customers;

describe mining_data
 Name                           Null?    Type
 ------------------------------ -------- ----------------
 CUST_ID                        NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH             NOT NULL NUMBER(4)
 CUST_CITY                      NOT NULL VARCHAR2(30)

BEGIN
  dbms_data_mining_transform.create_miss_cat(
      miss_table_name   => 'missc_tbl');
  dbms_data_mining_transform.insert_miss_cat_mode(
      miss_table_name   => 'missc_tbl',
      data_table_name   => 'mining_data');
END;
```

```
/

SELECT stats_mode(cust_city) FROM mining_data;

STATS_MODE(CUST_CITY)
------------------------------
Los Angeles

SELECT col, val
    from missc_tbl;

COL                            VAL
------------------------------ ------------------------------
CUST_CITY                      Los Angeles

DECLARE
    xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
    dbms_data_mining_transform.STACK_MISS_CAT (
        miss_table_name    => 'missc_tbl',
        xform_list        => xforms);
    dbms_data_mining_transform.XFORM_STACK (
         xform_list         => xforms,
         data_table_name    => 'mining_data',
         xform_view_name    => 'mining_data_stack');
END;
/

set long 3000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_YEAR_OF_BIRTH",NVL("CUST_CITY",'Los Angeles') "CUST_CITY"
 FROM mining_data
```

58.6.18 INSERT_MISS_NUM_MEAN Procedure

This procedure replaces missing numerical values with the average (the mean) and inserts the transformation definitions in a transformation definition table.

`INSERT_MISS_NUM_MEAN` replaces missing values in all `NUMBER` and `FLOAT` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_MISS_NUM_MEAN (
    miss_table_name    IN VARCHAR2,
    data_table_name    IN VARCHAR2,
    exclude_list       IN COLUMN_LIST DEFAULT NULL,
    round_num          IN PLS_INTEGER DEFAULT 6,
    miss_schema_name   IN VARCHAR2 DEFAULT NULL,
    data_schema_name   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-28 INSERT_MISS_NUM_MEAN Procedure Parameters

Parameter	Description

miss_table_name	Name of the transformation definition table for numerical missing value treatment. You can use the CREATE_MISS_NUM Procedure to create the definition table.
	The following columns are required by INSERT_MISS_NUM_MEAN:
	`COL VARCHAR2(30)`
	`VAL NUMBER`
	CREATE_MISS_NUM creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_MISS_NUM_MEAN.

data_table_name	Name of the table containing the data to be transformed
exclude_list	List of numerical columns to be excluded from missing value treatment. If you do not specify *exclude_list*, all numerical columns are transformed. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
round_num	The number of significant digits to use for the mean. The default number is 6.
miss_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

1. See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.
2. If you wish to replace numerical missing values with a value other than the mean, you can edit the transformation definition table.

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for information about default missing value treatment in Oracle Machine Learning for SQL

Example

In this example, INSERT_MISS_NUM_MEAN computes missing value treatment for cust_year_of_birth and inserts the transformation in a transformation definition table. The STACK_MISS_NUM Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in the view MINING_DATA_STACK. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
    SELECT cust_id, cust_year_of_birth, cust_city
    FROM sh.customers;

DESCRIBE mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                        NOT NULL NUMBER(4)
 CUST_CITY                                 NOT NULL VARCHAR2(30)

BEGIN
   dbms_data_mining_transform.create_miss_num(
       miss_table_name   => 'missn_tbl');
   dbms_data_mining_transform.insert_miss_num_mean(
       miss_table_name   => 'missn_tbl',
       data_table_name   => 'mining_data',
       exclude_list      => DBMS_DATA_MINING_TRANSFORM.COLUMN_LIST('cust_id'));
END;
/

set numwidth 4
column val off
SELECT col, val
  FROM missn_tbl;

COL                  VAL
-------------------- ----
CUST_YEAR_OF_BIRTH   1957

SELECT avg(cust_year_of_birth) FROM mining_data;

AVG(CUST_YEAR_OF_BIRTH)
-----------------------
                   1957

DECLARE
    xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
    dbms_data_mining_transform.STACK_MISS_NUM (
         miss_table_name    => 'missn_tbl',
         xform_list        => xforms);
    dbms_data_mining_transform.XFORM_STACK (
         xform_list         => xforms,
         data_table_name    => 'mining_data',
         xform_view_name    => 'mining_data_stack');
END;
/

set long 3000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID",NVL("CUST_YEAR_OF_BIRTH",1957.4) "CUST_YEAR_OF_BIRTH","CUST_CIT
Y" FROM mining_data
```

58.6.19 INSERT_NORM_LIN_MINMAX Procedure

This procedure performs linear normalization and inserts the transformation definitions in a transformation definition table.

INSERT_NORM_LIN_MINMAX computes the minimum and maximum values from the data and sets the value of *shift* and *scale* as follows:

```
shift = min
scale = max - min
```

Normalization is computed as:

```
x_new = (x_old - shift)/scale
```

INSERT_NORM_LIN_MINMAX rounds the value of *scale* to a specified number of significant digits before storing it in the transformation definition table.

INSERT_NORM_LIN_MINMAX normalizes all the NUMBER and FLOAT columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_NORM_LIN_MINMAX (
     norm_table_name     IN VARCHAR2,
     data_table_name     IN VARCHAR2,
     exclude_list        IN COLUMN_LIST DEFAULT NULL,
     round_num           IN PLS_INTEGER DEFAULT 6,
     norm_schema_name    IN VARCHAR2 DEFAULT NULL,
     data_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-29 *INSERT_NORM_LIN_MINMAX Procedure Parameters*

Parameter	Description

norm_table_name	Name of the transformation definition table for linear normalization. You can use the CREATE_NORM_LIN Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`SHIFT NUMBER`
	`SCALE NUMBER`
	CREATE_NORM_LIN creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_NORM_LIN_MINMAX.

data_table_name	Name of the table containing the data to be transformed
exclude_list	List of numerical columns to be excluded from normalization. If you do not specify *exclude_list*, all numerical columns are transformed. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
round_num	The number of significant digits to use for the minimum and maximum. The default number is 6.

Table 58-29 (Cont.) *INSERT_NORM_LIN_MINMAX Procedure Parameters*

Parameter	Description
norm_schema_name	Schema of *norm_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.

Examples

In this example, INSERT_NORM_LIN_MINMAX normalizes the cust_year_of_birth column and inserts the transformation in a transformation definition table. The STACK_NORM_LIN Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in the view MINING_DATA_STACK. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
         SELECT cust_id, cust_gender, cust_year_of_birth
         FROM sh.customers;

describe mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_GENDER                               NOT NULL CHAR(1)
 CUST_YEAR_OF_BIRTH                        NOT NULL NUMBER(4)

BEGIN
       dbms_data_mining_transform.CREATE_NORM_LIN(
         norm_table_name  => 'norm_tbl');
       dbms_data_mining_transform.INSERT_NORM_LIN_MINMAX(
         norm_table_name  => 'norm_tbl',
         data_table_name  => 'mining_data',
         exclude_list     => dbms_data_mining_transform.COLUMN_LIST( 'cust_id'),
         round_num        => 3);
END;
/

SELECT col, shift, scale FROM norm_tbl;

COL                              SHIFT      SCALE
-------------------------------- ---------- ----------
CUST_YEAR_OF_BIRTH                 1910         77

DECLARE
     xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
     dbms_data_mining_transform.STACK_NORM_LIN (
          norm_table_name    => 'norm_tbl',
          xform_list        => xforms);
     dbms_data_mining_transform.XFORM_STACK (
          xform_list         => xforms,
```

```
          data_table_name    => 'mining_data',
          xform_view_name    => 'mining_data_stack');
END;
/

set long 3000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_GENDER",("CUST_YEAR_OF_BIRTH"-1910)/77 "CUST_YEAR_OF_BIRT
H" FROM mining_data
```

58.6.20 INSERT_NORM_LIN_SCALE Procedure

This procedure performs linear normalization and inserts the transformation definitions in a transformation definition table.

`INSERT_NORM_LIN_SCALE` computes the minimum and maximum values from the data and sets the value of *`shift`* and *`scale`* as follows:

```
shift = 0
scale = max(abs(max), abs(min))
```

Normalization is computed as:

```
x_new = (x_old)/scale
```

`INSERT_NORM_LIN_SCALE` rounds the value of *`scale`* to a specified number of significant digits before storing it in the transformation definition table.

`INSERT_NORM_LIN_SCALE` normalizes all the `NUMBER` and `FLOAT` columns in the data source unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_NORM_LIN_SCALE (
     norm_table_name     IN VARCHAR2,
     data_table_name     IN VARCHAR2,
     exclude_list        IN COLUMN_LIST DEFAULT NULL,
     round_num           IN PLS_INTEGER DEFAULT 6,
     norm_schema_name    IN VARCHAR2 DEFAULT NULL,
     data_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-30 INSERT_NORM_LIN_SCALE Procedure Parameters

Parameter	Description

norm_table_name	Name of the transformation definition table for linear normalization. You can use the CREATE_NORM_LIN Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`SHIFT NUMBER`
	`SCALE NUMBER`
	CREATE_NORM_LIN creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_NORM_LIN_SCALE.

data_table_name	Name of the table containing the data to be transformed
exclude_list	List of numerical columns to be excluded from normalization. If you do not specify *exclude_list*, all numerical columns are transformed. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
round_num	The number of significant digits to use for *scale*. The default number is 6.
norm_schema_name	Schema of *norm_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.

Examples

In this example, INSERT_NORM_LIN_SCALE normalizes the cust_year_of_birth column and inserts the transformation in a transformation definition table. The STACK_NORM_LIN Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in the view MINING_DATA_STACK. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
      SELECT cust_id, cust_gender, cust_year_of_birth
      FROM sh.customers;

DESCRIBE mining_data
 Name                              Null?    Type
 --------------------------------- -------- ------------------
 CUST_ID                           NOT NULL NUMBER
 CUST_GENDER                       NOT NULL CHAR(1)
 CUST_YEAR_OF_BIRTH                NOT NULL NUMBER(4)

BEGIN
   dbms_data_mining_transform.CREATE_NORM_LIN(
```

```
       norm_table_name  => 'norm_tbl');
       dbms_data_mining_transform.INSERT_NORM_LIN_SCALE(
       norm_table_name  => 'norm_tbl',
       data_table_name  => 'mining_data',
       exclude_list     => dbms_data_mining_transform.COLUMN_LIST( 'cust_id'),
       round_num        => 3);
 END;
 /

SELECT col, shift, scale FROM norm_tbl;

COL                  SHIFT SCALE
-------------------- ----- -----
CUST_YEAR_OF_BIRTH       0  1990

DECLARE
    xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
    dbms_data_mining_transform.STACK_NORM_LIN (
        norm_table_name    => 'norm_tbl',
        xform_list        => xforms);
    dbms_data_mining_transform.XFORM_STACK (
        xform_list         => xforms,
        data_table_name    => 'mining_data',
        xform_view_name    => 'mining_data_stack');
END;
/

set long 3000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_GENDER",("CUST_YEAR_OF_BIRTH"-0)/1990 "CUST_YEAR_OF_BIRTH
" FROM mining_data
```

58.6.21 INSERT_NORM_LIN_ZSCORE Procedure

This procedure performs linear normalization and inserts the transformation definitions in a transformation definition table.

`INSERT_NORM_LIN_ZSCORE` computes the mean and the standard deviation from the data and sets the value of *`shift`* and *`scale`* as follows:

```
shift = mean
scale = stddev
```

Normalization is computed as:

```
x_new = (x_old - shift)/scale
```

`INSERT_NORM_LIN_ZSCORE` rounds the value of *`scale`* to a specified number of significant digits before storing it in the transformation definition table.

`INSERT_NORM_LIN_ZSCORE` normalizes all the `NUMBER` and `FLOAT` columns in the data unless you specify a list of columns to ignore.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.INSERT_NORM_LIN_ZSCORE (
     norm_table_name     IN VARCHAR2,
```

```
    data_table_name     IN VARCHAR2,
    exclude_list        IN COLUMN_LIST DEFAULT NULL,
    round_num           IN PLS_INTEGER DEFAULT 6,
    norm_schema_name    IN VARCHAR2 DEFAULT NULL,
    data_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-31 INSERT_NORM_LIN_ZSCORE Procedure Parameters

Parameter	Description

norm_table_name	Name of the transformation definition table for linear normalization. You can use the CREATE_NORM_LIN Procedure to create the definition table. The following columns are required:
	`COL VARCHAR2(30)`
	`SHIFT NUMBER`
	`SCALE NUMBER`
	CREATE_NORM_LIN creates an additional column, ATT, which may be used for specifying nested attributes. This column is not used by INSERT_NORM_LIN_ZSCORE.

data_table_name	Name of the table containing the data to be transformed
exclude_list	List of numerical columns to be excluded from normalization. If you do not specify *exclude_list*, all numerical columns are transformed. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
round_num	The number of significant digits to use for *scale*. The default number is 6.
norm_schema_name	Schema of *norm_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See *Oracle Machine Learning for SQL User's Guide* for details about numerical data.

Examples

In this example, INSERT_NORM_LIN_ZSCORE normalizes the cust_year_of_birth column and inserts the transformation in a transformation definition table. The STACK_NORM_LIN Procedure creates a transformation list from the contents of the definition table.

The SQL expression that computes the transformation is shown in the view MINING_DATA_STACK. The view is for display purposes only; it cannot be used to embed the transformations in a model.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, cust_gender, cust_year_of_birth
       FROM sh.customers;

DESCRIBE mining_data
 Name                              Null?    Type
 --------------------------------- -------- --------------------
 CUST_ID                           NOT NULL NUMBER
```

```
 CUST_GENDER                          NOT NULL CHAR(1)
 CUST_YEAR_OF_BIRTH                   NOT NULL NUMBER(4)

BEGIN
    dbms_data_mining_transform.CREATE_NORM_LIN(
       norm_table_name  => 'norm_tbl');
       dbms_data_mining_transform.INSERT_NORM_LIN_ZSCORE(
       norm_table_name  => 'norm_tbl',
       data_table_name  => 'mining_data',
       exclude_list     => dbms_data_mining_transform.COLUMN_LIST( 'cust_id'),
       round_num        => 3);
END;
/

SELECT col, shift, scale FROM norm_tbl;

COL                  SHIFT SCALE
-------------------- ----- -----
CUST_YEAR_OF_BIRTH    1960    15

DECLARE
    xforms      dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
    dbms_data_mining_transform.STACK_NORM_LIN (
        norm_table_name    => 'norm_tbl',
        xform_list        => xforms);
    dbms_data_mining_transform.XFORM_STACK (
        xform_list         => xforms,
        data_table_name    => 'mining_data',
        xform_view_name    => 'mining_data_stack');
END;
/

set long 3000
SQL> SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_STACK';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_GENDER",("CUST_YEAR_OF_BIRTH"-1960)/15 "CUST_YEAR_OF_BIRT
H" FROM mining_data
```

58.6.22 SET_EXPRESSION Procedure

This procedure appends a row to a `VARCHAR2` array that stores a SQL expression.

The array can be used for specifying a transformation expression that is too long to be used with the SET_TRANSFORM Procedure.

The GET_EXPRESSION Function returns a row in the array.

When you use `SET_EXPRESSION` to build a transformation expression, you must build a corresponding reverse transformation expression, create a transformation record, and add the transformation record to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.SET_EXPRESSION (
           expression    IN OUT NOCOPY EXPRESSION_REC,
           chunk                       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-32 SET_EXPRESSION Procedure Parameters

Parameter	Description
expression	An expression record (EXPRESSION_REC) that specifies a transformation expression or a reverse transformation expression for an attribute. Each expression record includes a VARCHAR2 array and index fields for specifying upper and lower boundaries within the array. There are two EXPRESSION_REC fields within a transformation record (TRANSFORM_REC): one for the transformation expression; the other for the reverse transformation expression. See Table 58-1 for a description of the EXPRESSION_REC type.
chunk	A VARCHAR2 chunk (row) to be appended to *expression*.

Notes

1. You can pass NULL in the *chunk* argument to SET_EXPRESSION to clear the previous chunk. The default value of *chunk* is NULL.
2. See "About Transformation Lists".
3. See "Operational Notes".

Examples

In this example, two calls to SET_EXPRESSION construct a transformation expression and two calls construct the reverse transformation.

> **Note:**
>
> This example is for illustration purposes only. It shows how SET_EXPRESSION appends the text provided in *chunk* to the text that already exists in *expression*. The SET_EXPRESSION procedure is meant for constructing very long transformation expressions that cannot be specified in a VARCHAR2 argument to SET_TRANSFORM.
>
> Similarly while transformation lists are intended for embedding in a model, the transformation list v_xlst is shown in an external view for illustration purposes.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, cust_year_of_birth, cust_postal_code, cust_credit_limit
       FROM sh.customers;

DECLARE
        v_expr dbms_data_mining_transform.EXPRESSION_REC;
        v_rexp dbms_data_mining_transform.EXPRESSION_REC;
        v_xrec dbms_data_mining_transform.TRANSFORM_REC;
        v_xlst dbms_data_mining_transform.TRANSFORM_LIST :=
                                dbms_data_mining_transform.TRANSFORM_LIST(NULL);
BEGIN
    dbms_data_mining_transform.SET_EXPRESSION(
         EXPRESSION  => v_expr,
         CHUNK       => '("CUST_YEAR_OF_BIRTH"-1910)');
    dbms_data_mining_transform.SET_EXPRESSION(
```

```
          EXPRESSION  => v_expr,
          CHUNK       => '/77');
    dbms_data_mining_transform.SET_EXPRESSION(
          EXPRESSION  => v_rexp,
          CHUNK       => '"CUST_YEAR_OF_BIRTH"*77');
    dbms_data_mining_transform.SET_EXPRESSION(
          EXPRESSION  => v_rexp,
          CHUNK       => '+1910');

    v_xrec := null;
    v_xrec.attribute_name := 'CUST_YEAR_OF_BIRTH';
    v_xrec.expression := v_expr;
    v_xrec.reverse_expression := v_rexp;
    v_xlst.TRIM;
    v_xlst.extend(1);
    v_xlst(1) := v_xrec;

    dbms_data_mining_transform.XFORM_STACK (
        xform_list           =>  v_xlst,
        data_table_name      => 'mining_data',
        xform_view_name      => 'v_xlst_view');

    dbms_output.put_line('====');
    FOR i IN 1..v_xlst.count LOOP
      dbms_output.put_line('ATTR: '||v_xlst(i).attribute_name);
      dbms_output.put_line('SUBN: '||v_xlst(i).attribute_subname);
      FOR j IN v_xlst(i).expression.lb..v_xlst(i).expression.ub LOOP
        dbms_output.put_line('EXPR: '||v_xlst(i).expression.lstmt(j));
      END LOOP;
      FOR j IN v_xlst(i).reverse_expression.lb..
               v_xlst(i).reverse_expression.ub LOOP
        dbms_output.put_line('REXP: '||v_xlst(i).reverse_expression.lstmt(j));
      END LOOP;
      dbms_output.put_line('====');
    END LOOP;
  END;
/
====
ATTR: CUST_YEAR_OF_BIRTH
SUBN:
EXPR: ("CUST_YEAR_OF_BIRTH"-1910)
EXPR: /77
REXP: "CUST_YEAR_OF_BIRTH"*77
REXP: +1910
====
```

58.6.23 SET_TRANSFORM Procedure

This procedure appends the transformation instructions for an attribute to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.SET_TRANSFORM (
          xform_list               IN OUT NOCOPY TRANSFORM_LIST,
          attribute_name           VARCHAR2,
          attribute_subname        VARCHAR2,
          expression               VARCHAR2,
          reverse_expression       VARCHAR2,
          attribute_spec           VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-33 SET_TRANSFORM Procedure Parameters

Parameter	Description
xform_list	A transformation list. See Table 58-1for a description of the TRANSFORM_LIST object type.
attribute_name	Name of the attribute to be transformed
attribute_subname	Name of the nested attribute if *attribute_name* is a nested column, otherwise NULL.
expression	A SQL expression that specifies the transformation of the attribute.
reverse_expression	A SQL expression that reverses the transformation for readability in model details and in the target of a supervised model (if the attribute is a target)
attribute_spec	One or more keywords that identify special treatment for the attribute during model build. Values are: • NOPREP — When ADP is on, prevents automatic transformation of the attribute. If ADP is not on, this value has no effect. • TEXT — Causes the attribute to be treated as unstructured text data • FORCE_IN — Forces the inclusion of the attribute in the model build. Applies only to GLM models with feature selection enabled (ftr_selection_enable = yes). Feature selection is disabled by default. If the model is not using GLM with feature selection, this value has no effect. See "Specifying Transformation Instructions for an Attribute" in *Oracle Machine Learning for SQL User's Guide*for more information about attribute_spec.

Usage Notes

1. See the following relevant sections in "Operational Notes":
   - About Transformation Lists
   - Nested Data Transformations
2. As shown in the following example, you can eliminate an attribute by specifying a null transformation expression and reverse expression. You can also use the STACK interface to remove a column (CREATE_COL_REM Procedure and STACK_COL_REM Procedure).

58.6.24 STACK_BIN_CAT Procedure

This procedure adds categorical binning transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_BIN_CAT (
     bin_table_name     IN            VARCHAR2,
     xform_list         IN OUT NOCOPY TRANSFORM_LIST,
     literal_flag       IN            BOOLEAN  DEFAULT FALSE,
     bin_schema_name    IN            VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-34 STACK_BIN_CAT Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table for categorical binning. You can use the CREATE_BIN_CAT Procedure to create the definition table. The table must be populated with transformation definitions before you call STACK_BIN_CAT. To populate the table, you can use one of the INSERT procedures for categorical binning or you can write your own SQL. See Table 58-4
xform_list	A transformation list. See Table 58-1 for a description of the TRANSFORM_LIST object type.
literal_flag	Indicates whether the values in the bin column in the transformation definition table are valid SQL literals. When *literal_flag* is FALSE (the default), the bin identifiers will be transformed to SQL literals by surrounding them with single quotes. Set *literal_flag* to TRUE if the bin identifiers are numbers that should have a numeric datatype, as is the case for an O-Cluster model. See "INSERT_BIN_NUM_EQWIDTH Procedure" for an example.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how a binning transformation for the categorical column cust_postal_code could be added to a stack called mining_data_stack.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. XFORM_STACK simply generates an external view of the transformed data. The actual purpose of the STACK procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to CREATE_MODEL in the xform_list parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
CREATE or REPLACE VIEW mining_data AS
    SELECT cust_id, cust_postal_code, cust_credit_limit
       FROM sh.customers
       WHERE cust_id BETWEEN 100050 AND 100100;
BEGIN
   dbms_data_mining_transform.CREATE_BIN_CAT ('bin_cat_tbl');
```

```
   dbms_data_mining_transform.INSERT_BIN_CAT_FREQ (
        bin_table_name   => 'bin_cat_tbl',
        data_table_name  => 'mining_data',
        bin_num          =>  3);
  END;
/
DECLARE
  MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
  dbms_data_mining_transform.STACK_BIN_CAT (
    bin_table_name       => 'bin_cat_tbl',
    xform_list           =>  mining_data_stack);
  dbms_data_mining_transform.XFORM_STACK (
    xform_list           =>  mining_data_stack,
    data_table_name      => 'mining_data',
    xform_view_name      => 'mining_data_stack_view');
  END;
/
-- Before transformation
column cust_postal_code format a16
SELECT * from mining_data
              WHERE cust_id BETWEEN 100050 AND 100053
              ORDER BY cust_id;

  CUST_ID CUST_POSTAL_CODE CUST_CREDIT_LIMIT
---------- ---------------- -----------------
    100050 76486                        1500
    100051 73216                        9000
    100052 69499                        5000
    100053 45704                        7000

-- After transformation
SELECT * FROM mining_data_stack_view
              WHERE cust_id BETWEEN 100050 AND 100053
              ORDER BY cust_id;

  CUST_ID CUST_POSTAL_CODE CUST_CREDIT_LIMIT
---------- ---------------- -----------------
    100050 4                            1500
    100051 1                            9000
    100052 4                            5000
    100053 4                            7000
```

58.6.25 STACK_BIN_NUM Procedure

This procedure adds numerical binning transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_BIN_NUM (
     bin_table_name     IN             VARCHAR2,
     xform_list         IN OUT  NOCOPY TRANSFORM_LIST,
     literal_flag       IN             BOOLEAN  DEFAULT FALSE,
     bin_schema_name    IN             VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-35 STACK_BIN_NUM Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table for numerical binning. You can use the CREATE_BIN_NUM Procedure to create the definition table. The table must be populated with transformation definitions before you call STACK_BIN_NUM. To populate the table, you can use one of the INSERT procedures for numerical binning or you can write your own SQL. See Table 58-6.
xform_list	A transformation list. See Table 58-1 for a description of the TRANSFORM_LIST object type.
literal_flag	Indicates whether the values in the bin column in the transformation definition table are valid SQL literals. When *literal_flag* is FALSE (the default), the bin identifiers will be transformed to SQL literals by surrounding them with single quotes. Set *literal_flag* to TRUE if the bin identifiers are numbers that should have a numeric datatype, as is the case for an O-Cluster model. See "INSERT_BIN_NUM_EQWIDTH Procedure" for an example.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how a binning transformation for the numerical column cust_credit_limit could be added to a stack called mining_data_stack.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. XFORM_STACK simply generates an external view of the transformed data. The actual purpose of the STACK procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to CREATE_MODEL in the xform_list parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
CREATE OR REPLACE VIEW mining_data AS
    SELECT cust_id, cust_postal_code, cust_credit_limit
      FROM sh.customers
      WHERE cust_id BETWEEN 100050 and 100100;
BEGIN
  dbms_data_mining_transform.create_bin_num ('bin_num_tbl');
```

```
  dbms_data_mining_transform.insert_bin_num_qtile (
  bin_table_name    => 'bin_num_tbl',
  data_table_name   => 'mining_data',
  bin_num              => 5,
  exclude_list      =>  dbms_data_mining_transform.COLUMN_LIST('cust_id'));
END;
/
DECLARE
   MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
   dbms_data_mining_transform.STACK_BIN_CAT (
      bin_table_name        => 'bin_num_tbl',
      xform_list            =>  mining_data_stack);
   dbms_data_mining_transform.XFORM_STACK (
      xform_list            =>  mining_data_stack,
      data_table_name   => 'mining_data',
      xform_view_name   => 'mining_data_stack_view');
END;
/
-- Before transformation
SELECT cust_id, cust_postal_code, ROUND(cust_credit_limit) FROM mining_data
   WHERE cust_id BETWEEN 100050 AND 100055
   ORDER BY cust_id;
CUST_ID   CUST_POSTAL_CODE   ROUND(CUST_CREDIT_LIMIT)
-------   -----------------  -------------------------
100050    76486                                 1500
100051    73216                                 9000
100052    69499                                 5000
100053    45704                                 7000
100055    74673                                11000
100055    74673                                11000

-- After transformation
SELECT cust_id, cust_postal_code, ROUND(cust_credit_limit)
   FROM mining_data_stack_view
   WHERE cust_id BETWEEN 100050 AND 100055
   ORDER BY cust_id;
CUST_ID   CUST_POSTAL_CODE   ROUND(CUST_CREDIT_LIMITT)
-------   ----------------   -------------------------
100050    76486
100051    73216                                    2
100052    69499                                    1
100053    45704
100054    88021                                    3
100055    74673                                    3
```

58.6.26 STACK_CLIP Procedure

This procedure adds clipping transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_CLIP (
          clip_table_name     IN            VARCHAR2,
          xform_list          IN OUT NOCOPY TRANSFORM_LIST,
          clip_schema_name    IN            VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-36 STACK_CLIP Procedure Parameters

Parameter	Description
clip_table_name	Name of the transformation definition table for clipping.You can use the CREATE_CLIP Procedure to create the definition table. The table must be populated with transformation definitions before you call STACK_CLIP. To populate the table, you can use one of the INSERT procedures for clipping or you can write your own SQL. See Table 58-8
xform_list	A transformation list. See Table 58-1 for a description of the TRANSFORM_LIST object type.
clip_schema_name	Schema of *clip_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See DBMS_DATA_MINING_TRANSFORM Operational Notes. The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how a clipping transformation for the numerical column cust_credit_limit could be added to a stack called mining_data_stack.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. XFORM_STACK simply generates an external view of the transformed data. The actual purpose of the STACK procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to CREATE_MODEL in the xform_list parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, cust_postal_code, cust_credit_limit
       FROM sh.customers
       WHERE cust_id BETWEEN 100050 AND 100100;
BEGIN
   dbms_data_mining_transform.create_clip ('clip_tbl');
   dbms_data_mining_transform.insert_clip_winsor_tail (
       clip_table_name   => 'clip_tbl',
       data_table_name   => 'mining_data',
       tail_frac         => 0.25,
       exclude_list      => dbms_data_mining_transform.COLUMN_LIST('cust_id'));
END;
/
```

```
DECLARE
      MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
      dbms_data_mining_transform.STACK_CLIP (
         clip_table_name     => 'clip_tbl',
         xform_list          =>  mining_data_stack);
      dbms_data_mining_transform.XFORM_STACK (
          xform_list        =>  mining_data_stack,
          data_table_name   => 'mining_data',
          xform_view_name   => 'mining_data_stack_view');
END;
/
-- Before transformation
SELECT cust_id, cust_postal_code, round(cust_credit_limit)
  FROM mining_data
    WHERE cust_id BETWEEN 100050 AND 100054
    ORDER BY cust_id;

CUST_ID   CUST_POSTAL_CODE   ROUND(CUST_CREDIT_LIMIT)
-------   ----------------   ------------------------
100050    76486                                 1500
100051    73216                                 9000
100052    69499                                 5000
100053    45704                                 7000
100054    88021                                11000

-- After transformation
SELECT cust_id, cust_postal_code, round(cust_credit_limit)
  FROM mining_data_stack_view
    WHERE cust_id BETWEEN 100050 AND 100054
    ORDER BY cust_id;

CUST_ID   CUST_POSTAL_CODE   ROUND(CUST_CREDIT_LIMIT)
-------   ----------------   ------------------------
100050    76486                                 5000
100051    73216                                 9000
100052    69499                                 5000
100053    45704                                 7000
100054    88021                                11000
```

58.6.27 STACK_COL_REM Procedure

This procedure adds column removal transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_COL_REM (
     rem_table_name     IN            VARCHAR2,
     xform_list         IN OUT NOCOPY TRANSFORM_LIST,
     rem_schema_name    IN            VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-37 STACK_COL_REM Procedure Parameters

Parameter	Description
rem_table_name	Name of the transformation definition table for column removal. You can use the CREATE_COL_REM Procedure to create the definition table. See Table 58-10. The table must be populated with column names before you call STACK_COL_REM. The INSERT_BIN_SUPER Procedure and the INSERT_AUTOBIN_NUM_EQWIDTH Procedure can optionally be used to populate the table. You can also use SQL INSERT statements.
xform_list	A transformation list. See Table 58-1 for a description of the TRANSFORM_LIST object type.
rem_schema_name	Schema of *rem_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how the column cust_credit_limit could be removed in a transformation list called mining_data_stack.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. XFORM_STACK simply generates an external view of the transformed data. The actual purpose of the STACK procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to CREATE_MODEL in the xform_list parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
CREATE OR REPLACE VIEW mining_data AS
     SELECT cust_id, country_id, cust_postal_code, cust_credit_limit
        FROM sh.customers;

BEGIN
    dbms_data_mining_transform.create_col_rem ('rem_tbl');
END;
/

INSERT into rem_tbl VALUES (upper('cust_postal_code'), null);

DECLARE
  MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
```

```
     dbms_data_mining_transform.stack_col_rem (
          rem_table_name     => 'rem_tbl',
          xform_list         =>  mining_data_stack);
      dbms_data_mining_transform.XFORM_STACK (
          xform_list         =>  mining_data_stack,
          data_table_name    => 'mining_data',
          xform_view_name    => 'mining_data_stack_view');
END;
/

SELECT *  FROM mining_data
  WHERE cust_id BETWEEN 100050 AND 100051
  ORDER BY cust_id;

CUST_ID   COUNTRY_ID   CUST_POSTAL_CODE   CUST_CREDIT_LIMIT
-------   ----------   ----------------   -----------------
100050         52773   76486                          1500
100051         52790   73216                          9000

SELECT *  FROM mining_data_stack_view
  WHERE cust_id BETWEEN 100050 AND 100051
  ORDER BY cust_id;

CUST_ID   COUNTRY_ID   CUST_CREDIT_LIMIT
-------   ----------   -----------------
100050         52773              1500
100051         52790              9000
```

58.6.28 STACK_MISS_CAT Procedure

This procedure adds categorical missing value transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_MISS_CAT (
     miss_table_name     IN       VARCHAR2,
     xform_list          IN OUT   NOCOPY TRANSFORM_LIST,
     miss_schema_name    IN       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-38 STACK_MISS_CAT Procedure Parameters

Parameter	Description
`miss_table_name`	Name of the transformation definition table for categorical missing value treatment. You can use the CREATE_MISS_CAT Procedure to create the definition table. The table must be populated with transformation definitions before you call `STACK_MISS_CAT`. To populate the table, you can use the INSERT_MISS_CAT_MODE Procedure or you can write your own SQL. See Table 58-12.
`xform_list`	A transformation list. See Table 58-1 for a description of the `TRANSFORM_LIST` object type.
`miss_schema_name`	Schema of *`miss_table_name`*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how the missing values in the column `cust_marital_status` could be replaced with the mode in a transformation list called `mining_data_stack`.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. `XFORM_STACK` simply generates an external view of the transformed data. The actual purpose of the `STACK` procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to `CREATE_MODEL` in the `xform_list` parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
CREATE OR REPLACE VIEW mining_data AS
      SELECT cust_id, country_id, cust_marital_status
         FROM sh.customers
         where cust_id BETWEEN 1 AND 10;

BEGIN
  dbms_data_mining_transform.create_miss_cat ('miss_cat_tbl');
  dbms_data_mining_transform.insert_miss_cat_mode ('miss_cat_tbl', 'mining_data');
END;
/

DECLARE
  MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
     dbms_data_mining_transform.stack_miss_cat (
          miss_table_name   => 'miss_cat_tbl',
          xform_list        => mining_data_stack);
      dbms_data_mining_transform.XFORM_STACK (
          xform_list        =>  mining_data_stack,
          data_table_name   => 'mining_data',
          xform_view_name   => 'mining_data_stack_view');
END;
/
SELECT * FROM mining_data
  ORDER BY cust_id;

CUST_ID   COUNTRY_ID   CUST_MARITAL_STATUS
-------   ----------   --------------------
      1        52789
      2        52778
      3        52770
      4        52770
      5        52789
      6        52769     single
      7        52790     single
      8        52790     married
      9        52770     divorced
     10        52790     widow
```

```
SELECT * FROM mining_data_stack_view
   ORDER By cust_id;

CUST_ID   COUNTRY_ID   CUST_MARITAL_STATUS
-------   ----------  --------------------
      1        52789   single
      2        52778   single
      3        52770   single
      4        52770   single
      5        52789   single
      6        52769   single
      7        52790   single
      8        52790   married
      9        52770   divorced
     10        52790   widow
```

58.6.29 STACK_MISS_NUM Procedure

This procedure adds numeric missing value transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_MISS_NUM (
     miss_table_name     IN       VARCHAR2,
     xform_list          IN OUT   NOCOPY TRANSFORM_LIST,
     miss_schema_name    IN       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-39 STACK_MISS_NUM Procedure Parameters

Parameter	Description
`miss_table_name`	Name of the transformation definition table for numerical missing value treatment. You can use the CREATE_MISS_NUM Procedure to create the definition table. The table must be populated with transformation definitions before you call `STACK_MISS_NUM`. To populate the table, you can use the INSERT_MISS_NUM_MEAN Procedure or you can write your own SQL. See Table 58-14.
`xform_list`	A transformation list. See Table 58-1 for a description of the `TRANSFORM_LIST` object type.
`miss_schema_name`	Schema of *`miss_table_name`*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how the missing values in the column `cust_credit_limit` could be replaced with the mean in a transformation list called `mining_data_stack`.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. `XFORM_STACK` simply generates an external view of the transformed data. The actual purpose of the `STACK` procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to `CREATE_MODEL` in the `xform_list` parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
describe mining_data
 Name                                                      Null?    Type
 --------------------------------------------------------- -------- -----
 CUST_ID                                                   NOT NULL NUMBER
 CUST_CREDIT_LIMIT                                                  NUMBER

BEGIN
   dbms_data_mining_transform.create_miss_num ('miss_num_tbl');
   dbms_data_mining_transform.insert_miss_num_mean ('miss_num_tbl','mining_data');
END;
/
SELECT * FROM miss_num_tbl;

COL                  ATT      VAL
-------------------- ----- ------
CUST_ID                        5.5
CUST_CREDIT_LIMIT          185.71

DECLARE
    MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
  BEGIN
    dbms_data_mining_transform.STACK_MISS_NUM (
          miss_table_name   => 'miss_num_tbl',
          xform_list        => mining_data_stack);
    dbms_data_mining_transform.XFORM_STACK (
          xform_list        =>  mining_data_stack,
          data_table_name   => 'mining_data',
          xform_view_name   => 'mining_data_stack_view');
END;
/
-- Before transformation
SELECT * FROM mining_data
  ORDER BY cust_id;
CUST_ID CUST_CREDIT_LIMIT
------- -----------------
      1               100
      2
      3               200
      4
      5               150
      6               400
      7               150
      8
      9               100
     10               200

-- After transformation
SELECT * FROM mining_data_stack_view
  ORDER BY cust_id;
```

```
CUST_ID CUST_CREDIT_LIMIT
------- -----------------
      1               100
      2            185.71
      3               200
      4            185.71
      5               150
      6               400
      7               150
      8            185.71
      9               100
     10               200
```

58.6.30 STACK_NORM_LIN Procedure

This procedure adds linear normalization transformations to a transformation list.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.STACK_NORM_LIN (
     norm_table_name     IN       VARCHAR2,
     xform_list          IN OUT   NOCOPY TRANSFORM_LIST,
     norm_schema_name    IN       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-40 STACK_NORM_LIN Procedure Parameters

Parameter	Description
norm_table_name	Name of the transformation definition table for linear normalization. You can use the CREATE_NORM_LIN Procedure to create the definition table. The table must be populated with transformation definitions before you call STACK_NORM_LIN.To populate the table, you can use one of the INSERT procedures for normalization or you can write your own SQL. See Table 58-16.
xform_list	A transformation list. See Table 58-1 for a description of the TRANSFORM_LIST object type.
norm_schema_name	Schema of *norm_table_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example shows how the column cust_credit_limit could be normalized in a transformation list called mining_data_stack.

> **Note:**
>
> This example invokes the XFORM_STACK Procedure to show how the data is transformed by the stack. `XFORM_STACK` simply generates an external view of the transformed data. The actual purpose of the `STACK` procedures is to assemble a list of transformations for embedding in a model. The transformations are passed to `CREATE_MODEL` in the `xform_list` parameter. See INSERT_BIN_NUM_EQWIDTH Procedure for an example.

```
CREATE OR REPLACE VIEW mining_data AS
       SELECT cust_id, country_id, cust_postal_code, cust_credit_limit
          FROM sh.customers;
BEGIN
   dbms_data_mining_transform.create_norm_lin ('norm_lin_tbl');
   dbms_data_mining_transform.insert_norm_lin_minmax (
       norm_table_name   => 'norm_lin_tbl',
       data_table_name   => 'mining_data',
       exclude_list      =>  dbms_data_mining_transform.COLUMN_LIST('cust_id',
                                                                    'country_id'));
END;
/
SELECT * FROM norm_lin_tbl;
COL                  ATT    SHIFT  SCALE
-------------------- ----- ------ ------
CUST_CREDIT_LIMIT            1500  13500

DECLARE
   MINING_DATA_STACK   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
   dbms_data_mining_transform.stack_norm_lin (
        norm_table_name   => 'norm_lin_tbl',
        xform_list        => mining_data_stack);
   dbms_data_mining_transform.XFORM_STACK (
        xform_list        =>  mining_data_stack,
        data_table_name   => 'mining_data',
        xform_view_name   => 'mining_data_stack_view');
END;
/
SELECT * FROM mining_data
  WHERE cust_id between 1 and 10
  ORDER BY cust_id;
CUST_ID COUNTRY_ID CUST_POSTAL_CODE     CUST_CREDIT_LIMIT
------- ---------- -------------------- -----------------
      1      52789 30828                             9000
      2      52778 86319                            10000
      3      52770 88666                             1500
      4      52770 87551                             1500
      5      52789 59200                             1500
      6      52769 77287                             1500
      7      52790 38763                             1500
      8      52790 58488                             3000
      9      52770 63033                             3000
     10      52790 52602                             3000

SELECT * FROM mining_data_stack_view
  WHERE cust_id between 1 and 10
  ORDER BY cust_id;
CUST_ID COUNTRY_ID CUST_POSTAL_CODE     CUST_CREDIT_LIMIT
```

```
------- ---------- ------------------- -----------------
      1      52789 30828                           .55556
      2      52778 86319                           .62963
      3      52770 88666                                0
      4      52770 87551                                0
      5      52789 59200                                0
      6      52769 77287                                0
      7      52790 38763                                0
      8      52790 58488                           .11111
      9      52770 63033                           .11111
     10      52790 52602                           .11111
```

58.6.31 XFORM_BIN_CAT Procedure

This procedure creates a view that implements the categorical binning transformations specified in a definition table. Only the columns that are specified in the definition table are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_BIN_CAT (
     bin_table_name       IN VARCHAR2,
     data_table_name      IN VARCHAR2,
     xform_view_name      IN VARCHAR2,
     literal_flag         IN BOOLEAN DEFAULT FALSE,
     bin_schema_name      IN VARCHAR2 DEFAULT NULL,
     data_schema_name     IN VARCHAR2 DEFAULT NULL,
     xform_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-41 XFORM_BIN_CAT Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table for categorical binning. You can use the CREATE_BIN_CAT Procedure to create the definition table. The table must be populated with transformation definitions before you call XFORM_BIN_CAT. To populate the table, you can use one of the INSERT procedures for categorical binning or you can write your own SQL. See Table 58-4.
data_table_name	Name of the table containing the data to be transformed.
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *bin_table_name*.
literal_flag	Indicates whether the values in the bin column in the transformation definition table are valid SQL literals. When *literal_flag* is FALSE (the default), the bin identifiers will be transformed to SQL literals by surrounding them with single quotes. Set *literal_flag* to TRUE if the bin identifiers are numbers that should have a numeric datatype, as is the case for an O-Cluster model. See "INSERT_BIN_NUM_EQWIDTH Procedure" for an example.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Table 58-41 (Cont.) XFORM_BIN_CAT Procedure Parameters

Parameter	Description
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes".

Examples

This example creates a view that bins the cust_postal_code column. The data source consists of three columns from sh.customer.

```
describe mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ------------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_POSTAL_CODE                          NOT NULL VARCHAR2(10)
 CUST_CREDIT_LIMIT                                  NUMBER

SELECT * FROM mining_data WHERE cust_id between 104066 and 104069;

   CUST_ID CUST_POSTAL_CODE
CUST_CREDIT_LIMIT
--------- --------------------
-----------------
    104066 69776
7000
    104067 52602
9000
    104068 55787
11000
    104069 55977
5000

BEGIN
  dbms_data_mining_transform.create_bin_cat(
     bin_table_name     => 'bin_cat_tbl');
  dbms_data_mining_transform.insert_bin_cat_freq(
     bin_table_name     => 'bin_cat_tbl',
     data_table_name    => 'mining_data',
     bin_num          => 10);
   dbms_data_mining_transform.xform_bin_cat(
     bin_table_name     => 'bin_cat_tbl',
     data_table_name    => 'mining_data',
     xform_view_name    => 'bin_cat_view');
END;
/

SELECT * FROM bin_cat_view WHERE cust_id between 104066 and 104069;

   CUST_ID CUST_POSTAL_CODE
CUST_CREDIT_LIMIT
---------- --------------------
-----------------
    104066 6
7000
```

```
    104067 11
9000
    104068 3
11000
    104069 11
5000

SELECT text FROM user_views WHERE view_name IN 'BIN_CAT_VIEW';

TEXT

--------------------------------------------------------------------------------

SELECT
"CUST_ID",DECODE("CUST_POSTAL_CODE",'38082','1','45704','9','48346','5','

55787','3','63736','2','67843','7','69776','6','72860','10','78558','4','80841',

'8',NULL,NULL,'11') "CUST_POSTAL_CODE","CUST_CREDIT_LIMIT" FROM
mining_data
```

58.6.32 XFORM_BIN_NUM Procedure

This procedure creates a view that implements the numerical binning transformations specified in a definition table. Only the columns that are specified in the definition table are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_BIN_NUM (
     bin_table_name     IN VARCHAR2,
     data_table_name    IN VARCHAR2,
     xform_view_name    IN VARCHAR2,
     literal_flag       IN BOOLEAN DEFAULT FALSE,
     bin_schema_name    IN VARCHAR2 DEFAULT NULL,
     data_schema_name   IN VARCHAR2 DEFAULT NULL,
     xform_schema_name  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-42 XFORM_BIN_NUM Procedure Parameters

Parameter	Description
bin_table_name	Name of the transformation definition table for numerical binning. You can use the CREATE_BIN_NUM Procedure to create the definition table. The table must be populated with transformation definitions before you call XFORM_BIN_NUM. To populate the table, you can use one of the INSERT procedures for numerical binning or you can write your own SQL. See "Table 58-6".
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *bin_table_name*.

Table 58-42 (Cont.) XFORM_BIN_NUM Procedure Parameters

Parameter	Description
literal_flag	Indicates whether the values in the bin column in the transformation definition table are valid SQL literals. When *literal_flag* is FALSE (the default), the bin identifiers will be transformed to SQL literals by surrounding them with single quotes. Set *literal_flag* to TRUE if the bin identifiers are numbers that should have a numeric datatype, as is the case for an O-Cluster model. See "INSERT_BIN_NUM_EQWIDTH Procedure" for an example.
bin_schema_name	Schema of *bin_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes".

Examples

This example creates a view that bins the cust_credit_limit column. The data source consists of three columns from sh.customer.

```
describe mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_POSTAL_CODE                          NOT NULL VARCHAR2(10)
 CUST_CREDIT_LIMIT                                  NUMBER

column cust_credit_limit off
SELECT * FROM mining_data WHERE cust_id between 104066 and 104069;

   CUST_ID CUST_POSTAL_CODE   CUST_CREDIT_LIMIT
---------  ------------------
--------------------
    104066 69776                          7000
    104067 52602                          9000
    104068 55787                         11000
    104069 55977                          5000

BEGIN
   dbms_data_mining_transform.create_bin_num(
           bin_table_name     => 'bin_num_tbl');
   dbms_data_mining_transform.insert_autobin_num_eqwidth(
           bin_table_name     => 'bin_num_tbl',
           data_table_name    => 'mining_data',
           bin_num                => 5,
           max_bin_num            => 10,
           exclude_list       => dbms_data_mining_transform.COLUMN_LIST('cust_id'));
  dbms_data_mining_transform.xform_bin_num(
          bin_table_name      => 'bin_num_tbl',
          data_table_name     => 'mining_data',
          xform_view_name     => 'mining_data_view');
```

```
END;
/
describe mining_data_view
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_POSTAL_CODE                          NOT NULL VARCHAR2(10)
 CUST_CREDIT_LIMIT                                  VARCHAR2(2)

col cust_credit_limit on
col cust_credit_limit format a25
SELECT * FROM mining_data_view WHERE cust_id between 104066 and 104069;

   CUST_ID CUST_POSTAL_CODE
CUST_CREDIT_LIMIT
---------- --------------------
-------------------------
    104066 69776
5
    104067 52602
6
    104068 55787
8
    104069 55977
3

set long 2000
SELECT text FROM user_views WHERE view_name IN 'MINING_DATA_VIEW';

TEXT

--------------------------------------------------------------------------------

SELECT "CUST_ID","CUST_POSTAL_CODE",CASE WHEN "CUST_CREDIT_LIMIT"<1500 THEN
NULL
 WHEN "CUST_CREDIT_LIMIT"<=2850 THEN '1' WHEN "CUST_CREDIT_LIMIT"<=4200 THEN
'2'
 WHEN "CUST_CREDIT_LIMIT"<=5550 THEN '3' WHEN "CUST_CREDIT_LIMIT"<=6900 THEN
'4'
 WHEN "CUST_CREDIT_LIMIT"<=8250 THEN '5' WHEN "CUST_CREDIT_LIMIT"<=9600 THEN
'6'
 WHEN "CUST_CREDIT_LIMIT"<=10950 THEN '7' WHEN "CUST_CREDIT_LIMIT"<=12300 THEN
'
8' WHEN "CUST_CREDIT_LIMIT"<=13650 THEN '9' WHEN "CUST_CREDIT_LIMIT"<=15000
THEN
 '10' END "CUST_CREDIT_LIMIT" FROM mining_data
```

58.6.33 XFORM_CLIP Procedure

This procedure creates a view that implements the clipping transformations specified in a definition table. Only the columns that are specified in the definition table are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_CLIP (
    clip_table_name       IN VARCHAR2,
    data_table_name       IN VARCHAR2,
    xform_view_name       IN VARCHAR2,
    clip_schema_name      IN VARCHAR2 DEFAULT NULL,
```

```
    data_schema_name      IN VARCHAR2,DEFAULT NULL,
    xform_schema_name     IN VARCHAR2,DEFAULT NULL);
```

Parameters

Table 58-43 XFORM_CLIP Procedure Parameters

Parameter	Description
clip_table_name	Name of the transformation definition table for clipping. You can use the CREATE_CLIP Procedure to create the definition table. The table must be populated with transformation definitions before you call XFORM_CLIP. To populate the table, you can use one of the INSERT procedures for clipping you can write your own SQL. See Table 58-8.
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *clip_table_name*.
clip_schema_name	Schema of *clip_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Examples

This example creates a view that clips the cust_credit_limit column. The data source consists of three columns from sh.customer.

```
describe mining_data
 Name                           Null?    Type
 ------------------------------ -------- ------------------------
 CUST_ID                        NOT NULL NUMBER
 CUST_POSTAL_CODE               NOT NULL VARCHAR2(10)
 CUST_CREDIT_LIMIT                       NUMBER

BEGIN
   dbms_data_mining_transform.create_clip(
      clip_table_name    => 'clip_tbl');
   dbms_data_mining_transform.insert_clip_trim_tail(
      clip_table_name   => 'clip_tbl',
      data_table_name   => 'mining_data',
      tail_frac         => 0.05,
      exclude_list      => dbms_data_mining_transform.COLUMN_LIST('cust_id'));
   dbms_data_mining_transform.xform_clip(
      clip_table_name     => 'clip_tbl',
      data_table_name    => 'mining_data',
      xform_view_name    => 'clip_view');
END;
/
describe clip_view
 Name                           Null?    Type
 ------------------------------ -------- ------------------------
 CUST_ID                        NOT NULL NUMBER
 CUST_POSTAL_CODE               NOT NULL VARCHAR2(10)
 CUST_CREDIT_LIMIT                       NUMBER
```

```
SELECT MIN(cust_credit_limit), MAX(cust_credit_limit) FROM mining_data;

MIN(CUST_CREDIT_LIMIT) MAX(CUST_CREDIT_LIMIT)
---------------------- ----------------------
                  1500                  15000

SELECT MIN(cust_credit_limit), MAX(cust_credit_limit) FROM clip_view;

MIN(CUST_CREDIT_LIMIT) MAX(CUST_CREDIT_LIMIT)
---------------------- ----------------------
                  1500                  11000

set long 2000
SELECT text FROM user_views WHERE view_name IN 'CLIP_VIEW';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID","CUST_POSTAL_CODE",CASE WHEN "CUST_CREDIT_LIMIT" < 1500 THEN NU
LL WHEN "CUST_CREDIT_LIMIT" > 11000 THEN NULL ELSE "CUST_CREDIT_LIMIT" END "CUST
_CREDIT_LIMIT" FROM mining_data
```

58.6.34 XFORM_COL_REM Procedure

This procedure creates a view that implements the column removal transformations specified in a definition table. Only the columns that are specified in the definition table are removed; the remaining columns from the data table are present in the view.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_COL_REM (
     rem_table_name     IN       VARCHAR2,
     data_table_name    IN       VARCHAR2,
     xform_view_name    IN       VARCHAR2,
     rem_schema_name    IN       VARCHAR2 DEFAULT NULL,
     data_schema_name   IN       VARCHAR2 DEFAULT NULL,
     xform_schema_name  IN       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-44 XFORM_COL_REM Procedure Parameters

Parameter	Description
rem_table_name	Name of the transformation definition table for column removal. You can use the CREATE_COL_REM Procedure to create the definition table. See Table 58-10. The table must be populated with column names before you call XFORM_COL_REM. The INSERT_BIN_SUPER Procedure and the INSERT_AUTOBIN_NUM_EQWIDTH Procedure can optionally be used to populate the table. You can also use SQL INSERT statements.
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents the columns in *data_table_name* that are not specified in *rem_table_name*.
rem_schema_name	Schema of *rem_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.

Table 58-44 (Cont.) XFORM_COL_REM Procedure Parameters

Parameter	Description
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes".

Examples

This example creates a view that includes all but one column from the table `customers` in the current schema.

```
describe customers
 Name                                    Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                 NOT NULL NUMBER
 CUST_MARITAL_STATUS                              VARCHAR2(20)
 OCCUPATION                                       VARCHAR2(21)
 AGE                                              NUMBER
 YRS_RESIDENCE                                    NUMBER

BEGIN
    DBMS_DATA_MINING_TRANSFORM.CREATE_COL_REM ('colrem_xtbl');
END;
 /
INSERT INTO colrem_xtbl VALUES('CUST_MARITAL_STATUS', null);

NOTE: This currently doesn't work. See bug 9310319

BEGIN
   DBMS_DATA_MINING_TRANSFORM.XFORM_COL_REM (
     rem_table_name        => 'colrem_xtbl',
     data_table_name       => 'customers',
     xform_view_name       => 'colrem_view');
END;
/
describe colrem_view

 Name                                    Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                 NOT NULL NUMBER
 OCCUPATION                                       VARCHAR2(21)
 AGE                                              NUMBER
 YRS_RESIDENCE                                    NUMBER
```

58.6.35 XFORM_EXPR_NUM Procedure

This procedure creates a view that implements the specified numeric transformations. Only the columns that you specify are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_EXPR_NUM (
     expr_pattern       IN       VARCHAR2,
```

```
data_table_name    IN       VARCHAR2,
xform_view_name    IN       VARCHAR2,
exclude_list       IN       COLUMN_LIST DEFAULT NULL,
include_list       IN       COLUMN_LIST DEFAULT NULL,
col_pattern        IN       VARCHAR2 DEFAULT ':col',
data_schema_name   IN       VARCHAR2 DEFAULT NULL,
xform_schema_name  IN       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-45 XFORM_EXPR_NUM Procedure Parameters

Parameter	Description
expr_pattern	A numeric transformation expression
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *expr_pattern* and *col_pattern*.
exclude_list	List of numerical columns to exclude. If NULL, no numerical columns are excluded. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
include_list	List of numeric columns to include. If NULL, all numeric columns are included. The format of *include_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
col_pattern	The value within *expr_pattern* that will be replaced with a column name. The value of *col_pattern* is case-sensitive. The default value of *col_pattern* is ':col'
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

1. The XFORM_EXPR_NUM procedure constructs numeric transformation expressions from the specified expression pattern (*expr_pattern*) by replacing every occurrence of the specified column pattern (*col_pattern*) with an actual column name.

   XFORM_EXPR_NUM uses the SQL REPLACE function to construct the transformation expressions.

   ```
   REPLACE (expr_pattern,col_pattern,'"column_name"') || '"column_name"'
   ```

   If there is a column match, then the replacement is made in the transformation expression; if there is not a match, then the column is used without transformation.

> **See:**
>
> *Oracle Database SQL Language Reference* for information about the `REPLACE` function

2. Because of the include and exclude list parameters, the `XFORM_EXPR_NUM` and `XFORM_EXPR_STR` procedures allow you to easily specify individual columns for transformation within large data sets. The other `XFORM_*` procedures support an exclude list only. In these procedures, you must enumerate every column that you do not want to transform.
3. See "Operational Notes"

Examples

This example creates a view that transforms the datatype of numeric columns.

```
describe customers
 Name                              Null?    Type
 --------------------------------- -------- ------------------------
 CUST_ID                           NOT NULL NUMBER
 CUST_MARITAL_STATUS                        VARCHAR2(20)
 OCCUPATION                                 VARCHAR2(21)
 AGE                                        NUMBER
 YRS_RESIDENCE                              NUMBER

BEGIN
  DBMS_DATA_MINING_TRANSFORM.XFORM_EXPR_NUM(
    expr_pattern         => 'to_char(:col)',
    data_table_name      => 'customers',
    xform_view_name      => 'cust_nonum_view',
    exclude_list         => dbms_data_mining_transform.COLUMN_LIST( 'cust_id'),
    include_list         => null,
    col_pattern          => ':col');
END;
/
describe cust_nonum_view
 Name                              Null?    Type
 --------------------------------- -------- ------------------------
 CUST_ID                           NOT NULL NUMBER
 CUST_MARITAL_STATUS                        VARCHAR2(20)
 OCCUPATION                                 VARCHAR2(21)
 AGE                                        VARCHAR2(40)
 YRS_RESIDENCE                              VARCHAR2(40)
```

58.6.36 XFORM_EXPR_STR Procedure

This procedure creates a view that implements the specified categorical transformations. Only the columns that you specify are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_EXPR_STR (
     expr_pattern       IN       VARCHAR2,
     data_table_name    IN       VARCHAR2,
     xform_view_name    IN       VARCHAR2,
     exclude_list       IN       COLUMN_LIST DEFAULT NULL,
     include_list       IN       COLUMN_LIST DEFAULT NULL,
```

```
    col_pattern        IN       VARCHAR2 DEFAULT ':col',
    data_schema_name   IN       VARCHAR2 DEFAULT NULL,
    xform_schema_name  IN       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-46 XFORM_EXPR_STR Procedure Parameters

Parameter	Description
expr_pattern	A character transformation expression
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *expr_pattern* and *col_pattern*.
exclude_list	List of categorical columns to exclude. If NULL, no categorical columns are excluded. The format of *exclude_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
include_list	List of character columns to include. If NULL, all character columns are included. The format of *include_list* is: `dbms_data_mining_transform.COLUMN_LIST('col1','col2', ...'coln')`
col_pattern	The value within *expr_pattern* that will be replaced with a column name. The value of *col_pattern* is case-sensitive. The default value of *col_pattern* is ':col'
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

1. The XFORM_EXPR_STR procedure constructs character transformation expressions from the specified expression pattern (*expr_pattern*) by replacing every occurrence of the specified column pattern (*col_pattern*) with an actual column name.

   XFORM_EXPR_STR uses the SQL REPLACE function to construct the transformation expressions.

   ```
   REPLACE (expr_pattern,col_pattern,'"column_name"') || '"column_name"'
   ```

   If there is a column match, then the replacement is made in the transformation expression; if there is not a match, then the column is used without transformation.

> **See:**
>
> *Oracle Database SQL Language Reference* for information about the `REPLACE` function

2. Because of the include and exclude list parameters, the `XFORM_EXPR_STR` and `XFORM_EXPR_NUM` procedures allow you to easily specify individual columns for transformation within large data sets. The other `XFORM_*` procedures support an exclude list only. In these procedures, you must enumerate every column that you do not want to transform.
3. See "Operational Notes"

Examples

This example creates a view that transforms character columns to upper case.

```
describe customers
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_MARITAL_STATUS                                VARCHAR2(20)
 OCCUPATION                                         VARCHAR2(21)
 AGE                                                NUMBER
 YRS_RESIDENCE                                      NUMBER

SELECT cust_id,  cust_marital_status, occupation FROM customers
    WHERE   cust_id > 102995
    ORDER BY cust_id desc;

CUST_ID CUST_MARITAL_STATUS  OCCUPATION
------- -------------------- --------------------
 103000 Divorc.              Cleric.
 102999 Married              Cleric.
 102998 Married              Exec.
 102997 Married              Exec.
 102996 NeverM               Other

BEGIN
  DBMS_DATA_MINING_TRANSFORM.XFORM_EXPR_STR(
       expr_pattern            => 'upper(:col)',
       data_table_name         => 'customers',
       xform_view_name         => 'cust_upcase_view');
END;
/
describe cust_upcase_view
 Name                           Null?    Type
 ------------------------------ -------- --------------------
 CUST_ID                        NOT NULL NUMBER
 CUST_MARITAL_STATUS                     VARCHAR2(20)
 OCCUPATION                              VARCHAR2(21)
 AGE                                     NUMBER
 YRS_RESIDENCE                           NUMBER

SELECT cust_id,  cust_marital_status, occupation FROM cust_upcase_view
   WHERE   cust_id > 102995
   ORDER BY cust_id desc;

CUST_ID CUST_MARITAL_STATUS  OCCUPATION
```

```
------- -------------------- --------------------
 103000 DIVORC.              CLERIC.
 102999 MARRIED              CLERIC.
 102998 MARRIED              EXEC.
 102997 MARRIED              EXEC.
 102996 NEVERM               OTHER
```

58.6.37 XFORM_MISS_CAT Procedure

This procedure creates a view that implements the categorical missing value treatment transformations specified in a definition table. Only the columns that are specified in the definition table are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_MISS_CAT (
    miss_table_name       IN VARCHAR2,
    data_table_name       IN VARCHAR2,
    xform_view_name       IN VARCHAR2,
    miss_schema_name      IN VARCHAR2 DEFAULT NULL,
    data_schema_name      IN VARCHAR2 DEFAULT NULL,
    xform_schema_name     IN VARCHAR2 DEFAULT NULL;
```

Parameters

Table 58-47 XFORM_MISS_CAT Procedure Parameters

Parameter	Description
miss_table_name	Name of the transformation definition table for categorical missing value treatment. You can use the CREATE_MISS_CAT Procedure to create the definition table. The table must be populated with transformation definitions before you call XFORM_MISS_CAT. To populate the table, you can use the INSERT_MISS_CAT_MODE Procedure or you can write your own SQL. See Table 58-12.
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *miss_table_name*.
miss_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes".

Examples

This example creates a view that replaces missing categorical values with the mode.

```
SELECT * FROM geog;

REG_ID REGION
```

```
------ ------------------------------
     1 NE
     2 SW
     3 SE
     4 SW
     5
     6 NE
     7 NW
     8 NW
     9
    10
    11 SE
    12 SE
    13 NW
    14 SE
    15 SE

SELECT STATS_MODE(region) FROM geog;

STATS_MODE(REGION)
------------------------------
SE

BEGIN
  DBMS_DATA_MINING_TRANSFORM.CREATE_MISS_CAT('misscat_xtbl');
  DBMS_DATA_MINING_TRANSFORM.INSERT_MISS_CAT_MODE (
    miss_table_name         => 'misscat_xtbl',
    data_table_name         => 'geog' );
END;
/

SELECT col, val FROM misscat_xtbl;

COL        VAL
---------- ----------
REGION     SE

BEGIN
  DBMS_DATA_MINING_TRANSFORM.XFORM_MISS_CAT (
     miss_table_name        => 'misscat_xtbl',
     data_table_name        => 'geog',
     xform_view_name        => 'geogxf_view');
END;
/

SELECT * FROM geogxf_view;

REG_ID REGION
------ ------------------------------
     1 NE
     2 SW
     3 SE
     4 SW
     5 SE
     6 NE
     7 NW
     8 NW
     9 SE
    10 SE
    11 SE
    12 SE
    13 NW
```

```
    14 SE
    15 SE
```

58.6.38 XFORM_MISS_NUM Procedure

This procedure creates a view that implements the numerical missing value treatment transformations specified in a definition table. Only the columns that are specified in the definition table are transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_MISS_NUM (
    miss_table_name       IN VARCHAR2,
    data_table_name       IN VARCHAR2,
    xform_view_name       IN VARCHAR2,
    miss_schema_name      IN VARCHAR2 DEFAULT NULL,
    data_schema_name      IN VARCHAR2 DEFAULT NULL,
    xform_schema_name     IN VARCHAR2 DEFAULT NULL;
```

Parameters

Table 58-48 XFORM_MISS_NUM Procedure Parameters

Parameter	Description
miss_table_name	Name of the transformation definition table for numerical missing value treatment. You can use the CREATE_MISS_NUM Procedure to create the definition table. The table must be populated with transformation definitions before you call XFORM_MISS_NUM. To populate the table, you can use the INSERT_MISS_NUM_MEAN Procedure or you can write your own SQL. See Table 58-14.
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *miss_table_name*.
miss_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes".

Examples

This example creates a view that replaces missing numerical values with the mean.

```
SELECT * FROM items;

ITEM_ID       QTY
---------- ------
aa            200
```

```
bb            200
cc            250
dd
ee
ff            100
gg            250
hh            200
ii
jj            200

SELECT AVG(qty) FROM items;

AVG(QTY)
--------
     200

BEGIN
  DBMS_DATA_MINING_TRANSFORM.CREATE_MISS_NUM('missnum_xtbl');
  DBMS_DATA_MINING_TRANSFORM.INSERT_MISS_NUM_MEAN (
     miss_table_name          => 'missnum_xtbl',
     data_table_name         => 'items' );
END;
/

SELECT col, val FROM missnum_xtbl;

COL           VAL
----------- ------
QTY           200

BEGIN
    DBMS_DATA_MINING_TRANSFORM.XFORM_MISS_NUM (
        miss_table_name        => 'missnum_xtbl',
        data_table_name        => 'items',
        xform_view_name        => 'items_view');
END;
/

SELECT * FROM items_view;

ITEM_ID       QTY
---------- ------
aa            200
bb            200
cc            250
dd            200
ee            200
ff            100
gg            250
hh            200
ii            200
jj            200
```

58.6.39 XFORM_NORM_LIN Procedure

This procedure creates a view that implements the linear normalization transformations specified in a definition table. Only the columns that are specified in the definition table are

transformed; the remaining columns from the data table are present in the view, but they are not changed.

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_NORM_LIN (
     norm_table_name      IN VARCHAR2,
     data_table_name      IN VARCHAR2,
     xform_view_name      IN VARCHAR2,
     norm_schema_name     IN VARCHAR2 DEFAULT NULL,
     data_schema_name     IN VARCHAR2 DEFAULT NULL,
     xform_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-49 XFORM_NORM_LIN Procedure Parameters

Parameter	Description
norm_table_name	Name of the transformation definition table for linear normalization. You can use the CREATE_NORM_LIN Procedure to create the definition table. The table must be populated with transformation definitions before you call XFORM_NORM_LIN. To populate the table, you can use one of the INSERT procedures for normalization or you can write your own SQL. See Table 58-12.
data_table_name	Name of the table containing the data to be transformed
xform_view_name	Name of the view to be created. The view presents columns in *data_table_name* with the transformations specified in *miss_table_name*.
norm_schema_name	Schema of *miss_table_name*. If no schema is specified, the current schema is used.
data_schema_name	Schema of *data_table_name*. If no schema is specified, the current schema is used.
xform_schema_name	Schema of *xform_view_name*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes".

Examples

This example creates a view that normalizes the cust_year_of_birth and cust_credit_limit columns. The data source consists of three columns from sh.customer.

```
CREATE OR REPLACE VIEW mining_data AS
     SELECT cust_id, cust_year_of_birth, cust_credit_limit
     FROM sh.customers;

describe mining_data
 Name                                      Null?    Type
 ----------------------------------------- -------- ----------------------------
 CUST_ID                                   NOT NULL NUMBER
 CUST_YEAR_OF_BIRTH                        NOT NULL NUMBER(4)
 CUST_CREDIT_LIMIT                                  NUMBER

SELECT * FROM mining_data WHERE cust_id > 104495
      ORDER BY cust_year_of_birth;
```

```
 CUST_ID CUST_YEAR_OF_BIRTH CUST_CREDIT_LIMIT
-------- ------------------ -----------------
  104496               1947              3000
  104498               1954             10000
  104500               1962             15000
  104499               1970              3000
  104497               1976              3000

BEGIN
  dbms_data_mining_transform.CREATE_NORM_LIN(
       norm_table_name        => 'normx_tbl');
 dbms_data_mining_transform.INSERT_NORM_LIN_MINMAX(
      norm_table_name     => 'normx_tbl',
      data_table_name     => 'mining_data',
      exclude_list     => dbms_data_mining_transform.COLUMN_LIST( 'cust_id'),
      round_num      => 3);
END;
/

SELECT col, shift, scale FROM normx_tbl;

COL                             SHIFT    SCALE
------------------------------ -------- --------
CUST_YEAR_OF_BIRTH                1910       77
CUST_CREDIT_LIMIT                 1500    13500

BEGIN
  DBMS_DATA_MINING_TRANSFORM.XFORM_NORM_LIN (
     norm_table_name     => 'normx_tbl',
     data_table_name     => 'mining_data',
     xform_view_name     => 'norm_view');
END;
/

SELECT * FROM norm_view WHERE cust_id > 104495
      ORDER BY cust_year_of_birth;

 CUST_ID CUST_YEAR_OF_BIRTH CUST_CREDIT_LIMIT
-------- ------------------ -----------------
  104496          .4805195          .1111111
  104498          .5714286          .6296296
  104500          .6753247                 1
  104499          .7792208          .1111111
  104497          .8571429          .1111111


set long 2000
SQL> SELECT text FROM user_views WHERE view_name IN 'NORM_VIEW';

TEXT
--------------------------------------------------------------------------------
SELECT "CUST_ID",("CUST_YEAR_OF_BIRTH"-1910)/77 "CUST_YEAR_OF_BIRTH",("CUST
_CREDIT_LIMIT"-1500)/13500 "CUST_CREDIT_LIMIT" FROM mining_data
```

58.6.40 XFORM_STACK Procedure

This procedure creates a view that implements the transformations specified by the stack. Only the columns and nested attributes that are specified in the stack are transformed. Any

remaining columns and nested attributes from the data table appear in the view without changes.

To create a list of objects that describe the transformed columns, use the DESCRIBE_STACK Procedure.

> **See Also:**
>
> "Overview"
>
> *Oracle Machine Learning for SQL User's Guide* for more information about machine learning attributes

Syntax

```
DBMS_DATA_MINING_TRANSFORM.XFORM_STACK (
     xform_list         IN     TRANSFORM_list,
     data_table_name    IN     VARCHAR2,
     xform_view_name    IN     VARCHAR2,
     data_schema_name   IN     VARCHAR2 DEFAULT NULL,
     xform_schema_name  IN     VARCHAR2 DEFAULT NULL);
```

Parameters

Table 58-50 XFORM_STACK Procedure Parameters

Parameter	Description
`xform_list`	The transformation list. See Table 58-1 for a description of the `TRANSFORM_LIST` object type.
`data_table_name`	Name of the table containing the data to be transformed
`xform_view_name`	Name of the view to be created. The view applies the transformations in *`xform_list`* to *`data_table_name`*.
`data_schema_name`	Schema of *`data_table_name`*. If no schema is specified, the current schema is used.
`xform_schema_name`	Schema of *`xform_view_name`*. If no schema is specified, the current schema is used.

Usage Notes

See "Operational Notes". The following sections are especially relevant:

- "About Transformation Lists"
- "About Stacking"
- "Nested Data Transformations"

Examples

This example applies a transformation list to the view `oml_user.cust_info` and shows how the data is transformed. The `CREATE` statement for `cust_info` is shown in "DESCRIBE_STACK Procedure".

```
BEGIN
  dbms_data_mining_transform.CREATE_BIN_NUM ('birth_yr_bins');
  dbms_data_mining_transform.INSERT_BIN_NUM_QTILE (
        bin_table_name   => 'birth_yr_bins',
        data_table_name  =>  'cust_info',
        bin_num          =>   6,
        exclude_list     =>  dbms_data_mining_transform.column_list(
                                 'cust_id','country_id'));
END;
/
SELECT * FROM birth_yr_bins;

COL                  ATT      VAL BIN
-------------------- ----- ------ ----------
CUST_YEAR_OF_BIRTH           1922
CUST_YEAR_OF_BIRTH           1951 1
CUST_YEAR_OF_BIRTH           1959 2
CUST_YEAR_OF_BIRTH           1966 3
CUST_YEAR_OF_BIRTH           1973 4
CUST_YEAR_OF_BIRTH           1979 5
CUST_YEAR_OF_BIRTH           1986 6

DECLARE
      cust_stack   dbms_data_mining_transform.TRANSFORM_LIST;
BEGIN
      dbms_data_mining_transform.SET_TRANSFORM (cust_stack,
          'country_id', NULL, 'country_id/10', 'country_id*10');
      dbms_data_mining_transform.STACK_BIN_NUM ('birth_yr_bins',
          cust_stack);
      dbms_data_mining_transform.SET_TRANSFORM (cust_stack,
          'custprods', 'Mouse Pad', 'value*100', 'value/100');
      dbms_data_mining_transform.XFORM_STACK(
          xform_list        => cust_stack,
          data_table_name   => 'cust_info',
          xform_view_name   =>  'cust_xform_view');
  END;
/

-- Two rows of data without transformations
SELECT * from cust_info WHERE cust_id BETWEEN 100010 AND 100011;

CUST_ID COUNTRY_ID CUST_YEAR_OF_BIRTH CUSTPRODS(ATTRIBUTE_NAME, VALUE)
------- ---------- ------------------ ----------------------------------------
 100010      52790               1975 DM_NESTED_NUMERICALS(
                                        DM_NESTED_NUMERICAL(
                                          '18" Flat Panel Graphics Monitor', 1),
                                        DM_NESTED_NUMERICAL(
                                          'SIMM- 16MB PCMCIAII card', 1))
 100011      52775               1972 DM_NESTED_NUMERICALS(
                                       DM_NESTED_NUMERICAL(
                                         'External 8X CD-ROM', 1),
                                       DM_NESTED_NUMERICAL(
                                         'Mouse Pad', 1),
                                       DM_NESTED_NUMERICAL(
                                        'SIMM- 16MB PCMCIAII card', 1),
                                       DM_NESTED_NUMERICAL(
                                         'Keyboard Wrist Rest', 1),
                                       DM_NESTED_NUMERICAL(
                                         '18" Flat Panel Graphics Monitor', 1),
                                       DM_NESTED_NUMERICAL(
                                         'O/S Documentation Set - English', 1))
```

```
-- Same two rows of data with transformations
SELECT * FROM cust_xform_view WHERE cust_id BETWEEN 100010 AND 100011;

CUST_ID  COUNTRY_ID  C  CUSTPRODS(ATTRIBUTE_NAME, VALUE)
-------  ----------  -  ------------------------------------------------------------
 100010        5279  5  DM_NESTED_NUMERICALS(
                           DM_NESTED_NUMERICAL(
                            '18" Flat Panel Graphics Monitor', 1),
                           DM_NESTED_NUMERICAL(
                             'SIMM- 16MB PCMCIAII card', 1))
 100011     5277.5  4  DM_NESTED_NUMERICALS(
                           DM_NESTED_NUMERICAL(
                             'External 8X CD-ROM', 1),
                           DM_NESTED_NUMERICAL(
                             'Mouse Pad', 100),
                           DM_NESTED_NUMERICAL(
                             'SIMM- 16MB PCMCIAII card', 1),
                           DM_NESTED_NUMERICAL(
                             'Keyboard Wrist Rest', 1),
                           DM_NESTED_NUMERICAL(
                             '18" Flat Panel Graphics Monitor', 1),
                           DM_NESTED_NUMERICAL(
                             'O/S Documentation Set - English', 1))
```

59
DBMS_DATAPUMP

The DBMS_DATAPUMP package is used to move all, or part of, a database between databases, including both data and metadata.

> **See Also:**
>
> *Oracle Database Utilities* for more information on the concepts behind the DBMS_DATAPUMP API, how it works, and how it is implemented in the Data Pump Export and Import utilities

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Data Structures
- Summary of DBMS_DATAPUMP Subprograms

59.1 DBMS_DATAPUMP Overview

DBMS_DATAPUMP provides the following support and functionality.

- The source and target databases can have different hardware, operating systems, character sets, time zones, and versions.
- All object types and datatypes existing in Oracle Database 11*g* and higher are supported.
- Data and metadata can be transferred between databases without using any intermediary files.
- A subset of a database can be moved based upon object type and names of objects.
- Schema names, data file names, tablespace names, and data can be transformed at import time.
- Previously aborted export and import jobs can be restarted without duplicating or omitting any data or metadata from the original job.
- The resources applied to an export or import job can be modified.
- Data in an Oracle proprietary format can be unloaded and loaded.

59.2 DBMS_DATAPUMP Security Model

Security for the DBMS_DATAPUMP package is implemented through roles.

The DATAPUMP_EXP_FULL_DATABASE and DATAPUMP_IMP_FULL_DATABASE roles enables privileged users to take full advantage of the API. The Oracle Data Pump API will use these roles to

determine whether privileged application roles should be assigned to the processes comprising the job.

- DATAPUMP_EXP_FULL_DATABASE

  The DATAPUMP_EXP_FULL_DATABASE role affects only Export operations. It allows users running these operations to do the following:

  – Perform the operation outside of the scope of their schema
  – Monitor jobs that were initiated by another user
  – Export objects (for example, TABLESPACE definitions) that unprivileged users cannot reference

  Although the SYS schema does not have the DATAPUMP_EXP_FULL_DATABASE role assigned to it, all security checks performed by Data Pump that require the DATAPUMP_EXP_FULL_DATABASE role will also grant access to the SYS schema.

- DATAPUMP_IMP_FULL_DATABASE

  The DATAPUMP_IMP_FULL_DATABASE role affects only Import and SQL_FILE operations. It allows users running these operations to do the following:

  – Perform the operation outside of the scope of their schema
  – Monitor jobs that were initiated by another user
  – Import objects (for example, DIRECTORY definitions) that unprivileged users cannot create

  Although the SYS schema does not have the DATAPUMP_IMP_FULL_DATABASE role assigned to it, all security checks performed by Oracle Data Pump that require the DATAPUMP_IMP_FULL_DATABASE role will also grant access to the SYS schema.

59.3 DBMS_DATAPUMP Constants

There are several public constants defined for use with the DBMS_DATAPUMP.GET_STATUS procedure. All such constants are defined as part of the DBMS_DATAPUMP package. Any references to these constants must be prefixed by DBMS_DATAPUMP. and followed by the symbols in the following lists:

Mask Bit Definitions

The following mask bit definitions are used for controlling the return of data through the DBMS_DATAPUMP.GET_STATUS procedure.

- KU$_STATUS_WIP CONSTANT BINARY_INTEGER := 1;
- KU$_STATUS_JOB_DESC CONSTANT BINARY_INTEGER := 2;
- KU$_STATUS_JOB_STATUS CONSTANT BINARY_INTEGER := 4;
- KU$_STATUS_JOB_ERROR CONSTANT BINARY_INTEGER := 8;

Dump File Type Definitions

The following definitions are used for identifying types of dump files returned through the DBMS_DATAPUMP.GET_STATUS procedure.

- KU$_DUMPFILE_TYPE_DISK CONSTANT BINARY_INTEGER := 0;
- KU$_DUMPFILE_TYPE_TEMPLATE CONSTANT BINARY_INTEGER := 3;

59.4 DBMS_DATAPUMP Data Structures

The DBMS_DATAPUMP package defines OBJECT types. The types described in this section are defined in the SYS schema for use by the GET_STATUS function. *The way in which these types are defined and used may be different than what you are accustomed to. Be sure to read this section carefully.*
The collection of types defined for use with the GET_STATUS procedure are version-specific and include version information in the names of the types. Once introduced, these types will always be provided and supported in future versions of Oracle Database and will not change. However, in future releases of Oracle Database, new versions of these types might be created that provide new or different information. The new versions of these types will have different version information embedded in the type names.

For example, in Oracle Database 12c Release 1 (12.1), there is a sys.ku$_Status1210 type, and in the next Oracle Database release, there can be a sys.ku$_Status1310 type defined. Both types can be used with the GET_STATUS procedure.

Public synonyms have been defined for each of the types used with the GET_STATUS procedure. This makes it easier to use the types and means that you do not have to be concerned with changes to the actual type names or schemas where they reside. Oracle recommends that you use these synonyms whenever possible.

For each of the types, there is a version-specific synonym and a generic synonym. For example, the version-specific synonym ku$_Status1210 is defined for the sys.ku$_Status1210 type.

The generic synonym always describes the latest version of that type. For example, in Oracle Database 12c, Release 1 (12.1), the generic synonym ku$_Status is defined as ku$_Status1210. In a future release, there might be a ku$_Status1310 synonym for sys.ku$Status1310. Because the ku$_Status generic synonym always points to the latest definition, it would then point to ku$_Status1310 rather than to ku$_Status1210.

The choice of whether to use version-specific synonyms or generic synonyms makes a significant difference in how you work. Using version-specific names protects your code from changes in future releases of Oracle Database because those types will continue to exist and be supported. However, access to new information will require code changes to use new synonym names for each of the types. Using the generic names implies that you always want the latest definition of the types and are prepared to deal with changes in different releases of Oracle Database.

When the version of Oracle Database that you are using changes, any C code that accesses types through generic synonym names will need to be recompiled.

> **Note:**
>
> Languages other than PL/SQL must ensure that their type definitions are properly aligned with the version-specific definitions.

> **See Also:**
>
> GET_STATUS Procedure for additional information about how types are used

Data Structures — Object Types

The DBMS_DATAPUMP package defines the following kinds of OBJECT types:

- Worker Status Types
- Log Entry and Error Types
- Job Status Types
- Job Description Types
- Status Types

Worker Status Types

The worker status types describe what each worker process in a job is doing. The schema, object name, and object type of an object being processed will be provided. For workers processing user data, the partition name for a partitioned table (if any), the number of bytes processed in the partition, and the number of rows processed in the partition are also returned. Workers processing metadata provide status on the last object that was processed. No status for idle threads is returned.

The percent_done refers to the amount completed for the current data item being processed. It is not updated for metadata objects.

The worker status types are defined as follows:

```
CREATE TYPE sys.ku$_WorkerStatus1010 AS OBJECT
        (
                worker_number     NUMBER,       -- Worker process identifier
                process_name      VARCHAR2(30), -- Worker process name
                state             VARCHAR2(30), -- Worker process state
                schema            VARCHAR2(30), -- Schema name
                name              VARCHAR2(4000),-- Object name
                object_type       VARCHAR2(200),-- Object type
                partition         VARCHAR2(30), -- Partition name
                completed_objects NUMBER,       -- Completed number of objects
                total_objects     NUMBER,       -- Total number of objects
                completed_rows    NUMBER,       -- Number of rows completed
                completed_bytes   NUMBER,       -- Number of bytes completed
                percent_done      NUMBER        -- Percent done current object
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatus1010
  FOR sys.ku$_WorkerStatus1010;

CREATE TYPE sys.ku$_WorkerStatus1020 AS OBJECT
        (
                worker_number     NUMBER,       -- Worker process identifier
                process_name      VARCHAR2(30), -- Worker process name
                state             VARCHAR2(30), -- Worker process state
                schema            VARCHAR2(30), -- Schema name
                name              VARCHAR2(4000),-- Object name
                object_type       VARCHAR2(200),-- Object type
                partition         VARCHAR2(30), -- Partition name
                completed_objects NUMBER,       -- Completed number of objects
                total_objects     NUMBER,       -- Total number of objects
                completed_rows    NUMBER,       -- Number of rows completed
                completed_bytes   NUMBER,       -- Number of bytes completed
                percent_done      NUMBER,       -- Percent done current object
```

```
                degree            NUMBER        -- Degree of parallelism
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatus1020
  FOR sys.ku$_WorkerStatus1020;

CREATE TYPE sys.ku$_WorkerStatus1120 AS OBJECT
        (
                worker_number     NUMBER,       -- Worker process identifier
                process_name      VARCHAR2(30), -- Worker process name
                state             VARCHAR2(30), -- Worker process state
                schema            VARCHAR2(30), -- Schema name
                name              VARCHAR2(4000),-- Object name
                object_type       VARCHAR2(200),-- Object type
                partition         VARCHAR2(30), -- Partition name
                completed_objects NUMBER,       -- Completed number of objects
                total_objects     NUMBER,       -- Total number of objects
                completed_rows    NUMBER,       -- Number of rows completed
                completed_bytes   NUMBER,       -- Number of bytes completed
                percent_done      NUMBER,       -- Percent done current object
                degree            NUMBER,       -- Degree of parallelism
                instance_id       NUMBER        -- Instance ID where running
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatus1120
  FOR sys.ku$_WorkerStatus1120;

CREATE TYPE sys.ku$_WorkerStatus1210 AS OBJECT
        (
                worker_number     NUMBER,       -- Worker process identifier
                process_name      VARCHAR2(30), -- Worker process name
                state             VARCHAR2(30), -- Worker process state
                schema            VARCHAR2(30), -- Schema name
                name              VARCHAR2(4000),-- Object name
                object_type       VARCHAR2(200),-- Object type
                partition         VARCHAR2(30), -- Partition name
                completed_objects NUMBER,       -- Completed number of objects
                total_objects     NUMBER,       -- Total number of objects
                completed_rows    NUMBER,       -- Number of rows completed
                completed_bytes   NUMBER,       -- Number of bytes completed
                percent_done      NUMBER,       -- Percent done current object
                degree            NUMBER,       -- Degree of parallelism
                instance_id       NUMBER,       -- Instance ID where running
                instance_name     VARCHAR2(60), -- Instance Name where running
                host_name         VARCHAR2(64)  -- Host name where running
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatus1210
  FOR sys.ku$_WorkerStatus1210;

CREATE TYPE sys.ku$_WorkerStatus1220 AS OBJECT
        (
                worker_number     NUMBER,       -- Worker process identifier
                process_name      VARCHAR2(128),-- Worker process name
                state             VARCHAR2(30), -- Worker process state
                schema            VARCHAR2(128),-- Schema name
                name              VARCHAR2(4000),-- Object name
                object_type       VARCHAR2(200),-- Object type
                partition         VARCHAR2(128),-- Partition name
                completed_objects NUMBER,       -- Completed number of objects
                total_objects     NUMBER,       -- Total number of objects
                completed_rows    NUMBER,       -- Number of rows completed
```

```
                completed_bytes   NUMBER,       -- Number of bytes completed
                percent_done      NUMBER,       -- Percent done current object
                degree            NUMBER,       -- Degree of parallelism
                instance_id       NUMBER,       -- Instance ID where running
                instance_name     VARCHAR2(60), -- Instance Name where running
                host_name         VARCHAR2(101), -- Host name where running
                access_method     VARCHAR2(16), -- Access Method of object
                obj_start_time    DATE,         -- Object start  time
                obj_status        DATE         -- Object status at current time

        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatus1220
  FOR sys.ku$_WorkerStatus1220;

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatus FOR ku$_WorkerStatus1220;

CREATE TYPE sys.ku$_WorkerStatusList1010 AS TABLE OF sys.ku$_WorkerStatus1010
CREATE TYPE sys.ku$_WorkerStatusList1020 AS TABLE OF sys.ku$_WorkerStatus1020
CREATE TYPE sys.ku$_WorkerStatusList1120 AS TABLE OF sys.ku$_WorkerStatus1120
CREATE TYPE sys.ku$_WorkerStatusList1210 AS TABLE OF sys.ku$_WorkerStatus1210
CREATE TYPE sys.ku$_WorkerStatusList1220 AS TABLE OF sys.ku$_WorkerStatus1220

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatusList1010
  FOR sys.ku$_WorkerStatusList1010;
CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatusList1020
  FOR sys.ku$_WorkerStatusList1020;
CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatusList1120
  FOR sys.ku$_WorkerStatusList1120;
CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatusList1210
  FOR sys.ku$_WorkerStatusList1210;
CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatusList1220
  FOR sys.ku$_WorkerStatusList1220;

CREATE OR REPLACE PUBLIC SYNONYM ku$_WorkerStatusList
  FOR ku$_WorkerStatusList1220;
```

Log Entry and Error Types

These types provide informational and error text to attached clients and the log stream. The `ku$LogLine.errorNumber` type is set to `NULL` for informational messages but is specified for error messages. Each log entry may contain several lines of text messages.

The log entry and error types are defined as follows:

```
CREATE TYPE sys.ku$_LogLine1010 AS OBJECT (
               logLineNumber   NUMBER,
               errorNumber     NUMBER,
               LogText         VARCHAR2(2000))

CREATE OR REPLACE PUBLIC SYNONYM ku$_LogLine1010 FOR sys.ku$_LogLine1010;
CREATE OR REPLACE PUBLIC SYNONYM ku$_LogLine1020 FOR sys.ku$_LogLine1010;
CREATE OR REPLACE PUBLIC SYNONYM ku$_LogLine FOR ku$_LogLine1010;
CREATE TYPE sys.ku$_LogEntry1010 AS TABLE OF sys.ku$_LogLine1010

CREATE OR REPLACE PUBLIC SYNONYM ku$_LogEntry1010 FOR sys.ku$_LogEntry1010;
CREATE OR REPLACE PUBLIC SYNONYM ku$_LogEntry1020 FOR sys.ku$_LogEntry1010;
CREATE OR REPLACE PUBLIC SYNONYM ku$_LogEntry FOR ku$_LogEntry1010;
```

Job Status Types

The job status type returns status about a job. Usually, the status concerns a running job, but it could also be about a stopped job when a client attaches. It is typically requested at attach time, when the client explicitly requests status from interactive mode and every *N* seconds when the client has requested status periodically.

The job status types are defined as follows (`percent_done` applies to data only):

```
CREATE TYPE sys.ku$_JobStatus1010 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- Name of the job
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                bytes_processed NUMBER,                 -- Bytes so far
                total_bytes     NUMBER,                 -- Total bytes for job
                percent_done    NUMBER,                 -- Percent done
                degree          NUMBER,                 -- Of job parallelism
                error_count     NUMBER,                 -- #errors so far
                state           VARCHAR2(30),           -- Current job state
                phase           NUMBER,                 -- Job phase
                restart_count   NUMBER,                 -- #Job restarts
                worker_status_list ku$_WorkerStatusList1010, -- For (non-idle)
                                                        -- job worker processes
                files           ku$_DumpFileSet1010     -- Dump file info
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobStatus1010 FOR sys.ku$_JobStatus1010;

CREATE TYPE sys.ku$_JobStatus1020 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- Name of the job
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                bytes_processed NUMBER,                 -- Bytes so far
                total_bytes     NUMBER,                 -- Total bytes for job
                percent_done    NUMBER,                 -- Percent done
                degree          NUMBER,                 -- Of job parallelism
                error_count     NUMBER,                 -- #errors so far
                state           VARCHAR2(30),           -- Current job state
                phase           NUMBER,                 -- Job phase
                restart_count   NUMBER,                 -- #Job restarts
                worker_status_list ku$_WorkerStatusList1020, -- For (non-idle)
                                                        -- job worker processes
                files           ku$_DumpFileSet1010     -- Dump file info
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobStatus1020 FOR sys.ku$_JobStatus1020;

CREATE TYPE sys.ku$_JobStatus1120 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- Name of the job
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                bytes_processed NUMBER,                 -- Bytes so far
                total_bytes     NUMBER,                 -- Total bytes for job
                percent_done    NUMBER,                 -- Percent done
                degree          NUMBER,                 -- Of job parallelism
                error_count     NUMBER,                 -- #errors so far
                state           VARCHAR2(30),           -- Current job state
                phase           NUMBER,                 -- Job phase
```

```
                restart_count   NUMBER,                 -- #Job restarts
                worker_status_list ku$_WorkerStatusList1120, -- For (non-idle)
                                                        -- job worker processes
                    files           ku$_DumpFileSet1010     -- Dump file info
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobStatus1120 FOR sys.ku$_JobStatus1120;

CREATE TYPE sys.ku$_JobStatus1210 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- Name of the job
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                bytes_processed NUMBER,                 -- Bytes so far
                total_bytes     NUMBER,                 -- Total bytes for job
                percent_done    NUMBER,                 -- Percent done
                degree          NUMBER,                 -- Of job parallelism
                error_count     NUMBER,                 -- #errors so far
                state           VARCHAR2(30),           -- Current job state
                phase           NUMBER,                 -- Job phase
                restart_count   NUMBER,                 -- #Job restarts
                worker_status_list ku$_WorkerStatusList1210, -- For (non-idle)
                                                        -- job worker processes
                files           ku$_DumpFileSet1010     -- Dump file info
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobStatus1210 FOR sys.ku$_JobStatus1210;

CREATE TYPE sys.ku$_JobStatus1220 IS OBJECT
        (
                job_name        VARCHAR2(128),          -- Name of the job
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                bytes_processed NUMBER,                 -- Bytes so far
                total_bytes     NUMBER,                 -- Total bytes for job
                percent_done    NUMBER,                 -- Percent done
                degree          NUMBER,                 -- Of job parallelism
                error_count     NUMBER,                 -- #errors so far
                state           VARCHAR2(30),           -- Current job state
                phase           NUMBER,                 -- Job phase
                restart_count   NUMBER,                 -- #Job restarts
                heartbeat       NUMBER,                 -- Job heartbeat
                worker_status_list ku$_WorkerStatusList1220, -- For (non-idle)
                                                        -- job worker processes
                files           ku$_DumpFileSet1010     -- Dump file info
        )



CREATE OR REPLACE PUBLIC SYNONYM ku$_JobStatus1220 FOR sys.ku$_JobStatus1220;

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobStatus FOR ku$_JobStatus1220;
```

Job Description Types

The job description type holds all the environmental information about the job such as parameter settings and dump file set members. There are a couple of subordinate types required as well.

The job description types are defined as follows:

```
CREATE TYPE sys.ku$_JobDesc1010 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- The job name
                guid            RAW(16),                -- The job GUID
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                remote_link     VARCHAR2(4000),         -- DB link, if any
                owner           VARCHAR2(30),           -- Job owner
                instance        VARCHAR2(16),           -- The instance name
                db_version      VARCHAR2(30),           -- Version of objects
                creator_privs   VARCHAR2(30),           -- Privs of job
                start_time      DATE,                   -- This job start time
                max_degree      NUMBER,                 -- Max. parallelism
                log_file        VARCHAR2(4000),         -- Log file name
                sql_file        VARCHAR2(4000),         -- SQL file name
                params          ku$_ParamValues1010     -- Parameter list
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobDesc1010 FOR sys.ku$_JobDesc1010;

CREATE TYPE sys.ku$_JobDesc1020 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- The job name
                guid            RAW(16),                -- The job GUID
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                remote_link     VARCHAR2(4000),         -- DB link, if any
                owner           VARCHAR2(30),           -- Job owner
                platform        VARCHAR2(101),          -- Current job platform
                exp_platform    VARCHAR2(101),          -- Export platform
                global_name     VARCHAR2(4000),         -- Current global name
                exp_global_name VARCHAR2(4000),         -- Export global name
                instance        VARCHAR2(16),           -- The instance name
                db_version      VARCHAR2(30),           -- Version of objects
                exp_db_version  VARCHAR2(30),           -- Export version
                scn             NUMBER,                 -- Job SCN
                creator_privs   VARCHAR2(30),           -- Privs of job
                start_time      DATE,                   -- This job start time
                exp_start_time  DATE,                   -- Export start time
                term_reason     NUMBER,                 -- Job termination code
                max_degree      NUMBER,                 -- Max. parallelism
                log_file        VARCHAR2(4000),         -- Log file name
                sql_file        VARCHAR2(4000),         -- SQL file name
                params          ku$_ParamValues1010     -- Parameter list
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobDesc1020 FOR sys.ku$_JobDesc1020;

CREATE TYPE sys.ku$_JobDesc1210 IS OBJECT
        (
                job_name        VARCHAR2(30),           -- The job name
                guid            RAW(16),                -- The job GUID
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                remote_link     VARCHAR2(4000),         -- DB link, if any
                owner           VARCHAR2(30),           -- Job owner
                platform        VARCHAR2(101),          -- Current job platform
                exp_platform    VARCHAR2(101),          -- Export platform
                global_name     VARCHAR2(4000),         -- Current global name
                exp_global_name VARCHAR2(4000),         -- Export global name
                instance        VARCHAR2(16),           -- The instance name
                db_version      VARCHAR2(30),           -- Cur. server software version
```

```
                exp_db_version  VARCHAR2(30),           -- Export svr. software version
                job_version     VARCHAR2(30),           -- Negotiated data version
                scn             NUMBER,                 -- Job SCN
                creator_privs   VARCHAR2(30),           -- Privs of job
                start_time      DATE,                   -- This job start time
                exp_start_time  DATE,                   -- Export start time
                term_reason     NUMBER,                 -- Job termination code
                max_degree      NUMBER,                 -- Max. parallelism
                timezone        VARCHAR2(64),           -- Cur. server timezone
                exp_timezone    VARCHAR2(64),           -- Exp. server timezone
                tstz_version    NUMBER,                 -- Cur. server timezone version
                exp_tstz_version NUMBER,                -- Exp. server timezone
                endianness      VARCHAR2(16),           -- Cur. platform's endianness
                exp_endianness  VARCHAR2(16),           -- Exp. platform's endianness
-- endianness is 'BIG' or 'LITTLE'
                charset         VARCHAR2(28),           -- Cur. server charset
                exp_charset     VARCHAR2(28),           -- Exp. server charset
                ncharset        VARCHAR2(28),           -- Cur. server national charset
                exp_ncharset    VARCHAR2(28),           -- Exp. server national charset
                log_file        VARCHAR2(4000),         -- Log file name
                sql_file        VARCHAR2(4000),         -- SQL file name
                params          ku$_ParamValues1010     -- Parameter list
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobDesc1210 FOR sys.ku$_JobDesc1210;

CREATE TYPE sys.ku$_JobDesc1220 IS OBJECT
        (
                job_name        VARCHAR2(128),          -- The job name
                guid            RAW(16),                -- The job GUID
                operation       VARCHAR2(30),           -- Current operation
                job_mode        VARCHAR2(30),           -- Current mode
                remote_link     VARCHAR2(4000),         -- DB link, if any
                owner           VARCHAR2(128),          -- Job owner
                platform        VARCHAR2(101),          -- Current job platform
                exp_platform    VARCHAR2(101),          -- Export platform
                global_name     VARCHAR2(4000),         -- Current global name
                exp_global_name VARCHAR2(4000),         -- Export global name
                instance        VARCHAR2(60),           -- The instance name
                db_version      VARCHAR2(60),           -- Cur. server software version
                exp_db_version  VARCHAR2(60),           -- Export svr. software version
                job_version     VARCHAR2(60),           -- Negotiated data version
                scn             NUMBER,                 -- Job SCN
                creator_privs   VARCHAR2(30),           -- Privs of job
                start_time      DATE,                   -- This job start time
                exp_start_time  DATE,                   -- Export start time
                term_reason     NUMBER,                 -- Job termination code
                max_degree      NUMBER,                 -- Max. parallelism
                timezone        VARCHAR2(64),           -- Cur. server timezone
                exp_timezone    VARCHAR2(64),           -- Exp. server timezone
                tstz_version    NUMBER,                 -- Cur. server timezone version
                exp_tstz_version NUMBER,                -- Exp. server timezone
                endianness      VARCHAR2(16),           -- Cur. platform's endianness
                exp_endianness  VARCHAR2(16),           -- Exp. platform's endianness
-- endianness is 'BIG' or 'LITTLE'
                charset         VARCHAR2(28),           -- Cur. server charset
                exp_charset     VARCHAR2(28),           -- Exp. server charset
                ncharset        VARCHAR2(28),           -- Cur. server national charset
                exp_ncharset    VARCHAR2(28),           -- Exp. server national charset
                log_file        VARCHAR2(4000),         -- Log file name
                sql_file        VARCHAR2(4000),         -- SQL file name
                params          ku$_ParamValues1010     -- Parameter list
```

```
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobDesc1220 FOR sys.ku$_JobDesc1220;

CREATE OR REPLACE PUBLIC SYNONYM ku$_JobDesc FOR ku$_JobDesc1220;
```

Status Types

The status type is an aggregate of some the previous types defined and is the return value for the `GET_STATUS` call. The mask attribute indicates which types of information are being returned to the caller. It is created by a client's shadow process from information it retrieves off the status queue or directly from the master table.

For errors, the `ku$_LogEntry` that is returned has already had its log lines ordered for proper output. That is, the original `ku$_LogEntry` objects have been ordered from outermost context to innermost.

The status types are defined as follows:

```
CREATE TYPE sys.ku$_Status1010 IS OBJECT
        (
                mask            NUMBER,           -- Status types present
                wip             ku$_LogEntry1010, -- Work in progress
                job_description ku$_JobDesc1010,  -- Complete job description
                job_status      ku$_JobStatus1010,-- Detailed job status
                error           ku$_LogEntry1010  -- Multi-level context errors
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_Status1010 FOR sys.ku$_Status1010;

CREATE TYPE sys.ku$_Status1020 IS OBJECT
        (
                mask            NUMBER,           -- Status types present
                wip             ku$_LogEntry1010, -- Work in progress
                job_description ku$_JobDesc1020,  -- Complete job description
                job_status      ku$_JobStatus1020,-- Detailed job status
                error           ku$_LogEntry1010  -- Multi-level context errors
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_Status1020 FOR sys.ku$_Status1020;

CREATE TYPE sys.ku$_Status1120 IS OBJECT
        (
                mask            NUMBER,           -- Status types present
                wip             ku$_LogEntry1010, -- Work in progress
                job_description ku$_JobDesc1020,  -- Complete job description
                job_status      ku$_JobStatus1120,-- Detailed job status
                error           ku$_LogEntry1010  -- Multi-level context errors
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_Status1120 FOR sys.ku$_Status1120;

CREATE TYPE sys.ku$_Status1210 IS OBJECT
        (
                mask            NUMBER,           -- Status types present
                wip             ku$_LogEntry1010, -- Work in progress
                job_description ku$_JobDesc1210,  -- Complete job description
                job_status      ku$_JobStatus1210,-- Detailed job status
                error           ku$_LogEntry1010  -- Multi-level context errors
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_Status1210 FOR sys.ku$_Status1210;
```

```
CREATE TYPE sys.ku$_Status1220 IS OBJECT
        (
                mask            NUMBER,           -- Status types present
                wip             ku$_LogEntry1010, -- Work in progress
                job_description ku$_JobDesc1220,  -- Complete job description
                job_status      ku$_JobStatus1220,-- Detailed job status
                error           ku$_LogEntry1010  -- Multi-level context errors
        )

CREATE OR REPLACE PUBLIC SYNONYM ku$_Status1220 FOR sys.ku$_Status1220;

CREATE OR REPLACE PUBLIC SYNONYM ku$_Status FOR ku$_Status1220;
```

59.5 Summary of DBMS_DATAPUMP Subprograms

This table lists the `DBMS_DATAPUMP` subprograms in alphabetical order and briefly describes them.

Table 59-1 DBMS_DATAPUMP Package Subprograms

Subprogram	Description
ADD_FILE Procedure	Adds dump files to the dump file set for an Export, Import, or SQL_FILE operation. In addition to dump files, other types of files can also be added by using the `FILETYPE` parameter provided with this procedure.
ATTACH Function	Used to gain access to a Data Pump job that is in the Defining, Executing, Idling, or Stopped state
DATA_FILTER Procedures	Specifies restrictions on the rows that are to be retrieved
DATA_REMAP Procedure	Specifies transformations to be applied to column data as it is exported from, or imported into, a database.
DETACH Procedure	Specifies that the user has no further interest in using the handle
GET_DUMPFILE_INFO Procedure	Retrieves information about a specified dump file
GET_STATUS Procedure	Monitors the status of a job or waits for the completion of a job or for more details on API errors
LOG_ENTRY Procedure	Inserts a message into the log file
METADATA_FILTER Procedure	Provides filters that allow you to restrict the items that are included in a job
METADATA_REMAP Procedure	Specifies a remapping to be applied to objects as they are processed in the specified job
METADATA_TRANSFORM Procedure	Specifies transformations to be applied to objects as they are processed in the specified job
OPEN Function	Declares a new job using the Data Pump API, the handle returned being used as a parameter for calls to all other procedures (but not to the `ATTACH` function)
SET_PARALLEL Procedure	Adjusts the degree of parallelism within a job
SET_PARAMETER Procedures	Specifies job-processing options
START_JOB Procedure	Begins or resumes execution of a job
STOP_JOB Procedure	Terminates a job, but optionally, preserves the state of the job
WAIT_FOR_JOB Procedure	Runs a job until it either completes normally or stops for some other reason

59.5.1 ADD_FILE Procedure

This procedure adds files to the dump file set for an Export, Import, or `SQL_FILE` operation, or specifies the log file or the output file for a `SQL_FILE` operation.

Syntax

```
DBMS_DATAPUMP.ADD_FILE (
   handle     IN NUMBER,
   filename   IN VARCHAR2,
   directory  IN VARCHAR2,
   filesize   IN VARCHAR2 DEFAULT NULL,
   filetype   IN NUMBER DEFAULT DBMS_DATAPUMP.KU$_FILE_TYPE_DUMP_FILE,
   reusefile  IN NUMBER DEFAULT NULL);
```

Parameters

Table 59-2 ADD_FILE Procedure Parameters

Parameter	Description
`handle`	The handle of a job. The current session must have previously attached to the handle through a call to either the `OPEN` or `ATTACH` function.
`filename`	The name of the file that is being added. The `filename` parameter must be a simple filename without any directory path information. For dump files, the `filename` can include a substitution variable. For a description of available substitution variables, see the following table. The file can be written to or read from either the local file system or the Oracle Object Store. If you are interfacing with the local file system, the `filename` parameter must contain a filename without directory path information. If you are interfacing with the Oracle Object Store, the `filename` parameter must contain a valid `URI` to the location of a bucket within your compartment in the Oracle Object Store. You must also specify a valid `CREDENTAL` in the directory parameter to access the bucket and the `DBMS_DATAPUMP.KU$_FILE_TYPE_URIDUMP_FILE` file type in the filetype parameter. For examples, see "Using the Oracle Data Pump API" in *Oracle Database Utilities*.
`directory`	If you are interfacing with the local file system, then the directory parameter specifies the name of a directory object within the database that is used to locate the filename. If you are interfacing with the Oracle Object Store, then the `directory` parameter must specify a valid `CREDENTIAL` to access the bucket where filename resides. Ensure that you specify a valid `directory` parameter.

Table 59-2 (Cont.) ADD_FILE Procedure Parameters

Parameter	Description
filesize	The size of the dump file that is being added. It can be specified as follows: • The number of bytes • The number of kilobytes (if followed by K) • The number of megabytes (if followed by M) • The number of gigabytes (if followed by G) • The number of terabytes (if followed by T) An Export operation will write no more than the specified number of bytes to the file. When the file is full, it will be closed. If there is insufficient space on the device to write the specified number of bytes, then the Export operation will fail, but it can be restarted. If not specified, then filesize defaults to an unlimited size. For Import and SQL_FILE operations, filesize is ignored. The minimum value for filesize is ten times the default Data Pump block size, which is 4 kilobytes. A filesize can only be specified for dump files.
filetype	The type of the file that you want to add. The supported values are as follows: • DBMS_DATAPUMP.KU$_FILE_TYPE_DUMP_FILE—Dump file for a job • DBMS_DATAPUMP.KU$_FILE_TYPE_LOG_FILE—Log file for a job • DBMS_DATAPUMP.KU$_FILE_TYPE_SQL_FILE—Output for SQL_FILE job • DBMS_DATAPUMP.KU$_FILE_TYPE_URIDUMP_FILE—Specifies a filename in the Oracle Object Store. For more information, see the filename and directory parameters described above.
reusefile	If the value is 0, then a preexisting file will cause an error. If the value is 1, then a preexisting file will be overwritten. If the value is NULL, then the default action for the file type will be applied (that is, dump files will not be overwritten). This parameter should only be non-NULL for dump files. The reusefile parameter is restricted to export jobs.

Substitution Variables

Table 59-3 Substitution Variables Available for the Filename Parameter on DBMS_DATAPUMP.ADD_FILE

Substitution Variable	Description
%U	The %U is expanded in the resulting file names into a two-character, fixed-width, incrementing integer starting at 01, and ending at 99. For example, the dump filename of export%U would cause export01, export02, export03, and so on, to be created depending on how many files are needed to perform the export up to export99. For filenames containing the % character, the % must be represented as %% to avoid ambiguity. Note: If you have more than 100 files, then Oracle recommends that you use the %L substitution variable.

Table 59-3 (Cont.) Substitution Variables Available for the Filename Parameter on DBMS_DATAPUMP.ADD_FILE

Substitution Variable	Description
%l, %L	Specifies a system-generated unique file name. The file names can contain a substitution variable (%L), which implies that multiple files may be generated. The substitution variable is expanded in the resulting file names into a 2-digit, fixed-width, incrementing integer starting at 01 and ending at 99 which is the same as (%U). In addition, the substitution variable is expanded in the resulting file names into a 3-digit to 10-digit, variable-width, incrementing integers starting at 100 and ending at 2147483646. The width field is determined by the number of digits in the integer. For example if the current integer was 1, exp%Laa%L.dmp would resolve to exp01aa01.dmp exp02aa02.dmp and so forth up until 99. Then, the next file name would have 3 digits substituted: exp100aa100.dmp exp101aa101.dmp and so forth up until 999 where the next file would have 4 digits substituted. The substitution will continue up to the largest number substitution allowed, which is 2147483646.
%d, %D	Specifies the current day of the month from the Gregorian calendar in format DD. Note: This substitution variable cannot be used in an import file name.
%m, %M	Specifies the month in the Gregorian calendar in format MM. Note: This substitution variable cannot be used in an import file name.
%t, %T	Specifies the year, month, and day in the Gregorian calendar in this format: YYYYMMDD. Note: This substitution variable cannot be used in an import file name.
%y, %Y	Specifies the year in this format: YYYY. Note: This substitution variable cannot be used in an import file name.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to an Oracle Data Pump job.
- INVALID_ARGVAL. An invalid value was supplied for an input parameter.
- INVALID_STATE. The job is completing, or the job is past the defining state for an import or SQL_FILE job, or is past the defining state for LOG and SQL files.
- INVALID_OPERATION. A dump file was specified for a Network Import or ESTIMATE_ONLY export operation.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- This procedure adds files to an Oracle Data Pump job. You can add the following types of files to a job:
  - Log files—To record the messages associated with an operation. The Log file overwrites the previously existing files.
  - SQL files—To record the output of a `SQL_FILE` operation. The SQL file overwrites the previously existing files.
  - Dump files—To contain the data that is being moved. The Dump files do not overwrite the existing files. However, an error is generated.
- Import and `SQL_FILE` operations require that that you specify all dump files during the definition phase of the job. For Export operations, dump files can be added at any time. For example, if the user ascertains that the file space is running low during an Export, additional dump files can be added through this API. If the specified dump file already exists for an Export operation and `reusefile` is not set to 1, an error will be returned.
- For Export operations, the parallelism setting should be less than or equal to the number of dump files in the dump file set. If there are not enough dump files, then the job will not be able to maximize parallelism to the degree specified by the `SET_PARALLEL` procedure.
- For Import operations, the parallelism setting should also be less than or equal to the number of dump files in the dump file set. If there are not enough dump files, the performance will not be optimal, as multiple threads of execution try to access the same dump file.
- If the substitution variable (`%U`) is included in a filename, then multiple dump files can be specified through a single call to `ADD_FILE`. For Export operations, the new dump files will be created as they are needed. Enough dump files will be created to allow all of the processes specified by the current `SET_PARALLEL` value to be active. If one of the dump files is full, then it is closed, and a new dump file (with a new generated name) is created to take its place. If multiple `ADD_FILE`s with substitution variables have been specified for dump files in a job, then they will be used to generate dump files in a round-robin fashion. For example, if `expa%U`, `expb%U` and `expc%U` were all specified for a job having a parallelism of 6, then the initial dump files created would appear as follows: `expa01`, `expb01`, `expc01`, `expa02`, `expb02`, and `expc02`.
- If presented with dump file specifications, `expa%U`, `expb%U` and `expc%U`, then an Import or `SQL_FILE` operation will begin by attempting to open the dump files, `expa01`, `expb01`, and `expc01`. If the dump file containing the master table is not found in this set, then the operation will expand its search for dump files by incrementing the substitution variable and looking up the new filenames (for example, `expa02`, `expb02`, and `expc02`). The Oracle Data Pump API will keep expanding the search until it locates the dump file containing the master table. If the Oracle Data Pump API determines that the dump file does not exist, or that it is not part of the current dump set at any iteration, then the Oracle Data Pump API stops incrementing the substitution variable for the dump file specification that was in error. After the master table is found, the master table is used to ascertain when all of dump files in the dump file set have been located.

Examples

> **Note:**
>
> The examples in this section assume that the credentials, network ACLs, database account, and object-store information are already set up.

Example 1

The following example performs a table mode export to the Oracle Object Store.

```
CONNECT user;
Enter password: password

SET SERVEROUTPUT ON
SET ECHO ON
SET FEEDBACK 1
SET NUMWIDTH 10
SET LINESIZE 80
SET TRIMSPOOL ON
SET TAB OFF
SET PAGESIZE 100

DECLARE
  hdl          NUMBER;         -- Datapump handle
  ind          NUMBER;         -- Loop index
  le           ku$_LogEntry;   -- For WIP and error messages
  js           ku$_JobStatus;  -- The job status from get_status
  jd           ku$_JobDesc;    -- The job description from get_status
  sts          ku$_Status;     -- The status object returned by get_status
  jobState     VARCHAR2(30);   -- To keep track of job state
  dumpFile     VARCHAR2(1024)  := 'https://example.oraclecloud.com/test/
den02ten_foo3b_split_%u.dat';
  dumpType     NUMBER          := dbms_datapump.ku$_file_type_uridump_file;
  credName     VARCHAR2(1024)  := 'BMCTEST';
  logFile      VARCHAR2(1024)  := 'tkopc_export3b_cdb2.log';
  logDir       VARCHAR2(9)     := 'WORK';
  logType      NUMBER          := dbms_datapump.ku$_file_type_log_file;

BEGIN

  --
  -- Open a schema-based export job and perform defining-phase initialization.
  --
  hdl := dbms_datapump.open('EXPORT', 'TABLE');
  dbms_datapump.set_parameter(hdl, 'COMPRESSION', 'ALL');
  dbms_datapump.set_parameter(hdl, 'CHECKSUM', 1);
  dbms_datapump.add_file(hdl, logfile, logdir, null, logType);
  dbms_datapump.add_file(hdl, dumpFile, credName, '3MB', dumpType, 1);
  dbms_datapump.data_filter(hdl, 'INCLUDE_ROWS', 1);
  dbms_datapump.metadata_filter(hdl, 'TABLE_FILTER', 'FOO', '');
  --
  -- Start the job.
```

```
--
dbms_datapump.start_job(hdl);

--
-- Now grab output from the job and write to standard out.
--
jobState := 'UNDEFINED';
WHILE (jobState != 'COMPLETED') AND (jobState != 'STOPPED')
LOOP
  dbms_datapump.get_status(hdl,
         dbms_datapump.ku$_status_job_error +
         dbms_datapump.ku$_status_job_status +
         dbms_datapump.ku$_status_wip, -1, jobState,sts);
  js := sts.job_status;

  --
  -- If we received any WIP or Error messages for the job, display them.
  --
  IF (BITAND(sts.mask,dbms_datapump.ku$_status_wip) != 0) THEN
    le := sts.wip;
  ELSE
    IF (bitand(sts.mask,dbms_datapump.ku$_status_job_error) != 0) THEN
      le := sts.error;
    ELSE
      le := NULL;
    END IF;
  END IF;

  IF le IS NOT NULL THEN
    ind := le.FIRST;
    WHILE ind IS NOT NULL LOOP
      dbms_output.put_line(le(ind).LogText);
      ind := le.NEXT(ind);
    END LOOP;
  END IF;
END LOOP;

--
-- Detach from job.
--
dbms_datapump.detach(hdl);

--
-- Any exceptions that propagated to this point will be captured.
-- The details are retrieved from get_status and displayed.
--
EXCEPTION
  WHEN OTHERS THEN
    BEGIN
      dbms_datapump.get_status(hdl, dbms_datapump.ku$_status_job_error, 0,
                               jobState, sts);
      IF (BITAND(sts.mask,dbms_datapump.ku$_status_job_error) != 0) THEN
        le := sts.error;
        IF le IS NOT NULL THEN
          ind := le.FIRST;
          WHILE ind IS NOT NULL LOOP
```

```
            dbms_output.put_line(le(ind).LogText);
            ind := le.NEXT(ind);
          END LOOP;
        END IF;
      END IF;

    BEGIN
      dbms_datapump.stop_job (hdl, 1, 0, 0);
    EXCEPTION
      WHEN OTHERS THEN NULL;
    END;

    EXCEPTION
    WHEN OTHERS THEN
      dbms_output.put_line('Unexpected exception while in exception ' ||
                           'handler. sqlcode = ' || TO_CHAR(SQLCODE));
    END;
END;
/
EXIT;
```

Example 2

The follwoing example performs a table mode import from the Oracle Object Store.

```
CONNECT user;
Enter password: password

SET SERVEROUTPUT ON
SET ECHO ON
SET FEEDBACK 1
SET NUMWIDTH 10
SET LINESIZE 80
SET TRIMSPOOL ON
SET TAB OFF
SET PAGESIZE 100

DECLARE
  hdl          NUMBER;         -- Datapump handle
  ind          NUMBER;         -- Loop index
  le           ku$_LogEntry;   -- For WIP and error messages
  js           ku$_JobStatus;  -- The job status from get_status
  jd           ku$_JobDesc;    -- The job description from get_status
  sts          ku$_Status;     -- The status object returned by get_status
  jobState     VARCHAR2(30);   -- To keep track of job state
  dumpFile     VARCHAR2(1024)  := 'https://example.oraclecloud.com/test/
den02ten_foo3b_split_%u.dat';
  dumpType     NUMBER          := dbms_datapump.ku$_file_type_uridump_file;
  credName     VARCHAR2(1024)  := 'BMCTEST';
  logFile      VARCHAR2(1024)  := 'tkopc_import3b_cdb2.log';
  logDir       VARCHAR2(9)     := 'WORK';
  logType      NUMBER          := dbms_datapump.ku$_file_type_log_file;

BEGIN

  --
```

```
-- Open a schema-based export job and perform defining-phase initialization.
--
hdl := dbms_datapump.open('IMPORT', 'TABLE', NULL, 'OSI');
dbms_datapump.add_file(hdl, logfile, logdir, null, logType);
dbms_datapump.add_file(hdl, dumpFile, credName, null, dumpType);
dbms_datapump.metadata_filter(hdl, 'TABLE_FILTER', 'FOO', '');
dbms_datapump.set_parameter(hdl, 'TABLE_EXISTS_ACTION', 'REPLACE');
dbms_datapump.set_parameter(hdl, 'VERIFY_CHECKSUM', 1);

--
-- Start the job.
--
dbms_datapump.start_job(hdl);

--
-- Now grab output from the job and write to standard out.
--
jobState := 'UNDEFINED';
WHILE (jobState != 'COMPLETED') AND (jobState != 'STOPPED')
LOOP
  dbms_datapump.get_status(hdl,
         dbms_datapump.ku$_status_job_error +
         dbms_datapump.ku$_status_job_status +
         dbms_datapump.ku$_status_wip, -1, jobState,sts);
  js := sts.job_status;

  --
  -- If we received any WIP or Error messages for the job, display them.
  --
  IF (BITAND(sts.mask,dbms_datapump.ku$_status_wip) != 0) THEN
    le := sts.wip;
  ELSE
    IF (bitand(sts.mask,dbms_datapump.ku$_status_job_error) != 0) THEN
      le := sts.error;
    ELSE
      le := NULL;
    END IF;
  END IF;

  IF le IS NOT NULL THEN
    ind := le.FIRST;
    WHILE ind IS NOT NULL LOOP
      dbms_output.put_line(le(ind).LogText);
      ind := le.NEXT(ind);
    END LOOP;
  END IF;
END LOOP;

--
-- Detach from job.
--
dbms_datapump.detach(hdl);

--
-- Any exceptions that propagated to this point will be captured.
-- The details are retrieved from get_status and displayed.
```

```
  --
  EXCEPTION
    WHEN OTHERS THEN
      BEGIN
        dbms_datapump.get_status(hdl, dbms_datapump.ku$_status_job_error, 0,
                                 jobState, sts);
        IF (BITAND(sts.mask,dbms_datapump.ku$_status_job_error) != 0) THEN
          le := sts.error;
          IF le IS NOT NULL THEN
            ind := le.FIRST;
            WHILE ind IS NOT NULL LOOP
              dbms_output.put_line(le(ind).LogText);
              ind := le.NEXT(ind);
            END LOOP;
          END IF;
        END IF;

      BEGIN
        dbms_datapump.stop_job (hdl, 1, 0, 0);
      EXCEPTION
        WHEN OTHERS THEN NULL;
      END;

      EXCEPTION
      WHEN OTHERS THEN
        dbms_output.put_line('Unexpected exception while in exception ' ||
                             'handler. sqlcode = ' || TO_CHAR(SQLCODE));
      END;
END;
/
EXIT;
```

59.5.2 ATTACH Function

This function provides access to a previously created job.

Syntax

```
DBMS_DATAPUMP.ATTACH(
   job_name    IN VARCHAR2 DEFAULT NULL,
   job_owner    IN VARCHAR2 DEFAULT NULL)
 RETURN NUMBER;
```

Parameters

Table 59-4 ATTACH Function Parameters

Parameter	Description
job_name	The name of the job. The default is the job name owned by the user who is specified in the job_owner parameter (assuming that user has only one job in the Defining, Executing, or Idling states).

Table 59-4 (Cont.) ATTACH Function Parameters

Parameter	Description
job_owner	The user who originally started the job. If NULL, then the value defaults to the owner of the current session. To specify a job owner other than yourself, you must have either the DATAPUMP_EXP_FULL_DATABASE role (for export operations) or the DATAPUMP_IMP_FULL_DATABASE role (for import and SQL_FILE operations). Being a privileged user allows you to monitor another user's job, but you cannot restart another user's job.

Return Values

An opaque handle for the job. This handle is used as input to the following procedures: ADD_FILE, DATA_FILTER, DETACH, GET_STATUS, LOG_ENTRY, METADATA_FILTER, METADATA_REMAP, METADATA_TRANSFORM, SET_PARALLEL, SET_PARAMETER, START_JOB, STOP_JOB, and WAIT_FOR_JOB.

Exceptions

- INVALID_ARGVAL. An invalid value was supplied for an input parameter.
- OBJECT_NOT_FOUND. The specified job no longer exists or the user specified a job owned by another schema, but the user did not have the DATAPUMP_EXP_FULL_DATABASE or DATAPUMP_IMP_FULL_DATABASE role.
- SUCCESS_WITH_INFO. The function succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- If the job was in the Stopped state, then the job is placed into the Idling state. After the ATTACH succeeds, you can monitor the progress of the job or control the job. The stream of KU$_STATUS_WIP and KU$_STATUS_JOB_ERROR messages returned through the GET_STATUS procedure will be returned to the newly attached job starting at the approximate time of the client's attachment. There will be no repeating of status and error messages that were processed before the client attached to a job.
- If you want to perform a second attach to a job, then you must do so from a different session.
- If the ATTACH fails, then use a null handle in a subsequent call to GET_STATUS for more information about the failure.

59.5.3 DATA_FILTER Procedures

This procedure specifies restrictions on the rows that are to be retrieved.

Syntax

```
DBMS_DATAPUMP.DATA_FILTER (
   handle      IN NUMBER,
   name        IN VARCHAR2,
   value       IN NUMBER,
   table_name  IN VARCHAR2 DEFAULT NULL,
   schema_name IN VARCHAR2 DEFAULT NULL);
```

```
DBMS_DATAPUMP.DATA_FILTER(
   handle      IN NUMBER,
   name        IN VARCHAR2,
   value       IN VARCHAR2,
   table_name  IN VARCHAR2 DEFAULT NULL,
   schema_name IN VARCHAR2 DEFAULT NULL);

DBMS_DATAPUMP.DATA_FILTER(
   handle      IN NUMBER,
   name        IN VARCHAR2,
   value       IN CLOB,
   table_name  IN VARCHAR2 DEFAULT NULL,
   schema_name IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 59-5 DATA_FILTER Procedure Parameters

Parameter	Description
handle	The handle that is returned from the OPEN function
name	The name of the filter
value	The value of the filter
table_name	The name of the table on which the data filter is applied. If no table name is supplied, the filter applies to all tables in the job.
schema_name	The name of the schema that owns the table on which the filter is applied. If no schema name is specified, the filter applies to all schemas in the job. If you supply a schema name you must also supply a table name.

Exceptions

- INVALID_ARGVAL. There can be several reasons for this message:
  - A bad filter name is specified
  - The mode is TRANSPORTABLE, which does not support data filters
  - The specified table does not exist
  - The filter has already been set for the specified values of schema_name and table_name
- INVALID_STATE. The user called DATA_FILTER when the job was not in the Defining state.
- INCONSISTENT_ARGS. The value parameter is missing or its datatype does not match the filter name. Or a schema name was supplied, but not a table name.
- PRIVILEGE_ERROR. A schema name was supplied, but the user did not have the DATAPUMP_EXP_FULL_DATABASE or DATAPUMP_IMP_FULL_DATABASE role.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- Each data filter can only appear once in each table (for example, you cannot supply multiple SUBQUERY filters to a table) or once in each job. If different filters using the same

name are applied to both a particular table and to the whole job, the filter parameter supplied for the specific table will take precedence.

With the exception of the `INCLUDE_ROWS` filter, data filters are not supported on tables having nested tables or domain indexes defined upon them. Data filters are not supported in jobs performed in Transportable Tablespace mode.

The available data filters are described in Table 59-6.

Table 59-6 Data Filters

Name	Datatype	Operations that Support Filter	Description
`INCLUDE_ROWS`	`NUMBER`	EXPORT, IMPORT	If nonzero, this filter specifies that user data for the specified table should be included in the job. The default is 1.
`PARTITION_EXPR` `PARTITION_LIST`	`TEXT`	EXPORT, IMPORT	**Note:** In this description, the information about partitions also applies to subpartions. For Export jobs, these filters specify which partitions are unloaded from the database. For Import jobs, they specify which table partitions are loaded into the database. Partition names are included in the job if their names satisfy the specified expression (for `PARTITION_EXPR`) or are included in the list (for `PARTITION_LIST`). Whereas the expression version of the filter offers more flexibility, the list version provides for full validation of the partition names. Double quotation marks around partition names are required only if the partition names contain special characters. `PARTITION_EXPR` is not supported on jobs across a network link. Default=All partitions are processed
`SAMPLE`	`NUMBER`	EXPORT, IMPORT	For Export jobs, specifies a percentage for sampling the data blocks to be moved. This filter allows subsets of large tables to be extracted for testing purposes.
`SUBQUERY`	`TEXT`	EXPORT, IMPORT	Specifies a subquery that is added to the end of the `SELECT` statement for the table. If you specify a `WHERE` clause in the subquery, you can restrict the rows that are selected. Specifying an `ORDER BY` clause orders the rows dumped in the export which improves performance when migrating from heap-organized tables to index-organized tables.

59.5.4 DATA_REMAP Procedure

This procedure specifies transformations to be applied to column data as it is exported from, or imported into, a database.

Syntax

```
DBMS_DATAPUMP.DATA_REMAP(
   handle          IN NUMBER,
   name            IN VARCHAR2,
   table_name      IN VARCHAR2,
   column          IN VARCHAR2,
   remap_function  IN VARCHAR2,
   schema          IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 59-7 DATA_REMAP Procedure Parameters

Parameter	Description
handle	The handle of the current job. The current session must have previously attached to the handle through a call to an OPEN function.
name	The name of the remap
table_name	The table containing the column to be remapped
column	The name of the column to be remapped
remap_function	The meaning of remap_function is dependent upon the value of name. See Table 59-8 for a list of possible names.
schema	The schema containing the column to be remapped. If NULL, the remapping applies to all schemas moved in the job that contain the specified table.

Exceptions

- INVALID_ARGVAL. The mode is transportable (which does not support data modifications) or it has specified that no data to be included in the job. An invalid remap name was supplied.
- INVALID_OPERATION. Data remaps are only supported for Export and Import operations.
- INVALID_STATE. The DATA_REMAP procedure was called after the job started (that is, it was not in the defining state).
- NO_SUCH_JOB. The job handle is no longer valid.

Usage Notes

- The DATA_REMAP procedure is only supported for Export and Import operations. It allows you to manipulate user data being exported or imported. The name of the remap determines the remap operation to be performed.
- For export operations, you might wish to define a data remap to obscure sensitive data such as credit card numbers from a dump file, but leave the remainder of the data so that it can be read. To accomplish this, the remapping should convert each unique source number into a distinct generated number. So that the mapping is consistent across the dump file set, the same function should be called for every column that contains the credit card number.

- For import operations, you might wish to define a data remap to reset the primary key when data is being merged into an existing table that contains colliding primary keys. A single remapping function should be provided for all columns defining or referencing the primary key to ensure that remapping is consistent.

> **Note:**
>
> If the called function uses package state variables, then to ensure that remapping is performed consistently across all tables, the job should be run with a SET_PARALLEL value of 1 and no restart operations should be performed.

The Data Remap functions are listed in Table 59-8.

Table 59-8 Names of Data Remap Functions

Name	Meaning of `remap_function`	Meaning
`COLUMN_FUNCTION`	String having the format: `[schema.]package.function`	The `name` parameter references a PL/SQL package function which is called to modify the data for the specified column. The function accepts a single parameter, which has the same datatype as the remapped column, and returns a value having the same datatype as the remapped column. Note that the default for the schema is the schema of the user performing the export.

59.5.5 DETACH Procedure

This procedure specifies that the user has no further interest in using the handle.

Syntax

```
DBMS_DATAPUMP.DETACH(
   handle  IN NUMBER);
```

Parameters

Table 59-9 DETACH Procedure Parameters

Parameter	Description
`handle`	The handle of the job. The current session must have previously attached to the handle through a call to either an `OPEN` or `ATTACH` function.

Exceptions

- `INVALID_HANDLE`. The specified handle is not attached to a Data Pump job.
- `SUCCESS_WITH_INFO`. The procedure succeeded, but further information is available through the `GET_STATUS` procedure.

- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- Through this call, you specify that you have no further interest in using the handle. Resources associated with a completed job cannot be reclaimed until all users are detached from the job. An implicit detach from a handle is performed when the user's session is exited or aborted. An implicit detach from a handle is also performed upon the expiration of the timeout associated with a STOP_JOB that was applied to the job referenced by the handle. All previously allocated DBMS_DATAPUMP handles are released when an instance is restarted.

59.5.6 GET_DUMPFILE_INFO Procedure

This procedure retrieves information about a specified dump file.

Syntax

```
DBMS_DATAPUMP.GET_DUMPFILE_INFO(
   filename    IN VARCHAR2,
   directory   IN VARCHAR2,
   info_table  OUT ku$_dumpfile_info,
   filetype    OUT NUMBER);
```

Parameters

Table 59-10 GET_DUMPFILE_INFO Procedure Parameters

Parameter	Description
filename	A simple filename with no directory path information
directory	A directory object that specifies where the file can be found
info_table	A PL/SQL table for storing information about the dump file
filetype	The type of file (Data Pump dump file, original Export dump file, external tables dump file, or unknown)

Exceptions

The GET_DUMPFILE_INFO procedure is a utility routine that operates outside the context of any Data Pump job. Exceptions are handled differently for this procedure than for procedures associated in some way with a Data Pump job. A full exception stack should be available directly, without the need to call the GET_STATUS procedure to retrieve the detailed information. The exception for this procedure is as follows:

- NO_DUMPFILE_INFO. Unable to retrieve dump file information as specified.

Usage Notes

You can use the GET_DUMPFILE_INFO procedure to request information about a specific file. If the file is not recognized as any type of dump file, then a filetype of 0 (zero) is returned and the dump file info_table remains empty.

A filetype value of 1 indicates a Data Pump dump file. A filetype value of 2 indicates an original Export dump file. A filetype value of 3 indicates an external tables dump file. In all cases, the dump file info_table will be populated with information retrieved from the dump file header. Rows of this table consist of item code and value pairs, where the item code indicates the type

of information and the value column is a VARCHAR2 containing the actual data (converted to a string in some cases). The table is defined as follows:

```
CREATE TYPE sys.ku$_dumpfile_item IS OBJECT (
               item_code       NUMBER,           -- Identifies header item
               value           VARCHAR2(2048)    -- Text string value)/

GRANT EXECUTE ON sys.ku$_dumpfile_item TO PUBLIC;
CREATE OR REPLACE PUBLIC SYNONYM ku$_dumpfile_item FOR sys.ku$_dumpfile_item;

CREATE TYPE sys.ku$_dumpfile_info AS TABLE OF sys.ku$_dumpfile_item/

GRANT EXECUTE ON sys.ku$_dumpfile_info TO PUBLIC;
CREATE OR REPLACE PUBLIC SYNONYM ku$_dumpfile_info FOR sys.ku$_dumpfile_info;
```

The item codes, which can easily be extended to provide more information as needed, are currently defined as shown in Table 59-11 (prepended with the package name, DBMS_DATAPUMP.). Assume the following with regard to these item codes:

- Unless otherwise stated, all item codes may be returned only for Oracle Data Pump and external tables dump files (filetypes 1 and 3).
- Unless otherwise stated, all item codes have been available since Oracle Database 10*g* Release 2 (10.2).

Table 59-11 Item Codes For the DBMS_DATAPUMP.GET_DUMPFILE_INFO Procedure

Item Code	Description
KU$_DFHDR_FILE_VERSION	The internal file version of the dump file.
KU$_DFHDR_MASTER_PRESENT	If the Data Pump master table is present in the dump file, then the value for this item code is 1; otherwise the value is 0. Returned only for filetype 1.
KU$_DFHDR_GUID	A unique identifier assigned to the Data Pump export job or the external tables unload job that produced the dump file. For a multifile dump set, each file in the set has the same value for this item code.
KU$_DFHDR_FILE_NUMBER	A numeric identifier assigned to the dump file. Each dump file in a multifile dump set has its own identifier, unique only within the dump set.
KU$_DFHDR_CHARSET_ID	A numeric code that represents the character set in use at the source system when the dump file was created. Returned for all filetypes.
KU$_DFHDR_CREATION_DATE	The date and time that the dump file was created.
KU$_DFHDR_FLAGS	Internal flag values.
KU$_DFHDR_JOB_NAME	The name assigned to the export job that created the dump file. Returned only for filetype 1.
KU$_DFHDR_PLATFORM	The operating system name of the source system on which the dump file was created.
KU$_DFHDR_INSTANCE	The instance name of the source system on which the dump file was created.
KU$_DFHDR_LANGUAGE	The language name that corresponds to the character set of the source system where the export dump file was created.
KU$_DFHDR_BLOCKSIZE	The blocksize, in bytes, of the dump file.
KU$_DFHDR_DIRPATH	If direct path mode was used when the dump file was created, then the value for this item code is 1, otherwise the value is 0. Returned only for filetype 2.

Table 59-11 (Cont.) Item Codes For the DBMS_DATAPUMP.GET_DUMPFILE_INFO Procedure

Item Code	Description
KU$_DFHDR_METADATA_COMPRESSED	If the system metadata is stored in the dump file in compressed format, then the value for this item code is 1, otherwise the value is 0. Returned only for filetype 1.
KU$_DFHDR_DB_VERSION	The database job version used to create the dump file. Returned for all filetypes.
KU$_DFHDR_MASTER_PIECE_COUNT	The Data Pump master table may be split into multiple pieces and written to multiple dump files in the set, one piece per file. The value returned for this item code indicates the number of dump files that contain pieces of the master table. The value for this item code is only meaningful if the Data Pump master table is present in the dump file, as indicated by the item code KU$_DFHDR_MASTER_PRESENT. Returned only for filetype 1. Only available since Oracle Database 11*g* Release 1 (11.1).
KU$_DFHDR_MASTER_PIECE_NUMBER	The Data Pump master table may be split into multiple pieces and written to multiple dump files in the set, one piece per file. The value returned for this item code indicates which master table piece is contained in the dump file. The value for this item code is only meaningful if the Data Pump master table is present in the dump file, as indicated by the item code KU$_DFHDR_MASTER_PRESENT. Returned only for filetype 1. Only available since Oracle Database 11*g* Release 1 (11.1).
KU$_DFHDR_DATA_COMPRESSED	If the table data is stored in the dump file in compressed format, then the value for this item code is 1, otherwise the value is 0. Only available since Oracle Database 11*g* Release 1 (11.1).
KU$_DFHDR_METADATA_ENCRYPTED	If the system metadata is stored in the dump file in encrypted format, then the value for this item code is 1, otherwise the value is 0. Returned only for filetype 1. Only available since Oracle Database 11*g* Release 1 (11.1).
KU$_DFHDR_DATA_ENCRYPTED	If the table data is stored in the dump file in encrypted format, then the value for this item code is 1, otherwise the value is 0. Only available since Oracle Database 11*g* Release 1 (11.1).
KU$_DFHDR_COLUMNS_ENCRYPTED	If encrypted column data is stored in the dump file in encrypted format, then the value for this item code is 1, otherwise the value is 0. Returned only for filetype 1. Only available since Oracle Database 11*g* Release 1 (11.1).

Table 59-11 (Cont.) Item Codes For the DBMS_DATAPUMP.GET_DUMPFILE_INFO Procedure

Item Code	Description
KU$_DFHDR_ENCRYPTION_MODE	The encryption mode indicates whether a user-provided password or the Oracle Encryption Wallet was used to encrypt data written to the dump file. The possible values returned for this item code are: • KU$_DFHDR_ENCMODE_NONE No data was written to the dump file in encrypted format. • KU$_DFHDR_ENCMODE_PASSWORD Data was written to the dump file in encrypted format using a provided password. • KU$_DFHDR_ENCMODE_DUAL Data was written to the dump file in encrypted format using both a provided password as well as an Oracle Encryption Wallet. • KU$_DFHDR_ENCMODE_TRANS Data was written to the dump file in encrypted format transparently using an Oracle Encryption Wallet. Only available since Oracle Database 11*g* Release 1 (11.1).
KU$_DFHDR_COMPRESSION_ALG	The compression algorithm used when writing system metadata and/or table data to the dump file in compressed format. The possible values returned for this item code are: • KU$_DFHDR_CMPALG_NONE No data was written to the dump file in compressed format. • KU$_DFHDR_CMPALG_BASIC Data was written to the dump file in compressed format using an internal algorithm. This is the default algorithm used since Oracle Database 10*g* Release 2 (10.2). • KU$_DFHDR_CMPALG_LOW Data was written to the dump file in compressed format using the LOW algorithm. • KU$_DFHDR_CMPALG_MEDIUM Data was written to the dump file in compressed format using the MEDIUM algorithm. • KU$_DFHDR_CMPALG_HIGH Data was written to the dump file in compressed format using the HIGH algorithm. Only available since Oracle Database 12*c* Release 1 (12.1).

59.5.7 GET_STATUS Procedure

This procedure monitors the status of a job or waits for the completion of a job.

Syntax

```
DBMS_DATAPUMP.GET_STATUS(
   handle    IN NUMBER,
   mask      IN BINARY_INTEGER,
   timeout   IN NUMBER DEFAULT NULL,
   job_state OUT VARCHAR2,
   status    OUT ku$_Status);
```

Parameters

Table 59-12 GET_STATUS Procedure Parameters

Parameter	Description
handle	The handle of a job. The current session must have previously attached to the handle through a call to either the OPEN or ATTACH function. A null handle can be used to retrieve error information after OPEN and ATTACH failures.
mask	A bit mask that indicates which of four types of information to return: • KU$_STATUS_WIP • KU$_STATUS_JOB_DESC • KU$_STATUS_JOB_STATUS • KU$_STATUS_JOB_ERROR Each status has a numerical value. You can request multiple types of information by adding together different combinations of values.
timeout	Maximum number of seconds to wait before returning to the user. A value of 0 requests an immediate return. A value of -1 requests an infinite wait. If KU$_STATUS_WIP or KU$_STATUS_JOB_ERROR information is requested and becomes available during the timeout period, then the procedure returns before the timeout period is over.
job_state	Current state of the job. If only the job state is needed, it is much more efficient to use this parameter than to retrieve the full ku$_Status structure.
status	A ku$_Status is returned. The ku$_Status mask indicates what kind of information is included. This could be none if only KU$_STATUS_WIP or KU$_STATUS_JOB_ERROR information is requested and the timeout period expires.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_VALUE. The mask or timeout contains an illegal value.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

The GET_STATUS procedure is used to monitor the progress of an ongoing job and to receive error notification. You can request various type of information using the mask parameter. The KU$_STATUS_JOB_DESC and KU$_STATUS_JOB_STATUS values are classified as synchronous information because the information resides in the master table. The KU$_STATUS_WIP and KU$_STATUS_JOB_ERROR values are classified as asynchronous because the messages that embody these types of information can be generated at any time by various layers in the Data Pump architecture.

- If synchronous information *only* is requested, the interface will ignore the timeout parameter and simply return the requested information.
- If asynchronous information is requested, the interface will wait a *maximum* of timeout seconds before returning to the client. If a message of the requested asynchronous information type is received, the call will complete prior to timeout seconds. If synchronous information was also requested, it will be returned whenever the procedure returns.

- If the job_state returned by GET_STATUS does not indicate a terminating job, it is possible that the job could still terminate before the next call to GET_STATUS. This would result in an INVALID_HANDLE exception. Alternatively, the job could terminate during the call to GET_STATUS, which would result in a NO_SUCH_JOB exception. Callers should be prepared to handle these cases.

Error Handling

There are two types of error scenarios that need to be handled using the GET_STATUS procedure:

- Errors resulting from other procedure calls: For example, the SET_PARAMETER procedure may produce an INCONSISTENT_ARGS exception. The client should immediately call GET_STATUS with mask=8 (errors) and timeout=0. The returned ku$_Status.error will contain a ku$_LogEntry that describes the inconsistency in more detail.
- Errors resulting from events asynchronous to the client(s): An example might be Table already exists when trying to create a table. The ku$_Status.error will contain a ku$_LogEntry with all error lines (from all processing layers that added context about the error) properly ordered.

After a job has begun, a client's main processing loop will typically consist of a call to GET_STATUS with an infinite timeout (-1) "listening" for KU$_STATUS_WIP and KU$_STATUS_JOB_ERROR messages. If status was requested, then JOB_STATUS information will also be in the request.

When the ku$_Status is interpreted, the following guidelines should be used:

- ku$_Status.ku$_JobStatus.percent_done refers only to the amount of data that has been processed in a job. Metadata is not considered in the calculation. It is determined using the following formulas:
  - EXPORT or network IMPORT--(bytes_processed/estimated_bytes) * 100
  - IMPORT--(bytes_processed/total_expected_bytes) * 100
  - SQL_FILE or estimate-only EXPORT--0.00 if not done or 100.00 if done

  The effects of the QUERY and PARTITION_EXPR data filters are not considered in computing percent_done.

  It is expected that the status returned will be transformed by the caller into more user-friendly status. For example, when percent done is not zero, an estimate of completion time could be produced using the following formula:

  ```
  ((SYSDATE - start time) / ku$_Status.ku$_JobStatus.percent_done) * 100
  ```
- The caller should not use ku$_Status.ku$_JobStatus.percent_done for determining whether the job has completed. Instead, the caller should only rely on the state of the job as found in job_state.

59.5.8 LOG_ENTRY Procedure

This procedure inserts a message into the log file.

Syntax

```
DBMS_DATAPUMP.LOG_ENTRY(
   handle         IN NUMBER,
   message        IN VARCHAR2
   log_file_only  IN NUMBER DEFAULT 0);
```

Parameters

Table 59-13 LOG_ENTRY Procedure Parameters

Parameter	Description
handle	The handle of a job. The current session must have previously attached to the handle through a call to either the OPEN or ATTACH function.
message	A text line to be added to the log file
log_file_only	Specified text should be written only to the log file. It should not be returned in GET_STATUS work-in-progress (KU$_STATUS_WIP) messages.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

The message is added to the log file. If log_file_only is zero (the default), the message is also broadcast as a KU$_STATUS_WIP message through the GET_STATUS procedure to all users attached to the job.

The LOG_ENTRY procedure allows applications to tailor the log stream to match the abstractions provided by the application. For example, the command-line interface supports INCLUDE and EXCLUDE parameters defined by the user. Identifying these values as calls to the underlying METADATA_FILTER procedure would be confusing to users. Instead, the command-line interface can enter text into the log describing the settings for the INCLUDE and EXCLUDE parameters.

Lines entered in the log stream from LOG_ENTRY are prefixed by the string, ";;; "

59.5.9 METADATA_FILTER Procedure

This procedure provides filters that allow you to restrict the items that are included in a job.

Syntax

```
DBMS_DATAPUMP.METADATA_FILTER(
   handle       IN NUMBER,
   name         IN VARCHAR2,
   value        IN VARCHAR2,
   object_path  IN VARCHAR2 DEFAULT NULL);

DBMS_DATAPUMP.METADATA_FILTER(
   handle       IN NUMBER,
   name         IN VARCHAR2,
   value        IN CLOB,
   object_path  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 59-14 METADATA_FILTER Procedure Parameters

Parameter	Description
handle	The handle returned from the OPEN function
name	The name of the filter. See Table 59-15 for descriptions of the available filters.
value	The value of the filter
object_path	The object path to which the filter applies. If the default is used, the filter applies to all applicable objects. Lists of the object paths supported for each mode are contained in the catalog views for DATABASE_EXPORT_OBJECTS, SCHEMA_EXPORT_OBJECTS, and TABLE_EXPORT_OBJECTS. (Note that the TABLE_EXPORT_OBJECTS view is applicable to both Table and Tablespace mode because their object paths are the same.) For an import operation, object paths reference the mode used to create the dump file rather than the mode being used for the import.

Table 59-15 describes the name, the object type, and the meaning of the filters available with the METADATA_FILTER procedure. The datatype for all the filters is a text expression. All operations support all filters.

Table 59-15 Filters Provided by METADATA_FILTER Procedure

Name	Object Type	Meaning
NAME_EXPR NAME_LIST	Named objects	Defines which object names are included in the job. You use the object type parameter to limit the filter to a particular object type. For Table mode, identifies which tables are to be processed.
SCHEMA_EXPR SCHEMA_LIST	Schema objects	Restricts the job to objects whose owning schema name is satisfied by the expression. For Table mode, only a single SCHEMA_EXPR filter is supported. If specified, it must only specify a single schema (for example, 'IN (''SCOTT'')'). For Schema mode, identifies which users are to be processed.
TABLESPACE_EXPR TABLESPACE_LIST	TABLE, CLUSTER, INDEX, ROLLBACK_SEGMENT	Restricts the job to objects stored in a tablespace whose name is satisfied by the expression. For Tablespace mode, identifies which tablespaces are to be processed. If a partition of an object is stored in the tablespace, the entire object is added to the job. For Transportable mode, identifies which tablespaces are to be processed. If a table has a single partition in the tablespace set, all partitions must be in the tablespace set. An index is not included within the tablespace set unless all of its partitions are in the tablespace set. A domain index is not included in the tablespace set unless all of its secondary objects are included in the tablespace set.

Table 59-15 (Cont.) Filters Provided by METADATA_FILTER Procedure

Name	Object Type	Meaning
INCLUDE_PATH_EXPR INCLUDE_PATH_LIST EXCLUDE_PATH_EXPR EXCLUDE_PATH_LIST	All	Defines which object paths are included in, or excluded from, the job. You use these filters to select only certain object types from the database or dump file set. Objects of paths satisfying the condition are included (INCLUDE_PATH_*) or excluded (EXCLUDE_PATH_*) from the operation. The object_path parameter is not supported for these filters.
EXCLUDE_TABLES	TABLE_EXPORT	Specifies that no tables are to be exported. The EXCLUDE_TABLES parameter is needed only when VIEWS_AS_TABLES is specified, and when there are no tables in a table mode export.
VIEWS_AS_TABLES	TABLE_EXPORT	A comma-separated list of views that you want to be exported as tables: [*schema_name*.]*view_name*[:*table_name*] The filter can be called multiple times with multiple values. All values are added to a list. All views on the list are exported as tables.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_ARGVAL. This exception can indicate any of the following conditions:
  - An object_path was specified for an INCLUDE_PATH_EXPR or EXCLUDE_PATH_EXPR filter.
  - The specified object_path is not supported for the current mode.
  - The SCHEMA_EXPR filter specified multiple schemas for a Table mode job.
- INVALID_STATE. The user called the METADATA_FILTER procedure after the job left the defining state.
- INCONSISTENT_ARGS. The filter value is of the wrong datatype, or the filter is missing.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- Metadata filters identify a set of objects that you want to be included or excluded from an Oracle Data Pump operation. Except for EXCLUDE_PATH_EXPR and INCLUDE_PATH_EXPR, dependent objects of an identified object will be processed along with the identified object. For example, if an index is identified for inclusion by a filter, then grants on that index will also be included by the filter. Likewise, if a table is excluded by a filter, then indexes, constraints, grants and triggers on the table will also be excluded by the filter.
- Two versions of each filter are supported: SQL expression and List. The SQL expression version of the filters offer maximum flexibility for identifying objects (for example the use of LIKE to support use of wild cards). The names of the expression filters are as follows:

  - NAME_EXPR
  - SCHEMA_EXPR
  - TABLESPACE_EXPR
  - INCLUDE_PATH_EXPR
  - EXCLUDE_PATH_EXPR

  The list version of the filters allow maximum validation of the filter. An error will be reported if one of the elements in the filter is not found within the source database (for Export and network-based jobs) or is not found within the dump file (for file-based Import and SQLFILE jobs). The names of the list filters are as follows:

  - NAME_LIST
  - SCHEMA_LIST
  - TABLESPACE_LIST
  - INCLUDE_PATH_LIST
  - EXCLUDE_PATH_LIST

- Filters allow a user to restrict the items that are included in a job. For example, a user could request a full export, but without Package Specifications or Package Bodies.
- If multiple filters are specified for a object type, they are implicitly 'ANDed' together (that is, objects participating in the job must pass all of the filters applied to their object types).
- The same filter name can be specified multiple times within a job. For example, specifying `NAME_EXPR as '!=''EMP'''` and `NAME_EXPR as '!=''DEPT'''` on a Table mode export would produce a file set containing all of the tables except for `EMP` and `DEPT`.
- `EXCLUDE_TABLES` is only used in concert with `VIEWS_AS_TABLES`, so that you can exclude tables from an export while enabling exports of views as tables. For example, if the only thing you want to do is to export the contents of a view as a table, then you can use `VIEWS_AS_TABLES` to export the views, but not the tables. If you want to export some actual tables, and you also want export other views as tables, then you would select `VIEWS_AS_TABLES`, but not select `EXCLUDE_TABLES`.

DBMS_DATAPUMP Examples

In the following example, we use the metadata filters `EXCLUDE_TABLES` with `VIEWS_AS_TABLES` to export a view as a table, while excluding the table from which the view is created.

Suppose you have a view called `SCOTT.EMPV` that is defined for a specific set of columns in the table `SCOTT.EMP` (for example, only the columns `empno,ename` and `hiredate`). You want to export that view as a table to a target database using the Oracle Data Pump API, so that you only expose that subset of columns from `SCOTT.EMP` in your export. In this case, you can run the export as follows:

```
BEGIN
  h := dbms_datapump.open('EXPORT' ,'TABLE');
  dbms_datapump.metadata_filter(h, 'EXCLUDE_TABLES', 'Y');
  dbms_datapump.metadata_filter(h, 'VIEWS_AS_TABLES','SCOTT.EMPV');
:
:
:
 dbms_datapump.start_job(h, 0, 0);
```

:
:

If you don't use the EXCLUDE_TABLES filter, then the job fails, because there are no tables included in a table-model export. However, because we only want to export the view SCOTT.EMPV as if it was a table, and we specifically do not want to export the table SCOTT.EMP, we set EXCLUDE_TABLES. The tables in the source are excluded, and the view SCOTT.EMPV is exported to the target as a table.

59.5.10 METADATA_REMAP Procedure

This procedure specifies a remapping to be applied to objects as they are processed in the specified job.

Syntax

```
DBMS_DATAPUMP.METADATA_REMAP (
   handle      IN NUMBER,
   name        IN VARCHAR2,
   old_value   IN VARCHAR2,
   value       IN VARCHAR2,
   object_type IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 59-16 METADATA_REMAP Procedure Parameters

Parameter	Description
handle	The handle for the current job. The current session must have previously attached to the handle through a call to the OPEN function.
name	The name of the remap. See Table 59-17 for descriptions of the available remaps.
old_value	Specifies which value in the dump file set should be reset to value
value	The value of the parameter for the remap. This signifies the new value that old_value should be translated into.
object_type	Designates the object type to which the remap applies. The list of object types supported for each mode are contained in the DATABASE_EXPORT_OBJECTS, SCHEMA_EXPORT_OBJECTS, TABLE_EXPORT_OBJECTS, and TABLESPACE_EXPORT_OBJECTS catalog views. By default, the remap applies to all applicable objects within the job. The object_type parameter allows a caller to specify different parameters for different object types within a job. Remaps that explicitly specify an object type override remaps that apply to all object types.

Table 59-17 describes the remaps provided by the METADATA_REMAP procedure.

Table 59-17 Remaps Provided by the METADATA_REMAP Procedure

Name	Datatype	Object Type	Meaning
REMAP_SCHEMA	TEXT	Schema objects	Any schema object in the job that matches the object_type parameter and was located in the old_value schema will be moved to the value schema. Privileged users can perform unrestricted schema remaps. Nonprivileged users can perform schema remaps only if their schema is the target schema of the remap. For example, SCOTT can remap his BLAKE's objects to SCOTT, but SCOTT cannot remap SCOTT's objects to BLAKE.
REMAP_TABLESPACE	TEXT	TABLE, INDEX, ROLLBACK_SEGMENT, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG, TABLE_SPACE	Any storage segment in the job that matches the object_type parameter and was located in the old_value tablespace will be relocated to the value tablespace.
REMAP_DATAFILE	TEXT	LIBRARY, TABLESPACE, DIRECTORY	If old_value and value are both full file specifications, then any data file reference in the job that matches the object_type parameter and that referenced the old_value data file will be redefined to use the value data file. If old_value and value are both directory paths, then any data file reference whose object path matches old_value will have its path substituted with value.
REMAP_TABLE	TEXT	TABLE	Any reference to a table in the job that matches the old_value table name will be replaced with the value table name. The old_value parameter may refer to a partition such as employees.low. This allows names for tables constructed the by PARTITION_OPTIONS=DEPARTITION parameter to be specified by the user.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_ARGVAL. This message can indicate any of the following:

  - The job's mode does not include the specified `object_type`.
  - The remap has already been specified for the specified `old_value` and `object_type`.
- `INVALID_OPERATION`. Remaps are only supported for SQL_FILE and Import operations. The job's operation was Export, which does not support the use of metadata remaps.
- `INVALID_STATE`. The user called `METADATA_REMAP` after the job had started (that is, the job was not in the defining state).
- `INCONSISTENT_ARGS`. There was no `value` supplied or it was of the wrong datatype for the remap.
- `PRIVILEGE_ERROR`. A nonprivileged user attempted to do a `REMAP_SCHEMA` to a different user's schema or a `REMAP_DATAFILE`.
- `SUCCESS_WITH_INFO`. The procedure succeeded, but further information is available through the `GET_STATUS` procedure.
- `NO_SUCH_JOB`. The specified job does not exist.

Usage Notes

- The `METADATA_REMAP` procedure is only supported for Import and SQL_FILE operations. It enables you to apply commonly desired, predefined remappings to the definition of objects as part of the transfer. If you need remaps that are not supported within this procedure, you should do a preliminary SQL_FILE operation to produce a SQL script corresponding to the dump file set. By editing the DDL directly and then executing it, you can produce any remappings that you need.
- Transforms for the DataPump API are a subset of the remaps implemented by the `DBMS_METADATA.SET_TRANSFORM_PARAMETER` API. Multiple remaps can be defined for a single job. However, each remap defined must be unique according its parameters. That is, two remaps cannot specify conflicting or redundant remaps.

59.5.11 METADATA_TRANSFORM Procedure

This procedure specifies transformations to be applied to objects as they are processed in the specified job.

Syntax

```
DBMS_DATAPUMP.METADATA_TRANSFORM (
   handle      IN NUMBER,
   name        IN VARCHAR2,
   value       IN VARCHAR2,
   object_type IN VARCHAR2 DEFAULT NULL);

DBMS_DATAPUMP.METADATA_TRANSFORM (
   handle      IN NUMBER,
   name        IN VARCHAR2,
   value       IN NUMBER,
   object_type IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 59-18 METADATA_TRANSFORM Procedure Parameters

Parameter	Description
handle	The handle for the current job. The current session must have previously attached to the handle through a call to the OPEN function.
name	The name of the transformation.
value	The value of the parameter for the transform
object_type	Designates the object type to which the transform applies. The list of object types supported for each mode are contained in the DATABASE_EXPORT_OBJECTS, SCHEMA_EXPORT_OBJECTS, TABLE_EXPORT_OBJECTS, and TABLESPACE_EXPORT_OBJECTS catalog views. By default, the transform applies to all applicable objects within the job. The object_type parameter allows a caller to specify different transform parameters for different object types within a job. Transforms that explicitly specify an object type override transforms that apply to all object types.

The following table describes the transforms provided by the METADATA_TRANSFORM procedure.

Table 59-19 Transforms Provided by the METADATA_TRANFORM Procedure

Name	Datatype	Object Type	Meaning
DISABLE_ARCHIVE_LOGGING	NUMBER	TABLE INDEX	Specifies whether to disable archive logging for specified object types during import. A value of zero (FALSE) is the default. It specifies that archive logging will take place. This is the default behavior if this transform parameter is not specified. A non-zero (TRUE) value disables archive logging for the specified object types before data is loaded during import. If no object type is specified, then archive logging is disabled for both TABLE and INDEX object types. All other object types processed during data pump import are logged. **Note:** If the database is in FORCE LOGGING mode, then the DISABLE_ARCHIVE_LOGGING transform does not disable logging when indexes and tables are created.

Table 59-19 (Cont.) Transforms Provided by the METADATA_TRANFORM Procedure

Name	Datatype	Object Type	Meaning
INMEMORY	NUMBER	TABLE TABLESPACE MATERIALIZED VIEWS	The INMEMORY transform is related to the In-Memory Column Store (IM column store), an optional area in the SGA that stores whole tables, table partitions, and individual columns in a compressed columnar format. If a non-zero value (TRUE) is specified on import, then Data Pump generates an IM clause that preserves the IM setting those objects had at export time. This is the default. If a value of zero (FALSE) is specified on import, then Data Pump does not include an IM clause for any objects. **Note**: The INMEMORY transform is available only in Oracle Database 12c Release 1 (12.1.0.2) or higher.
INMEMORY_CLAUSE	TEXT	TABLE TABLESPACE MATERIALIZED VIEWS	The INMEMORY_CLAUSE transform is related to the In-Memory Column Store (IM column store), an optional area in the SGA that stores whole tables, table partitions, and individual columns in a compressed columnar format. When you specify this transform, Data Pump uses the contents of the string as the IM clause for all objects being imported that have an IM clause in their DDL. This transform is useful when you want to override the IM clause for an object in the dump file. **Note**: The INMEMORY_CLAUSE transform is available only in Oracle Database 12c Release (12.1.0.2) or higher.
LOB_STORAGE	TEXT	TABLE	Specifies the storage type to use for LOB segments. The options are as follows: • SECUREFILE - LOB storage is returned as SECUREFILE • BASICFILE - LOB storage is returned as BASICFILE • DEFAULT - The keyword (SECUREFILE or BASICFILE) is omitted in the LOB STORE AS clause. • NO_CHANGE - LOB segments are created with the same storage they had in the source database. This is the default. Specifying this transform changes the LOB storage for all tables in the job, including tables that provide storage for materialized views.
OID	NUMBER	TYPE TABLE	If zero, inhibits the assignment of the exported OID during type or table creation. Instead, a new OID will be assigned. Use of this transform on Object Tables will cause breakage in REF columns that point to the table. Defaults to 1.

Table 59-19 (Cont.) Transforms Provided by the METADATA_TRANFORM Procedure

Name	Datatype	Object Type	Meaning
OMIT_ACDR_METADATA	NUMBER	TABLE	Used with Oracle GoldenGate. When set to 1 (true), excludes invisible columns from importing replicated tables, deletes tombstone tables, and deletes all the automatic conflict detection and resolution (ACDR) instance procedural actions.
PCTSPACE	NUMBER	TABLE INDEX TABLESPACE	Specifies a percentage multiplier used to alter extent allocations and data file sizes. Used to shrink large tablespaces for testing purposes. Defaults to 100.
SEGMENT_ATTRIBUTES	NUMBER	TABLE, INDEX	If nonzero (TRUE), emit storage segment parameters. Defaults to 1.
SEGMENT_CREATION	NUMBER	TABLE	If nonzero (TRUE), the SQL SEGMENT CREATION clause is added to the CREATE TABLE statement. That is, the CREATE TABLE statement will explicitly say either SEGMENT CREATION DEFERRED or SEGMENT CREATION IMMEDIATE. If the value is FALSE, then the SEGMENT CREATION clause is omitted from the CREATE TABLE statement. Set this parameter to FALSE to use the default segment creation attributes for the table(s) being loaded. Defaults to nonzero (TRUE).
STORAGE	NUMBER	TABLE	If nonzero (TRUE), emit storage clause. (Ignored if SEGMENT_ATTRIBUTES is zero.) Defaults to nonzero (TRUE).
TABLE_COMPRESSION_CLAUSE	TEXT	TABLE	Specifies a table compression clause (for example, COMPRESS BASIC) to use when the table is created. Specify NONE to omit the table compression clause. The table will have the default compression for the tablespace. Specifying this transform changes the compression type for all tables in the job, including tables that provide storage for materialized views.

Table 59-19 (Cont.) Transforms Provided by the METADATA_TRANFORM Procedure

Name	Datatype	Object Type	Meaning
XMLTYPE_STORAGE_CLAUSE	TEXT	TABLE	Specifies to use the TRANSPORTABLE BINARY XML XMLType. This new type is the recommended storage type for Oracle Database 23ai to store the data in a self-contained binary format or the non-transportable compact BINARY XML storage XMLType for XML data storage. There is no default. If the transform is not used, then the source datatype in the dumpfile is the datatype defined on the target. syntax: dbms_datapump.metadata_transform (handle => h1, name => 'XMLTYPE_STORAGE_CLAUSE', value => 'transform param value');

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_ARGVAL. This message can indicate any of the following:
  - The mode is transportable, which doesn't support transforms.
  - The job's mode does not include the specified object_type.
  - The transform has already been specified for the specified value and object_type.
- INVALID_OPERATION. Transforms are only supported for SQL_FILE and Import operations. The job's operation was Export which does not support the use of metadata transforms.
- INVALID_STATE. The user called METADATA_TRANSFORM after the job had started (that is, the job was not in the defining state).
- INCONSISTENT_ARGS. There was no value supplied or it was of the wrong datatype for the transform.
- PRIVILEGE_ERROR. A nonprivileged user attempted to do a REMAP_SCHEMA to a different user's schema or a REMAP_DATAFILE.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- The METADATA_TRANSFORM procedure is only supported for Import and SQL_FILE operations. It enables you to apply commonly desired, predefined transformations to the definition of objects as part of the transfer. If you need transforms that are not supported within this procedure, you should do a preliminary SQL_FILE operation to produce a SQL script corresponding to the dump file set. By editing the DDL directly and then executing it, you can produce any transformations that you need.
- Transforms for the Oracle Data Pump API are a subset of the transforms implemented by the DBMS_METADATA.SET_TRANSFORM_PARAMETER API. Multiple transforms can be defined for

a single job. However, each transform defined must be unique according its parameters. That is, two transforms cannot specify conflicting or redundant transformations.

- In Oracle Database 23ai, Oracle recommends that you use the `TRANSPORTABLE BINARY XML XMLType` to store the data in a self-contained binary format. In Oracle Database 19c and 21c, you can use this XMLType in Cloud deployments. This XMLType does not store the metadata used to encode or decode XML data in a central table, which simplifies the XML data storage and makes it easier to transport. For more information about this XMLType, see *Oracle XML DB Developer's Guide*.

Related Topics

- Export/Import Limitations for Oracle XML DB Repository in *Oracle XML DB Developer's Guide*

59.5.12 OPEN Function

This function is used to declare a new job using the Data Pump API.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

The handle that is returned is used as a parameter for calls to all other procedures (but not to the `ATTACH` function).

Syntax

```
DBMS_DATAPUMP.OPEN (
   operation    IN VARCHAR2,
   job_mode         IN VARCHAR2,
   remote_link  IN VARCHAR2 DEFAULT NULL,
   job_name     IN VARCHAR2 DEFAULT NULL,
   version      IN VARCHAR2 DEFAULT 'COMPATIBLE')
 RETURN NUMBER;
```

Parameters

Table 59-20 OPEN Function Parameters

Parameter	Meaning
`operation`	The type of operation to be performed. Table 59-21 contains descriptions of valid operation types.
`job_mode`	The scope of the operation to be performed. Table 59-22 contains descriptions of valid modes. Specifying NULL generates an error.
`remote_link`	If the value of this parameter is non-null, it provides the name of a database link to the remote database that will be the source of data and metadata for the current job.

Table 59-20 (Cont.) OPEN Function Parameters

Parameter	Meaning
job_name	The name of the job. The name is limited to 128 characters; it will be truncated if more than 128 characters are used. It may consist of printable characters and spaces. It is implicitly qualified by the schema of the user executing the OPEN function and must be unique to that schema (that is, there cannot be other Data Pump jobs using the same name). The name is used to identify the job both within the API and with other database components such as identifying the job in the DBA_RESUMABLE view if the job becomes suspended through lack of resources. If no name is supplied, a system generated name will be provided for the job in the following format: "SYS_<OPERATION>_<MODE>_%N". The default job name is formed where %N expands to a two-digit incrementing integer starting at '01' (for example, "SYS_IMPORT_FULL_03"). The name supplied for the job will also be used to name the master table and other resources associated with the job.
version	The version of database objects to be extracted. This option is only valid for Export, network Import, and SQL_FILE operations. Database objects or attributes that are incompatible with the version will not be extracted. Legal values for this parameter are as follows: • COMPATIBLE - (default) the version of the metadata corresponds to the database compatibility level and the compatibility release level for feature (as given in the V$COMPATIBILITY view). Database compatibility must be set to 9.2 or higher. • LATEST - the version of the metadata corresponds to the database version. • A specific database version, for example, '11.0.0'. Specify a value of 12 to allow all existing database features, components, and options to be exported from Oracle Database 11*g* release 2 (11.2.0.3) or later into an Oracle Database 12 *c* Release 1 (12.1) (either a multitenant container database (CDB) or a non-CDB).

Table 59-21 describes the valid operation types for the OPEN function.

Table 59-21 Valid Operation Types for the OPEN Function

Operation	Description
EXPORT	Saves data and metadata to a dump file set or obtains an estimate of the size of the data for an operation.
IMPORT	Restores data and metadata from a dump file set or across a database link.
SQL_FILE	Displays the metadata within a dump file set, or from across a network link, as a SQL script. The location of the SQL script is specified through the ADD_FILE procedure.

Table 59-22 describes the valid modes for the OPEN function.

Table 59-22 Valid Modes for the OPEN Function

Mode	Description
FULL	Operates on the full database or full dump file set except for Oracle Database internal schemas. (Some tables from Oracle Database internal schemas may be registered to be exported and imported in full operations in order to provide consistent metadata during import.) The TRANSPORTABLE parameter can be set to ALWAYS during a full database export in order to move data via transportable tablespaces rather than in the Data Pump dump file.
SCHEMA	Operates on a set of selected schemas. Defaults to the schema of the current user. All objects in the selected schemas are processed. In SCHEMA mode, you cannot specify Oracle-internal schemas (for example, SYS, XDB, ORDSYS, MDSYS, CTXSYS, ORDPLUGINS, or LBACSYS).
TABLE	Operates on a set of selected tables. Defaults to all of the tables in the current user's schema. Only tables and their dependent objects are processed.
TABLESPACE	Operates on a set of selected tablespaces. No defaulting is performed. Tables that have storage in the specified tablespaces are processed in the same manner as in Table mode.
TRANSPORTABLE	Operates on metadata for tables (and their dependent objects) within a set of selected tablespaces to perform a transportable tablespace export/import.

Return Values

- An opaque handle for the job. This handle is used as input to the following procedures: ADD_FILE, CREATE_JOB_VIEW, DATA_FILTER, DETACH, GET_STATUS, LOG_ENTRY, LOG_ERROR,METADATA_FILTER, METADATA_REMAP, METADATA_TRANSFORM, SET_PARALLEL,SET_PARAMETER, START_JOB,STOP_JOB, and WAIT_FOR_JOB

Exceptions

- INVALID_ARGVAL. An invalid operation or mode was specified. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- JOB_EXISTS. A table already exists with the specified job name.
- PRIVILEGE_ERROR. The user does not have the necessary privileges or roles to use the specified mode.
- INTERNAL_ERROR. The job was created under the wrong schema or the master table was of the wrong format.
- SUCCESS_WITH_INFO. The function succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- When the job is created, a master table is created for the job under the caller's schema within the caller's default tablespace. A handle referencing the job is returned that attaches the current session to the job. Once attached, the handle remains valid until either an explicit or implicit detach occurs. The handle is only valid in the caller's session. Other

handles can be attached to the same job from a different session by using the ATTACH function.

- If the call to the OPEN function fails, call the GET_STATUS procedure with a null handle to retrieve additional information about the failure.

59.5.13 SET_PARALLEL Procedure

This procedure adjusts the degree of parallelism within a job.

Syntax

```
DBMS_DATAPUMP.SET_PARALLEL(
   handle      IN NUMBER,
   degree      IN NUMBER);
```

Parameters

Table 59-23 SET_PARALLEL Procedure Parameters

Parameter	Description
handle	The handle of a job. The current session must have previously attached to the handle through a call to either the OPEN or ATTACH function.
degree	The maximum number of worker processes that can be used for the job. You use this parameter to adjust the amount of resources used for a job.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_OPERATION. The SET_PARALLEL procedure is only valid for export and import operations.
- INVALID_ARGVAL. An invalid value was supplied for an input parameter.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- The SET_PARALLEL procedure is only available in the Enterprise Edition of the Oracle database.
- The SET_PARALLEL procedure can be executed by any session attached to a job. The job must be in one of the following states: Defining, Idling, or Executing.
- The effect of decreasing the degree of parallelism may be delayed because ongoing work needs to find an orderly completion point before SET_PARALLEL can take effect.
- Decreasing the parallelism will not result in fewer worker processes associated with the job. It will only decrease the number of worker processes that will be executing at any given time.
- Increasing the parallelism will take effect immediately if there is work that can be performed in parallel.

- The degree of parallelism requested by a user may be decreased based upon settings in the resource manager or through limitations introduced by the PROCESSES or SESSIONS initialization parameters in the init.ora file.
- To parallelize an Export job to a degree of *n*, the user should supply *n* files in the dump file set or specify a substitution variable in a file specification. Otherwise, some of the worker processes will be idle while waiting for files.
- SQL_FILE operations always operate with a degree of 1. Jobs running in the Transportable mode always operate with a degree of 1.

59.5.14 SET_PARAMETER Procedures

The DBMS_DATAPUMP procedure SET_PARAMETER is used to specify job-processing options.

Syntax

```
DBMS_DATAPUMP.SET_PARAMETER(
   handle       IN NUMBER,
   name         IN VARCHAR2,
   value        IN VARCHAR2);

DBMS_DATAPUMP.SET_PARAMETER (
   handle      IN NUMBER,
   name        IN VARCHAR2,
   value       IN NUMBER);
```

Parameters

Table 59-24 SET_PARAMETER Procedure Parameters

Parameter	Description
handle	The handle of a job. The current session must have previously attached to the handle through a call to the OPEN function.
name	The name of the parameter.
value	The value for the specified parameter

The following table describes the valid options for the name parameter of the SET_PARAMETER procedure.

Table 59-25 Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
CLIENT_COMMAND	TEXT	All	An opaque string used to describe the current operation from the client's perspective. The command-line procedures will use this string to store the original command used to invoke the job.

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
COMPRESSION	TEXT	Export	Allows you to trade off the size of the dump file set versus the time it takes to perform export and import operations. The DATA_ONLY option compresses only user data in the dump file set. The METADATA_ONLY option compresses only metadata in the dump file set. The ALL option compresses both user data and metadata. The NONE option stores the dump file set in an uncompressed format. The METADATA_ONLY and NONE options require a job version of 10.2 or later. All other options require a job version of 11.1 or later. Default=METADATA_ONLY
COMPRESSION_ALGORITHM	TEXT	Export	Indicates the compression algorithm is to be used when compressing dump file data. The choices are as follows: • BASIC--Offers a good combination of compression ratios and speed; the algorithm used is the same as in previous versions of Oracle Data Pump. • LOW---Least impact on backup throughput and suited for environments where CPU resources are the limiting factor. • MEDIUM---Recommended for most environments. This option, like the BASIC option, provides a good combination of compression ratios and speed, but it uses a different algorithm than BASIC. • HIGH--Best suited for exports over slower networks where the limiting factor is network speed. To use this feature, the COMPATIBLE initialization parameter must be set to at least 12.0.0. This feature requires that the Oracle Advanced Compression option is enabled.
DATA_ACCESS_METHOD	TEXT	Export and Import	Allows you to specify an alternative method for Oracle Data Pump Export (unloading) or Oracle Data Pump Import (loading) to unload or load data if the default method does not work for some reason. For Export, the options are AUTOMATIC, DIRECT_PATH, or EXTERNAL_TABLE. For Import, the options are AUTOMATIC, DIRECT_PATH, EXTERNAL_TABLE, CONVENTIONAL, or (network imports only) INSERT_AS_SELECT. The INSERT_AS_SELECT option is valid only if you are importing directly over a database link. If you are importing directly over a database link, then the allowed options are AUTOMATIC, DIRECT_PATH, and INSERT_AS_SELECT. **Note**: For both Export and Import, Oracle recommends that you use the default option (AUTOMATIC) whenever possible because it allows Oracle Data Pump to automatically select the most efficient method.

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
DATA_OPTIONS	Number	Export and Import	A bitmask to supply special options for processing the job. The possible values are as follows: KU$_DATAOPT_SKIP_CONST_ERR KU$_DATAOPT_XMLTYPE_CLOB KU$_DATAOPT_DISABL_APPEND_HINT KU$_DATAOPT_GRP_PART_TAB KU$_DATAOPT_TRUST_EXIST_TB_PAR KU$_DATAOPT_CONT_LOAD_ON_FMT_ERR KU$_DATAOPT_REJECT_ROWS_REPCHR KU$_DATAOPT_ENABLE_NET_COMP KU$_DATAOPT_VALIDATE_TBL_DATA KU$_DATAOPT_VERIFY_STREAM_FORM Export supports the value KU$_DATAOPT_XMLTYPE_CLOB. This option stores compressed XMLType columns in the dump file in CLOB format rather than as XML documents. **Note**: XMLType stored as CLOB is deprecated as of Oracle Database 12c Release 1 (12.1). Import supports the value KU$_DATAOPT_SKIP_CONST_ERR. This option specifies that if constraint violations occur while data is being imported into user tables, the rows that cause the violations will be rejected and the load will continue. If this option is not set, a constraint error will abort the loading of the entire partition (or table for unpartitioned tables). Setting this option may affect performance, especially for pre-existing tables with unique indexes or constraints. Import also supports the value KU$_DATAOPT_DISABL_APPEND_HINT. This option prevents the append hint from being applied to the data load. Disabling the APPEND hint can be useful if there is a small set of data objects to load that already exist in the database and some other application may be concurrently accessing one or more of the data objects. Oracle Data Pump Export supports the value KU$_DATAOPT_GRP_PART_TAB. This option tells Data Pump to unload all table data in one operation rather than unload each table partition as a separate operation. As a result, the definition of the table will not matter at import time because Import will see one partition of data that will be loaded into the entire table. Oracle Data Pump Import supports the value KU$_DATAOPT_TRUST_EXIST_TB_PAR. This option tells Data Pump to load partition data in parallel into existing tables. Use of the DATA_OPTIONS parameter requires that the version on the OPEN function be set to 11.1 or later. Default=0 Oracle Data Pump Import supports the value KU$_DATAOPT_CONT_LOAD_ON_FMT_ERR. This option tells Oracle Data Pump to skip forward to the start of the next granule if a stream format error is encountered while loading table data. Most stream format errors are caused by corrupt dump files. This value can be used if Oracle Data Pump encounters a stream format error and the original export database is not available to export the table data

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
			again. If Oracle Data Pump skips over data, not all data from the source database are imported, which potentially can mean skipping hundreds or thousands of rows.
ENCRYPTION	TEXT	Export	Specifies what to encrypt in the dump file set, as follows: ALL enables encryption for all data and metadata in the export operation. DATA_ONLY specifies that only data is written to the dump file set in encrypted format. ENCRYPTED_COLUMNS_ONLY specifies that only encrypted columns are written to the dump file set in encrypted format. METADATA_ONLY specifies that only metadata is written to the dump file set in encrypted format. NONE specifies that no data is written to the dump file set in encrypted format. This parameter requires a job version of 11.1 or later. The default value depends upon the combination of encryption-related parameters that are used. To enable encryption, either ENCRYPTION or ENCRYPTION_PASSWORD or both, must be specified. If only ENCRYPTION_PASSWORD is specified, then ENCRYPTION defaults to ALL. If neither ENCRYPTION nor ENCRYPTION_PASSWORD is specified, then ENCRYPTION defaults to NONE. To specify ALL, DATA_ONLY, or METADATA_ONLY, the COMPATIBLE initialization parameter must be set to at least 11.1. **NOTE**: If the data being exported includes SecureFiles that you want to be encrypted, then you must specify ENCRYPTION=ALL to encrypt the entire dump file set. Encryption of the entire dump file set is the only way to achieve encryption security for SecureFiles during an Oracle Data Pump export operation.
ENCRYPTION_ALGORITHM	TEXT	Export	Identifies which cryptographic algorithm should be used to perform encryption. Possible values are AES128, AES192, and AES256. The ENCRYPTION_ALGORITHM parameter requires that you also specify either ENCRYPTION or ENCRYPTION_PASSWORD; otherwise an error is returned. For the most current information, see: Supported Encryption and Integrity Algorithms in *Oracle Database Transparent Data Encryption Guide*. This parameter requires a job version of 1.1 or later. Default=AES128

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
ENCRYPTION_MODE	TEXT	Export	Identifies the types of security used for encryption and decryption. The values are as follows: PASSWORD requires that you provide a password when creating encrypted dump file sets. You will need to provide the same password when you import the dump file set. PASSWORD mode requires that you also specify the ENCRYPTION_PASSWORD parameter. The PASSWORD mode is best suited for cases in which the dump file set will be imported into a different or remote database, but which must remain secure in transit. TRANSPARENT allows an encrypted dump file set to be created without any intervention from a database administrator (DBA), provided the required Oracle Encryption Wallet is available. Therefore, the ENCRYPTION_PASSWORD parameter is not required, and will in fact, cause an error if it is used in TRANSPARENT mode. This encryption mode is best suited for cases in which the dump file set will be imported into the same database from which it was exported. DUAL creates a dump file set that can later be imported using either the Oracle Encryption Wallet or the password that was specified with the ENCRYPTION_PASSWORD parameter. DUAL mode is best suited for cases in which the dump file set will be imported on-site using the Oracle Encryption Wallet, but which may also need to be imported offsite where the Oracle Encryption Wallet is not available. When you use the ENCRYPTION_MODE parameter, you must also use either the ENCRYPTION or ENCRYPTION_PASSWORD parameter. Otherwise, an error is returned. To use DUAL or TRANSPARENT mode, the COMPATIBLE initialization parameter must be set to at least 11.1. The default mode depends on which other encryption-related parameters are used. If only ENCRYPTION is specified, then the default mode is TRANSPARENT. If ENCRYPTION_PASSWORD is specified and the Oracle Encryption Wallet is open, then the default is DUAL. If ENCRYPTION_PASSWORD is specified and the Oracle Encryption Wallet is closed, then the default is PASSWORD.
ENCRYPTION_PASSWORD	TEXT	Export and Import	For export operations, this parameter is required if ENCRYPTION_MODE is set to either PASSWORD or DUAL. It is also required for transportable export/import operations (job mode=FULL and TRANSPORTABLE=ALWAYS) when the database includes either encrypted tablespaces or tables with encrypted columns.
ESTIMATE	TEXT	Export and Import	Specifies that the estimate method for the size of the tables should be performed before starting the job. If BLOCKS, a size estimate for the user tables is calculated using the count of blocks allocated to the user tables. If STATISTICS, a size estimate for the user tables is calculated using the statistics associated with each table. If no statistics are available for a table, the size of the table is estimated using BLOCKS. The ESTIMATE parameter cannot be used in Transportable Tablespace mode. Default=STATISTICS

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
ESTIMATE_ONLY	NUMBER	Export	Specifies that only the estimation portion of an export job should be performed. This option is useful for estimating the size of dump files when the size of the export is unknown.
FLASHBACK_SCN	NUMBER	Export and network Import	System change number (SCN) to serve as transactionally consistent point for reading user data. If neither FLASHBACK_SCN nor FLASHBACK_TIME is specified, there will be no transactional consistency between partitions, except for logical standby databases and Streams targets. FLASHBACK_SCN is not supported in Transportable mode. For FLASHBACK_SCN, Oracle Data Pump supports the 8–byte big SCNs used in Oracle Database 12c release 2 (12.2) and later releases.
FLASHBACK_TIME	TEXT	Export and network Import	Either the date and time used to determine a consistent point for reading user data or a string of the form TO_TIMESTAMP(...). If neither FLASHBACK_SCN nor FLASHBACK_TIME is specified, there will be no transactional consistency between partitions. FLASHBACK_SCN and FLASHBACK_TIME cannot both be specified for the same job. FLASHBACK_TIME is not supported in Transportable mode.
INCLUDE_METADATA	NUMBER	Export and Import	If nonzero, then metadata for objects will be moved in addition to user table data. If zero, then metadata for objects will not moved. This parameter converts an Export operation into an unload of user data and an Import operation into a load of user data. INCLUDE_METADATA is not supported in Transportable mode. Default=1
INDEX_THRESHOLD	TEXT	Import	Sets a size threshold for creating large indexes with a degree of parallelism (DOP) greater than 1 in conjunction with the ONESTEP_INDEX parameter. Indexes equal to and above the threshold can be created with a degre of parallism DOP greater than 1 if the parameter ONESTEP_INDEX=FALSE. The INDEX_THRESHOLD value must be specified as a text string, such as 1000B, 100k, 200kb, 100M, 200mb, 100G, 200gb, 100t, 200TB. An invalid threshold will produce an error. Oracle recommends that you use the default in most cases. Default=150M
KEEP_MASTER	NUMBER	Export and Import	Specifies whether the master table should be deleted or retained at the end of a Data Pump job that completes successfully. The master table is automatically retained for jobs that do not complete successfully. Default=0
LOGTIME	TEXT	Export and Import	Specifies that messages displayed during export and import operations be timestamped. Valid options are as follows: • NONE--No timestamps on status or log file messages (this is the default) • STATUS--Timestamps on status messages only • LOGFILE--Timestamps on log file messages only • ALL--Timestamps on both status and log file messages

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
MASTER_ONLY	NUMBER	Import	Indicates whether to import just the master table and then stop the job so that the contents of the master table can be examined. Default=0
METRICS	NUMBER	Export and Import	Indicates whether additional information about the job should be reported to the Data Pump log file. Default=0
ONESTEP_INDEX	NUMBER	Import	Attempts to optimize index creation concurrency and balance it with job parallelism in conjunction with the INDEX_THRESHOLD parameter. After a method is selected, it cannot be changed for the job, including a job restart. Valid options: 0 (False), and 1 (True). In most cases, 0 will yield better performance than 1: ONESTEP_INDEX=0, in conjunction with the INDEX_THRESHOLD parameter starts a two-step process to balance parallelism for creating very large indexes with parallelism for importing all objects in the overall import job, up to the PARALLEL setting for the job. Oracle Data Pump attempts to create all the indexes in the shortest amount of time, given the constraints of the job. Large index creation is prioritized first because they take longer to create. ONESTEP_INDEX=1 uses a single step process and a DOP of 1 for each index created during import. However, multiple indexes can be created in parallel up to the DOP setting for the import job. 1 (true) can be optimal if all indexes are relatively small, below the INDEX_THRESHOLD default value. 1 is also the default for jobs with PARALLEL=1.
PARTITION_OPTIONS	TEXT	Import	Specifies how partitioned tables should be handled during an import operation. The options are as follows: NONE means that partitioning is reproduced on the target database as it existed in the source database. DEPARTITION means that each partition or subpartition that contains storage in the job is reproduced as a separate unpartitioned table. Intermediate partitions that are subpartitioned are not re-created (although their subpartitions are converted into tables). The names of the resulting tables are system-generated from the original table names and partition names unless the name is overridden by the REMAP_TABLE metadata transform. MERGE means that each partitioned table is re-created in the target database as an unpartitioned table. The data from all of the source partitions is merged into a single storage segment. This option is not supported for transportable jobs or when the TRANSPORTABLE parameter is set to ALWAYS. This parameter requires a job version of 11.1 or later. Default=NONE
REUSE_DATAFILES	NUMBER	Import	Specifies whether the import job should reuse existing data files for tablespace creation. Default=0
SKIP_UNUSABLE_INDEXES	NUMBER	Import	If nonzero, rows will be inserted into tables having unusable indexes. SKIP_UNUSABLE_INDEXES is not supported in Transportable mode. Default=1

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
SOURCE_EDITION	TEXT	Export and network Import	The application edition that will be used for determining the objects that will be unloaded for export and for network import.
STREAMS_CONFIGURATION	NUMBER	Import	Specifies whether to import any Streams metadata that may be present in the export dump file. Default=1
TABLE_EXISTS_ACTION	TEXT	Import	Specifies the action to be performed when data is loaded into a preexisting table. The possible actions are: TRUNCATE, REPLACE, APPEND, and SKIP. If INCLUDE_METADATA=0, only TRUNCATE and APPEND are supported. If TRUNCATE, rows are removed from a preexisting table before inserting rows from the Import. Note that if TRUNCATE is specified on tables referenced by foreign key constraints, the TRUNCATE will be modified into a REPLACE. If REPLACE, preexisting tables are replaced with new definitions. Before creating the new table, the old table is dropped. If APPEND, new rows are added to the existing rows in the table. If SKIP, the preexisting table is left unchanged. TABLE_EXISTS_ACTION is not supported in Transportable mode. The default is SKIP if metadata is included in the import. The default is APPEND if INCLUDE_METADATA is set to 0.
TABLESPACE_DATAFILE	TEXT	Import	Specifies the full file specification for a data file in the transportable tablespace set. TABLESPACE_DATAFILE is only valid for transportable mode imports. TABLESPACE_DATAFILE can be specified multiple times, but the value specified for each occurrence must be different.
TARGET_EDITION	TEXT	Import	The application edition that will be used for determining where the objects will be loaded for import and for network import.

Table 59-25 (Cont.) Valid Options for the name Parameter in the SET_PARAMETER Procedure

Parameter Name	Datatype	Supported Operations	Meaning
TRANSPORTABLE	TEXT	Export (and network import or full-mode import)	This option is for export operations done in table mode, and also for full-mode imports and network imports. It allows the data to be moved using transportable tablespaces. In table-mode storage segments in the moved tablespaces that are not associated with the parent schemas (tables) will be reclaimed at import time. If individual partitions are selected in a table-mode job, only the tablespaces referenced by those partitions will be moved. During import, the moved partitions can only be reconstituted as tables by using the PARTITION_OPTIONS=DEPARTITION parameter. Use of the TRANSPORTABLE parameter prohibits the subsequent import of the dump file into a database at a lower version or using different character sets. Additionally, the data files may need to be converted if the target database is on a different platform. In table-mode, the TRANSPORTABLE parameter is not allowed if a network link is supplied on the OPEN call. The possible values for this parameter are as follows: ALWAYS - data is always moved by moving data files. This option is valid only for table mode and full mode. NEVER - data files are never used for copying user data This parameter requires a job version of 11.1 or later This parameter requires a job version of 12.1 or later when the job mode is FULL. Default=NEVER
TTS_FULL_CHECK	NUMBER	Export	If nonzero, verifies that a transportable tablespace set has no dependencies (specifically, IN pointers) on objects outside the set, and vice-versa. Only valid for Transportable mode Exports. Default=0
USER_METADATA	NUMBER	Export and network Import	For schema-mode operations, specifies that the metadata to re-create the users' schemas (for example, privilege grants to the exported schemas) should also be part of the operation if set to nonzero. Users must be privileged to explicitly set this parameter. The USER_METADATA parameter cannot be used in Table, Tablespace, or Transportable Tablespace mode. Default=1 if user has DATAPUMP_EXP_FULL_DATABASE role; 0 otherwise.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_ARGVAL. This exception could be due to any of the following causes:
  - An invalid name was supplied for an input parameter
  - The wrong datatype was used for value
  - A value was not supplied
  - The supplied value was not allowed for the specified parameter name
  - A flashback parameter had been established after a different flashback parameter had already been established

  - A parameter was specified that did not support duplicate definitions
- INVALID_OPERATION. The operation specified is invalid in this context.
- INVALID_STATE. The specified job is not in the Defining state.
- INCONSISTENT_ARGS. Either the specified parameter is not supported for the current operation type or it is not supported for the current mode.
- PRIVILEGE_ERROR. The user does not have the DATAPUMP_EXP_FULL_DATABASE or DATAPUMP_IMP_FULL_DATABASE role required for the specified parameter.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- The SET_PARAMETER procedure is used to specify optional features for the current job. For a list of supported options, see the preceding table, "Valid Options for the name Parameter in the SET_PARAMETER Procedure".

59.5.15 START_JOB Procedure

This procedure begins or resumes job execution.

Syntax

```
DBMS_DATAPUMP.START_JOB (
   handle       IN NUMBER,
   skip_current    IN  NUMBER DEFAULT 0,
   abort_step      IN  NUMBER DEFAULT 0,
   cluster_ok      IN  NUMBER DEFAULT 1,
   service_name    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 59-26 START_JOB Procedure Parameters

Parameter	Description
handle	The handle of a job. The current session must have previously attached to the handle through a call to either the OPEN or ATTACH function.
skip_current	If nonzero, causes actions that were 'in progress' on a previous execution of the job to be skipped when the job restarts. The skip will only be honored for Import jobs. This mechanism allows the user to skip actions that trigger fatal bugs and cause the premature termination of a job. Multiple actions can be skipped on a restart. The log file will identify which actions are skipped. If a domain index was being processed, all pieces of the domain index are skipped even if the error occurred in only a subcomponent of the domain index. A description of the actions skipped is entered into the log file. skip_current is ignored for the initial START_JOB in a job. If zero, no data or metadata is lost upon a restart.
abort_step	Value must be 0. Inserting values other than 0 into this argument will have unintended consequences.
cluster_ok	If = 0, all workers are started on the current instance. Otherwise, workers are started on instances usable by the job.

Table 59-26 (Cont.) START_JOB Procedure Parameters

Parameter	Description
service_name	If specified, indicates a service name used to constrain the job to specific instances or to a specific resource group.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID_STATE. The causes of this exception can be any of the following:
  - No files have been defined for an Export, non-network Import, or SQL_FILE job
  - An ADD_FILE procedure has not been called to define the output for a SQL_FILE job
  - A TABLESPACE_DATAFILE parameter has not been defined for a Transportable Import job
  - A TABLESPACE_EXPR metadata filter has not been defined for a Transportable or Tablespace mode Export or Network job
  - The dump file set on an Import or SQL_FILE job was either incomplete or missing a master table specification
- INVALID_OPERATION. Unable to restore master table from a dump file set.
- INTERNAL_ERROR. An inconsistency was detected when the job was started. Additional information may be available through the GET_STATUS procedure.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- When this procedure is called to request that the corresponding job be started or restarted, the state of the job is changed from either the Defining or Idling state to the Executing state.
- If the SET_PARALLEL procedure was not called prior to the START_JOB procedure, the initial level of parallelism used in the job will be 1. If SET_PARALLEL was called prior to the job starting, the degree specified by the last SET_PARALLEL call determines the parallelism for the job. On restarts, the parallelism is determined by the previous parallel setting for the job, unless it is overridden by another SET_PARALLEL call.
- To restart a stopped job, an ATTACH function must be performed prior to executing the START_JOB procedure.

59.5.16 STOP_JOB Procedure

This procedure terminates a job, but optionally, preserves the state of the job.

Syntax

```
DBMS_DATAPUMP.STOP_JOB (
   handle      IN NUMBER,
   immediate   IN NUMBER DEFAULT 0,
```

```
keep_master IN NUMBER DEFAULT NULL,
delay       IN NUMBER DEFAULT 60);
```

Parameters

Table 59-27 STOP_JOB Procedure Parameters

Parameter	Description
handle	The handle of a job. The current session must have previously attached to the handle through a call to either the OPEN or ATTACH function. At the end of the procedure, the user is detached from the handle.
immediate	If nonzero, the worker processes are aborted immediately. This halts the job quickly, but parts of the job will have to be rerun if the job is ever restarted. If zero, the worker processes are allowed to complete their current work item (either metadata or table data) before they are terminated. The job is placed in a Stop Pending state while the workers finish their current work.
keep_master	If nonzero, the master table is retained when the job is stopped. If zero, the master table is dropped when the job is stopped. If the master table is dropped, the job will not be restartable. If the master table is dropped during an export job, the created dump files are deleted.
delay	The number of seconds to wait until other attached sessions are forcibly detached. The delay allows other sessions attached to the job to be notified that a stop has been performed. The job keeps running until either all clients have detached or the delay has been satisfied. If no delay is specified, then the default delay is 60 seconds. If a shorter delay is used, clients might not be able to retrieve the final messages for the job through the GET_STATUS procedure.

Exceptions

- INVALID_HANDLE. The specified handle is not attached to a Data Pump job.
- INVALID STATE. The job is already in the process of being stopped or completed.
- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS procedure.
- NO_SUCH_JOB. The specified job does not exist.

Usage Notes

- This procedure is used to request that the corresponding job stop executing.
- The termination of a job that is in an Executing state may take several minutes to complete in an orderly fashion.
- For jobs in the Defining, Idling, or Completing states, this procedure is functionally equivalent to the DETACH procedure.
- Once a job is stopped, it can be restarted using the ATTACH function and START_JOB procedures, provided the master table and the dump file set are left intact.
- If the KEEP_MASTER parameter is not specified, and the job is in the Defining state or has a mode of Transportable, the master table is dropped. Otherwise, the master table is retained.

59.5.17 WAIT_FOR_JOB Procedure

This procedure runs a job until it either completes normally or stops for some other reason.

Syntax

```
DBMS_DATAPUMP.WAIT_FOR_JOB (
  handle      IN   NUMBER,
  job_state   OUT  VARCHAR2);
```

Parameters

Table 59-28 WAIT_FOR_JOB Procedure Parameters

Parameter	Description
handle	The handle of the job. The current session must have previously attached to the handle through a call to either the OPEN or ATTACH function. At the end of the procedure, the user is detached from the handle.
job_state	The state of the job when it has stopped executing; either STOPPED or COMPLETED.

Exceptions

- SUCCESS_WITH_INFO. The procedure succeeded, but further information is available through the GET_STATUS API.
- INVALID_HANDLE. The job handle is no longer valid.

Usage Notes

This procedure provides the simplest mechanism for waiting for the completion of a Data Pump job. The job should be started before calling WAIT_FOR_JOB. When WAIT_FOR_JOB returns, the job will no longer be executing. If the job completed normally, the final status will be COMPLETED. If the job stopped executing because of a STOP_JOB request or an internal error, the final status will be STOPPED.

60
DBMS_DB_VERSION

The `DBMS_DB_VERSION` package specifies the Oracle version numbers and other information useful for simple conditional compilation selections based on Oracle versions.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* regarding conditional compilation

This package contains the following topics:

- Overview
- Constants
- Examples

60.1 DBMS_DB_VERSION Overview

The `DBMS_DB_VERSION` package specifies the Oracle version numbers and other information useful for simple conditional compilation selections based on Oracle versions.

The package for the Oracle Database 21c version is shown below.

```
PACKAGE DBMS_DB_VERSION IS
   VERSION CONSTANT PLS_INTEGER := 21; -- RDBMS version number
   RELEASE CONSTANT PLS_INTEGER := 0;  -- RDBMS release number
   ver_le_9_1    CONSTANT BOOLEAN := FALSE;
   ver_le_9_2    CONSTANT BOOLEAN := FALSE;
   ver_le_9      CONSTANT BOOLEAN := FALSE;
   ver_le_10_1   CONSTANT BOOLEAN := FALSE;
   ver_le_10_2   CONSTANT BOOLEAN := FALSE;
   ver_le_10     CONSTANT BOOLEAN := FALSE;
   ver_le_11_1   CONSTANT BOOLEAN := FALSE;
   ver_le_11_2   CONSTANT BOOLEAN := FALSE;
   ver_le_11     CONSTANT BOOLEAN := FALSE;
   ver_le_12_1   CONSTANT BOOLEAN := FALSE;
   ver_le_12_2   CONSTANT BOOLEAN := FALSE;
   ver_le_12     CONSTANT BOOLEAN := FALSE;
   ver_le_18     CONSTANT BOOLEAN := FALSE;
   ver_le_19     CONSTANT BOOLEAN := FALSE;
   ver_le_20     CONSTANT BOOLEAN := FALSE;
   ver_le_21     CONSTANT BOOLEAN := TRUE;
END DBMS_DB_VERSION;
```

The boolean constants follow a naming convention. Each constant gives a name for a boolean expression. For example:

- `VER_LE_9_1` represents version <= 9 and release <= 1
- `VER_LE_10_2` represents version <= 10 and release <= 2

- `VER_LE_10` represents version <= 10

A typical usage of these boolean constants is:

```
$IF DBMS_DB_VERSION.VER_LE_10 $THEN
   version 10 and earlier code
$ELSIF DBMS_DB_VERSION.VER_LE_11 $THEN
   version 11 code
$ELSE
   version 12 and later code
$END
```

This code structure will protect any reference to the code for version 12. It also prevents the controlling package constant `DBMS_DB_VERSION.VER_LE_11` from being referenced when the program is compiled under version 10. A similar observation applies to version 11. This scheme works even though the static constant `VER_LE_11` is not defined in version 10 database because conditional compilation protects the `$ELSIF` from evaluation if `DBMS_DB_VERSION.VER_LE_10` is `TRUE`.

60.2 DBMS_DB_VERSION Constants

The `DBMS_DB_VERSION` package contains different constants for different Oracle Database releases.

The Oracle Database 21*c* version of the `DBMS_DB_VERSION` package uses the constants shown in the following table.

Table 60-1 DBMS_DB_VERSION Constants

Name	Type	Value	Description
VERSION	PLS_INTEGER	21	Current version
RELEASE	PLS_INTEGER	0	Current release
VER_LE_9	BOOLEAN	FALSE	Version <= 9
VER_LE_9_1	BOOLEAN	FALSE	Version <= 9 and release <= 1
VER_LE_9_2	BOOLEAN	FALSE	Version <= 9 and release <= 2
VER_LE_10	BOOLEAN	FALSE	Version <= 10
VER_LE_10_1	BOOLEAN	FALSE	Version <= 10 and release <= 1
VER_LE_10_2	BOOLEAN	FALSE	Version <=10 and release <= 2
VER_LE_11	BOOLEAN	FALSE	Version <= 11
VER_LE_11_1	BOOLEAN	FALSE	Version <=11 and release <= 1
VER_LE_11_2	BOOLEAN	FALSE	Version <=11 and release <= 2
VER_LE_12	BOOLEAN	FALSE	Version <=12
VER_LE_12_1	BOOLEAN	FALSE	Version <=12 and release <= 1
VER_LE_12_2	BOOLEAN	FALSE	Version <=12 and release <= 2
VER_LE_18	BOOLEAN	FALSE	Version <=18
VER_LE_19	BOOLEAN	FALSE	Version <=19
VER_LE_20	BOOLEAN	FALSE	Version <=20
VER_LE_21	BOOLEAN	TRUE	Version <=21

60.3 DBMS_DB_VERSION Examples

This example uses conditional compilation to guard new features.

```
CREATE OR REPLACE PROCEDURE whetstone IS

 -- Notice that conditional compilation constructs
 -- can interrupt a regular PL/SQL statement.
 -- You can locate a conditional compilation directive anywhere
 -- there is whitespace in the regular statement.

 SUBTYPE my_real IS
    $IF DBMS_DB_VERSION.VER_LE_9 $THEN NUMBER
                                 $ELSE BINARY_DOUBLE
    $END;

 t  CONSTANT my_real := $IF DBMS_DB_VERSION.VER_LE_9 $THEN 0.499975
                                                     $ELSE 0.499975d
                        $END;

 t2 CONSTANT my_real := $if DBMS_DB_VERSION.VER_LE_9 $THEN 2.0
                                                     $ELSE 2.0d
                        $END;

 x  CONSTANT my_real := $IF DBMS_DB_VERSION.VER_LE_9 $THEN 1.0
                                                     $ELSE 1.0d
                        $END;

 y  CONSTANT my_real := $IF DBMS_DB_VERSION.VER_LE_9 $THEN 1.0
                                                     $ELSE 1.0d
                        $END;

 z  MY_REAL;

 PROCEDURE P(x IN my_real, y IN my_real, z OUT NOCOPY my_real) IS
   x1 my_real;
   y1 my_real;
 BEGIN
   x1 := x;
   y1 := y;
   x1 := t * (x1 + y1);
   y1 := t * (x1 + y1);
   z := (x1 + y1)/t2;
 END P;
BEGIN
 P(x, y, z);
 DBMS_OUTPUT.PUT_LINE ('z = '|| z);
END whetstone;
/
```

61
DBMS_DBCOMP

The DBMS_DBCOMP package performs block comparison to detect lost writes or database inconsistencies between a primary database and one or more physical standby databases. It contains one procedure, DBCOMP, which can be executed at any time. (It does not require that the DB_LOST_WRITE_PROTECT initialization parameter be enabled.)

This chapter contains the following sections:

- Using DBMS_DBCOMP
- DBMS_DBCOMP Security Model
- DBCOMP Procedure

Related Topics

- *Oracle Data Guard Concepts and Administration*

61.1 Using DBMS_DBCOMP

The DBMS_DBCOMP.DBCOMP procedure assumes that there is one primary database and one or more physical standby databases. The databases should be at least mounted before block comparison.

> **Note:**
>
> Logical standby databases, far sync instances, and cascaded standbys cannot be the target database for the DBMS_DBCOMP.DBCOMP procedure.

While this procedure is running, you can monitor progress of the compare operation by querying the V$SESSION_LONGOPS view. For example:

```
SQL> SELECT target_desc, sofar, totalwork FROM V$SESSION_LONGOPS WHERE opname
= 'BlockCompare';
```

The output might look as follows:

```
TARGET_DESC                    SO FAR     TOTALWORK
------------------------------ ---------  ----------
Compared Blocks                23914      340639
Lost Writes                    0          0
```

The results can be interpreted as follows:

- The Compared Blocks value is the number of blocks that were scanned (including those that did not use compare callback functions). The number is the sum of all secondary processes (that is, each secondary process is responsible for a connection between

primary and standby). For example, if there are two secondary processes and each secondary process has compared 1000 blocks, then the value of SOFAR should be 2000.

- The Lost Writes value is the number of lost writes found at both the primary and the standby. Also, it is the sum of all secondary processes. For Lost Writes, the value of SOFAR should be always equal to the value of TOTALWORK.
- The value of SOFAR should be always less than or equal to the value of TOTALWORK.
- When the block compare secondary processes finish, the value of SOFAR should be equal to the value of TOTALWORK.

61.2 DBMS_DBCOMP Security Model

The DBMS_DBCOMP package runs with invoker's rights and requires the SYSDBA or SYSDG or SYSBACKUP privilege.

61.3 Summary of the DBMS_DBCOMP Subprogram

The DBMS_DBCOMP package contains one procedure, DBCOMP.

Subprogram	Description
DBCOMP procedure	The DBMS_DBCOMP.DBCOMP procedure performs block comparison to detect lost writes or database inconsistencies between a primary database and one or more physical standby databases. The DBCOMP procedure can be executed on the primary or on the physical standby while the databases are mounted or open.

61.3.1 DBCOMP Procedure

Use DBCOMP procedure to compare blocks on a primary database and one or more physical standby databases.

Syntax

```
DBMS_DBCOMP.DBCOMP (
  datafile IN varchar2,
  outputfile IN varchar2,
  block_dump IN boolean);
```

Parameters

Parameter	Description
datafile	This can be a data file name or a data file number. Specify 'ALL' to compare all data files.
outputfile	This value is used as a prefix in the name of the output file. By default, output files are placed in the dbs/ directory, but the output file argument can contain a relative directory path or an absolute directory path. See Usage Notes for descriptions of the information contained in the output file.

Parameter	Description
block_dump	This is a Boolean parameter to specify whether block content should be dumped into output files if a pair of blocks of primary and standby databases is not identical. This parameter does not guarantee that all different block pairs are dumped. The default value is FALSE.

Usage Notes

- The output file contains the following columns:
  - **ID**: The block type. For instance, type 02 is undo block type, and 06 is KTB-managed data block type.
  - **TOTAL**: The total number of blocks in the specific block type.
  - **CORR**: The number of corrupted blocks.
  - **SKIPPED**: The number of block pairs that are skipped (ignored) without comparison, such as direct loaded blocks, RMAN optimized blocks, and flashback optimized blocks
  - **DIFFV**: The number of blocks that have different version numbers between primary and standby.
  - **SAMEV**: The number of block pairs that have the same version number between primary and standby.
  - **SAMEV&C**: The number of block pairs that have the same version and the same checksum.
  - **ENCERR**: The number of block pairs that have encryption related errors (usually because the Wallet is not open).
  - **LWLOC**: The number of lost writes on the local database.
  - **LWRMT**: The number of lost writes on the remote databases.
  - **DIFFPAIR** : The number of block pairs that have the same version but different checksum, and which the block compare callback function has concluded are different.

Example 61-1 Example

The following is an example of using the DBMS_DBCOMP.DBCOMP procedure.

```
declare
  Datafile_Name_or_Number varchar2(1000);
  Output_File varchar2(1000);
begin
  dbms_output.enable(1000000);
  Datafile_Name_or_Number:= 'all' ;
  Output_File:='BlockCompareAll_';
  sys.dbms_dbcomp.DBCOMP(Datafile_Name_or_Number, Output_File, true);
end;
```

62
DBMS_DBFS_CONTENT

The DBMS_DBFS_CONTENT package provides an interface comprising a file system-like abstraction backed by one or more Store Providers.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Exceptions
- Operational Notes
- Data Structures
- Summary of DBMS_DBFS_CONTENT Subprograms

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide*

62.1 DBMS_DBFS_CONTENT Overview

The DBMS_DBFS_CONTENT package provides an interface that connects the Oracle database to the DBFS client-side.

In the server, the DBMS_DBFS_CONTENT package is backed by the DBMS_DBFS_CONTENT_SPI package, which includes descriptions but not actual implementations of DBFS stores.

> **See Also:**
>
> - Oracle Database SecureFiles and Large Objects Developer's Guide for a description of DBMS_DBFS_CONTENT and the DBFS architecture
> - Oracle Database SecureFiles and Large Objects Developer's Guide for conceptual information about the DBMS_DBFS_CONTENT package

62.2 DBMS_DBFS_CONTENT Security Model

The DBMS_DBFS_CONTENT package runs under AUTHID CURRENT_USER.

62.3 DBMS_DBFS_CONTENT Constants

The `DBMS_DBFS_CONTENT` package uses various types of constants.

These are shown in the following tables:

- Table 62-1
- Table 62-2
- Table 62-3
- Table 62-4
- Table 62-5
- Table 62-6
- Table 62-7
- Table 62-8
- Table 62-9
- Table 62-10

Path Name Constants and Types

The following constants are useful for declaring paths and item names. Paths are limited to 1024 characters and item names are limited to 256 characters.

Table 62-1 DBMS_DBFS_CONTENT Constants - Path Names

Constant	Type	Value	Description
NAME_MAX	PLS_INTEGER	256	Maximum length of an absolute path name visible to clients
NAME_T	VARCHAR2(256)	NAME_MAX	Portable alias for string that can represent component names
PATH_MAX	PLS_INTEGER	1024	Maximum length of any individual component of an absolute path name visible to clients
PATH_T	VARCHAR2(1024)	PATH_MAX	Portable alias for string that can represent path names

ContentID Constants

Stores may expose to the user a unique identifier that represents a particular path item in the store. These identifiers are limited to 128 characters.

Table 62-2 DBMS_DBFS_CONTENT Constants - ContentID

Constant	Type	Value	Description
CONTENT_ID_MAX	PLS_INTEGER	128	Maximum length of a store-specific provider-generated contentID that identifies a file-type content item
CONTENT_ID_T	RAW(128)	CONTENT_ID_MAX	Portable alias for raw buffers that can represent contentID values

Path Properties Constants

Every path name in a store is associated with a set of properties. Each property is identified by a string "name", has a string "value" (which might be NULL if unset or undefined or unsupported by a specific store implementation) and a value "typecode" (a numeric discriminant for the actual type of value held in the "value" string.)

Table 62-3 DBMS_DBFS_CONTENT Constants - Properties

Constant	Type	Value	Description
PROPNAME_MAX	PLS_INTEGER	32	Maximum length of a property name
PROPNAME_T	VARCHAR2(32)	PROPNAME_MAX	Portable alias for string that can represent property names
PROPVAL_MAX	PLS_INTEGER	1024	Maximum length of the string value of a property
PROPVAL_T	VARCHAR2(1024)	PATH_MAX	Portable alias for string that can represent property values

Path Name Type Constants

Path items in a store have a item type associated with them. These types represent the kind of entry the item represents in the store.

Table 62-4 DBMS_DBFS_CONTENT Constants - Path Name Types

Constant	Type	Value	Description
TYPE_FILE	PLS_INTEGER	1	A regular file storing content (a logically linear sequence of bytes accessed as a BLOB
TYPE_DIRECTORY	PLS_INTEGER	2	A container of other path name types, including file types
TYPE_LINK	PLS_INTEGER	3	A symbolic link (that is, an uninterpreted string value associated with a path name). Since symbolic links may represent path names that fall outside the scope of any given store (or even the entire aggregation of stores managed by the DBMS_DBFS_CONTENT interface), or may not even represent path names, clients must be careful in creating symbolic links, and stores must be careful in trying to resolve these links internally.
TYPE_REFERENCE	PLS_INTEGER	4	A hard link which is always a valid path name alias to content
TYPE_SOCKET	PLS_INTEGER	5	UNIX domain socket created using socket interface defined as in socket.h with domain defined as AF_UNIX or AF_LOCAL.

Store Feature Constants

The DBFS content API allows different store providers (and different stores) to describe themselves through a *feature set* (a bitmask indicating which features they support and which ones they do not).

Table 62-5 DBMS_DBFS_CONTENT Constants - Store Features

Constant	Type	Value	Description
FEATURE_FOLDERS	PLS_INTEGER	1	Set if the store supports folders (or directories) as part of hierarchical path names
FEATURE_FOIAT	PLS_INTEGER	2	Set if implicit folder operations within the store (performed as part of a client-requested operation) runs inside autonomous transactions.In general, the use of autonomous transactions is a compromise between (a) simplicity in the implementation and client-controlled transaction scope for all operations, at the cost of greatly reduced concurrency (FEATURE_FOIAT not set), versus (b) more complex implementation and smaller client-controlled transaction scope, at the benefit of greatly increased concurrency (FEATURE_FOIAT set).
FEATURE_NOWAIT	PLS_INTEGER	4	Set if the store allows nowait gets of path elements. The default behavior is to wait for row locks; if nowait gets are implemented, the get operation raises an ORA-54 exception if the path element is already locked by another transaction.
FEATURE_ACLS	PLS_INTEGER	8	Set if the store supports Access Control Lists (ACLs) and internal authorization or checking based on these ACLs. ACLs are standard properties but a store may do nothing more than store and retrieve the ACLs without interpreting them in any way.
FEATURE_LINKS	PLS_INTEGER	16	Set if the store supports symbolic links, and if certain types of symbolic links (specifically non-absolute path names) can be internally resolved by the store itself
FEATURE_LINK_DEREF	PLS_INTEGER	32	Set if the store supports symbolic links, and if certain types of symbolic links (specifically non-absolute path names) can be internally resolved by the store itself
FEATURE_REFERENCES	PLS_INTEGER	64	Set if the store supports hard links

Table 62-5 (Cont.) DBMS_DBFS_CONTENT Constants - Store Features

Constant	Type	Value	Description
FEATURE_LOCKING	PLS_INTEGER	128	Set if the store supports user-level locks (read-only, write-only, read-write) that can be applied on various items of the store, and if the store uses these lock settings to control various types of access to the locked items. User-level locks are orthogonal to transaction locks and persist beyond the scope of any specific transaction, session, or connection — this implies that the store itself may not be able to clean up after dangling locks, and client-applications need to perform any garbage collection.
FEATURE_LOCK_HIERARCHY	PLS_INTEGER	256	Set if the store allows a user-lock to control access to the entire sub-tree under the locked path name.
FEATURE_LOCK_CONVERT	PLS_INTEGER	512	Set if the store supports upgrade or downgrade of locks from one mode to another
FEATURE_VERSIONING	PLS_INTEGER	1024	Set if the store supports at least a linear versioning and version management. Different versions of the same path name are identified by monotonic version numbers, with a version-nonqualified path name representing the latest version.
FEATURE_VERSION_PATH	PLS_INTEGER	2048	Set if the store supports a hierarchical namespace for different versions of a path name
FEATURE_SOFT_DELETES	PLS_INTEGER	4096	Set if the store supports a "soft-delete", that is, the ability to delete a path name and make it invisible to normal operations, but retain the ability to restore the path name later (as long as it has not been overwritten by a new create operation). The store also supports purging soft-deleted path names (making them truly deleted), and navigation modes that show soft-deleted items.
FEATURE_HASHING	PLS_INTEGER	8192	Set if the store automatically computes and maintains some type of a secure hash of the contents of a path name (typically a TYPE_FILE path).
FEATURE_HASH_LOOKUP	PLS_INTEGER	16384	Set if the store allows "content-based addressing", that is, the ability to locate a content item based, not on its path name, but on its content hash.

Table 62-5 (Cont.) DBMS_DBFS_CONTENT Constants - Store Features

Constant	Type	Value	Description

FEATURE_FILTERING	PLS_INTEGER	32768	Set if the store allows clients to pass a filter function (a PL/SQL function conforming to the signature below) that returns a logical boolean indicating if a given store item satisfies a selection predicate. Stores that support filtering may be able to more efficiently perform item listing, directory navigation, and deletions by embedding the filtering logic inside their implementation. If filtering is not supported, clients can retrieve more items than necessary and perform the filtering checks themselves, albeit less efficiently.
			A filter predicate is a function with the following signature:
			`function filterFunction(`
			` path IN VARCHAR2,`
			` store_name IN VARCHAR2,`
			` opcode IN INTEGER,`
			` item_type IN INTEGER,`
			` properties IN DBMS_DBFS_CONTENT_PROPERTIES_T,`
			` content IN BLOB)`
			`  RETURN INTEGER;`
			Any PL/SQL function conforming to this signature can examine the contents and properties of a store item, and determine if the item satisfies the selection criterion for the current operation. Any nonzero return value results in the DBMS_DBFS_CONTENT interface processing the item as part of the current operation; a return value that is zero or NULL results in the item being skipped from processing.

FEATURE_SEARCHING	PLS_INTEGER	65536	Set if the store allows clients to pass a text-search filter query to locate type_file path names based on their content. Stores that support searching may use indexes to accelerate such searches; otherwise, clients need to build their own indexes, or else search a potentially larger set of items to locate the ones of interest for the current search.
FEATURE_ASOF	PLS_INTEGER	131072	Set if the store allows clients to use a flashback timestamp in query operations (non-mutating GETPATH Procedures and LIST Function).

Table 62-5 (Cont.) DBMS_DBFS_CONTENT Constants - Store Features

Constant	Type	Value	Description
FEATURE_PROVIDER_PROPS	PLS_INTEGER	262144	Set if the store allows per-operation properties (that control the behavior of the store with regard to the current operation, as opposed to properties associated with individual items).
FEATURE_SNAPSHOTS	PLS_INTEGER	524288	Set if the store allows the use of named, read-only snapshots of its contents. It is up to the provider to implement snapshots using any suitable means (including creating immediate copies of the content, or using copy-on-write) and managing dependencies between snapshots and its parent content view.
FEATURE_CLONES	PLS_INTEGER	1048576	Set if the store allows the use of named, writable clones of its contents. It is up to the provider to implement clones using any suitable means (including creating immediate copies of the content, or using copy-on-write) and managing dependencies between clones and its parent content view.
FEATURE_LOCATOR	PLS_INTEGER	2097152	Set if the store allows direct access to file contents through a LOB locator. Stores that internally manipulate the file contents, perhaps by shredding or reassembling them in separate pieces, performing other transformations, and so on, cannot transparently give out a LOB locator to clients. The file contents of these stores should be accessed using the buffer-based interfaces.
FEATURE_CONTENT_ID	PLS_INTEGER	4194304	Set if the store allows a "pathless", contentID-based access to files (there is no notion of a directory, link, or reference in this model)
FEATURE_LAZY_PATH	PLS_INTEGER	8388608	Set if the store allows a lazy binding of a path name to file content elements that are otherwise identified by a contentID; this feature makes sense only in conjunction with FEATURE_CONTENT_ID

Lock Type Constants

Stores that support locking should implement 3 types of locks: LOCK_READ_ONLY, LOCK_WRITE_ONLY, and LOCK_READ_WRITE.

Table 62-6 DBMS_DBFS_CONTENT Constants - Lock Types

Constant	Type	Value	Description
LOCK_READ_ONLY	PLS_INTEGER	1	Locks as read-only

Table 62-6 (Cont.) DBMS_DBFS_CONTENT Constants - Lock Types

Constant	Type	Value	Description
LOCK_WRITE_ONLY	PLS_INTEGER	2	Locks as write-only
LOCK_READ_WRITE	PLS_INTEGER	3	Locks as read-write

Standard Property Constants

Standard properties are well-defined, mandatory properties associated with all content path names that all stores should support (in the manner described by the content interface), with some exceptions. For example, a read-only store need not implement a `modification_time` or `creation_time`.

All standard properties informally use the `STD` namespace, which clients and stores should avoid using.

Table 62-7 DBMS_DBFS_CONTENT Constants - Standard Properties

Constant	Type	Value	Description
STD_ACCESS_TIME	VARCHAR2(32)	'std:access_time'	TYPECODE_TIMESTAMP in UTC: The time of last access of a path name's contents
STD_ACL	VARCHAR2(32)	'std:acl'	TYPECODE_VARCHAR2: The access control list (in standard ACL syntax) associated with the path name
STD_CANONICAL_PATH	VARCHAR2(32)	'std:canonical_path'	TYPECODE_VARCHAR2: The canonical store-specific path name of an item.
STD_CHANGE_TIME	VARCHAR2(32)	'std:change_time'	TYPECODE_TIMESTAMP in UTC: The time of last change to the metadata of a path name
STD_CHILDREN	VARCHAR2(32)	'std:children'	TYPECODE_NUMBER: The number of child directories/folders a directory/folder path has (this property should be available in providers that support the FEATURE_FOLDERS feature)
STD_CONTENT_TYPE	VARCHAR2 (32)	'std:content_type'	TYPECODE_NUMBER: The client-supplied mime-type(s) (in standard RFC syntax) describing the (typically type_file) path name. The content type is not necessarily interpreted by the store.
STD_CREATION_TIME	VARCHAR2(32)	'std:creation_time'	TYPECODE_TIMESTAMP in UTC: The time at which the item was created (once set, this value never changes for the lifetime of the path name)
STD_DELETED	VARCHAR2(32)	'std:deleted'	TYPECODE_NUMBER as a BOOLEAN: Set to a nonzero number if the path name has been soft-deleted but not yet purged.

Table 62-7 (Cont.) DBMS_DBFS_CONTENT Constants - Standard Properties

Constant	Type	Value	Description
STD_GUID	VARCHAR2(32)	'std:guid'	TYPECODE_NUMBER: A store-specific unique identifier for a path name. Clients must not depend on the GUID being unique across different stores, but a given (*store-name*, *store-specific-path name*) has a stable and unique GUID for its lifetime.
STD_LENGTH	VARCHAR2(32)	'std:length'	TYPECODE_NUMBER: The length of the content (BLOB) of a TYPE_FILE/ TYPE_REFERENCE path, or the length of the referent of a TYPE_LINK symbolic link. Directories do not have a well-defined length and stores are free to set this property to zero, NULL, or any other value.
STD_MODIFICATION_TIME	VARCHAR2(32)	'std:modification_time'	TYPECODE_TIMESTAMP in UTC: The time of last change to the data associated with a path name. Change to the content of a TYPE_FILE/ TYPE_REFERENCE path, the referent of the TYPE_LINK path, and addition or deletion of immediate children in a TYPE_DIRECTORY path, all constitute data changes.
STD_OWNER	VARCHAR2(32)	'std:owner'	TYPECODE_VARCHAR2: A client-supplied (or implicit) owner name for the path name. The owner name may be used (along with the current "principal") for access checks by stores that support ACLs, locking, or both.
STD_PARENT_GUID	VARCHAR2(32)	'std:parent_guid'	TYPECODE_NUMBER: A store-specific unique identifier for the parent of a path name. Clients must not depend on the GUID being unique across different stores, but a given (*store-name*, *store-specific-path name*) has a stable and unique GUID for its lifetime. STD_PARENT_GUID(path name) == STD_GUID(parent(path name))
STD_REFERENT	VARCHAR2(32)	'std:referent'	TYPECODE_VARCHAR2: The content of the symbolic link of a TYPE_LINK path; NULL otherwise. As mentioned, the STD_REFERENT can be an arbitrary string and must not necessarily be interpreted as path name by clients (or such interpretation should be done with great care).

Optional Property Constants

Optional properties are well-defined properties (not mandatory) associated with all content path names that all stores are free to support (but only in the manner described by the DBFS content API).

All optional properties informally use the `opt:` namespace, which clients and stores should avoid using.

Table 62-8 DBMS_DBFS_CONTENT Constants - Optional Properties

Constant	Type	Value	Description
OPT_HASH_TYPE	VARCHAR2(32)	'opt:hash_type'	TYPECODE_NUMBER: The type of hash provided in the opt_hash_value property; see DBMS_CRYPTO for possible options.
OPT_HASH_VALUE	VARCHAR2(32)	'opt:hash_value'	TYPECODE_NUMBER: The hash value of type OPT_HASH_TYPE describing the content of the path name.
OPT_LOCK_COUNT	VARCHAR2(32)	'opt:lock_count'	TYPECODE_NUMBER: The number of (compatible) locks placed on a path name. If different principals are allowed to place compatible (read) locks on a path, the opt_locker must specify all lockers (with repeats so that lock counts can be correctly maintained).
OPT_LOCK_DATA	VARCHAR2(32)	'opt:lock_data'	TYPECODE_NUMBER: The client-supplied user-data associated with a user-lock, uninterpreted by the store.
OPT_LOCKER	VARCHAR2(32)	'opt:locker'	TYPECODE_NUMBER: One or more implicit or client-specified principals that applied a user-lock on a path name.
OPT_LOCK_STATUS	VARCHAR2(32)	'opt:lock_status'	TYPECODE_NUMBER: One of the LOCK_READ_ONLY, LOCK_WRITE_ONLY, LOCK_READ_WRITE values describing the type of lock currently applied on a path name.
OPT_VERSION	VARCHAR2(32)	'opt:version'	TYPECODE_NUMBER: A sequence number for linear versioning of a path name.
OPT_VERSION_PATH	VARCHAR2(32)	'opt:version_path'	TYPECODE_NUMBER: A version-path name for hierarchical versioning of a path name.
OPT_CONTENT_ID	VARCHAR2(32)	'opt:content_id'	TYPECODE_NUMBER: A provider-generated store-specific unique contentID in the form of a string for a file content element (that may optionally not be associated with a path; see FEATURE_CONTENT_ID and FEATURE_LAZY_PATH).

Property Access Flag Constants

Content interface methods to get or set properties can use combinations of property access flags to fetch properties from different name spaces in a single interface call.

Table 62-9 DBMS_DBFS_CONTENT Constants - Property Access Flags

Constant	Type	Value	Description
PROP_NONE	PLS_INTEGER	0	None: used when the client is not interested in any properties, and is invoking the content access method for other reasons (path name existence or lockability validation, data access, and so on)
PROP_STD	PLS_INTEGER	1	Mandatory: used when the client is interested in the standard properties; all standard properties are retrieved if this flag is specified.
PROP_OPT	PLS_INTEGER	2	Optional: used when the client is interested in the optional properties; all optional properties are retrieved if this flag is specified.
PROP_USR	PLS_INTEGER	3	User-defined: used when the client is interested in the user-defined properties; all user-defined properties are retrieved if this flag is specified.
PROP_ALL	PLS_INTEGER	PROP_STD + PROP_OPT + PROP_USR ;	All: an alias for the combination of all standard, optional, and user-defined properties
PROP_DATA	PLS_INTEGER	8	Content: used when the client is interested only in data access, and does not care about properties
PROP_SPC	PLS_INTEGER	16	Specific: used when the client is interested in a mix-and-match of different subsets of various property name spaces; the names of the specific properties to fetch are passed into the content interface method call as arguments, and only these property values are fetched and returned to the client. This is useful in cases where there are a very large number of properties potentially accessible, but the client is interested in only a small number of them (and knows the names of these "interesting" properties beforehand). PROP_SPC is applicable only to the various GETPATH operations. Other operations that specify properties ignore PROP_SPC specifications.

Operation Code Constants

All of the operations in the DBFS content API are represented as abstract opcodes.

Clients can use these opcodes to directly and explicitly by invoking the CHECKACCESS Function to verify if a particular operation can be invoked by a given principal on a particular path name.

Table 62-10 DBMS_DBFS_CONTENT Constants - Operation Codes

Constant	Type	Value	Description
OP_CREATE	PLS_INTEGER	1	Create a path item
OP_CREATEFILE	PLS_INTEGER	OP_CREATE	Create a file
OP_CREATELINK	PLS_INTEGER	OP_CREATE	Create a soft link
OP_CREATEREFERENCE	PLS_INTEGER	OP_CREATE	Create a reference (hard link)
OP_DELETE	PLS_INTEGER	2	Soft-deletion, purge, and restore operations are all represented by OP_DELETE
OP_DELETEFILE	PLS_INTEGER	OP_DELETE	Delete a file
OP_DELETEDIRECTORY	PLS_INTEGER	OP_DELETE	Delete a directory
OP_RESTORE	PLS_INTEGER	OP_DELETE	Restore a soft-deleted path item
OP_PURGE	PLS_INTEGER	OP_DELETE	Purge a soft-deleted path item
OP_READ	PLS_INTEGER	3	Read from a path item
OP_GET	PLS_INTEGER	OP_READ	Get a path item for either read or update operations
OP_WRITE	PLS_INTEGER	4	Write a path item
OP_PUT	PLS_INTEGER	OP_WRITE	Put (write) to a path item
OP_RENAME	PLS_INTEGER	5	Rename a path item
OP_RENAMEFROM	PLS_INTEGER	OP_RENAME	Operations performed on the source of a rename
OP_RENAMETO	PLS_INTEGER	OP_RENAME	Operations performed on the destination of a rename
OP_SETPATH	PLS_INTEGER	OP_RENAME	Set a path item name
OP_LIST	PLS_INTEGER	6	Perform a path listing
OP_SEARCH	PLS_INTEGER	7	Perform a search
OP_LOCK	PLS_INTEGER	8	Lock a path item
OP_UNLOCK	PLS_INTEGER	9	Unlock a path item

Table 62-10 (Cont.) DBMS_DBFS_CONTENT Constants - Operation Codes

Constant	Type	Value	Description
OP_ACL	PLS_INTEGER	10	An implicit operation invoked during an OP_CREATE or OP_PUT that specifies a STD_ACL property; the operation tests to see if the principal is allowed to set or change the ACL of a store item
OP_STORE	PLS_INTEGER	11	A catch-all category for miscellaneous store operations that do not fall under any of the other operational interfaces

62.4 Exceptions

DBFS content API operations can raise any one of these top-level exceptions.

Table 62-11 DBMS_DBFS_CONTENT Exceptions

Exception	Code	Description
PATH_EXISTS	64000	A specified path name already exists
INVALID_PARENT	64001	Parent of a specified path name does not exist
INVALID_PATH	64002	Specified path name does not exist, or is not valid
UNSUPPORTED_OPERATION	64003	An operation unsupported by a store was invoked
INVALID_ARGUMENTS	64004	An operation was invoked with invalid arguments
INVALID_ACCESS	64005	Access control checks failed for the current operation
LOCK_CONFLICT	64006	Current operation failed lock conflict check
INVALID_STORE	64007	An invalid store name was specified
INVALID_MOUNT	64008	An invalid mount point was specified
INVALID_PROVIDER	64009	An invalid provider-package was specified
READONLY_PATH	64010	A mutating operation was invoked on a read-only mount or store

62.5 DBMS_DBFS_CONTENT Operational Notes

This topic lists operational notes for DBMS_DBFS_CONTENT implementation, path names, and other operations.

- Implementation
- Path Names
- Other DBMS_DBFS_CONTENT Operations

Implementation

Since the interconnection of the DBMS_DBFS_CONTENT interface and the provider SPI is a 1-to-many pluggable architecture, the interface uses dynamic SQL to invoke methods in the provider SPI, this can lead to runtime errors.

There are no explicit `INIT` or `FINI` methods to indicate when the `DBMS_DBFS_CONTENT` interface plugs or unplugs a particular provider SPI. Provider SPIs must be willing to auto-initialize themselves at any SPI entry-point.

All operations performed by a store provider are "stateless" in that they are complete operations unto themselves. If state is necessary to be maintained for some reason, then the state must be maintained in data structures such as auxiliary tables that can be queried as needed.

Path Names

All path names used in the provider SPI are store-qualified in pair form (`store_name`, `pathname`) where the path name is rooted within the store namespace.

Stores and their providers that support contentID-based access (see `FEATURE_CONTENT_ID` in Table 62-5) also support a form of addressing that is not based on path names. Content items are identified by an explicit store name, a `NULL` path name, and possibly a contentID specified as a parameter or by way of the `OPT_CONTENT_ID` (see Table 62-8) property.

Not all operations are supported with contentID-based access, and applications should depend only on the simplest create or delete functionality being available.

This table lists other operations and provides links to related discussions.

Other DBMS_DBFS_CONTENT Operations

Table 62-12 Other DBMS_DBFS_CONTENT Operations

Other Operations	See ...
Creation	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on creation operations
Deletion	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on deletion operations
Get (Retrieve) and Put (Insert)	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Get and Put operations
Rename and Move	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Rename and Move operations
Directory Navigation and Search	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Navigation and Search operations
Locking	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Locking operations
Access Check	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Access Check operations

62.6 DBMS_DBFS_CONTENT Data Structures

The `DBMS_DBFS_CONTENT` package defines `RECORD` types and `TABLE` types.

RECORD Types

- FEATURE_T Record Type
- MOUNT_T Record Type
- PATH_ITEM_T Record Type

- PROP_ITEM_T Record Type
- PROPERTY_T Record Type
- STORE_T Record Type

TABLE Types

- FEATURES_T Table Type
- MOUNTS_T Table Type
- PATH_ITEMS_T Table Type
- PROP_ITEMS_T Table Type
- PROPERTIES_T Table Type
- STORES_T Table Type

Usage Notes

There is an approximate correspondence between `DBMS_DBFS_CONTENT_PROPERTY_T` and `PROPERTY_T` — the former is a SQL object type that describes the full property tuple, while the latter is a PL/SQL record type that describes only the property value component.

Likewise, there is an approximate correspondence between `DBMS_DBFS_CONTENT_PROPERTIES_T` and `PROPERTIES_T` — the former is a SQL nested table type, while the latter is a PL/SQL hash table type.

Dynamic SQL calling conventions force the use of SQL types, but PL/SQL code may be implemented more conveniently in terms of the hash-table types.

The `DBMS_DBFS_CONTENT` interface provides convenient utility functions to convert between `DBMS_DBFS_CONTENT_PROPERTIES_T` and `PROPERTIES_T` (see propertiesT2H and propertiesH2T).

Clients can query the `DBMS_DBFS_CONTENT` interface for the list of available stores, determine which store is to handle access to a given path name, and determine the feature set for the store.

62.6.1 DBMS_DBFS_CONTENT FEATURE_T Record Type

This type describes a store mount point and its properties.

Syntax

```
TYPE feature_t IS RECORD (
   feature_name    VARCHAR2(32),
   feature_mask    INTEGER,
   feature_state   VARCHAR2(3));
```

Fields

Table 62-13 MOUNT_T Fields

Field	Description
`feature_name`	Name of feature
`feature_mask`	Value used to mask off all other bits other than this feature in the feature value

Table 62-13 (Cont.) MOUNT_T Fields

Field	Description
feature_state	'YES' or 'NO' depending on whether the feature is supported on this store

62.6.2 DBMS_DBFS_CONTENT MOUNT_T Record Type

This type describes a store mount point and its properties.

Syntax

```
TYPE mount_t IS RECORD (
   store_name          VARCHAR2(32),
   store_id            NUMBER,
   provider_name       VARCHAR2(32),
   provider_pkg        VARCHAR2(32),
   provider_id         NUMBER,
   provider_version    VARCHAR2(32),
   store_features      INTEGER,
   store_guid          NUMBER,
   store_mount         NAME_T,
   mount_properties    DBMS_DBFS_CONTENT_PROPERTIES_T);
```

Fields

Table 62-14 MOUNT_T Fields

Field	Description
store_name	Name of store
store_id	ID of store
provider_name	Name of the content store
provider_pkg	PL/SQL package name for the content store
provider_id	Unique identifier for the content store
provider_version	Version number for the content store
respos_features	Features supported by this content store
store_guid	Unique ID for this instance of the store
store_mount	Location at which this store instance is mounted
mount_properties	Properties for this mount point (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)

62.6.3 DBMS_DBFS_CONTENT PATH_ITEM_T Record Type

A PATH_ITEM_T is a tuple describing a (store, mount) qualified path in a store, with all standard and optional properties associated with it.

Syntax

```
TYPE path_item_t IS RECORD (
   store                    NAME_T,
```

```
  mount                   NAME_T,
  pathname                PATH_T,
  pathtype                VARCHAR2(32),
  filedata                BLOB,
  std_access_time         TIMESTAMP,
  std_acl                 VARCHAR2(1024),
  std_change_time         TIMESTAMP,
  std_children            NUMBER,
  std_content_type        VARCHAR2(1024),
  std_creation_time       TIMESTAMP,
  std_deleted             INTEGER,
  std_guid                INTEGER,
  std_modification_time   TIMESTAMP,
  std_owner               VARCHAR2(32),
  std_parent_guid         INTEGER,
  std_referent            VARCHAR2(1024),
  opt_hash_type           VARCHAR2(32),
  opt_hash_value          VARCHAR2(128),
  opt_lock_count          INTEGER,
  opt_lock_data           VARCHAR2(128),
  opt_locker              VARCHAR2(128),
  opt_lock_status         INTEGER,
  opt_version             INTEGER,
  opt_version_path        PATH_T,
  opt_content_id          CONTENT_ID_T);
```

Fields

Table 62-15 PATH_ITEM_T Fields

Field	Description
`store`	Name of store
`mount`	Location at which instance of store is mounted
`pathname`	Name of path to item
`pathtype`	Type of object path (see Table 62-4)
`filedata`	`BLOB` locator `that can be` used to access data in the path item
`std_access_time`	Time of last access of a pathname's contents
`std_acl`	Access Control List (in standard ACL syntax)
`std_change_time`	Time of last change to the metadata of a path name
`std_children`	Number of child directories or folders a directory or folder path (this property should be available in providers that support the `feature_folders` feature).
`std_content_type`	One or more client-supplied mime-types (in standard RFC syntax) describing the path name which is typically of `type_file`. The content type s not necessarily interpreted by the store.
`std_creation_time`	Time at which the item was created. Once set, this value remains the same for the lifetime of the path name.
`std_deleted`	Set to a nonzero number if the path name has been soft-deleted but not yet purged (see Table 62-5)
`std_guid`	Store-specific unique identifier for a path name. Clients must not depend on the GUID being unique across different stores, but a given *store-name*, *store-specific-pathname* has a stable and unique GUID for its lifetime.

Table 62-15 (Cont.) PATH_ITEM_T Fields

Field	Description
std_modification_time	Time of last change to the data associated with a path name. Changes to the content of a type_file or type_reference path, the referent of the type_link path, and addition or deletion of immediate children in a type_directory path, all constitute data changes.
std_owner	Client-supplied (or implicit) owner name for the path name
std_parent_guid	Store-specific unique identifier for the parent of a path name. Clients must not depend on the GUID being unique across different stores, but a given *store-name*, *store-specific-pathname* has a stable and unique GUID for its lifetime. std_parent_guid(pathname) == std_guid(parent(pathname))
std_referent	Content of the symbolic link of a type_link path, otherwise NULL. As mentioned before, the std_referent can be an arbitrary string and must not necessarily be interpreted as pathname by clients (or such interpretation should be done with great care).
opt_hash_type	Type of hash provided in the opt_hash_value property (see DBMS_CRYPTO for possible options)
opt_hash_value	Hash value of type opt_hash_type describing the content of the path name
opt_lock_count	Number of compatible locks placed on a path name. If different principals are allowed to place compatible (read) locks on a path, the opt_locker must specify all lockers with repeats so that lock counts can be correctly maintained.
opt_lock_data	Client-supplied user-data associated with a user-lock, uninterpreted by the store
opt_locker	One or more implicit or client-specified principals that applied a user-lock on a path name
opt_lock_status	One of the lock_read_only, lock_write_only, lock_read_write values describing the type of lock currently applied on a path name
opt_version	Sequence number for linear versioning of a path name
opt_version_path	Version path name for hierarchical versioning of a path name
opt_content_id	Stringified provider-generated store-specific unique contentID for a file element (that may optionally not be associated with a path (see FEATURE_CONTENT_ID and FEATURE_LAZY_PATH in Table 62-5)

62.6.4 DBMS_DBFS_CONTENT PROP_ITEM_T Record Type

A PROP_ITEM_T is a tuple describing a (store, mount) qualified path in a store, with all user-defined properties associated with it, expanded out into individual (name, value, type) tuples.

Syntax

```
TYPE prop_item_t IS RECORD (
   store               NAME_T,
   mount                    NAME_T,
   pathname                 PATH_T,
   property_name            PROPNAME_T,
```

```
   property_value          PROPVAL_T,
   property_type           INTEGER);
```

Fields

Table 62-16 PROP_ITEM_T Fields

Field	Description
store	Name of store
mount	Location at which instance of store is mounted
pathname	Name of path to item
property_name	Name of the property
property_value	Value of the property
property_type	PL/SQL typecode for the property value

62.6.5 PROPERTY_T Record Type

This type describes a single (value, typecode) property value tuple; the property name is implied.

See PROPERTIES_T Table Type for more information.

Syntax

```
TYPE property_t  IS RECORD (
   propvalue    PROPVAL_T,
   typecode     INTEGER);
```

Fields

Table 62-17 PROPERTY_T Fields

Field	Description
propvalue	Value of property
typecode	Typecode

62.6.6 DBMS_DBFS_CONTENT STORE_T Record Type

This type describes a store registered with and managed by the DBMS_DBFS_CONTENT interface.

Syntax

```
TYPE store_t IS RECORD (
   store_name          VARCHAR2(32),
   store_id            NUMBER,
   provider_name       VARCHAR2(32),
   provider_pkg        VARCHAR2(32),
   provider_id         NUMBER,
   provider_version    VARCHAR2(32),
   store_features      INTEGER,
   store_guid          NUMBER);
```

Fields

Table 62-18 STORET_T Fields

Field	Description
store_name	Name of store
store_name	ID of store
provider_name	Name of the content store
provider_pkg	PL/SQL package name for the content store
provider_id	Unique identifier for the content store
provider_version	Version number for the content store
respos_features	Features supported by this content store
store_guid	Unique ID for this instance of the store

62.6.7 DBMS_DBFS_CONTENT FEATURES_T Table Type

A table type of FEATURE_T Record Type.

Syntax

```
TYPE features_t IS TABLE OF feature_t;
```

Related Topics

- DBMS_DBFS_CONTENT FEATURE_T Record Type
  This type describes a store mount point and its properties.

62.6.8 MOUNTS_T Table Type

A table type of MOUNT_T Record Type.

Syntax

```
TYPE mounts_t IS TABLE OF mount_t;
```

Related Topics

- DBMS_DBFS_CONTENT MOUNT_T Record Type
  This type describes a store mount point and its properties.

62.6.9 DBMS_DBFS_CONTENT PATH_ITEMS_T Table Type

A table type of PATH_ITEM_T Record Type

Syntax

```
TYPE path_items_t IS TABLE OF path_item_t;
```

Related Topics

- DBMS_DBFS_CONTENT PATH_ITEM_T Record Type
  A `PATH_ITEM_T` is a tuple describing a (store, mount) qualified path in a store, with all standard and optional properties associated with it.

62.6.10 DBMS_DBFS_CONTENT PROP_ITEMS_T Table Type

A table type of PATH_ITEM_T Record Type.

Syntax

```
TYPE prop_items_t IS TABLE OF prop_item_t;
```

Related Topics

- DBMS_DBFS_CONTENT PATH_ITEM_T Record Type
  A `PATH_ITEM_T` is a tuple describing a (store, mount) qualified path in a store, with all standard and optional properties associated with it.

62.6.11 DBMS_DBFS_CONTENT PROPERTIES_T Table Type

This is a name-indexed hash table of property tuples. The implicit hash-table association between the index and the value allows the client to build up the full `DBMS_DBFS_CONTENT_PROPERTY_T` tuples for a `PROPERTIES_T`.

Syntax

```
TYPE properties_t IS TABLE OF property_t INDEX BY propname_t;
```

62.6.12 STORES_T Table Type

This type describes a store registered with and managed by the `DBMS_DBFS_CONTENT` interface.

Syntax

```
TYPE stores_t IS TABLE OF store_t;
```

62.7 Summary of DBMS_DBFS_CONTENT Subprograms

This table lists and describes the subprograms used in the `DBMS_DBFS_CONTENT` Package.

Table 62-19 DBMS_DBFS_CONTENT Package Subprograms

Subprogram	Description
CHECKACCESS Function	Reports if the user (`principal`) can perform the specified operation on the given path
CHECKSPI Functions and Procedures	Checks if a user-provided package implements all of the `DBMS_DBFS_CONTENT_SPI` subprograms with the proper signatures, and reports on the conformance.
CREATEDIRECTORY Procedures	Creates a directory
CREATEFILE Procedures	Creates a file
CREATELINK Procedures	Creates a new reference to the source file system element
CREATEREFERENCE Procedures	Creates a physical link to an already existing file system element
DECODEFEATURES Function	Given a feature bit set integer value, returns a `FEATURES_T` table of the feature bits as `FEATURE_T` records
DELETECONTENT Procedure	Deletes the file specified by the given contentID

Table 62-19 (Cont.) DBMS_DBFS_CONTENT Package Subprograms

Subprogram	Description
DELETEDIRECTORY Procedure	Deletes a directory
DELETEFILE Procedure	Deletes a file
FEATURENAME Function	Given a feature bit, returns a `VARCHAR2` of that feature's name
FLUSHSTATS Function	Flushes `DBMS_DBFS_CONTENT` statistics to disk
GETDEFAULTACL Procedure	Returns the ACL parameter of the default context
GETDEFAULTASOF Procedure	Returns the `asof` parameter of the default context
GETTDEFAULTCONTEXT Procedure	Returns the default context
GETDEFAULTOWNER Procedure	Returns the `owner` parameter of the default context
GETDEFAULTPRINCIPAL Procedure	Returns the `principal` parameter of the default context
GETFEATURESBYMOUNT Function	Returns features of a store by mount point
GETFEATURESBYNAME Function	Returns features of a store by store name
GETFEATURESBYPATH Function	Returns features of a store by path
GETPATHBYMOUNTID Function	Returns the full absolute path name
GETPATH Procedures	Returns existing path items (such as files and directories)
GETPATHBYSTOREID Function	If the underlying GUID is found in the underlying store, returns the store-qualified path name
GETPATHNOWAIT Procedures	Implies that the operation is for an update, and, if implemented, allows providers to return an exception (`ORA-00054`) rather than wait for row locks.
GETSTOREBYMOUNT Function	Returns a store by way of its mount point
GETSTOREBYNAME Function	Returns a store by way of its name
GETSTOREBYPATH Function	Returns a store by way of its path
GETSTATS Procedure	Returns information about `DBMS_DBFS_CONTENT` statistics collection
GETTRACE Function	Returns whether or not `DBMS_DBFS_CONTENT` tracing is turned on
GETVERSION Function	Returns the version of the `DBMS_DBFS_CONTENT` interface in a standardized format associated with a store
LIST Function	Lists the path items in the specified path meeting the specified filter and other criteria
LISTALLCONTENT Function	Lists all path items in all mounts
LISTALLPROPERTIES Function	Returns a table of all properties for all path items in all mounts
LISTMOUNTS Function	Lists all available mount points, their backing stores, and the store features
LISTSTORES Function	Lists all available stores and their features
LOCKPATH Procedure	Applies user-level locks to the given valid path name
MOUNTSTORE Procedure	Mounts a previously registered store and binds it to the mount point
NORMALIZEPATH Functions	Converts a store-specific or full-absolute path name into normalized form

Table 62-19 (Cont.) DBMS_DBFS_CONTENT Package Subprograms

Subprogram	Description
PROPANY Functions	Provides constructors that take one of a variety of types and return a PROPERTY_T
PROPERTIESH2T Function	Converts a `PROPERTY_T` hash to a `DBMS_DBFS_CONTENT_PROPERTIES_T` table
PROPERTIEST2H Function	Converts a `DBMS_DBFS_CONTENT_PROPERTIES_T` table to a `PROPERTY_T` hash
PROPNUMBER Function	Is a constructor that takes a `NUMBER` and returns a `PROPERTY_T`
PROPRAW Function	Is a constructor that takes a `RAW` and returns a `PROPERTY_T`
PROPTIMESTAMP Function	Is a constructor that takes a `TIMESTAMP` and returns a `PROPERTY_T`
PROPVARCHAR2 Function	Is a constructor that takes a `VARCAHR2` and returns a `PROPERTY_T`
PURGEALL Procedure	Purges all soft-deleted entries matching the path and optional filter criteria
PURGEPATH Procedure	Purges any soft-deleted versions of the given path item
PUTPATH Procedures	Creates a new path item
REGISTERSTORE Procedure	Registers a new store
RENAMEPATH Procedures	Renames or moves a path
RESTOREALL Procedure	Restores all soft-deleted path items meeting the path and filter criteria
RESTOREPATH Procedure	Restores all soft-deleted path items that match the given path and filter criteria
SETDEFAULTACL Procedure	Sets the ACL parameter of the default context
SETDEFAULTASOF Procedure	Sets the "as of" parameter of the default context
SETDEFAULTCONTEXT Procedure	Sets the default context
SETDEFAULTOWNER Procedure	Sets the "owner" parameter of the default context
SETDEFAULTPRINCIPAL Procedure	Sets the "principal" parameter of the default context
SETPATH Procedures	Assigns a path name to a path item represented by contentID
SETSTATS Procedure	Enables and disables statistics collection
SETTRACE Procedure	Sets `DBMS_DBFS_CONTENT` tracing on or off
SPACEUSAGE Procedure	Queries file system space usage statistics
TRACE Procedure	Returns a `CLOB` that contains the evaluation results
TRACEENABLED Function	Determines if the current trace "severity" set by the SETTRACE Procedure is at least as high as the given trace level
UNLOCKPATH Procedure	Unlocks path items that were previously locked with the LOCKPATH Procedure
UNMOUNTSTORE Procedure	Unmounts a registered store
UNREGISTERSTORE Procedure	Unregisters a store

62.7.1 CHECKACCESS Function

This function reports if the user (`principal`) can perform the specified operation on the given path. This enables verifying the validity of an operation without attempting to perform the operation. If `CHECKACCESS` returns 0, then the subprogram invoked to implement that operation should fail with an error.

Syntax

```
DBMS_DBFS_CONTENT.CHECKACCESS (
   path           IN     VARCHAR2,
   pathtype       IN     INTEGER,
   operation      IN     VARCHAR2,
   principal      IN     VARCHAR2,
   store_name     IN     VARCHAR2  DEFAULT NULL)
  RETURN  BOOLEAN;
```

Parameters

Table 62-20 CHECKACCESS Procedure Parameters

Parameter	Description
`path`	Name of path to check for access
`pathtype`	Type of object `path` represents (see Table 62-4)
`operation`	Operation to be checked (see Table 62-8)
`principal`	File system user for whom the access check is made
`store_name`	Name of store

Usage Notes

Whether or not the user invokes this function, a store that supports access control internally performs these checks to guarantee security.

62.7.2 CHECKSPI Functions and Procedures

Given the name of a putative `DBMS_DBFS_CONTENT_SPI` conforming package, this function or procedure checks whether the package implements all of the provider subprograms with the proper signatures, and reports on the conformance.

Syntax

```
DBMS_DBFS_CONTENT.CHECKSPI (
   package_name        IN              VARCHAR2)
  RETURN  CLOB;

DBMS_DBFS_CONTENT.CHECKSPI (
   schema_name         IN              VARCHAR2,
   package_name        IN              VARCHAR2)
  return  clob;

DBMS_DBFS_CONTENT.CHECKSPI (
   package_name        IN              VARCHAR2,
   chk                 IN OUT NOCOPY   CLOB);
```

```
DBMS_DBFS_CONTENT.CHECKSPI (
   schema_name         in              VARCHAR2,
   package_name        in              VARCHAR2,
   chk                 in out nocopy   CLOB);
```

Parameters

Table 62-21 CHECKSPI Procedure Parameters

Parameter	Description
package_name	Name of package
schema_name	Name of schema
chk	CLOB that contains the evaluation results

Usage Notes

- The functional form returns a cached temporary LOB of session duration with the results of the analysis. The caller is expected to manage the lifetime of this LOB, as needed.
- The procedural form generates the results of the analysis into the chk LOB parameter; if the value passed in is NULL, the results are written to the foreground trace file provided that DBMS_DBFS_CONTENT interface tracing is enabled. If neither tracing is enabled nor a valid LOB passed in, the checker does not provide any useful indication of the analysis (other than raise exceptions if it encounters a serious error).
- If schema_name is NULL, standard name resolution rules (current schema, private synonym, public synonym) are used to try and locate a suitable package to analyze.

62.7.3 CREATEDIRECTORY Procedures

This procedure creates a directory.

Syntax

```
DBMS_DBFS_CONTENT.CREATEDIRECTORY (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   prop_flags  IN              INTEGER     DEFAULT PROP_STD,
   recurse     IN              BOOLEAN     DEFAULT FALSE,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.CREATEDIRECTORY (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   prop_flags  IN              INTEGER     DEFAULT PROP_STD,
   recurse     IN              BOOLEAN     DEFAULT FALSE,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-22 CREATEDIRECTORY Procedure Parameters

Parameter	Description
path	Name of path to the directory

Table 62-22 (Cont.) CREATEDIRECTORY Procedure Parameters

Parameter	Description
properties	One or more properties and their values to be set, returned, or both, depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting PROP_SPC (see Table 62-9), and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
recurse	If 0, do not execute recursively; otherwise, recursively create the directories above the given directory
store_name	Name of store
principal	File system user for whom the access check is made

62.7.4 CREATEFILE Procedures

This procedure creates a file.

Syntax

```
DBMS_DBFS_CONTENT.CREATEFILE (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   content     IN OUT NOCOPY   BLOB,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD + PROP_DATA),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.CREATEFILE (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   content     IN OUT NOCOPY   BLOB,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD + PROP_DATA),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-23 CREATEFILE Procedure Parameters

Parameter	Description
path	Name of path to the file
properties	One or more properties and their values to be set, returned, or both, depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
content	BLOB holding data with which to populate the file (optional)

Table 62-23 (Cont.) CREATEFILE Procedure Parameters

Parameter	Description
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
store_name	Name of store
principal	File system user for whom the access check is made

62.7.5 CREATELINK Procedures

This procedure creates a new link element srcPath with the value of dstPath. The value of dstPath is not validated or interpreted in any way by this procedure. This is analogous to a UNIX file system symbolic link.

Syntax

```
DBMS_DBFS_CONTENT.CREATELINK (
   srcPath     IN              VARCHAR2,
   dstPath     IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   prop_flags  IN              INTEGER     DEFAULT PROP_STD,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.CREATELINK (
   srcPath     IN              VARCHAR2,
   dstPath     IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   prop_flags  IN              INTEGER     DEFAULT PROP_STD,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-24 CREATELINK Procedure Parameters

Parameter	Description
srcPath	File system entry to create.
dstPath	Value to associate with srcPath.
properties	One or more properties and their values to be set, returned depending, or both, on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
store_name	Name of store
principal	File system user for whom the access check is made

62.7.6 CREATEREFERENCE Procedures

This procedure creates a physical link, `srcPath`, to an already existing file system element, `dstPath` (such as file or directory). The resulting entry shares the same metadata structures as the value of the `dstPath` parameter, and so is similar to incrementing a reference count on the file system element. This is analogous to a UNIX file system hard link.

Syntax

```
DBMS_DBFS_CONTENT.CREATEREFERENCE (
   srcPath     IN              VARCHAR2,
   dstPath     IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   prop_flags  IN              INTEGER     DEFAULT PROP_STD,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.CREATEREFERENCE (
   srcPath     IN              VARCHAR2,
   dstPath     IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   prop_flags  IN              INTEGER     DEFAULT PROP_STD,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-25 CREATEREFERENCE Procedure Parameters

Parameter	Description
`srcPath`	File system entry to create.
`dstPath`	Path that is the reference to `srcPath`.
`properties`	One or more properties and their values to be set, returned, or both, depending on `prop_flags` (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
`prop_flags`	Determines which properties are set, returned. Default is `PROP_STD`. Specify properties to be returned by setting `prop_spec`, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
`store_name`	Name of store
`principal`	File system user for whom the access check is made

62.7.7 DECODEFEATURES Function

Given a feature bit set integer value, this function returns a `FEATURES_T` table of the feature bits as `FEATURE_T` records.

Syntax

```
DBMS_DBFS_CONTENT.DECODEFEATURES (
   featureSet          IN      INTEGER)
  RETURN FEATURES_T DETERMINISTIC PIPELINED;
```

Parameters

Table 62-26 DECODEFEATURES Function Parameters

Parameter	Description
featureSet	Feature set

Return Values

FEATURES_T Table Type

62.7.8 DELETECONTENT Procedure

This procedure deletes the file specified by the given contentID.

Syntax

```
DBMS_DBFS_CONTENT.DELETECONTENT (
   store_name     IN     VARCHAR2    DEFAULT NULL,
   contentID      IN     RAW,
   filter         IN     VARCHAR2    DEFAULT NULL,
   soft_delete    IN     BOOLEAN     DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-27 DELETECONTENT Procedure Parameters

Parameter	Description
store_name	Name of store
contentID	Unique identifier for the file to be deleted
filter	A filter, if any, to be applied
soft_delete	If 0, execute a hard (permanent) delete. For any value other than 0, perform a soft delete (see *Oracle Database SecureFiles and Large Objects Developer's Guide*, Deletion Operations).
principal	File system user for whom the access check is made

62.7.9 DELETEDIRECTORY Procedure

This procedure deletes a directory.

If recurse is nonzero, it recursively deletes all elements of the directory. A filter, if supplied, determines which elements of the directory are deleted.

Syntax

```
DBMS_DBFS_CONTENT.DELETEDIRECTORY (
   path           IN     VARCHAR2,
   filter         IN     VARCHAR2    DEFAULT NULL,
   soft_delete    IN     BOOLEAN     DEFAULT NULL,
   recurse        IN     BOOLEAN     DEFAULT FALSE,
   store_name     IN     VARCHAR2    DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-28 DELETEDIRECTORY Procedure Parameters

Parameter	Description
path	Name of path to the directory
filter	A filter, if any, to be applied
soft_delete	If 0, execute a hard (permanent) delete. For any value other than 0, perform a soft delete see *Oracle Database SecureFiles and Large Objects Developer's Guide*, Deletion Operations.
recurse	If 0, do not execute recursively. Otherwise, recursively delete the directories and files below the given directory.
store_name	Name of store
principal	File system user for whom the access check is made

62.7.10 DELETEFILE Procedure

This procedure deletes the specified file.

Syntax

```
DBMS_DBFS_CONTENT.DELETEFILE (
   path           IN     VARCHAR2,
   filter         IN     VARCHAR2    DEFAULT NULL,
   soft_delete    IN     BOOLEAN     DEFAULT NULL,
   store_name     IN     VARCHAR2    DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-29 DELETEFILE Procedure Parameters

Parameter	Description
path	Name of path to the file
filter	A filter, if any, to be applied
soft_delete	If 0, execute a hard (permanent) delete. For any value other than 0, perform a soft delete (see *Oracle Database SecureFiles and Large Objects Developer's Guide*, Deletion Operations).
store_name	Name of store
principal	File system user for whom the access check is made

62.7.11 FEATURENAME Function

Given a feature bit, this function returns a `VARCHAR2` of that feature's name.

Syntax

```
DBMS_DBFS_CONTENT.FEATURENAME (
   featureBit          IN      INTEGER)
  RETURN VARCHAR2 DETERMINISTIC;
```

Parameters

Table 62-30 FEATURENAME Function Parameters

Parameter	Description
featureBit	Bit representation of the feature (see Table 62-5)

Return Values

Name of the feature

62.7.12 FLUSHSTATS Function

This procedure flushes DBMS_DBFS_CONTENT statistics to disk.

Syntax

```
DBMS_DBFS_CONTENT.FLUSHSTATS;
```

62.7.13 GETDEFAULTACL Procedure

This procedure returns the ACL parameter of the default context. This information can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.GETDEFAULTACL (
   acl    OUT NOCOPY     VARCHAR2);
```

Parameters

Table 62-31 GETDEFAULTACL Procedure Parameters

Parameter	Description
acl	ACL for all new elements created (implicitly or explicitly) by the current operation

62.7.14 GETDEFAULTASOF Procedure

This procedure returns the "as of" parameter of the default context. This information can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.GETDEFAULTASOF (
  asof    OUT NOCOPY     TIMESTAMP);
```

Parameters

Table 62-32 GETDEFAULTASOF Procedure Parameters

Parameter	Description
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes

62.7.15 GETDEFAULTCONTEXT Procedure

This procedure returns the default context. The information contained in the context can be inserted explicitly by way of arguments to the various method calls, allowing for fine-grained control over individual operations.

Syntax

```
DBMS_DBFS_CONTENT.GETDEFAULTCONTEXT (
   principal    OUT NOCOPY     VARCHAR2,
   owner        OUT NOCOPY     VARCHAR2,
   acl          OUT NOCOPY     VARCHAR2,
   asof         OUT NOCOPY     TIMESTAMP);
```

Parameters

Table 62-33 GETDEFAULTCONTEXT Procedure Parameters

Parameter	Description
principal	Agent (principal) invoking the current operation
owner	Owner for new elements created (implicitly or explicitly) by the current operation
acl	ACL for all new elements created (implicitly or explicitly) by the current operation
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes

62.7.16 GETDEFAULTOWNER Procedure

This procedure returns the "owner" parameter of the default context. This information can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.GETDEFAULTOWNER (
   principal    IN     VARCHAR2);
```

Parameters

Table 62-34 GETDEFAULTOWNER Procedure Parameters

Parameter	Description
owner	Owner for new elements created (implicitly or explicitly) by the current operation

62.7.17 GETDEFAULTPRINCIPAL Procedure

This procedure returns the "principal" parameter of the default context. This information contained can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.GETDEFAULTPRINCIPAL (
   principal    OUT NOCOPY     VARCHAR2);
```

Parameters

Table 62-35 GETDEFAULTPRINCIPAL Procedure Parameters

Parameter	Description
principal	Agent (principal) invoking the current operation

62.7.18 GETFEATURESBYMOUNT Function

This function returns features of a store by mount point.

Syntax

```
DBMS_DBFS_CONTENT.GETFEATURESBYMOUNT (
   store_mount       IN      VARCHAR2)
  RETURN INTEGER;
```

Parameters

Table 62-36 GETFEATURESBYMOUNT Function Parameters

Parameter	Description
store_mount	Mount point

Return Values

A bit mask of supported features (see FEATURES_T Table Type)

62.7.19 GETFEATURESBYNAME Function

This function returns features of a store by store name.

Syntax

```
DBMS_DBFS_CONTENT.GETFEATURESBYNAME (
   store_name       IN      VARCHAR2)
  RETURN INTEGER;
```

Parameters

Table 62-37 GETFEATURESBYNAME Function Parameters

Parameter	Description
store_name	Name of store

Return Values

A bit mask of supported features (see FEATURES_T Table Type)

62.7.20 GETFEATURESBYPATH Function

This function returns features of a store by path.

Syntax

```
DBMS_DBFS_CONTENT.GETFEATURESBYPATH (
   path       IN      PATH_T)
  RETURN INTEGER;
```

Parameters

Table 62-38 GETFEATURESBYPATH Function Parameters

Parameter	Description
path	PATH_T

Return Values

A bit mask of supported features (see FEATURES_T Table Type)

62.7.21 GETPATH Procedures

This procedure returns existing path items (such as files and directories). This includes both data and metadata (properties).

The client can request (using `prop_flags`) that specific properties be returned. File path names can be read either by specifying a `BLOB` locator using the `prop_data` bitmask in `prop_flags` (see Table 62-9) or by passing one or more `RAW` buffers.

When `forUpdate` is 0, this procedure also accepts a valid `asof` timestamp parameter as part of `ctx` that can be used by stores to implement "as of" style flashback queries. Mutating versions of the `GETPATH` Procedures do not support these modes of operation.

Syntax

```
DBMS_DBFS_CONTENT.GETPATH (
        path        IN              VARCHAR2,
        properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
        content     OUT    NOCOPY   BLOB,
        item_type   OUT             INTEGER,
        prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                         PROP_OPT +
                                                         PROP_DATA),
        asof        IN              TIMESTAMP   DEFAULT NULL,
        forUpdate   IN              BOOLEAN     DEFAULT FALSE,
        deref       IN              BOOLEAN     DEFAULT FALSE,
        store_name  IN              VARCHAR2    DEFAULT NULL,
        principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.GETPATH (
        path        IN              VARCHAR2,
        properties  IN OUT NOCOPY   PROPERTIES_T,
        content     OUT    NOCOPY   BLOB,
        item_type   OUT             INTEGER,
        prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                         PROP_OPT +
                                                         PROP_DATA),
        asof        IN              TIMESTAMP   DEFAULT NULL,
        forUpdate   IN              BOOLEAN     DEFAULT FALSE,
        deref       IN              BOOLEAN     DEFAULT FALSE,
        store_name  IN              VARCHAR2    DEFAULT NULL,
        principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.GETPATH (
        path        IN              VARCHAR2,
        properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
        amount      IN OUT          NUMBER,
        offset      IN              NUMBER,
        buffers     OUT    NOCOPY   RAW,
        prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                         PROP_OPT +
                                                         PROP_DATA),
        asof        IN              TIMESTAMP   DEFAULT NULL,
        store_name  IN              VARCHAR2    DEFAULT NULL,
        principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.GETPATH (
        path        IN              VARCHAR2,
        properties  IN OUT NOCOPY   PROPERTIES_T,
        amount      IN OUT          NUMBER,
        offset      IN              NUMBER,
        buffers     OUT    NOCOPY   RAW,
        prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                         PROP_OPT +
                                                         PROP_DATA),
        asof        IN              TIMESTAMP   DEFAULT NULL,
        store_name  IN              VARCHAR2    DEFAULT NULL,
        principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.GETPATH (
        path        IN              VARCHAR2,
```

```
        properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
        amount      IN OUT          NUMBER,
        offset      IN              NUMBER,
        buffers     OUT    NOCOPY   DBMS_DBFS_CONTENT_RAW_T,
        prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                         PROP_OPT +
                                                         PROP_DATA),
        asof        IN              TIMESTAMP   DEFAULT NULL,
        store_name  IN              VARCHAR2    DEFAULT NULL,
        principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.GETPATH (
        path        IN              VARCHAR2,
        properties  IN OUT NOCOPY   PROPERTIES_T,
        amount      IN OUT          NUMBER,
        offset      IN              NUMBER,
        buffers     OUT    NOCOPY   DBMS_DBFS_CONTENT_RAW_T,
        prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                         PROP_OPT +
                                                         PROP_DATA),
        asof        IN              TIMESTAMP   DEFAULT NULL,
        store_name  IN              VARCHAR2    DEFAULT NULL,
        principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-39 GETPATH Procedure Parameters

Parameter	Description
path	Name of path to path items
properties	One or more properties and their values to be returned depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
content	BLOB holding data which populates the file (optional)
item_type	Type of the path item specified (see Table 62-4)
amount	On input, number of bytes to be read. On output, number of bytes read
offset	Byte offset from which to begin reading
buffer	Buffer to which to write
buffers	Buffers to which to write
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes
forUpdate	Specifies that a lock should be taken to signify exclusive write access to the path item
deref	If nonzero, attempts to resolve the given path item to actual data provided it is a reference
store_name	Name of store
principal	Agent (principal) invoking the current operation

62.7.22 GETPATHBYMOUNTID Function

If the underlying GUID is found in the underlying store, this function returns the full absolute path name.

Syntax

```
DBMS_DBFS_CONTENT.GETPATHBYMOUNTID (
   store_mount          IN      VARCHAR2,
   guid                 IN      INTEGER)
  RETURN VARCHAR2;
```

Parameters

Table 62-40 GETPATHBYMOUNTID Function Parameters

Parameter	Description
`store_mount`	Mount point in which the path item with `guid` resides
`guid`	Unique ID for the path item

Usage Notes

If the GUID is unknown, a `NULL` value is returned. Clients are expected to handle this as appropriate.

Return Values

Path of the path item represented by GUID in `store_mount`

62.7.23 GETPATHBYSTOREID Function

If the underlying GUID is found in the underlying store, this function returns the store-qualified path name.

Syntax

```
DBMS_DBFS_CONTENT.GETPATHBYSTOREID (
   store_name           IN      VARCHAR2,
   guid                 IN      INTEGER)
  RETURN VARCHAR2;
```

Parameters

Table 62-41 GETPATHBYSTOREID Function Parameters

Parameter	Description
`store_name`	Name of store
`guid`	Unique ID representing the desired path item

Usage Notes

If the GUID is unknown, a `NULL` value is returned. Clients are expected to handle this as appropriate.

Return Values

Store-qualified path name represented by the GUID

62.7.24 GETPATHNOWAIT Procedures

This procedure implies that the operation is for an update, and, if implemented, allows providers to return an exception (`ORA-00054`) rather than wait for row locks.

See `FEATURE_NOWAIT` in Table 62-5 for more information.

Syntax

```
DBMS_DBFS_CONTENT.GETPATHNOWAIT (
       path        IN              VARCHAR2,
       properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
       content     OUT    NOCOPY   BLOB,
       item_type   OUT             INTEGER,
       prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                        PROP_OPT +
                                                        PROP_DATA),
       deref       IN              BOOLEAN     DEFAULT FALSE,
       store_name  IN              VARCHAR2    DEFAULT NULL,
       principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.GETPATHNOWAIT (
       path        IN              VARCHAR2,
       properties  IN OUT NOCOPY   PROPERTIES_T,
       content     OUT    NOCOPY   BLOB,
       item_type   OUT             INTEGER,
       prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                        PROP_OPT +
                                                        PROP_DATA),
       deref       IN              BOOLEAN     DEFAULT FALSE,
       store_name  IN              VARCHAR2    DEFAULT NULL,
       principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-42 GETPATHNOWAIT Procedure Parameters

Parameter	Description
`path`	Name of path to path items
`properties`	One or more properties and their values to be returned depending on `prop_flags` (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
`content`	`BLOB` holding data which populates the file (optional)
`item_type`	Type of the path item specified (see Table 62-4)
`prop_flags`	Determines which properties are returned. Default is `PROP_STD`. Specify properties to be returned by setting `prop_spec`, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
`asof`	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes

Table 62-42 (Cont.) GETPATHNOWAIT Procedure Parameters

Parameter	Description
deref	If nonzero, attempts to resolve the given path item to actual data provided it is a reference
store_name	Name of store
principal	Agent (principal) invoking the current operation

62.7.25 GETSTOREBYMOUNT Function

This function returns a store by way of its name.

Syntax

```
DBMS_DBFS_CONTENT.GETSTOREBYMOUNT (
   store_mount       IN      VARCHAR2)
  RETURN STORE_T;
```

Parameters

Table 62-43 GETSTOREBYMOUNT Function Parameters

Parameter	Description
store_mount	Location at which the store instance is mounted

Return Values

STORE_T Record Type

62.7.26 GETSTOREBYNAME Function

This function returns a store by way of its name.

Syntax

```
DBMS_DBFS_CONTENT.GETSTOREBYNAME (
   store_name       IN      VARCHAR2)
  RETURN STORE_T;
```

Parameters

Table 62-44 GETSTOREBYNAME Function Parameters

Parameter	Description
store_name	Name of store

Return Values

STORE_T Record Type

62.7.27 GETSTOREBYPATH Function

This function returns a store by way of its path.

Syntax

```
DBMS_DBFS_CONTENT.GETSTOREBYPATH (
   path       IN      PATH_T)
  RETURN STORE_T;
```

Parameters

Table 62-45 GETSTOREBYPATH Function Parameters

Parameter	Description
path	PATH_T s

Return Values

STORE_T Record Type

62.7.28 GETSTATS Procedure

This procedure returns information about DBMS_DBFS_CONTENT statistics collection.

Syntax

```
DBMS_DBFS_CONTENT.GETSTATS (
   enabled         OUT    BOOLEAN,
   flush_time      OUT    INTEGER,
   flush_count     OUT    INTEGER);
```

Parameters

Table 62-46 GETSTATS Procedure Parameters

Parameter	Description
enabled	Whether statistics collection is enabled
flush_time	How often to flush the statistics to disk in centiseconds
flush_count	Number of operations to allow between statistics flushes

62.7.29 GETTRACE Function

This function returns whether DBMS_DBFS_CONTENT tracing is turned on or not.

Syntax

```
DBMS_DBFS_CONTENT.GETTRACE
  RETURN INTEGER.
```

Return Values

Returns zero if tracing is off, non-zero if tracing is on.

62.7.30 GETVERSION Function

This function marks each version of the `DBMS_DBFS_CONTENT` interface.

Syntax

```
DBMS_DBFS_CONTENT.GETVERSION (
  RETURN VARCHAR2;
```

Return Values

A string enumerating the version of the `DBMS_DBFS_CONTENT` interface in standard naming convention: string: *a.b.c* corresponding to *major*, *minor*, and *patch* components.

62.7.31 ISPATHLOCKED Procedure

This procedure checks if any user-level locks are applied on a given path.

Syntax

```
DBMS_DBFS_CONTENT.ISPATHLOCKED (
   path           IN     VARCHAR2,
   who            IN     VARCHAR2,
   lock_type      IN OUT INTEGER,
   store_name     IN     VARCHAR2    DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-47 ISPATHLOCKED Procedure Parameters

Parameter	Description
`path`	Name of path to file items
`who`	Transaction identifier that has locked the path
`lock_type`	One of the available lock types (see Table 62-6)
`store_name`	Name of store
`principal`	Agent (principal) invoking the current operation

62.7.32 LIST Function

This function lists the path items in the specified path meeting the specified filter and other criteria.

Syntax

```
DBMS_DBFS_CONTENT.LIST (
   path           IN     VARCHAR2,
   filter         IN     VARCHAR2    DEFAULT NULL,
   recurse        IN     INTEGER     DEFAULT 0,
   asof           IN     TIMESTAMP   DEFAULT NULL,
   store_name     IN     VARCHAR2    DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL)
  RETURN  DBMS_DBFS_CONTENT_LIST_ITEMS_T PIPELINED;
```

Parameters

Table 62-48 LIST Function Parameters

Parameter	Description
path	Name of path to directories
filter	A filter, if any, to be applied
recurse	If 0, do not execute recursively. Otherwise, recursively list the contents of directories and files below the given directory.
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes
store_name	Name of repository
principal	Agent (principal) invoking the current operation

Return Values

DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type

62.7.33 LISTALLPROPERTIES Function

This function returns a table of all properties for all path items in all mounts.

Syntax

```
DBMS_DBFS_CONTENT.LISTALLPROPERTIES
  RETURN  PROP_ITEMS_T PIPELINED;
```

Return Values

PROP_ITEMS_T Table Type

62.7.34 LISTALLCONTENT Function

This function lists all path items in all mounts.

Syntax

```
DBMS_DBFS_CONTENT.LISTALLCONTENT
  RETURN  PATH_ITEMS_T PIPELINED;
```

Return Values

PATH_ITEMS_T Table Type

62.7.35 LISTMOUNTS Function

This function lists all available mount points, their backing stores, and the store features.

Syntax

```
DBMS_DBFS_CONTENT.LISTMOUNTS
  RETURN MOUNTS_T PIPELINED;
```

Return Values

MOUNTS_T Table Type

Usage Notes

A single mount results in a single returned row, with its `store_mount` field of the returned records set to `NULL`.

62.7.36 LISTSTORES Function

This function lists all available stores and their features.

Syntax

```
DBMS_DBFS_CONTENT.LISTSTORES
  RETURN STORES_T PIPELINED;
```

Return Values

STORES_T Table Type

Usage Notes

The `store_mount` field of the returned records is set to `NULL` (since mount-points are separate from stores themselves).

62.7.37 LOCKPATH Procedure

This procedure applies user-level locks to the given valid path name (subject to store feature support), and optionally associates user-data with the lock.

Syntax

```
DBMS_DBFS_CONTENT.LOCKPATH (
   path           IN     VARCHAR2,
   who            IN     VARCHAR2,
   lock_type      IN     INTEGER,
   waitForRowLock IN     INTEGER     DEFAULT 1,
   store_name     IN     VARCHAR2    DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-49 LOCKPATH Procedure Parameters

Parameter	Description
`path`	Name of path to file items
`who`	Transaction identifier that has locked the path
`lock_type`	One of the available lock types (see Table 62-6)
`waitForRowLock`	Determines if a row is locked by a transaction or not
`store_name`	Name of store
`principal`	Agent (principal) invoking the current operation

62.7.38 MOUNTSTORE Procedure

This procedure mounts a previously registered store and binds it to the mount point.

Syntax

```
DBMS_DBFS_CONTENT.MOUNTSTORE (
   store_mount     in      VARCHAR2   DEFAULT NULL,
   singleton       in      BOOLEAN    DEFAULT FALSE,
   principal       in      VARCHAR2   DEFAULT NULL,
   owner           in      VARCHAR2   DEFAULT NULL,
   acl             in      VARCHAR2   DEFAULT NULL,
   asof            in      TIMESTAMP  DEFAULT NULL,
   read_only       in      BOOLEAN    DEFAULT FALSE);
```

Parameters

Table 62-50 MOUNTSTORE Procedure Parameters

Parameter	Description
store_mount	Path name to use to mount this store
singleton	Whether the mount is a single backend store on the system
principal	Agent (principal) invoking the current operation
owner	Owner for new elements created (implicitly or explicitly) by the current operation
acl	ACL for all new elements created (implicitly or explicitly) by the current operation
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes
read_only	Whether the mount is read-only

Usage Notes

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for information on mounting a registered store

62.7.39 NORMALIZEPATH Functions

This function converts a store-specific or full-absolute path name into normalized form.

It does the following:

- verifies that the path name is absolute, and so starts with "/"
- collapses multiple consecutive "/" into a single "/"
- strips trailing "/"
- breaks up a store-specific normalized path name into 2 components - parent pathname, trailing component name

- breaks up a full-absolute normalized path name into 3 components - store name, parent pathname, trailing component name

Syntax

```
DBMS_DBFS_CONTENT.NORMALIZEPATH (
   path        IN              VARCHAR2,
   parent      OUT NOCOPY      VARCHAR2,
   tpath       OUT NOCOPY      VARCHAR2)
  RETURN VARCHAR2;

DBMS_DBFS_CONTENT.NORMALIZEPATH (
   path        IN              VARCHAR2,
   store_name  OUT NOCOPY      VARCHAR2,
   parent      OUT NOCOPY      VARCHAR2,
   tpath       OUT NOCOPY      VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 62-51 NORMALIZEPATH Function Parameters

Parameter	Description
path	Name of path to file items
store_name	Name of store
parent	Parent path name
tpath	Name of trailing path item

Return Values

The completely normalized store-specific or full-absolute path name

62.7.40 PROPANY Functions

This function provides constructors that take one of a variety of types and return a PROPERTY_T.

Syntax

```
DBMS_DBFS_CONTENT.PROPANY (
   val      IN      NUMBER)
RETURN PROPERTY_T;

DBMS_DBFS_CONTENT.PROPANY (
   val      IN      VARCHAR2)
RETURN PROPERTY_T;

DBMS_DBFS_CONTENT.PROPANY (
   val      IN      TIMESTAMP)
RETURN PROPERTY_T;

DBMS_DBFS_CONTENT.PROPANY (
   val      IN      RAW)
RETURN PROPERTY_T;
```

Parameters

Table 62-52 PROPANY Function Parameters

Parameter	Description
val	Value

Return Values

PROPERTY_T Record Type

62.7.41 PROPERTIESH2T Function

This function converts a PROPERTY_T hash to a DBMS_DBFS_CONTENT_PROPERTIES_T table.

Syntax

```
DBMS_DBFS_CONTENT.PROPERTIEST2H (
   pprops      IN      PROPERTIES_T)
RETURN DBMS_DBFS_CONTENT_PROPERTIES_T;
```

Parameters

Table 62-53 PROPERTIEST2H Function Parameters

Parameter	Description
sprops	A PROPERTIES_T hash

Return Values

DBMS_DBFS_CONTENT_PROPERTIES_T Table Type

62.7.42 PROPERTIEST2H Function

This function converts a DBMS_DBFS_CONTENT_PROPERTIES_T table to a PROPERTY_T hash.

Syntax

```
DBMS_DBFS_CONTENT.PROPERTIEST2H (
   sprops      IN      DBMS_DBFS_CONTENT_PROPERTIES_T)
RETURN properties_t;
```

Parameters

Table 62-54 PROPERTIEST2H Function Parameters

Parameter	Description
sprops	A DBMS_DBFS_CONTENT_PROPERTIES_T table

Return Values

PROPERTIES_T Table Type

62.7.43 PROPNUMBER Function

This function is a constructor that takes a number and returns a `PROPERTY_T`.

Syntax

```
DBMS_DBFS_CONTENT.PROPNUMBER (
   val      IN      NUMBER)
RETURN PROPERTY_T;
```

Parameters

Table 62-55 PROPNUMBER Function Parameters

Parameter	Description
`val`	Value

Return Values

PROPERTY_T Record Type

62.7.44 PROPRAW Function

This function is a constructor that takes a `RAW` and returns a `PROPERTY_T`.

Syntax

```
DBMS_DBFS_CONTENT.PROPRAW (
   val      IN      RAW)
RETURN PROPERTY_T;
```

Parameters

Table 62-56 PROPRAW Function Parameters

Parameter	Description
`val`	Value

Return Values

PROPERTY_T Record Type

62.7.45 PROPTIMESTAMP Function

This function is a constructor that takes a `TIMESTAMP` and returns a `PROPERTY_T`.

Syntax

```
DBMS_DBFS_CONTENT.PROPTIMESTAMP (
   val      IN      TIMESTAMP)
RETURN PROPERTY_T;
```

Parameters

Table 62-57 PROPTIMESTAMP Function Parameters

Parameter	Description
val	Value

Return Values

PROPERTY_T Record Type

62.7.46 PROPVARCHAR2 Function

This function is a constructor that takes a VARCHAR2 and returns a PROPERTY_T.

Syntax

```
DBMS_DBFS_CONTENT.PROPVARCHAR2 (
   val      IN      VARCHAR2)
RETURN PROPERTY_T;
```

Parameters

Table 62-58 PROPNUMBER Function Parameters

Parameter	Description
val	Value

Return Values

PROPERTY_T Record Type

62.7.47 PURGEALL Procedure

This procedure purges all soft-deleted entries matching the path and optional filter criteria.

Syntax

```
DBMS_DBFS_CONTENT.PURGEALL (
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2    DEFAULT NULL,
   store_name     IN      VARCHAR2    DEFAULT NULL,
   principal      IN      VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-59 PURGEALL Procedure Parameters

Parameter	Description
path	Name of path to file items
filter	A filter, if any, to be applied based on specified criteria
store_name	Name of store

Table 62-59 (Cont.) PURGEALL Procedure Parameters

Parameter	Description
principal	Agent (principal) invoking the current operation

62.7.48 PURGEPATH Procedure

This procedure purges any soft-deleted versions of the given path item.

Syntax

```
DBMS_DBFS_CONTENT.PURGEPATH (
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2    DEFAULT NULL,
   store_name     IN      VARCHAR2    DEFAULT NULL,
   principal      IN      VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-60 PURGEPATH Procedure Parameters

Parameter	Description
path	Name of path to file items
filter	A filter, if any, to be applied
store_name	Name of store
principal	Agent (principal) invoking the current operation

62.7.49 PUTPATH Procedures

This procedure creates a new path item.

Syntax

```
DBMS_DBFS_CONTENT.PUTPATH (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   content     IN OUT NOCOPY   BLOB,
   item_type   OUT             INTEGER,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                    PROP_OPT +
                                                    PROP_DATA),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.PUTPATH (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   content     IN OUT NOCOPY   BLOB,
   item_type   OUT             INTEGER,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                    PROP_OPT +
                                                    PROP_DATA),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

```
DBMS_DBFS_CONTENT.PUTPATH (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   amount      IN              NUMBER,
   offset      IN              NUMBER,
   buffer      IN              RAW,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                    PROP_OPT),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.PUTPATH (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   amount      IN              NUMBER,
   offset      IN              NUMBER,
   buffer      IN              RAW,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                    PROP_OPT),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.PUTPATH (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   written     OUT             NUMBER,
   offset      IN              NUMBER,
   buffers     IN             DBMS_DBFS_CONTENT_RAW_T,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                    PROP_OPT),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.PUTPATH (
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   written     OUT             NUMBER,
   offset      IN              NUMBER,
   buffers     IN              DBMS_DBFS_CONTENT_RAW_T,
   prop_flags  IN              INTEGER     DEFAULT (PROP_STD +
                                                    PROP_OPT),
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-61 PUTPATH Procedure Parameters

Parameter	Description
`path`	Name of path to file items
`properties`	One or more properties and their values to be set depending on `prop_flags` (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
`content`	`BLOB` holding data which populates the file (optional)
`item_type`	Type of the path item specified (see Table 62-4)
`amount`	Number of bytes to be read

Table 62-61 (Cont.) PUTPATH Procedure Parameters

Parameter	Description
offset	Byte offset from which to begin reading
buffer	Buffer to which to write
buffers	Buffers to which to write
prop_flags	Determines which properties are set. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
store_name	Name of store
principal	Agent (principal) invoking the current operation

62.7.50 REGISTERSTORE Procedure

This procedure registers a new store backed by a provider that uses a store provider (conforming to the DBMS_DBFS_CONTENT_SPI package signature).

This method is to be used primarily by store providers after they have created a new store.

Syntax

```
DBMS_DBFS_CONTENT.REGISTERSTORE (
   store_name          IN      VARCHAR2,
   provider_name       IN      VARCHAR2,
   provider_package    IN      VARCHAR2);
```

Parameters

Table 62-62 REGISTERSTORE Procedure Parameters

Parameter	Description
store_name	Name of store, must be unique
provider_name	Name of provider
provider_package	Store provider

62.7.51 RENAMEPATH Procedures

This procedure renames or moves a path. This operation can be performed across directory hierarchies and mount-points as long as it is within the same store.

> **Note:**
>
> See *Oracle Database SecureFiles and Large Objects Developer's Guide* for Rename and Move operations

Syntax

```
DBMS_DBFS_CONTENT.RENAMEPATH (
   oldPath     IN              VARCHAR2,
   newPath     IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.RENAMEPATH (
   oldPath     IN              VARCHAR2,
   newPath     IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   store_name  IN              VARCHAR2    DEFAULT NULL,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-63 RENAMEPATH Procedure Parameters

Parameter	Description
oldPath	Name of path prior to renaming
newPath	Name of path after renaming
properties	One or more properties and their values to be set depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
store_name	Name of store, must be unique
principal	Agent (principal) invoking the current operation

62.7.52 RESTOREALL Procedure

This procedure restores all soft-deleted path items meeting the path and optional filter criteria.

Syntax

```
DBMS_DBFS_CONTENT.RESTOREALL (
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2    DEFAULT NULL,
   store_name     IN      VARCHAR2    DEFAULT NULL,
   principal      IN      VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-64 RESTOREALL Procedure Parameters

Parameter	Description
path	Name of path to file items
filter	A filter, if any, to be applied
store_name	Name of store
principal	Agent (principal) invoking the current operation

62.7.53 RESTOREPATH Procedure

This procedure restores all soft-deleted path items that match the given path and optional filter criteria.

Syntax

```
DBMS_DBFS_CONTENT.RESTOREPATH (
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2    DEFAULT NULL,
   store_name     IN      VARCHAR2    DEFAULT NULL,
   principal      IN      VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-65 RESTOREPATH Procedure Parameters

Parameter	Description
`path`	Name of path to file items
`filter`	A filter, if any, to be applied
`store_name`	Name of store
`principal`	Agent (principal) invoking the current operation

62.7.54 SETDEFAULTACL Procedure

This procedure sets the ACL parameter of the default context.

This information can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.SETDEFAULTACL (
   acl    IN     VARCHAR2);
```

Parameters

Table 62-66 SETDEFAULTACL Procedure Parameters

Parameter	Description
`acl`	ACL for all new elements created (implicitly or explicitly) by the current operation

Usage Notes

- `NULL` by default, this parameter be can be cleared by setting it to `NULL`.
- The parameters, once set, remain as a default for the duration of the session, and is inherited by all operations for which the default is not explicitly overridden.

62.7.55 SETDEFAULTASOF Procedure

This procedure sets the "as of" parameter of the default context. This information can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.SETDEFAULTASOF (
  asof     IN     TIMESTAMP);
```

Parameters

Table 62-67 SETDEFAULTASOF Procedure Parameters

Parameter	Description
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes

Usage Notes

- NULL by default, this parameter be can be cleared by setting it to NULL.
- The parameters, once set, remain as a default for the duration of the session, and is inherited by all operations for which the default is not explicitly overridden.

62.7.56 SETDEFAULTCONTEXT Procedure

This procedure sets the default context. The information contained in the context can be inserted explicitly by way of arguments to the various method calls, allowing for fine-grained control over individual operations.

Syntax

```
DBMS_DBFS_CONTENT.SETDEFAULTCONTEXT (
   principal     IN     VARCHAR2,
   owner         IN     VARCHAR2,
   acl           IN     VARCHAR2,
   asof          IN     TIMESTAMP);
```

Parameters

Table 62-68 SETDEFAULTCONTEXT Procedure Parameters

Parameter	Description
principal	Agent (principal) invoking the current operation
owner	Owner for new elements created (implicitly or explicitly) by the current operation
acl	ACL for all new elements created (implicitly or explicitly) by the current operation
asof	The "as of" timestamp at which the underlying read-only operation (or its read-only sub-components) executes

Usage Notes

- All of the context parameters are `NULL` by default, and be can be cleared by setting them to `NULL`.
- The context parameters, once set, remain as defaults for the duration of the session, and are inherited by all operations for which the defaults are not explicitly overridden.

62.7.57 SETDEFAULTOWNER Procedure

This procedure sets the "owner" parameter of the default context. This information can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.SETDEFAULTOWNER (
   principal    IN     VARCHAR2);
```

Parameters

Table 62-69 SETDEFAULTOWNER Procedure Parameters

Parameter	Description
`owner`	Owner for new elements created (implicitly or explicitly) by the current operation

Usage Notes

- `NULL` by default, this parameter be can be cleared by setting it to `NULL`.
- The parameters, once set, remain as a default for the duration of the session, and is inherited by all operations for which the default is not explicitly overridden.

62.7.58 SETDEFAULTPRINCIPAL Procedure

This procedure sets the "principal" parameter of the default context. This information contained can be inserted explicitly by way of argument into other method calls, allowing for a more fine-grained control.

Syntax

```
DBMS_DBFS_CONTENT.SETDEFAULTPRINCIPAL (
   principal    IN     VARCHAR2);
```

Parameters

Table 62-70 SETDEFAULTPRINCIPAL Procedure Parameters

Parameter	Description
`principal`	Agent (principal) invoking the current operation

Usage Notes

- `NULL` by default, this parameter be can be cleared by setting it to `NULL`.

- The parameters, once set, remain as a default for the duration of the session, and is inherited by all operations for which the default is not explicitly overridden.

62.7.59 SETPATH Procedures

This procedure assigns a path name to a path item represented by contentID.

Stores and their providers that support contentID-based access and lazy path name binding also support the `SETPATH Procedure` that associates an existing `contentID` with a new path.

> **Note:**
>
> See *Oracle Database SecureFiles and Large Objects Developer's Guide* for Rename and Move operations

Syntax

```
DBMS_DBFS_CONTENT.SETPATH (
   store_name  IN              VARCHAR2,
   contentID   IN              RAW,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   principal   IN              VARCHAR2    DEFAULT NULL);

DBMS_DBFS_CONTENT.SETPATH (
   store_name  IN              VARCHAR2,
   contentID   IN              RAW,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   PROPERTIES_T,
   principal   IN              VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-71 SETPATH Procedure Parameters

Parameter	Description
`store_name`	Name of the store
`contentID`	Unique identifier for the item to be associated
`path`	Name of path to path item
`properties`	One or more properties and their values to be set depending on `prop_flags` (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
`principal`	Agent (principal) invoking the current operation

62.7.60 SETSTATS Procedure

This procedure enables and disables statistics collection.

The client can optionally control the flush settings by specifying non-`NULL` values for the time, count or both parameters.

Syntax

```
DBMS_DBFS_CONTENT.SETSTATS (
   enable          IN    BOOLEAN,
   flush_time      IN    INTEGER,
   flush_count     IN    INTEGER);
```

Parameters

Table 62-72 SETSTATS Procedure Parameters

Parameter	Description
`enable`	If `TRUE`, enable statistics collection. If `FALSE`, disable statistics collection.
`flush_time`	How often to flush the statistics to disk in centiseconds
`flush_count`	Number of operations to allow between statistics flushes

Usage Notes

The `SETSTATS Procedure` buffers statistics in-memory for a maximum of `flush_time` centiseconds or a maximum of `flush_count` operations (whichever limit is reached first), or both, at which time the buffers are implicitly flushed to disk.

62.7.61 SETTRACE Procedure

This procedure sets the `DBMS_DBFS_CONTENT` tracing severity to the given level, `0` being "off".

Syntax

```
DBMS_DBFS_CONTENT.SETTRACE
     trclvl      IN         INTEGER);
```

Parameters

Table 62-73 SETTRACE Procedure Parameters

Parameter	Description
`trclvl`	Level of the tracing, higher values implying more tracing

62.7.62 SPACEUSAGE Procedure

This procedure queries file system space usage statistics.

Providers are expected to support this subprogram for their stores (and to make a best effort determination of space usage, especially if the store consists of multiple tables, indexes, LOBs, and so on).

Syntax

```
DBMS_DBFS_CONTENT.SPACEUSAGE (
   path          IN        VARCHAR2,
   blksize       OUT       INTEGER,
   tbytes        OUT       INTEGER,
```

```
   fbytes        OUT       INTEGER,
   nfile         OUT       INTEGER,
   ndir          OUT       INTEGER,
   nlink         OUT       INTEGER,
   nref          OUT       INTEGER,
   store_name    IN        VARCHAR2 DEFAULT NULL);
```

Parameters

Table 62-74 SPACEUSAGE Procedure Parameters

Parameter	Description
path	Name of path to file items
blksize	Natural tablespace blocksize that holds the store. If multiple tablespaces with different blocksizes are used, any valid blocksize is acceptable.
tbytes	Total size of the store in bytes computed over all segments that comprise the store
fbytes	Free or unused size of the store in bytes computed over all segments that comprise the store
nfile	Number of currently available files in the store
ndir	Number of currently available directories in the store
nlink	Number of currently available links in the store
nref	Number of currently available references in the store
store_name	Name of store

Usage Notes

- A space usage query on the top-level root directory returns a combined summary of the space usage of all available distinct stores under it (if the same store is mounted multiple times, is still counted only once).
- Since database objects are dynamically expandable, it is not easy to estimate the division between "free" space and "used" space.

62.7.63 TRACE Procedure

This procedure outputs tracing to the current foreground trace file.

Syntax

```
DBMS_DBFS_CONTENT.TRACE
   sev           IN              INTEGER,
   msg0          IN              VARCHAR2,
   msg1          IN              VARCHAR     DEFAULT '',
   msg2          IN              VARCHAR     DEFAULT '',
   msg3          IN              VARCHAR     DEFAULT '',
   msg4          IN              VARCHAR     DEFAULT '',
   msg5          IN              VARCHAR     DEFAULT '',
   msg6          IN              VARCHAR     DEFAULT '',
   msg7          IN              VARCHAR     DEFAULT '',
   msg8          IN              VARCHAR     DEFAULT '',
   msg9          IN              VARCHAR     DEFAULT '',
   msg10         IN              VARCHAR     DEFAULT '');
```

Parameters

Table 62-75 TRACE Procedure Parameters

Parameter	Description
sev	Severity at which trace message is output
msg*	One or more message strings to be output. If more than one message is specified, all are output.

Usage Notes

- Trace information is written to the foreground trace file, with varying levels of detail as specified by the trace level arguments.
- The global trace level consists of 2 components: "severity" and "detail". These can be thought of as additive bitmasks.

  The "severity" allows the separation of top level as compared to low-level tracing of different components, and allows the amount of tracing to be increased as needed. There are no semantics associated with different levels, and users are free to set trace at any severity they choose, although a good rule of thumb would use severity "1" for top level API entry and exit traces, "2" for internal operations, and "3" or greater for very low-level traces.

  The "detail" controls how much additional information: timestamps, short-stack, etc. is dumped along with each trace record.

62.7.64 TRACEENABLED Function

This function determines if the current trace "severity" set by the SETTRACE Procedure is at least as high as the given trace level.

Syntax

```
DBMS_DBFS_CONTENT.TRACEENABLED(
   sev          IN               INTEGER)
  RETURN  INTEGER;
```

Parameters

Table 62-76 TRACEENABLED Procedure Parameters

Parameter	Description
sev	Severity at which trace message is output

Return Values

Returns 0 if the requested severity level is lower than the currently set trace severity level; 1 otherwise.

Related Topics

- SETTRACE Procedure
  This procedure sets the DBMS_DBFS_CONTENT tracing severity to the given level, 0 being "off".

62.7.65 UNLOCKPATH Procedure

This procedure unlocks path items that were previously locked with the LOCKPATH Procedure.

Syntax

```
DBMS_DBFS_CONTENT.UNLOCKPATH (
   path           IN     VARCHAR2,
   who            IN     VARCHAR2,
   waitForRowLock IN     INTEGER     DEFAULT 1,
   store_name     IN     VARCHAR2    DEFAULT NULL,
   principal      IN     VARCHAR2    DEFAULT NULL);
```

Parameters

Table 62-77 UNLOCKPATH Procedure Parameters

Parameter	Description
`path`	Name of path to file items
`who`	Transaction identifier that has locked the path
`waitForRowLock`	Determines if a row is locked by a transaction or not
`store_name`	Name of store
`principal`	Agent (principal) invoking the current operation

Related Topics

- LOCKPATH Procedure
  This procedure applies user-level locks to the given valid path name (subject to store feature support), and optionally associates user-data with the lock.

62.7.66 UNMOUNTSTORE Procedure

This procedure unmounts a registered store, either by name or by mount point.

Syntax

```
DBMS_DBFS_CONTENT.UNMOUNTSTORE (
   store_name       IN      VARCHAR2   DEFAULT NULL,
   store_mount      IN      VARCHAR2   DEFAULT NULL,
   ignore_unknown   IN      BOOLEAN    DEFAULT FALSE);
```

Parameters

Table 62-78 UNMOUNTSTORE Procedure Parameters

Parameter	Description
`store_name`	Name of store
`store_mount`	Location at which the store instance is mounted
`ignore_unknown`	If `TRUE`, attempts to unregister unknown stores will not raise an exception.

Usage Notes

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on unmounting a previously unmounted store

62.7.67 UNREGISTERSTORE Procedure

This procedure unregisters a previously registered store (invalidating all mount points associated with it).

Syntax

```
DBMS_DBFS_CONTENT.UNREGISTERSTORE (
   store_name          IN      VARCHAR2,
   ignore_unknown      IN      BOOLEAN DEFAULT FALSE);
```

Parameters

Table 62-79 UNREGISTERSTORE Procedure Parameters

Parameter	Description
`store_name`	Name of store
`ignore_unknown`	If `TRUE`, attempts to unregister unknown stores will not raise an exception.

Usage Notes

- Once unregistered all access to the store (and its mount points) are not guaranteed to work
- If the `ignore_unknown` argument is `TRUE`, attempts to unregister unknown stores do not raise an exception.

63
DBMS_DBFS_CONTENT_SPI

The DBMS_DBFS_CONTENT_SPI package is a specification for DBMS_DBFS_CONTENT store providers, which must be implemented. Application designers can create PL/SQL packages conforming to this specification to extend DBMS_DBFS_CONTENT to use custom store providers.

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Summary of DBMS_DBFS_CONTENT_SPI Subprograms

Related Topics

- DBMS_DBFS_CONTENT
  The DBMS_DBFS_CONTENT package provides an interface comprising a file system-like abstraction backed by one or more Store Providers.

> **See Also:**
>
> - *Oracle Database SecureFiles and Large Objects Developer's Guide*

63.1 ODBMS_DBFS_CONTENT_SPI Overview

The DBMS_DBFS_CONTENT_SPI package describes an internal contract between the implementation of the DBMS_DBFS_CONTENT interface and individual store providers, and whichever package contains their code.

Since PL/SQL does not allow a compile-time, declarative type-conformation between package signatures, store providers should informally conform to the SPI, which is to say, they should implement the SPI by means of a package that contains all of the methods specified in package DBMS_DBFS_CONTENT_SPI, with the same method signatures and semantics.

Obviously, these provider packages can implement other methods and expose other interfaces, however, these interfaces are not to be used by the DBMS_DBFS_CONTENT interface itself.

Since the provider SPI is merely a contract specification, there is no package body for DBMS_DBFS_CONTENT_SPI, and it is not possible to actually invoke any methods using this package.

The SPI references various elements (constants, types, exceptions) defined by the DBMS_DBFS_CONTENT interface.

Additionally, there is an almost one-to-one correspondence between the client API exported by the DBMS_DBFS_CONTENT interface and the provider interface that the DBMS_DBFS_CONTENT interface itself expects to work against.

The main distinction in the method naming conventions is that all path name references are always store-qualified. That is, the notion of mount-points and full-absolute path names have been normalized and converted to store-qualified path names by the DBMS_DBFS_CONTENT interface before it invokes any of the provider SPI methods.

Since the interconnection of the DBMS_DBFS_CONTENT interface and the provider SPI is a 1-to-many pluggable architecture, and the interface uses dynamic SQL to invoke methods in the provider SPI, this can lead to runtime errors.

Related Topics

- DBMS_DBFS_CONTENT
  The DBMS_DBFS_CONTENT package provides an interface comprising a file system-like abstraction backed by one or more Store Providers.

63.2 DBMS_DBFS_CONTENT_SPI Security Model

Implementations of the DBMS_DBFS_CONTENT_SPI package should be created as AUTHID CURRENT_USER.

63.3 DBMS_DBFS_CONTENT_SPI Operational Notes

This topic lists operational notes for DBMS_DBFS_CONTENT_SPI implementation, path names, and other operations.

- Implementation
- Path Names
- Other DBMS_DBFS_CONTENT Operations

Implementation

Since the interconnection of the DBMS_DBFS_CONTENT interface and the provider SPI is a 1-to-many pluggable architecture, the interface uses dynamic SQL to invoke methods in the provider SPI, this can lead to runtime errors.

There are no explicit INIT or FINI methods to indicate when the DBMS_DBFS_CONTENT interface plugs or unplugs a particular provider SPI. Provider SPIs must be willing to auto-initialize themselves at any SPI entry-point.

All operations performed by a store provider are "stateless" in that they are complete operations unto themselves. If state is necessary to be maintained for some reason, then the state must be maintained in data structures such as auxiliary tables that can be queried as needed.

Path Names

All path names used in the provider SPI are store-qualified in pair form (store_name, pathname) where the path name is rooted within the store namespace.

Stores and their providers that support contentID-based access (see FEATURE_CONTENT_ID in Table 62-5) also support a form of addressing that is not based on path names. Content items

are identified by an explicit store name, a `NULL` path name, and possibly a contentID specified as a parameter or by way of the `OPT_CONTENT_ID` (seeTable 62-8) property.

Not all operations are supported with contentID-based access, and applications should depend only on the simplest create or delete functionality being available.

Other DBMS_DBFS_CONTENT Operations

This table lists other operations and provides links to related discussions.

Table 63-1 Other DBMS_DBFS_CONTENT Operations

Other Operations	See ...
Creation	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on creation operations
Deletion	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on deletion operations
Get (Retrieve) and Put (Insert)	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Get and Put operations
Rename and Move	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Rename and Move operations
Directory Navigation and Search	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Navigation and Search operations
Locking	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Locking operations
Access Check	*Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Access Check operations

63.4 Summary of DBMS_DBFS_CONTENT_SPI Subprograms

This table lists and describes the subprograms used in the `DBMS_DBFS_CONTENT_SPI` Package.

Table 63-2 DBMS_DBFS_CONTENT_SPI Package Subprograms

Subprogram	Description
CHECKACCESS Function	Reports if the user (`principal`) can perform the specified operation on the given path
CREATEDIRECTORY Procedure	Creates a directory
CREATEFILE Procedure	Creates a file
CREATELINK Procedure	Creates a physical link to an already existing file system element
CREATEREFERENCE Procedure	Creates a new reference to the source file system element
DELETECONTENT Procedure	Deletes the file specified by the given contentID
DELETEDIRECTORY Procedure	Deletes a directory
DELETEFILE Procedure	Deletes a file
GETFEATURES Function	Returns the features of a store
GETPATH Procedures	Returns existing path items (such as files and directories)
GETPATHBYSTOREID Function	If the underlying GUID is found in the underlying store, returns the store-qualified path name

Table 63-2 (Cont.) DBMS_DBFS_CONTENT_SPI Package Subprograms

Subprogram	Description
GETPATHNOWAIT Procedure	Implies that the operation is for an update, and, if implemented, allows providers to return an exception (ORA-00054) rather than wait for row locks.
GETSTOREID Function	Returns the ID of a store
GETVERSION Function	Returns the version associated with a store
LIST Function	Lists the contents of a directory path name
LOCKPATH Procedure	Applies user-level locks to the given valid path name
PURGEALL Procedure	Purges all soft-deleted entries matching the path and optional filter criteria
PURGEPATH Procedure	Purges any soft-deleted versions of the given path item
PUTPATH Procedures	Creates a new path item
RENAMEPATH Procedure	Renames or moves a path
RESTOREALL Procedure	Restores all soft-deleted path items meeting the path and filter criteria
RESTOREPATH Procedure	Restores all soft-deleted path items that match the given path and filter criteria
SEARCH Function	Searches for path items matching the given path and filter criteria
SETPATH Procedure	Assigns a path name to a path item represented by contentID
SPACEUSAGE Procedure	Queries file system space usage statistics
UNLOCKPATH Procedure	Unlocks path items that were previously locked with the LOCKPATH Procedure

63.4.1 CHECKACCESS Function

This function reports if the user (principal) can perform the specified operation on the given path. This enables verifying the validity of an operation without attempting to perform the operation. If CHECKACCESS returns 0, then the subprogram invoked to implement that operation should fail with an error.

Syntax

```
DBMS_DBFS_CONTENT_SPI.CHECKACCESS (
   store_name      IN     VARCHAR2  DEFAULT NULL,
   path            IN     VARCHAR2,
   pathtype        IN     INTEGER,
   operation       IN     VARCHAR2,
   principal       IN     VARCHAR2)
  RETURN  INTEGER;
```

Parameters

Table 63-3 CHECKACCESS Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to check for access

Table 63-3 (Cont.) CHECKACCESS Procedure Parameters

Parameter	Description
pathtype	Type of object path represents (see Table 62-4)
operation	Operation to be checked (see Table 62-8)
principal	File system user for whom the access check is made

Usage Notes

Whether or not the user invokes this function, a store that supports access control internally performs these checks to guarantee security.

63.4.2 CREATEDIRECTORY Procedure

This procedure creates a directory.

Syntax

```
DBMS_DBFS_CONTENT_SPI.CREATEDIRECTORY (
   store_name  IN              VARCHAR2,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   prop_flags  IN              INTEGER,
   recurse     IN              INTEGER,
   ctx         IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-4 CREATEDIRECTORY Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to the directory
properties	One or more properties and their values to be set, returned, or both, depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting PROP_SPC (see Table 62-9), and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
recurse	If 0, do not execute recursively; otherwise, recursively create the directories above the given directory
ctx	Context with which to create the directory (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.3 CREATEFILE Procedure

This procedure creates a file.

Syntax

```
DBMS_DBFS_CONTENT_SPI.CREATEFILE (
   store_name     IN              VARCHAR2,
   path           IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   content        IN OUT NOCOPY   BLOB,
   prop_flags     IN              INTEGER,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-5 CREATEFILE Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to the file
properties	One or more properties and their values to be set, returned or both depending, or both on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
content	BLOB holding data with which to populate the file (optional)
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
ctx	Context with which to create the file (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.4 CREATELINK Procedure

This procedure creates a physical link to an already existing file system element (such as file or directory). The resulting entry shares the same metadata structures as the value of the srcPath parameter, and so is similar to incrementing a reference count on the file system element. This is analogous to a UNIX file system hard link.

Syntax

```
DBMS_DBFS_CONTENT_SPI.CREATELINK (
   store_name     IN              VARCHAR2,
   srcPath        IN              VARCHAR2,
   dstPath        IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   prop_flags     IN              INTEGER,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-6 CREATELINK Procedure Parameters

Parameter	Description
store_name	Name of store
srcPath	File system entry with which to link
dstPath	Path of the new link element to be created
properties	One or more properties and their values to be set, returned, or both, depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
ctx	Context with which to create the link (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.5 CREATEREFERENCE Procedure

This procedure creates a new reference to the source file system element (such as a file, or directory). The resulting reference points to the source element but does not directly share metadata with the source element. This is analogous to a UNIX file system symbolic link.

Syntax

```
DBMS_DBFS_CONTENT_SPI.CREATEREFERENCE (
   srcPath       IN              VARCHAR2,
   dstPath       IN              VARCHAR2,
   properties    IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   prop_flags    IN              INTEGER,
   store_name    IN              VARCHAR2,
   ctx           IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-7 CREATEREFERENCE Procedure Parameters

Parameter	Description
store_name	Name of store
srcPath	File system entry with which to link
dstPath	Path of the new link element to be created
properties	One or more properties and their values to be set, returned, or both, depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.

Table 63-7 (Cont.) CREATEREFERENCE Procedure Parameters

Parameter	Description
ctx	Context with which to create the reference (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.6 DELETECONTENT Procedure

This procedure deletes the file specified by the given contentID.

Syntax

```
DBMS_DBFS_CONTENT_SPI.DELETECONTENT (
   store_name     IN     VARCHAR2,
   contentID      IN     RAW,
   filter         IN     VARCHAR2,
   soft_delete    IN     INTEGER,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-8 DELETECONTENT Procedure Parameters

Parameter	Description
store_name	Name of store
contentID	Unique identifier for the file to be deleted
filter	A filter, if any, to be applied
soft_delete	If 0, execute a hard (permanent) delete. For any value other than 0, perform a soft delete see *Oracle Database SecureFiles and Large Objects Developer's Guide*, Deletion Operations).
ctx	Context with which to delete the file (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.7 DELETEDIRECTORY Procedure

This procedure deletes a directory.

If recurse is nonzero, it recursively deletes all elements of the directory. A filter, if supplied, determines which elements of the directory are deleted.

Syntax

```
DBMS_DBFS_CONTENT_SPI.DELETEDIRECTORY (
   store_name     IN     VARCHAR2,
   path           IN     VARCHAR2,
   filter         IN     VARCHAR2,
   soft_delete    IN     INTEGER,
   recurse        IN     INTEGER,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-9 DELETEDIRECTORY Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to the directory
filter	A filter, if any, to be applied
soft_delete	If 0, execute a hard (permanent) delete. For any value other than 0, perform a soft delete see *Oracle Database SecureFiles and Large Objects Developer's Guide*, Deletion Operations).
recurse	If 0, do not execute recursively. Otherwise, recursively delete the directories and files below the given directory.
ctx	Context with which to delete the directory (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.8 DELETEFILE Procedure

This procedure deletes the specified file.

Syntax

```
DBMS_DBFS_CONTENT_SPI.DELETEFILE (
   store_name     IN     VARCHAR2,
   path           IN     VARCHAR2,
   filter         IN     VARCHAR2,
   soft_delete    IN     BOOLEAN,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-10 DELETEFILE Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to the file
filter	A filter, if any, to be applied
soft_delete	If 0, execute a hard (permanent) delete. For any value other than 0, perform a soft delete see *Oracle Database SecureFiles and Large Objects Developer's Guide*, Deletion Operations).
ctx	Context with which to delete the file (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.9 GETFEATURES Function

This function returns the features of a store.

Syntax

```
DBMS_DBFS_CONTENT_SPI.GETFEATURES (
   store_name          IN      VARCHAR2)
  RETURN  INTEGER;
```

Parameters

Table 63-11 GETFEATURES Function Parameters

Parameter	Description
store_name	Name of store

Return Values

DBMS_DBFS_CONTENT.FEATURE_* features supported by the Store Provider

63.4.10 GETPATH Procedures

This procedure returns existing path items (such as files and directories). This includes both data and metadata (properties).

The client can request (using prop_flags) that specific properties be returned. File path names can be read either by specifying a BLOB locator using the prop_data bitmask in prop_flags (see Table 62-9) or by passing one or more RAW buffers.

When forUpdate is 0, this procedure also accepts a valid "as of" timestamp parameter as part of ctx that can be used by stores to implement "as of" style flashback queries. Mutating versions of the GETPATH Procedures do not support these modes of operation.

Syntax

```
DBMS_DBFS_CONTENT_SPI.GETPATH (
   store_name  IN              VARCHAR2,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   content     OUT    NOCOPY   BLOB,
   item_type   OUT             INTEGER,
   prop_flags  IN              INTEGER,
   forUpdate   IN              INTEGER,
   deref       IN              INTEGER,
   ctx         IN              DBMS_DBFS_CONTENT_CONTEXT_T);

DBMS_DBFS_CONTENT_SPI.GETPATH (
   store_name  IN              VARCHAR2,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   amount      IN OUT          NUMBER,
   offset      IN              NUMBER,
   buffer      OUT    NOCOPY   RAW,
   prop_flags  IN              INTEGER,
   ctx         IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

```
DBMS_DBFS_CONTENT_SPI.GETPATH (
   store_name  IN              VARCHAR2,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   amount      IN OUT          NUMBER,
   offset      IN              NUMBER,
   buffers     OUT    NOCOPY   DBMS_DBFS_CONTENT_RAW_T,
   prop_flags  IN              INTEGER,
   ctx         IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-12 GETPATH Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to path items
properties	One or more properties and their values to be returned depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
content	BLOB holding data which populates the file (optional)
item_type	Type of the path item specified (see Table 62-4)
amount	On input, number of bytes to be read. On output, number of bytes read
offset	Byte offset from which to begin reading
buffer	Buffer to which to write
buffers	Buffers to which to write
prop_flags	Determines which properties are set, returned, or both. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
forUpdate	Specifies that a lock should be taken to signify exclusive write access to the path item
deref	If nonzero, attempts to resolve the given path item to actual data provided it is a reference (symbolic link)
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.11 GETPATHBYSTOREID Function

If the underlying GUID is found in the underlying store, this function returns the store-qualified path name.

Syntax

```
DBMS_DBFS_CONTENT_SPI.GETPATHBYSTOREID (
   store_name          IN      VARCHAR2,
   guid                IN      INTEGER)
  RETURN VARCHAR2;
```

Parameters

Table 63-13 GETPATHBYSTOREID Function Parameters

Parameter	Description
store_name	Name of store
guid	Unique ID representing the desired path item

Return Values

Store-qualified path name represented by the GUID

Usage Notes

If the STD_GUID is unknown, a NULL value is returned. Clients are expected to handle this as appropriate.

63.4.12 GETPATHNOWAIT Procedure

This procedure implies that the operation is for an update, and, if implemented, allows providers to return an exception (ORA-00054) rather than wait for row locks.

See FEATURE_NOWAIT in Table 62-5 for more information.

Syntax

```
DBMS_DBFS_CONTENT_SPI.GETPATHNOWAIT (
   store_name  IN              VARCHAR2,
   path        IN              VARCHAR2,
   properties  IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   content     OUT    NOCOPY   BLOB,
   item_type   OUT             INTEGER,
   prop_flags  IN              INTEGER,
   deref       IN              INTEGER,
   ctx         IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-14 GETPATHNOWAIT Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to path items
properties	One or more properties and their values to be returned depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
content	BLOB holding data which populates the file (optional)
item_type	Type of the path item specified (see Table 62-4)
prop_flags	Determines which properties are returned. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.

Table 63-14 (Cont.) GETPATHNOWAIT Procedure Parameters

Parameter	Description
deref	If nonzero, attempts to resolve the given path item to actual data provided it is a reference (symbolic link)
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.13 GETSTOREID Function

This function returns the ID of a store.

Syntax

```
DBMS_DBFS_CONTENT_SPI.GETSTOREID (
   store_name          IN      VARCHAR2)
  RETURN  NUMBER;
```

Parameters

Table 63-15 GETSTOREID Function Parameters

Parameter	Description
store_name	Name of store

Return Values

ID of the Store

Usage Notes

A store ID identifies a provider-specific store, across registrations and mounts, but independent of changes to the store contents. For this reason, changes to the store table or tables should be reflected in the store ID, but re-initialization of the same store table or tables should preserve the store ID.

63.4.14 GETVERSION Function

This function returns the version associated with a store.

Syntax

```
DBMS_DBFS_CONTENT_SPI.GETVERSION (
   store_name          IN      VARCHAR2)
  RETURN  VARCHAR2;
```

Parameters

Table 63-16 GETVERSION Function Parameters

Parameter	Description
store_name	Name of store

Return Values

A "version" (either specific to a provider package, or to an individual store) based on a standard *a.b.c* naming convention (for *major*, *minor*, and *patch* components)

63.4.15 ISPATHLOCKED Procedure

This procedure checks if any user-level locks are applied on a given path.

Syntax

```
DBMS_DBFS_CONTENT.ISPATHLOCKED (
   store_name     IN     VARCHAR2,
   path           IN     VARCHAR2,
   who            IN     VARCHAR2,
   lock_type      IN OUT INTEGER,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-17 ISPATHLOCKED Procedure Parameters

Parameter	Description
store_name	Name of store
path	Path name of items to be locked
who	Transaction identifier that has locked the path
lock_type	One of the available lock types (see Table 62-6)
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.16 LIST Function

This function lists the contents of a directory path name.

The invoker of the subprogram has the option to investigate recursively into sub-directories, to make soft-deleted items visible, to use a flashback "as of" a specified timestamp, and to filter items within the store based on list predicates.

Syntax

```
DBMS_DBFS_CONTENT_SPI.LIST (
   store_name     IN     VARCHAR2,
   path           IN     VARCHAR2,
   filter         IN     VARCHAR2,
   recurse        IN     INTEGER,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T)
  RETURN  DBMS_DBFS_CONTENT_LIST_ITEMS_T PIPELINED;
```

Parameters

Table 63-18 LIST Function Parameters

Parameter	Description
store_name	Name of repository
path	Name of path to directories
filter	A filter, if any, to be applied
recurse	If 0, do not execute recursively. Otherwise, recursively list the contents of directories and files below the given directory.
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

Return Values

Path items found that match the path, filter and criteria for executing recursively (see DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type)

Usage Notes

This function returns only list items; the client is expected to explicitly use one of the GETPATH Procedures to access the properties or content associated with an item.

63.4.17 LOCKPATH Procedure

This procedure applies user-level locks to the given valid path name (subject to store feature support), and optionally associates user-data with the lock.

Syntax

```
DBMS_DBFS_CONTENT_SPI.LOCKPATH (
   store_name     IN     VARCHAR2,
   path           IN     VARCHAR2,
   who            IN     VARCHAR2,
   lock_type      IN     INTEGER,
   waitForRowLock IN     INTEGER,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-19 LOCKPATH Procedure Parameters

Parameter	Description
store_name	Name of store
path	Path name of items to be locked
who	Transaction identifier that has locked the path
lock_type	One of the available lock types (see Table 62-6)
waitForRowLock	Determines if a row is locked by a transaction or not
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

Usage Notes

- It is the responsibility of the store and its providers (assuming it supports user-defined lock checking) to ensure that lock and unlock operations are performed in a consistent manner.
- The status of locked items is available by means of various optional properties (see `OPT_LOCK*` in Table 62-8).

63.4.18 PURGEALL Procedure

This procedure purges all soft-deleted entries matching the path and optional filter criteria.

Syntax

```
DBMS_DBFS_CONTENT_SPI.PURGEALL (
   store_name     IN      VARCHAR2,
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2,
   ctx            IN      DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-20 PURGEALL Procedure Parameters

Parameter	Description
`store_name`	Name of store
`path`	Name of path to file items
`filter`	A filter, if any, to be applied based on specified criteria
`ctx`	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.19 PURGEPATH Procedure

This procedure purges any soft-deleted versions of the given path item.

Syntax

```
DBMS_DBFS_CONTENT_SPI.PURGEPATH (
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2,
   store_name     IN      VARCHAR2,
   ctx            IN      DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-21 PURGEPATH Procedure Parameters

Parameter	Description
`store_name`	Name of store
`path`	Name of path to file items
`filter`	A filter, if any, to be applied

Table 63-21 (Cont.) PURGEPATH Procedure Parameters

Parameter	Description
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.20 PUTPATH Procedures

This procedure creates a new path item.

Syntax

```
DBMS_DBFS_CONTENT_SPI.PUTPATH (
   store_name     IN              VARCHAR2,
   path           IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   content        IN OUT NOCOPY   BLOB,
   item_type      OUT             INTEGER,
   prop_flags     IN              INTEGER,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);

DBMS_DBFS_CONTENT_SPI.PUTPATH (
   store_name     IN              VARCHAR2,
   path           IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   amount         IN              NUMBER,
   offset         IN              NUMBER,
   buffer         IN              RAW,
   prop_flags     IN              INTEGER,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);

DBMS_DBFS_CONTENT_SPI.PUTPATH (
   store_name     IN              VARCHAR2,
   path           IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   written        OUT             NUMBER,
   offset         IN              NUMBER,
   buffers        IN              DBMS_DBFS_CONTENT_RAW_T,
   prop_flags     IN              INTEGER,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-22 PUTPATH Procedure Parameters

Parameter	Description
store_name	Name of store
path	Path name of item to be put
properties	One or more properties and their values to be set depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
content	BLOB holding data which populates the file (optional)
item_type	Type of the path item specified (see Table 62-4)
amount	Number of bytes to be read

Table 63-22 (Cont.) PUTPATH Procedure Parameters

Parameter	Description
written	Number of bytes written
offset	Byte offset from which to begin reading
buffer	Buffer to which to write
buffers	Buffers to which to write
prop_flags	Determines which properties are set. Default is PROP_STD. Specify properties to be returned by setting prop_spec, and providing an instance of the DBMS_DBFS_CONTENT_PROPERTIES_T Table Type with properties whose values are of interest.
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

Usage Notes

- All path names allow their metadata (properties) to be read and modified. On completion of the call, the client can access specific properties using prop_flags (see Table 62-9).
- On completion of the call, the client can request a new BLOB locator that can be used to continue data access using the prop_data bitmask in prop_flags (see Table 62-9).
- Files can also be written without using BLOB locators, by explicitly specifying logical offsets or buffer-amounts, and a suitably sized buffer.

63.4.21 RENAMEPATH Procedure

This procedure renames or moves a path. This operation can be performed across directory hierarchies and mount-points as long as it is within the same store.

> **Note:**
>
> See *Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Rename and Move operations

Syntax

```
DBMS_DBFS_CONTENT_SPI.RENAMEPATH (
   store_name     IN              VARCHAR2,
   oldPath        IN              VARCHAR2,
   newPath        IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-23 RENAMEPATH Procedure Parameters

Parameter	Description
store_name	Name of store, must be unique

Table 63-23 (Cont.) RENAMEPATH Procedure Parameters

Parameter	Description
oldPath	Name of path prior to renaming
newPath	Name of path after renaming
properties	One or more properties and their values to be set depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.22 RESTOREALL Procedure

This procedure restores all soft-deleted path items meeting the path and optional filter criteria.

Syntax

```
DBMS_DBFS_CONTENT_SPI.RESTOREALL (
   store_name     IN      VARCHAR2,
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2,
   ctx            IN      DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-24 RESTOREALL Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to path items
filter	A filter, if any, to be applied
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.23 RESTOREPATH Procedure

This procedure restores all soft-deleted path items that match the given path and optional filter criteria.

Syntax

```
DBMS_DBFS_CONTENT_SPI.RESTOREPATH (
   store_name     IN      VARCHAR2,
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2,
   ctx            IN      DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-25 RESTOREPATH Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to path items
filter	A filter, if any, to be applied
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.24 SEARCH Function

This function searches for path items matching the given path and filter criteria.

Syntax

```
DBMS_DBFS_CONTENT_SPI.SEARCH (
   store_name     IN      VARCHAR2,
   path           IN      VARCHAR2,
   filter         IN      VARCHAR2,
   recurse        IN      INTEGER,
   ctx            IN      DBMS_DBFS_CONTENT_CONTEXT_T)
 RETURN  DBMS_DBFS_CONTENT_LIST_ITEMS_T PIPELINED;
```

Parameters

Table 63-26 LIST Function Parameters

Parameter	Description
store_name	Name of store
path	Name of path to the path items
filter	A filter, if any, to be applied
recurse	If 0, do not execute recursively. Otherwise, recursively search the contents of directories and files below the given directory.
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

Return Values

Path items matching the given path and filter criteria (see DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type)

63.4.25 SETPATH Procedure

This procedure assigns a path name to a path item represented by contentID.

Stores and their providers that support contentID-based access and lazy path name binding also support the `SETPATH Procedure` that associates an existing `contentID` with a new path.

> **Note:**
>
> See *Oracle Database SecureFiles and Large Objects Developer's Guide* for further information on Rename and Move operations

Syntax

```
DBMS_DBFS_CONTENT_SPI.SETPATH (
   store_name     IN              VARCHAR2,
   contentID      IN              RAW,
   path           IN              VARCHAR2,
   properties     IN OUT NOCOPY   DBMS_DBFS_CONTENT_PROPERTIES_T,
   ctx            IN              DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-27 SETPATH Procedure Parameters

Parameter	Description
store_name	Name of the store
contentID	Unique identifier for the item to be associated
path	Name of path to path item
properties	One or more properties and their values to be set depending on prop_flags (see DBMS_DBFS_CONTENT_PROPERTIES_T Table Type)
ctx	Context with which to access the path items (seE DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

63.4.26 SPACEUSAGE Procedure

This procedure queries file system space usage statistics.

Providers are expected to support this subprogram for their stores and to make a best effort determination of space usage, especially if the store consists of multiple tables, indexes, LOBs, and so on.

Syntax

```
DBMS_DBFS_CONTENT_SPI.SPACEUSAGE (
   store_name     IN        VARCHAR2,
   blksize        OUT       INTEGER,
   tbytes         OUT       INTEGER,
   fbytes         OUT       INTEGER,
   nfile          OUT       INTEGER,
   ndir           OUT       INTEGER,
   nlink          OUT       INTEGER,
   nref           OUT       INTEGER);
```

Parameters

Table 63-28 SPACEUSAGE Procedure Parameters

Parameter	Description
store_name	Name of store
blksize	Natural tablespace blocksize that holds the store. If multiple tablespaces with different blocksizes are used, any valid blocksize is acceptable.
tbytes	Total size of the store in bytes computed over all segments that comprise the store
fbytes	Free or unused size of the store in bytes computed over all segments that comprise the store
nfile	Number of currently available files in the store
ndir	Number of currently available directories in the store
nlink	Number of currently available links in the store
nref	Number of currently available references in the store

Usage Notes

- A space usage query on the top-level root directory returns a combined summary of the space usage of all available distinct stores under it (if the same store is mounted multiple times, it is still counted only once).
- Since database objects are dynamically expandable, it is not easy to estimate the division between "free" space and "used" space.

63.4.27 UNLOCKPATH Procedure

This procedure unlocks path items that were previously locked with the LOCKPATH Procedure.

Syntax

```
DBMS_DBFS_CONTENT_SPI.UNLOCKPATH (
   store_name     IN     VARCHAR2,
   path           IN     VARCHAR2,
   who            IN     VARCHAR2,
   waitForRowLock IN     INTEGER,
   ctx            IN     DBMS_DBFS_CONTENT_CONTEXT_T);
```

Parameters

Table 63-29 UNLOCKPATH Procedure Parameters

Parameter	Description
store_name	Name of store
path	Name of path to the path items
who	Transaction identifier that has locked the path
waitForRowLock	Determines if a row is locked by a transaction or not
ctx	Context with which to access the path items (see DBMS_DBFS_CONTENT_CONTEXT_T Object Type)

Related Topics

- LOCKPATH Procedure
  This procedure applies user-level locks to the given valid path name (subject to store feature support), and optionally associates user-data with the lock.

64
DBMS_DBFS_HS

The Oracle Database File System Hierarchical Store is implemented in the `DBMS_DBFS_HS` package. This package provides users the ability to use tape or Amazon S3 Web service as a storage tier when doing Information Lifecycle Management for their database tables.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operational Notes
- Summary of DBMS_DBFS_HS Subprograms

> **See Also:**
>
> - *Oracle Database SecureFiles and Large Objects Developer's Guide*

64.1 DBMS_DBFS_HS Overview

The `DBMS_DBFS_HS` package is a service provider underneath the `DBMS_DBFS_CONTENT` package that enables use of tape or Amazon S3 Web service as storage for data.

The data on tape or Amazon S3 Web service is part of the Oracle Database and can be accessed through all standard interfaces, but only through the database. The package allows users to use tape or Amazon S3 Web service as a storage tier when doing Information Lifecycle Management of their content.

The package initially stores all content files in level-1 cache. As the level-1 cache fills up, content files are moved to level-2 cache and then to an external storage device using bulk writes.

64.2 DBMS_DBFS_HS Security Model

The `DBMS_DBFS_HS` package runs with invoker's rights.

64.3 DBMS_DBFS_HS Constants

The `DBMS_DBFS_HS` package uses the constants shown in the tables in this topic.

Table 64-1 DBMS_DBFS_HS Constants - Used by the CREATESTORE Procedure

Constant	Type	Value	Description
STORETYPE_TAPE	VARCHAR2(50)	'HS_TAPE'	Use tape as a storage tier
STORETYPE_AMAZON S3	VARCHAR2(50)	'HS_S3'	Use Amazon S3 Web service as a storage tier

Table 64-2 DBMS_DBFS_HS Constants - Used by the SETSTOREPROPERTY Procedure and the GETSTOREPROPERTY Function

Constant	Type	Value	Description
PROPNAME_BUCKET	VARCHAR2(50)	'BUCKET'	Specifies the AWS bucket to be used as a storage tier by the Hierarchical Store. Restrictions on bucket name are: 1) Bucket names can only contain lowercase letters, numbers, periods (.) and dashes(-). Note that underscores (_) are invalid. 2) Bucket names must start with a number or letter. 3) Bucket names cannot be in an IP address style (192.168.5.4). 4) Bucket names must be between 3 and 63 characters long. 5) Bucket names should not end with a dash. 6) Dashes cannot appear next to periods. For example, my-.bucket.com is invalid.
PROPNAME_CACHESIZE	VARCHAR2(50)	'CACHE_SIZE'	Specifies the cumulative cache size used for the Hierarchical Store. This property is set by the CREATESTORE Procedure and can be modified by the RECONFIGCACHE Procedure. It cannot be modified by the SETSTOREPROPERTY Procedure, though its value can be queried by the GETSTOREPROPERTY Function.
PROPNAME_COMPRESSLEVEL	VARCHAR2(50)	'COMPRESSION_LEVEL'	Use to enable compression of files stored in the DBFS hierarchical store. It specifies the compression level to be used for compressing the files

Table 64-2 (Cont.) DBMS_DBFS_HS Constants - Used by the SETSTOREPROPERTY Procedure and the GETSTOREPROPERTY Function

Constant	Type	Value	Description
PROPNAME_ENABLECLEANUPONDELETE	VARCHAR2(50)	'ENABLE_CLEANUP_ON_DELETE'	If this property is set to 'TRUE', whenever the user invokes the DELETEFILE Procedure in the DBMS_DBFS_CONTENT interface on a file residing in the DBMS_DBFS_HS store, the DBMS_DBFS_HS removes the file on the external storage that contains this user file, provided that the file has no other useful data.By default, the property is set to 'TRUE' for STORETYPE_AMAZONS3 and 'FALSE' for STORETYPE_TAPE.
PROPNAME_HTTPPROXY	VARCHAR2(50)	'HTTP_PROXY'	Specifies the DNS name of the HTTP proxy, if any, that is needed to access the Amazon S3 storage service
PROPNAME_LICENSEID	VARCHAR2(50)	'LICENSE_ID'	Specifies the license ID associated with the library libosbws11.so.
PROPNAME_LOBCACHE_QUOTA	VARCHAR2(50)	'LOBCACHE_QUOTA'	Specifies fraction of the cache_size which is allocated for level 1 cache. The default value of this parameter is NULL which means that 0.8 (= 80%) of the cache_size is used for level 1 cache. This property cannot be modified by the SETSTOREPROPERTY Procedurethough its value can be queried by the GETSTOREPROPERTY Function. Its value is set by CREATESTORE Procedure and can be modified by the RECONFIGCACHE Procedure.
PROPNAME_MEDIAPOOL	VARCHAR2(50)	'MEDIA_POOL'	Specifies the media pool number to use for storing the content
PROPVAL_COMPLVL_NONE	VARCHAR2(50)	'NONE'	Indicates no compression
PROPVAL_COMPLVL_LOW	VARCHAR2(50)	'LOW'	Use to set the compression level to LOW. This is expected to have the best performance while still providing a good compression ratio.
PROPVAL_COMPLVL_MEDIUM	VARCHAR2(50)	'MEDIUM'	Use to set the compression level to MEDIUM. This compression level is expected to provide better compression ratio than LOW but the time required for compression will be higher than compression level LOW.
PROPVAL_COMPLVL_HIGH	VARCHAR2(50)	'HIGH'	Use to set the compression level to HIGH. This compression level is expected to provide the best compression ratio but compression time will in general be highest among the 3 compression levels.

Table 64-2 (Cont.) DBMS_DBFS_HS Constants - Used by the SETSTOREPROPERTY Procedure and the GETSTOREPROPERTY Function

Constant	Type	Value	Description
PROPNAME_OPTTARBALLSIZE	VARCHAR2(50)	'OPTIMAL_TARBALL_SIZE'	Specifies optimal_tarball_size as the maximum possible size of an archive file. Multiple content files are bundled together into one archive file and then the archive file is transferred to tape or Amazon S3. This is because creating one file on tape or Amazon S3 for every content file in the store is a prohibitively expensive operation. This property cannot be modified by the SETSTOREPROPERTY Procedurethough its value can be queried by the GETSTOREPROPERTY Function. Its value is set by CREATESTORE Procedure and can be modified by the RECONFIGCACHE Procedure.
PROPNAME_READCHUNKSIZE	VARCHAR2(50)	'READ_CHUNK_SIZE'	Specifies the size used by the SBT protocol to transfer data from tape or S3.This chunk is allocated in memory per transaction for retrieval of content files from an archive store, so the value of this property should be conservative. The default size of 1MB is typically good for most users.
PROPNAME_S3HOST	VARCHAR2(50)	'S3_HOST'	Specifies the HOST name of the Amazon S3 storage service. It must be s3.amazonaws.com.
PROPNAME_SBT_LIBRARY	VARCHAR2(50)	'SBT_LIBRARY'	Specifies the path of the shared library used by RMAN to communicate with Amazon S3. It is named libosbws11.so and is available in rdbms/lib directory.
PROPNAME_STREAMABLE	VARCHAR2(50)	'STREAMABLE'	Indicates whether buffer-based PUT or GET should be done on this store. Valid values for are TRUE and FALSE. The default value of this property is TRUE.

Table 64-2 (Cont.) DBMS_DBFS_HS Constants - Used by the SETSTOREPROPERTY Procedure and the GETSTOREPROPERTY Function

Constant	Type	Value	Description
PROPNAME_WALLET	VARCHAR2(50)	'WALLET'	The value of this property should be of the form: LOCATION=file:*filename* CREDENTIAL_ALIAS=*access/secret_alias* PROXY_AUTH_ALIAS=*proxyusername/password alias* Defines the Oracle Wallet which contains the credentials of the Amazon S3 account associated with the store under consideration. LOCATION: The directory path that contains the Oracle wallet. The format is file:*directory-path* The format of wallet_path in Windows is, for example: file:c:\WINNT\Profiles*username*\WALLETS In UNIX or Linux it is, for example: file:/home/*username*/wallets When the package is executed in the Oracle database server, the wallet is accessed from the database server. PASSWORD: Defines the wallet password. If auto-login is enabled in wallet, this parameter does not have to be specified. By default, the mkstore utility enables auto-login. CREDENTIAL_ALIAS: Defines the credential alias for ACCESS_KEY and SECRET_KEY
PROPNAME_WRITECHUNKSIZ	VARCHAR2(50)	'WRITE_CHUNK_SIZE'	Specifies the size used by the SBT protocol to transfer data to tape or S3. This chunk is allocated in memory per transaction for PUT of Content Files to an archive store so the value should be conservative. The default size of 1MB is typically good for most users.

Table 64-3 DBMS_DBFS_HS Constants - Used by the REGISTERSTORECOMMAND Function

Constant	Type	Value	Description
BEFORE_PUT	NUMBER	'1'	Specified operation must be performed before writing a SECUREFILE to the remote store

Table 64-3 (Cont.) DBMS_DBFS_HS Constants - Used by the REGISTERSTORECOMMAND Function

Constant	Type	Value	Description
BEFORE_GET	NUMBER	'2'	Specified operation must be performed before a retrieval operation such as reading a SECUREFILE from the remote device

Table 64-4 DBMS_DBFS_HS Constants - Failure/Success/Error

Constant	Type	Value	Description
FAIL	NUMBER	'0'	Procedure or function did not execute successfully
SUCCESS	NUMBER	'1'	Procedure or function completed successfully
ERROR	NUMBER	'2'	Procedure or function returned an error

64.4 DBMS_DBFS_HS Operational Notes

When the DBMS_DBFS_HS package is executed in the Oracle database server, the wallet is accessed from the database server.

64.5 Summary of DBMS_DBFS_HS Subprograms

This table lists and describes the subprograms in the DBMS_DBFS_HS package.

Table 64-5 DBMS_DBFS_HS Package Subprograms

Subprogram	Description
CLEANUPUNUSEDBACKUPFILES Procedure	Removes files created on the external storage device that hold no currently used data
CREATEBUCKET Procedure	Creates an AWS bucket, associated with a store of type STORETYPE_AMAZONS3 into which the Hierarchical Store can then move data
CREATESTORE Procedure	Creates a new hierarchical store
DEREGSTORECOMMAND Function	Removes a command that had been previously associated with a store through the RECONFIGCACHE Procedure
DROPSTORE Procedure	Deletes a previously created hierarchical store
FLUSHCACHE Procedure	Flushes (writes out) dirty contents from the level-1 cache.
GETSTOREPROPERTY Function	Retrieves the values of a property of a store
RECONFIGCACHE Procedure	Reconfigures the parameters of the database cache used by the store
REGISTERSTORECOMMAND Procedure	Registers commands for a store with the Hierarchical Store to be sent to the Media Manager for the external storage device associated with the store

Table 64-5 (Cont.) DBMS_DBFS_HS Package Subprograms

Subprogram	Description
SENDCOMMAND Procedures	Sends a command to be executed on the external storage device's Media Manager
SETSTOREPROPERTY Procedure	Stores properties of a store in the database
STOREPUSH Procedure	Pushes locally staged data to the remote storage

64.5.1 CLEANUPUNUSEDBACKUPFILES Procedure

This procedure removes files created on the external storage device that hold no currently used data in them.

Syntax

```
DBMS_DBFS_HS.CLEANUPUNUSEDBACKUPFILES (
   store_name      IN      VARCHAR2);
```

Parameters

Table 64-6 CLEANUPUNUSEDBACKUPFILES Procedure Parameters

Parameter	Description
store_name	Name of store

Usage Notes

- The action of removing files from external storage device can not be rolled back.
- This method can be executed periodically to clear space on the external storage device. Asynchronously deleting content from the external storage device is useful because it has minimal impact on the OLTP performance. The periodic scheduling can be accomplished using the DBMS_SCHEDULER package.

64.5.2 CREATEBUCKET Procedure

This procedure creates an AWS bucket, associated with a store of type `STORETYPE_AMAZONS3` into which the Hierarchical Store can then move data.

Syntax

```
DBMS_DBFS_HS.CREATEBUCKET   (
   store_name     IN      VARCHAR2);
```

Parameters

Table 64-7 CREATEBUCKET Procedure Parameters

Parameter	Description
store_name	Name of store

Usage Notes

- The PROPNAME_BUCKET property of the store should be set before this subprogram is called.
- Once this procedure has successfully created a bucket in Amazon S3, the bucket can only be deleted using out-of-band methods, such as logging-in to S3 and deleting data (directories, files, and other items) for the bucket.

64.5.3 CREATESTORE Procedure

This procedure creates a new hierarchical store store_name of type STORE_TYPE (STORETYPE_TAPE or STORETYPE_AMAZONS3) in schema schema_name (defaulting to the current schema) under the ownership of the invoking session user.

Syntax

```
DBMS_DBFS_HS.CREATESTORE  (
   store_name              IN     VARCHAR2,
   store_type              IN     VARCHAR2,
   tbl_name                IN     VARCHAR2,
   tbs_name                IN     VARCHAR2,
   cache_size              IN     NUMBER,
   lob_cache_quota         IN     NUMBER DEFAULT NULL,
   optimal_tarball_size    IN     NUMBER DEFAULT NULL,
   schema_name             IN     VARCHAR2 DEFAULT NULL);
```

Parameters

Table 64-8 CREATESTORE Procedure Parameters

Parameter	Description
store_name	Name of store
store_type	STORETYPE_TAPE or STORETYPE_AMAZONS3
tbl_name	Table for store entries
tbs_name	Tablespace for the store
cache_size	Amount of space used by the store to cache content in given tablespace
lob_cache_quota	Fraction of the cache_size which is allocated for level 1 cache. The default value of this parameter is NULL which means that 0.8 (= 80%) of the cache_size is used for level 1 cache.
optimal_tarball_size	Maximum possible size of the archive file. Multiple content files are bundled together into one archive file, and then the archive file is transferred to tape or Amazon S3. This is because creating one file on tape or Amazon S3 for every content file in the store is a prohibitively expensive operation. The value of is set by default to 10GB for tape and to 100MB for Amazon S3.
schema_name	Schema for the store

Usage Notes

CREATESTORE() sets certain properties of the store to default values. The user can use the methods SETSTOREPROPERTY() and RECONFIGCACHE() to appropriately change the property values and to set other properties of the store.

- Store names must be unique for an owner. The same store names can be used for different stores owned by different owners.
- Once a table space has been specified to store the store's content in a database, it cannot be changed later.
- This subprogram will execute like a DDL statement, performing an automatic `COMMIT` before and after execution.
- Stores using `DBMS_DBFS_HS` must not use singleton mount. This means that the singleton parameter should be `FALSE` and the `store_mount` parameter should have a non-`NULL` value in a call to the `DBMS_DBFS_CONTENT.MOUNTSTORE` procedure.

64.5.4 DEREGSTORECOMMAND Function

This procedure removes a command that had been previously associated with a store through the REGISTERSTORECOMMAND Procedure.

Syntax

```
DBMS_DBFS_HS.DEREGSTORECOMMAND (
   store_name       IN     VARCHAR2,
   message          IN     VARCHAR2);
```

Parameters

Table 64-9 DEREGSTORECOMMAND Procedure Parameters

Parameter	Description
`store_name`	Name of store
`message`	Message to be deregistered

Usage Notes

If this subprogram successfully executes, its actions cannot be rolled back by the user. If the user wants to restore the previous state, the user must call the REGISTERSTORECOMMAND Procedure.

Related Topics

- REGISTERSTORECOMMAND Procedure
  This procedure registers commands for a store with the Hierarchical Store. These commands are sent to the Media Manager for the external storage device associated with the store.

64.5.5 DROPSTORE Procedure

This procedure deletes a previously created hierarchical store specified by name and owned by the invoking session user.

Syntax

```
DBMS_DBFS_HS.DROPSTORE  (
   store_name    IN     VARCHAR2,
   opt_flags     IN     INTEGER DEFAULT 0);
```

Parameters

Table 64-10 DROPSTORE Procedure Parameters

Parameter	Description
store_name	Name of store owned by the invoking session user
opt_flags	User can specify optional flags. If DISABLE_CLEANUPBACKUPFILES is specified as one of the optional flags, the call to the CLEANUPUNUSEDBACKUPFILES Procedure is not issued. By default, when this flag is not set, the procedure implicitly cleans-up all unused backup files.

Usage Notes

- The procedure executes like a DDL in that it auto-commits before and after its execution.
- If CLEANUPBACKUPFILES is disabled during the procedure, the user must resort to out-of-band techniques to cleanup unused backup files. No further invocations of CLEANUPBACKFILES for a dropped store are possible through hierarchical store.
- This subprogram will un-register the store from DBMS_DBFS_CONTENT package. All files in the given store are deleted from the store (Tape or Amazon S3 Web service). The database table holding the store's entries in the database, is also dropped by this subprogram.

64.5.6 FLUSHCACHE Procedure

This procedure flushes out dirty contents from level-1 cache, which can be locked, to level-2 cache, thereby freeing-up space in level 1 cache.

Syntax

```
DBMS_DBFS_HS.FLUSHCACHE  (
   store_name    IN     VARCHAR2);
```

Parameters

Table 64-11 FLUSHCACHE Procedure Parameters

Parameter	Description
store_name	Name of store

64.5.7 GETSTOREPROPERTY Function

This function retrieves the values of a property.

Syntax

```
DBMS_DBFS_HS.GETSTOREPROPERTY  (
   store_name      IN     VARCHAR2,
   property_name   IN     VARCHAR2,
   noexcp          IN     BOOLEAN DEFAULT FALSE) RETURN VARCHAR2;
```

Parameters

Table 64-12 GETSTOREPROPERTY Function Parameters

Parameter	Description
store_name	Name of store
property_name	Name of property
noexcp	If set to FALSE, raises an exception if the property does not exist in the database. If noexcp is set to TRUE, returns NULL if the property does not exist.

Return Values

The values of a property.

Usage Notes

The specified store must already have been created.

64.5.8 RECONFIGCACHE Procedure

This procedure reconfigures the parameters of the database cache being used by the store.

Syntax

```
DBMS_DBFS_HS.RECONFIGCACHE (
   store_name              IN    VARCHAR2,
   cache_size              IN    NUMBER DEFAULT NULL,
   lobcache_quota          IN    NUMBER DEFAULT NULL,
   optimal_tarball_size    IN    NUMBER DEFAULT NULL);
```

Parameters

Table 64-13 RECONFIGCACHE Procedure Parameters

Parameter	Description
store_name	Name of store
cache_size	Cumulative cache size used for the Hierarchical Store
lobcache_quota	Fraction of the cache size that are assigned to level 1 cache
optimal_tarball_size	Maximum possible size of an archive file. Since creating one file for every content file in the store is a prohibitively expensive operation, multiple content files are bundled together into one archive file for transfer to tape or Amazon S3.

Usage Notes

- The specified store must already have been created before reconfiguration.
- The Hierarchical Store uses a level 1 cache and a level 2 cache. The level 1 cache subsumes most of the working set and the level 2 cache is used to perform bulk writes to the backend device.

- If any of the last 3 parameters is `NULL`, its value specified during store creation is used. If the parameter was `NULL` when the call to the CREATESTORE Procedure was issued, the `DBMS_DBFS_HS` package assigns a default value.

  The `DBMS_DBFS_HS` package optimistically tries to allocate more than 1 tarball's worth of size for level 2 cache to facilitate concurrency, though a minimum of 1 tarball size is necessary for level 2 cache.

  The values for cumulative cache size and LOB cache quota decide allocation of space for the two caches. If values are not provided, a user might see an `INSUFFICIENT_CACHE` exception. In that case, it is better to revise the cache parameters in order to have a working store.
- If this subprogram successfully executes, its actions cannot be rolled back by the user. In that case, the user should call `RECONFIGCACHE` again with new or modified parameters.

64.5.9 REGISTERSTORECOMMAND Procedure

This procedure registers commands for a store with the Hierarchical Store. These commands are sent to the Media Manager for the external storage device associated with the store.

Syntax

```
DBMS_DBFS_HS.REGISTERSTORECOMMAND (
   store_name       IN     VARCHAR2,
   message          IN     VARCHAR2,
   flags            IN     NUMBER);
```

Parameters

Table 64-14 REGISTERSTORECOMMAND Procedure Parameters

Parameter	Description
`store_name`	Name of store
`message`	Message to be sent to the Media Manager of the external store
`flags`	Valid values: • `BEFORE_PUT CONSTANT NUMBER := 1;` • `BEFORE_GET CONSTANT NUMBER := 2;`

Usage Notes

- These commands are sent before the next read or write of content. When the Hierarchical Store wants to push (or get) data to (or from) the storage device, it begins a session (to communicate with the device). After beginning the session, it sends all registered commands for the to the relevant device before writing (or getting) any data.
- If this method successfully executes, its actions cannot be rolled back by the user. To restore the previous state the user must call the DEREGSTORECOMMAND Function.

64.5.10 SENDCOMMAND Procedure

This procedure sends a command to be executed on the external storage device's Media Manager.

Syntax

```
DBMS_DBFS_HS.SENDCOMMAND (
   store_name     IN       VARCHAR2,
   message        IN       VARCHAR2);
```

Parameters

Table 64-15 SENDCOMMAND Procedure Parameters

Parameter	Description
`store_name`	Name of store
`message`	Message string to be executed

64.5.11 SETSTOREPROPERTY Procedure

This procedure stores properties of a store in the database as name-value pairs.

Syntax

```
DBMS_DBFS_HS.SETSTOREPROPERTY  (
   store_name      IN     VARCHAR2,
   property_name   IN     VARCHAR2,
   property_value  IN     VARCHAR2);
```

Parameters

Table 64-16 SETSTOREPROPERTY Procedure Parameters

Parameter	Description
`store_name`	Name of store
`property_name`	For a store using Tape device, there are three properties whose values must be set by the user, and four properties that have default values. Stores of type `STORETYPE_AMAZONS3` have properties with default values. The various options for both types of stores are detailed under `property_value`.

Table 64-16 (Cont.) SETSTOREPROPERTY Procedure Parameters

Parameter	Description
property_value	**Stores using a Tape Device** The values for the following properties must be set by the user: • PROPNAME_SBTLIBRARY - This should point to the shared library used by RMAN to communicate with the external tape device. It is usually named libobk.so. • PROPNAME_MEDIAPOOL - Media pool number for storing content • PROPNAME_CACHE_SIZE - Amount of space, in bytes, used for the cache of this store The following properties, which have default values assigned to them when a store is created, benefit from tuning: • PROPNAME_READCHUNKSIZE and PROPNAME_WRITECHUNKSIZE - These are the sizes used by the SBT protocol to transfer data to and from the tapes. These chunks are allocated in memory per transaction, so the values should be conservative. The default size is 1MB. • PROPNAME_STREAMABLE – Indicates whether DBFS_LINKs can perform read operations (for example SELECT or DBMS_LOB.READ) directly from the store, or if the data must be copied back into the database before it can be read • PROPNAME_ENABLECLEANUPONDELETE - Indicates if DBMS_DBFS_HS should delete unused files on the external storage. Valid values for this property are 'FALSE' for STORETYPE_TAPE. • PROPNAME_COMPRESSLEVEL - Describes how files written to Tape should be compressed. It can be set to PROPVAL_COMPLVL_NONE, PROPVAL_COMPLVL_LOW, PROPVAL_COMPLVL_MEDIUM or PROPVAL_COMPLVL_HIGH. By default it is set to PROPVAL_COMPLVL_NONE.

Table 64-16 (Cont.) SETSTOREPROPERTY Procedure Parameters

Parameter	Description
(cont) property_value	**Stores of type STORETYPE_AMAZONS3** It is mandatory that the following properties have assigned values, and default values are provided: • PROPNAME_SBTLIBRARY - Specifies the path of the shared library used by RMAN to communicate with Amazon S3. It is named libosbws11.so and is available in rdbms/lib directory. • PROPNAME_S3HOST - Defines the HOST name of the Amazon S3 storage service. It must be s3.amazonaws.com. • PROPNAME_BUCKET - Defines the AWS bucket used as a storage tier by the Hierarchical Store. Restrictions on bucket names are: -- Bucket names can only contain lowercase letters, numbers, periods (.) and dashes (-). Use of an underscore (_) is invalid. -- Bucket names must start with a number or letter -- Bucket names cannot be in an IP address style ("192.168.5.4") -- Bucket names must be between 3 and 63 characters in length -- Bucket names should not end with a dash -- Dashes cannot appear next to periods. For example, "my-.bucket.com" is invalid. • PROPNAME_LICENSEID - Specifies the license ID associated with the library libosbws11.so. • PROPNAME_WALLET - Has the form: 'LOCATION=file:<filename> CREDENTIAL_ALIAS=<access/secret_alias> PROXY_AUTH_ALIAS=<proxyusername/password alias>' -- LOCATION - Directory path that contains the Oracle wallet. The format is file:*directory-path*. See Examples for variations in format. -- PASSWORD - Defines the wallet password. If auto-login is enabled in the wallet (this can be changed using the user's own utility), and does not have to be specified. By default, the mkstore utility enables auto-login. -- CREDENTIAL_ALIAS - Defines the credential alias for ACCESS_KEY and SECRET_KEY -- PROXY_AUTH_ALIAS - Defines authentication credentials for the proxy server, if applicable.

Table 64-16 (Cont.) SETSTOREPROPERTY Procedure Parameters

Parameter	Description
(property_value (contd.)	The following properties are optional: • PROPNAME_HTTPPROXY - Defines the DNS name of the HTTP proxy, if any, that is needed to access the Amazon S3 storage service. • PROPNAME_STREAMABLE – Indicates whether buffer-based PUT or GET operation should be done on this store. Valid values for this property are TRUE (default) and FALSE. • PROPNAME_ENABLECLEANUPONDELETE - Indicates if DBMS_DBFS_HS should delete unused files on the external storage device. Default values for this property are FALSE for STORETYPE_TAPE and TRUE for STORETYPE_AMAZONS3. • PROPNAME_COMPRESSLEVEL - Describes how files written to tape should be compressed. It can be set to PROPVAL_COMPLVL_NONE, PROPVAL_COMPLVL_LOW, PROPVAL_COMPLVL_MEDIUM or PROPVAL_COMPLVL_HIGH. By default it is set to PROPVAL_COMPLVL_NONE.

Usage Notes

- The specified store must already have been created.
- If this subprogram successfully executes, its actions cannot be rolled back by the user.
- The same property can be set multiple times to the same or different values using this subprogram
- Regarding PROPNAME_ENABLECLEANUPONDELETE behavior, a job is created for each store by the DBMS_DBFS_HS to remove the unused files from the external storage. By default, the job is enabled for STORETYPE_AMAZONS3 and is disabled for STORETYPE_TAPE. If the ENABLECLEANUPONDELETE property is set to TRUE, the job is enabled; if the property is set to FALSE, the job is disabled. If enabled, the job runs at an interval of one hour by default. The DBMS_SCHEDULER package can be used to modify the schedule. The name of the job can be obtained by querying USER_DBFS_HS_FIXED_PROPERTIES for prop_name = 'DELJOB_NAME'.

Examples

Format

The format of wallet_path in Windows is, for example:

```
file:c:\WINNT\Profiles\<username>\WALLETS
```

The format of wallet_path in UNIX or Linux is, for example:

```
file:/home/username/wallets
```

64.5.12 STOREPUSH Procedure

This procedure pushes locally staged data to the remote storage.

Syntax

```
DBMS_DBFS_HS.STOREPUSH (
   store_name   IN        VARCHAR2,
   path         IN        VARCHAR2 DEFAULT NULL);
```

Parameters

Table 64-17 STOREPUSH Procedure Parameters

Parameter	Description
`store_name`	Name of store whose content the client writes from local cache to the external store
`path`	A non-mount qualified (without mount point) path within the store. By default, its value is `NULL` which corresponds to the root path of the store.

Usage Notes

- The Hierarchical Store caches the content files locally in database tables. When enough content is amassed in the cache to make it efficient to write to the external storage device (or the cache is completely filled), the Hierarchical Store creates a tarball out of the local content and writes these tarballs as files on the external device. The size of the tarball created by the Hierarchical Store is controlled by the store property `PROPNAME_OPTTARBALLSIZE`.
- When the amount of free space in the cache is such that the caching of a content file will push the space used above `cache_size`, the Hierarchical Store will internally call `STOREPUSH`. The `STOREPUSH Procedure` creates tarball(s) out of the existing dirty or modified content files in the cache and writes them out to the external device. A `STOREPUSH` call is not guaranteed to write all the dirty content from local cache to the external storage, since some files may be locked by other sessions.
- `STOREPUSH` has a built-in ability feature allowing it to automatically resume operation. If a `STOREPUSH` call is interrupted (say by a network outage) after it has transferred some tarballs to the external device, it can be restarted after the outage and will then resume transferring data from the point it was interrupted. In other words, work done before the outage is not lost. `STOREPUSH` can safely be restarted and the effect is such as if the outage never occurred.
- If this method successfully executes, its actions cannot be rolled back by the user.
- By default, when `path` is `NULL`, all files in the store are candidates for `STOREPUSH`. If `path` has a valid input value, all files which are under the namespace of given path are written from the local cache to the external store. If a given path is an existing file, it is pushed out again to the remote store.

65
DBMS_DBFS_SFS

The `DBMS_DBFS_SFS` package provides an interface to operate a SecureFile-based store (SFS) for the content interface described in the `DBMS_DBFS_CONTENT` package.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_DBFS_SFS Subprograms

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide*

65.1 DBMS_DBFS_SFS Overview

The `DBMS_DBFS_SFS` package is a sample implementation of a package that implements and extends the `DBMS_DBFS_CONTENT_SPI` interface. It provides a POSIX-compliant file system stored in the RDBMS.

65.2 DBMS_DBFS_SFS Security Model

The `DBMS_DBFS_SFS` package runs with `AUTHID CURRENT_USER`.

65.3 DBMS_DBFS_SFS Constants

The `DBMS_DBFS_SFS` package uses the constants shown in the following tables.

> **Note:**
>
> Oracle has deprecated the older encryptions and hashing algorithms. The deprecated algorithms for `DBMS_CRYPTO` and native network encryption include MD4, MD5, DES, 3DES, and RC4-related algorithms as well as 3DES for Transparent Data Encryption (TDE). Removing older, less secure cryptography algorithms prevents accidental use of these algorithms. To meet your security requirements, Oracle recommends that you use more modern cryptography algorithms, such as the Advanced Encryption Standard (AES).

- Table 65-1

- Table 65-2
- Table 65-3
- Table 65-4

Table 65-1 DBMS_DBFS_SFS Constants - Compression Levels

Constant	Type	Value	Description
COMPRESSION_DEFAULT	VARCHAR2(32)	''	Use the default SecureFile compression level
COMPRESSION_LOW	VARCHAR2(32)	'LOW'	Use compression level 'LOW'
COMPRESSION_MEDIUM	VARCHAR2(32)	'MEDIUM'	Use compression level 'MEDIUM'
COMPRESSION_HIGH	VARCHAR2(32)	'HIGH'	Use compression level 'HIGH'

Table 65-2 DBMS_DBFS_SFS Constants - Used by the encryption Parameter

Constant	Type	Value	Description
ENCRYPTION_DEFAULT	VARCHAR2(32)	''	Use the default SecureFile encryption algorithm
ENCRYPTION_AES128	VARCHAR2(32)	'AES128'	Use encryption AES 128 bit
ENCRYPTION_AES192	VARCHAR2(32)	'AES192'	Use encryption AES 192 bit
ENCRYPTION_AES256	VARCHAR2(32)	'AES256'	Use encryption AES 256 bit

Table 65-3 DBMS_DBFS_SFS Constants - Used by the npartitions Parameter

Constant	Type	Value	Description
DEFAULT_PARTITIONS	INTEGER	16	Default to 16 partitions

Table 65-4 DBMS_DBFS_SFS Constants - Used by the partition_key Parameter

Constant	Type	Value	Description
PARTITION_BY_ITEM	INTEGER	1	Use a hash of the item name for the partition key
PARTITION_BY_PATH	INTEGER	2	Use a hash of the path name for the partition key
PARTITION_BY_GUID	INTEGER	3	Use a hash of the GUID as the partition key

65.4 Summary of DBMS_DBFS_SFS Subprograms

This table lists and describes the DBMS_DBFS_SFS Package subprograms.

Table 65-5 DBMS_DBFS_SFS Package Subprograms

Subprogram	Description
CREATEFILESYSTEM Procedure	Creates a file system store
CREATESTORE Procedure	Creates a new DBFS SFS store
DROPFILESYSTEM Procedures	Drops the DBFS SFS store
INITFS Procedure	Initializes a POSIX file system store

65.4.1 CREATEFILESYSTEM Procedure

This procedure creates a file system store.

Syntax

```
DBMS_DBFS_SFS.CREATEFILESYSTEM (
   store_name         IN     VARCHAR2,
   schema_name        IN     VARCHAR2    DEFAULT NULL,
   tbl_name           IN     VARCHAR2    DEFAULT NULL,
   tbl_tbs            IN     VARCHAR2    DEFAULT NULL,
   lob_tbs            IN     VARCHAR2    DEFAULT NULL,
   use_bf             IN     BOOLEAN     DEFAULT FALSE,
   properties         IN     DBMS_DBFS_CONTENT_PROPERTIES_T DEFAULT NULL,
   create_only        IN     BOOLEAN     FALSE,
   use_objects        IN     BOOLEAN     DEFAULT FALSE,
   with_grants        IN     BOOLEAN     DEFAULT FALSE,
   do_dedup           IN     BOOLEAN     DEFAULT FALSE,
   do_compress        IN     BOOLEAN     DEFAULT FALSE
   compression        IN     VARCHAR2    DEFAULT COMPRESSION_DEFAULT,
   do_encrypt         IN     BOOLEAN     DEFAULT FALSE,
   encryption         IN     VARCHAR2    DEFAULT ENCRYPTION_DEFAULT,
   do_partition       IN     BOOLEAN     DEFAULT FALSE,
   npartitions        IN     NUMBER      DEFAULTDEFAULT_PARTITIONS,
   partition_key      IN     NUMBER      DEFAULT PARTITION_BY_ITEM,
   partition_guidi    IN     BOOLEAN     DEFAULT FALSE,
   partition_pathi    IN     BOOLEAN     DEFAULT FALSE,
   partition_prop     IN     BOOLEAN     DEFAULT TRUE);
```

Parameters

Table 65-6 CREATEFILESYSTEM Procedure Parameters

Parameter	Description
store_name	Name of store
schema_name	Schema for the store, defaulting to the current schema
tbl_name	Table for store entries. If not specified, an internally generated name is used.
tbl_tb	Tablespace for the store, defaulting to the schema's default tablespace
lob_tbs	Tablespace in which to create the LOB segment. It defaults to the user's default tablespace.
use_bf	If TRUE, a BasicFile LOB is used; otherwise a SecureFile LOB is used.
properties	Table of (name, value, typecode) tuples used to configure the store properties. Currently no such properties are defined or used.

Table 65-6 (Cont.) CREATEFILESYSTEM Procedure Parameters

Parameter	Description
create_only	If TRUE, the file system is created, but not registered with the current user
use_objects	If TRUE, a single base-table with an object-type column (using a nested table) is created to backup the new file system. Otherwise, a pair of (parent, child) tables is used to backup the file system. In both cases, the object type nested table or the child table is used only for user-defined properties.
with_grants	If TRUE, DML and query access permissions are granted to the DBFS_ROLE as part of creating the file system. Otherwise, explicit grants (or existing permissions) are required to access the file system.
do_dedup	If TRUE, do deduplication the underlying SecureFile column
do_compress	If TRUE, do compression the underlying SecureFile column
compression	Compression algorithm to use (see Table 65-1)
do_encrypt	If TRUE, encrypt the underlying SecureFile column
encryption	encryption algorithm to use (see Table 65-2)
do_partition	If TRUE, partition the table used for storage
npartitions	Number of partitions to create for the table (see Table 65-3).
partition_key	How to partition the table: by item name, by path name, or by GUID (see Table 65-4).
partition_guidi	If TRUE, build an index on GUID
partition_pathi	If TRUE, build an index on path name
partition_prop	If TRUE, partition the properties table

Usage Notes

The procedure executes like a DDL in that it auto-commits before and after its execution.

65.4.2 CREATESTORE Procedure

This procedure creates a new DBFS SFS store owned by the invoking session user.

Syntax

```
DBMS_DBFS_SFS.CREATESTORE  (
   store_name     IN     VARCHAR2,
   tbl_name       IN     VARCHAR2  DEFAULT NULL,
   tbs_name       in     VARCHAR2  DEFAULT NULL,
   use_bf         in     BOOLEAN   DEFAULT FALSE,
   stgopts        in     VARCHAR2 DEFAULT '');
```

Parameters

Table 65-7 CREATESTORE Procedure Parameters

Parameter	Description
store_name	Name of store

Table 65-7 (Cont.) CREATESTORE Procedure Parameters

Parameter	Description
store_type	STORETYPE_TAPE or STORETYPE_AMAZONS3
tbl_name	Placeholder for the store content cached in database
tbs_name	Named tablespace
use_bf	If TRUE, a BasicFile LOB is used; otherwise a SecureFile LOB is used.
stgopts	Currently non-operational, reserved for future use

65.4.3 DROPFILESYSTEM Procedures

This procedure drops the DBFS SFS store, purging all dictionary information associated with the store, and dropping the underlying file system table.

Syntax

```
DBMS_DBFS_SFS.DROPFILESYSTEM  (
   schema_name     IN      VARCHAR2 DEFAULT NULL,
   tbl_name        IN      INTEGER);

DBMS_DBFS_SFS.DROPFILESYSTEM  (
   store_name      IN      VARCHAR2);
```

Parameters

Table 65-8 DROPFILESYSTEM Procedure Parameters

Parameter	Description
schema_name	Name of schema
tbl_name	Name of tablespace
store_name	Name of store path

Usage Notes

- If the specified store table is registered by the current user, it will be unregistered from the content interface described in the DBMS_DBFS_CONTENT package and the POSIX metadata tables.
- Subsequent to unregistration, an attempt will be made to store table(s). This operation may fail if other users are currently using this store table.
- The user attempting a drop of the tables underlying the store must actually have the privileges to complete the drop operation (either as the owner of the tables, or as a sufficiently privileged user for cross-schema operations).
- The procedure executes like a DDL in that it auto-commits before and after its execution.

65.4.4 INITFS Procedure

This procedure initialize a POSIX file system store. The table associated with the POSIX file system store `store_name` is truncated and reinitialized with a single "root" directory entry.

Syntax

```
DBMS_DBFS_SFS.INITFS (
   store_name      IN      VARCHAR2);
```

Parameters

Table 65-9 INITFS Procedure Parameters

Parameter	Description
`store_name`	Name of store

Usage Notes

The procedure executes like a DDL in that it auto-commits before and after its execution.

66
DBMS_DDL

This package provides access to some SQL data definition language (DDL) statements from stored procedures. It also provides special administration operations that are not available as Data Definition Language statements (DDLs).

This chapter contains the following topics:

- Deprecated Subprograms
- Security Model
- Operational Notes
- Summary of DBMS_DDL Subprograms

66.1 DBMS_DDL Deprecated Subprograms

Oracle recommends that you do not use deprecated subprograms in new applications. Support for deprecated features is for backward compatibility only

The following subprograms are deprecated with release Oracle Database 10*g*:

- ALTER_COMPILE Procedure

66.2 DBMS_DDL Security Model

This package runs with the privileges of the calling user, rather than the package owner `SYS`.

66.3 **DBMS_DDL** Operational Notes

The `ALTER_COMPILE` procedure commits the current transaction, performs the operation, and then commits again.

66.4 Summary of DBMS_DDL Subprograms

This table lists the `DDBMS_DDL` subprograms in alphabetical order and briefly describes them

Table 66-1 ***DBMS_DDL Package Subprograms***

Subprogram	Description
ALTER_COMPILE Procedure	Compiles the PL/SQL object
ALTER_TABLE_NOT_REFERENCEABLE Procedure	Reorganizes object tables
ALTER_TABLE_REFERENCEABLE Procedure	Reorganizes object tables

Table 66-1 (Cont.) *DBMS_DDL Package Subprograms*

Subprogram	Description
CREATE_WRAPPED Procedures	Takes as input a single `CREATE OR REPLACE` statement that specifies creation of a PL/SQL package specification, package body, function, procedure, type specification or type body, generates a `CREATE OR REPLACE` statement with the PL/SQL source text obfuscated and executes the generated statement
IS_TRIGGER_FIRE_ONCE Function	Returns `TRUE` if the specified DML or DDL trigger is set to fire once. Otherwise, returns `FALSE`
SET_TRIGGER_FIRING_PROPERTY Procedures	Sets the specified DML or DDL trigger's firing property
WRAP Functions	Takes as input a `CREATE OR REPLACE` statement that specifies creation of a PL/SQL package specification, package body, function, procedure, type specification or type body and returns a `CREATE OR REPLACE` statement where the text of the PL/SQL unit has been obfuscated

66.4.1 ALTER_COMPILE Procedure

This procedure is equivalent to the SQL statement: `ALTER PROCEDURE|FUNCTION|PACKAGE [<schema>.] <name> COMPILE [BODY]`

> **Note:**
>
> This procedure is deprecated in Oracle Database 10*g* Release 2 (10.2) While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the DDL equivalent in a dynamic SQL statement.

Syntax

```
DBMS_DDL.ALTER_COMPILE (
   type             VARCHAR2,
   schema           VARCHAR2,
   name             VARCHAR2
   reuse_settings   BOOLEAN := FALSE);
```

Parameters

Table 66-2 ALTER_COMPILE Procedure Parameters

Parameter	Description
`type`	Must be either `PROCEDURE`, `FUNCTION`, `PACKAGE`, `PACKAGE BODY` or `TRIGGER`
`schema`	Schema name If `NULL`, then use current schema (case-sensitive)
`name`	Name of the object (case-sensitive)
`reuse_settings`	Indicates whether the session settings in the objects should be reused, or whether the current session settings should be adopted instead

Exceptions

Table 66-3 ALTER_COMPILE Procedure Exceptions

Exception	Description
ORA-20000:	Insufficient privileges or object does not exist
ORA-20001:	Remote object, cannot compile
ORA-20002:	Bad value for object type: should be either PACKAGE, PACKAGE BODY, PROCEDURE, FUNCTION, or TRIGGER

66.4.2 ALTER_TABLE_NOT_REFERENCEABLE Procedure

This procedure alters the given object table table_schema.table_name so it becomes not the default referenceable table for the schema affected_schema.

This is equivalent to SQL:

```
ALTER TABLE [<table_schema>.]<table_name> NOT REFERENCEABLE FOR <affected_schema>
```

which is currently not supported or available as a DDL statement.

Syntax

```
DBMS_DDL.ALTER_TABLE_NOT_REFERENCEABLE (
   table_name        IN           VARCHAR2,
   table_schema      IN  DEFAULT  NULL,
   affected_schema   IN  DEFAULT  NULL);
```

Parameters

Table 66-4 ALTER_TABLE_NOT_REFERENCEABLE Procedure Parameters

Parameter	Description
table_name	Name of the table to be altered. Cannot be a synonym. Must not be NULL. Case sensitive.
table_schema	Name of the schema owning the table to be altered. If NULL then the current schema is used. Case sensitive.
affected_schema	Name of the schema affected by this alteration. If NULL then the current schema is used. Case sensitive.

Usage Notes

This procedure simply reverts for the affected schema to the default table referenceable for PUBLIC; that is., it simply undoes the previous ALTER_TABLE_REFERENCEABLE call for this specific schema. The affected schema must a particular schema (cannot be PUBLIC).

The user that executes this procedure must own the table (that is, the schema is the same as the user), and the affected schema must be the same as the user.

If the user executing this procedure has ALTER ANY TABLE and SELECT ANY TABLE and DROP ANY TABLE privileges, the user doesn't have to own the table and the affected schema can be any valid schema.

66.4.3 ALTER_TABLE_REFERENCEABLE Procedure

This procedure alters the given object table `table_schema.table_name` so it becomes the referenceable table for the given schema `affected_schema`.

This is equivalent to SQL:

```
ALTER TABLE [<table_schema>.]<table_name>  REFERENCEABLE FOR <affected_schema>
```

which is currently not supported or available as a DDL statement.

Syntax

```
DBMS_DDL.ALTER_TABLE_REFERENCEABLE
   table_name       IN  VARCHAR2,
   table_schema     IN  DEFAULT  NULL,
   affected_schema  IN  DEFAULT  NULL);
```

Parameters

Table 66-5 ALTER_TABLE_REFERENCEABLE Procedure Parameters

Parameter	Description
`table_name`	Name of the table to be altered. Cannot be a synonym. Must not be `NULL`. Case sensitive.
`table_schema`	Name of the schema owning the table to be altered. If `NULL` then the current schema is used. Case sensitive.
`affected_schema`	Name of the schema affected by this alteration. If NULL then the current schema is used. Case sensitive.

Usage Notes

When you create an object table, it automatically becomes referenceable, unless you use the `OID AS` clause when creating the table. The `OID AS` clause makes it possible for you to create an object table and to assign to the new table the same EOID as another object table of the same type. After you create a new table using the `OID AS` clause, you end up with two object table with the same `EOID`; the new table is not referenceable, the original one is. All references that used to point to the objects in the original table still reference the same objects in the same original table.

If you execute this procedure on the new table, it makes the new table the referenceable table replacing the original one; thus, those references now point to the objects in the new table instead of the original table.

66.4.4 CREATE_WRAPPED Procedures

The procedure takes as input a single `CREATE OR REPLACE` statement that specifies creation of a PL/SQL package specification, package body, function, procedure, type specification or type body. It then generates a `CREATE OR REPLACE` statement with the PL/SQL source text

obfuscated and executes the generated statement. In effect, this procedure bundles together the operations of wrapping the text and creating the PL/SQL unit.

> **See Also:**
>
> WRAP Functions

This procedure has 3 overloads. Each of the three functions provides better performance than using a combination of individual WRAP Functions and `DBMS_SQL.PARSE` (or `EXECUTE IMMEDIATE`) calls. The different functionality of each form of syntax is presented with the definition.

Syntax

Is a shortcut for `EXECUTE IMMEDIATE SYS.DBMS_DDL.WRAP`(*ddl*):

```
DBMS_DDL.CREATE_WRAPPED (
   ddl      VARCHAR2);
```

Is a shortcut for `DBMS_SQL.PARSE`(`cursor`, `SYS.DBMS_DDL.WRAP` (`input`, `lb`, `ub`)):

```
DBMS_DDL.CREATE_WRAPPED(
   ddl      DBMS_SQL.VARCHAR2A,
   lb       PLS_INTEGER,
   ub       PLS_INTEGER);
```

Is a shortcut for `DBMS_SQL.PARSE`(`cursor`, `SYS.DBMS_DDL.WRAP` (`input`, `lb`, `ub`)):

```
DBMS_DDL.CREATE_WRAPPED(
   ddl      DBMS_SQL.VARCHAR2S,
   lb       PLS_INTEGER,
   ub       PLS_INTEGER);
```

Parameters

Table 66-6 CREATE_WRAPPED Procedure Parameters

Parameter	Description
`ddl`	A `CREATE OR REPLACE` statement that specifies creation of a PL/SQL package specification, package body, function, procedure, type specification or type body
`lb`	Lower bound for indices in the string table that specify the `CREATE OR REPLACE` statement
`ub`	Upper bound for indices in the string table that specify the `CREATE OR REPLACE` statement.

Usage Notes

- The `CREATE OR REPLACE` statement is executed with the privileges of the user invoking `DBMS_DDL.CREATE_WRAPPED`.
- Any PL/SQL code that attempts to call these interfaces should use the fully qualified package name `SYS.DBMS_DDL` to avoid the possibility that the name `DBMS_DDL` is captured by a locally-defined unit or by redefining the `DBMS_DDL` public synonym.

- Each invocation of any accepts only a single PL/SQL unit. By contrast, the PL/SQL `wrap` utility accepts a entire SQL*Plus file and obfuscates the PL/SQL units within the file leaving all other text as-is. These interfaces are intended to be used in conjunction with or as a replacement for PL/SQL's dynamic SQL interfaces (`EXECUTE IMMEDIATE` and `DBMS_SQL.PARSE`). Since these dynamic SQL interfaces only accept a single unit at a time (and do not understand the SQL*Plus "`/`" termination character), both the CREATE_WRAPPED Procedures and the WRAP Functions require input to be a single unit.

Exceptions

`ORA-24230`: If the input is not a `CREATE OR REPLACE` statement specifying a PL/SQL unit, exception `DBMS_DDL.MALFORMED_WRAP_INPUT` is raised.

Examples

```
DECLARE
    ddl VARCHAR2(32767);
BEGIN
    ddl := GENERATE_PACKAGE(...);
    SYS.DBMS_DDL.CREATE_WRAPPED(ddl); -- Instead of EXECUTE IMMEDIATE ddl
END;
```

66.4.5 IS_TRIGGER_FIRE_ONCE Function

This function returns `TRUE` if the specified DML or DDL trigger is set to fire once. Otherwise, it returns `FALSE`.

A fire once trigger fires in a user session but does not fire in the following cases:

- For changes made by a Streams apply process
- For changes made by executing one or more Streams apply errors using the `EXECUTE_ERROR` or `EXECUTE_ALL_ERRORS` procedure in the `DBMS_APPLY_ADM` package
- For changes made by a Logical Standby apply process

> **Note:**
>
> Only DML and DDL triggers can be fire once. All other types of triggers always fire.

> **See Also:**
>
> "SET_TRIGGER_FIRING_PROPERTY Procedures"

Syntax

```
DBMS_DDL.IS_TRIGGER_FIRE_ONCE
   trig_owner         IN  VARCHAR2,
   trig_name          IN  VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 66-7 IS_TRIGGER_FIRE_ONCE Function Parameters

Parameter	Description
trig_owner	Schema of trigger
trig_name	Name of trigger

66.4.6 SET_TRIGGER_FIRING_PROPERTY Procedures

This procedure sets the specified DML or DDL trigger's firing property whether or not the property is set for the trigger.

Use this procedure to control a DML or DDL trigger's firing property for changes:

- Applied by a Streams apply process
- Made by executing one or more Streams apply errors using the EXECUTE_ERROR or EXECUTE_ALL_ERRORS procedure in the DBMS_APPLY_ADM package.
- Applied by a Logical Standby apply process

Syntax

```
DBMS_DDL.SET_TRIGGER_FIRING_PROPERTY (
   trig_owner         IN  VARCHAR2,
   trig_name          IN  VARCHAR2,
   fire_once          IN  BOOLEAN);

DBMS_DDL.SET_TRIGGER_FIRING_PROPERTY (
   trig_owner         IN  VARCHAR2,
   trig_name          IN  VARCHAR2,
   property           IN INTEGER,
   setting            IN BOOLEAN);
```

Parameters

Table 66-8 SET_TRIGGER_FIRING_PROPERTY Procedure Parameters

Parameter	Description
trig_owner	Schema of the trigger to set
trig_name	Name of the trigger to set
fire_once	• If TRUE, the trigger is set to fire once. By default, the fire_once parameter is set to TRUE for DML and DDL triggers. • If FALSE, the trigger is set to always fire unless apply_server_only property is set to TRUE, which overrides fire_once property setting.
property	• DBMS_DDL.fire_once to set the fire_once property of the trigger • DBMS_DDL.apply_server_only to indicate whether trigger fires only in the context of SQL apply processes maintaining a logical standby database or Streams apply processes
setting	Value of property being set

Usage Notes

DML triggers created on a table have their fire-once property set to `TRUE`. In this case, the triggers only fire when the table is modified by an user process, and they are automatically disabled inside Oracle processes maintaining either a logical standby database (SQL Apply) or Oracle processes doing replication (Streams Apply) processes, and thus do not fire when a SQL Apply or a Streams Apply process modifies the table. There are two ways for a user to fire a trigger as a result of SQL Apply or a Streams Apply process making a change to a maintained table: (a) setting the fire-once property of a trigger to `FALSE`, which allows it fire both in the context of a user process or a SQL or Streams Apply process, or (b) by setting the apply-server-only property to `TRUE` and thus making the trigger fire only in the context of a SQL Apply or a Streams Apply process and not in the context of a user process.

- `FIRE_ONCE=TRUE`, `APPLY_SERVER_ONLY=FALSE`

  This is the default property setting for a DML trigger. The trigger only fires when user process modifies the base table.

- `FIRE_ONCE=TRUE` or `FALSE`, `APPLY_SERVER_ONLY=TRUE`

  The trigger only fires when SQL Apply or Streams Apply process modifies the base table. The trigger does not fire when a user process modifies the base table.Thus the apply-server-only property overrides the fire-once property of a trigger.

> **Note:**
>
> - If you dequeue an error transaction from the error queue and execute it without using the `DBMS_APPLY_ADM` package, then relevant changes resulting from this execution cause a trigger to fire, regardless of the trigger firing property.
> - Only DML and DDL triggers can be fire once. All other types of triggers always fire.

66.4.7 WRAP Functions

This function takes as input a single `CREATE OR REPLACE` statement that specifies creation of a PL/SQL package specification, package body, function, procedure, type specification or type body and returns a `CREATE OR REPLACE` statement where the text of the PL/SQL unit has been obfuscated.

The function has 3 overloads to allow for the different ways in which DDL statements can be generated dynamically and presented to `DBMS_SQL` or `EXECUTE IMMEDIATE`. The different functionality of each form of syntax is presented with the definition.

> **See Also:**
>
> CREATE_WRAPPED Procedures

Syntax

Provides basic functionality:

```
DBMS_DDL.WRAP(
   ddl      VARCHAR2)
  RETURN VARCHAR2;
```

Provides the same functionality as the first form, but allows for larger inputs. This function is intended to be used with the PARSE Procedures in the DBMS_SQL package and its argument list follows the convention of DBMS_SQL.PARSE:

```
DBMS_DDL.WRAP(
   ddl      DBMS_SQL.VARCHAR2S,
   lb       PLS_INTEGER,
   ub       PLS_INTEGER)
  RETURN DBMS_SQL.VARCHAR2S;
```

Provides the same functionality as the second form and is provided for compatibility with multiple forms of the PARSE Procedures in the DBMS_SQL package:

```
DBMS_DDL.WRAP(
   ddl      DBMS_SQL.VARCHAR2A,
   lb       PLS_INTEGER,
   ub       PLS_INTEGER)
  RETURN DBMS_SQL.VARCHAR2A;
```

Parameters

Table 66-9 WRAP Function Parameters

Parameter	Description
ddl	A CREATE OR REPLACE statement that specifies creation of a PL/SQL package specification, package body, function, procedure, type specification or type body
lb	Lower bound for indices in the string table that specify the CREATE OR REPLACE statement
ub	Upper bound for indices in the string table that specify the CREATE OR REPLACE statement.

Return Values

A CREATE OR REPLACE statement with the text obfuscated. In the case of the second and third form, the return value is a table of strings that need to be concatenated in order to construct the CREATE OR REPLACE string containing obfuscated source text.

Usage Notes

- Any PL/SQL code that attempts to call these interfaces should use the fully qualified package name SYS.DBMS_DDL to avoid the possibility that the name DBMS_DDL is captured by a locally-defined unit or by redefining the DBMS_DDL public synonym.
- Each invocation of any accepts only a single PL/SQL unit. By contrast, the PL/SQL wrap utility accepts a full SQL file and obfuscates the PL/SQL units within the file leaving all other text as-is. These interfaces are intended to be used in conjunction with or as a replacement for PL/SQL's dynamic SQL interfaces (EXECUTE IMMEDIATE and DBMS_SQL.PARSE). Since these dynamic SQL interfaces only accept a single unit at a time (and do not understand the SQL*Plus "/" termination character), both the CREATE_WRAPPED Procedures and the WRAP Functions require input to be a single unit.

Exceptions

ORA-24230: If the input is not a CREATE OR REPLACE statement specifying a PL/SQL unit, exception DBMS_DDL.MALFORMED_WRAP_INPUT is raised.

Examples

```
DECLARE
   ddl VARCHAR2(32767);
BEGIN
   ddl := GENERATE_PACKAGE(...);
EXECUTE IMMEDIATE SYS.DBMS_DDL.WRAP(ddl); -- Instead of EXECUTE IMMEDIATE ddl
END;
```

67
DBMS_DEBUG

`DBMS_DEBUG` is deprecated. Use `DBMS_DEBUG_JDWP` instead.

See DBMS_DEBUG_JDWP for more information.

`DBMS_DEBUG` is a PL/SQL interface to the PL/SQL debugger layer, Probe, in the Oracle server.

This package is primarily intended to implement server-side debuggers and it provides a way to debug server-side PL/SQL program units.

> **Note:**
>
> The term *program unit* refers to a PL/SQL program of any type (procedure, function, package, package body, trigger, anonymous block, object type, or object type body).

This chapter contains the following topics:

- Overview
- Constants
- Variables
- Exceptions
- Operational Notes
- Data Structures
- Summary of DBMS_DEBUG Subprograms

67.1 DBMS_DEBUG Overview

To debug server-side code, you must have two database sessions: one session to run the code in debug mode (the target session), and a second session to supervise the target session (the debug session).

The target session becomes available for debugging by making initializing calls with `DBMS_DEBUG`. This marks the session so that the PL/SQL interpreter runs in debug mode and generates debug events. As debug events are generated, they are posted from the session. In most cases, debug events require return notification: the interpreter pauses awaiting a reply.

Meanwhile, the debug session must also initialize itself using `DBMS_DEBUG`: This tells it which target session to supervise. The debug session may then call entry points in `DBMS_DEBUG` to read events that were posted from the target session and to communicate with the target session.

The following subprograms are run in the target session (the session that is to be debugged):

- SYNCHRONIZE Function
- DEBUG_ON Procedure

- DEBUG_OFF Procedure

DBMS_DEBUG does not provide an interface to the PL/SQL compiler, but it does depend on debug information optionally generated by the compiler. Without debug information, it is not possible to examine or modify the values of parameters or variables.

67.2 DBMS_DEBUG Constants

A breakpoint status may have the following value: breakpoint_status_unused—breakpoint is not in use.

Otherwise, the status is a mask of the following values:

- breakpoint_status_active—a line breakpoint
- breakpoint_status_disabled—breakpoint is currently disabled
- breakpoint_status_remote—a shadow breakpoint (a local representation of a remote breakpoint)

67.3 DBMS_DEBUG Variables

The DBMS_DEBUG uses the variables shown in the following table.

Table 67-1 DBMS_DEBUG Variables

Variable	Description
default_timeout	The timeout value (used by both sessions).The smallest possible timeout is 1 second. If this value is set to 0, then a large value (3600) is used.

67.4 DBMS_DEBUG Exceptions

These values are returned by the various functions called in the debug session (SYNCHRONIZE, CONTINUE, SET_BREAKPOINT, and so on). If PL/SQL exceptions worked across client/server and server/server boundaries, then these would all be exceptions rather than error codes.

Table 67-2 DBMS_DEBUG Exceptions

Status	Description
success	Normal termination

Statuses returned by GET_VALUE and SET_VALUE:

Table 67-3 DBMS_DEBUG Exceptions Returned by GET_VALUE and SET_VALUE

Status	Description
error_bogus_frame	No such entrypoint on the stack
error_no_debug_info	Program was compiled without debug symbols
error_no_such_object	No such variable or parameter
error_unknown_type	Debug information is unreadable

Table 67-3 (Cont.) DBMS_DEBUG Exceptions Returned by GET_VALUE and SET_VALUE

Status	Description
error_indexed_table	Returned by GET_VALUE if the object is a table, but no index was provided
error_illegal_index	No such element exists in the collection
error_nullcollection	Table is atomically NULL
error_nullvalue	Value is NULL

Statuses returned by SET_VALUE:

Table 67-4 DBMS_DEBUG Exceptions Returned by SET_VALUE

Status	Description
error_illegal_value	Constraint violation
error_illegal_null	Constraint violation
error_value_malformed	Unable to decipher the given value
error_other	Some other error
error_name_incomplete	Name did not resolve to a scalar

Statuses returned by the breakpoint functions:

Table 67-5 Statuses Returned by the Breakpoint Functions

Status	Description
error_no_such_breakpt	No such breakpoint
error_idle_breakpt	Cannot enable or disable an unused breakpoint
error_bad_handle	Unable to set breakpoint in given program (nonexistent or security violation)

General error codes (returned by many of the DBMS_DEBUG subprograms):

Table 67-6 DBMS_DEBUG Subprograms Error Codes

Status	Description
error_unimplemented	Functionality is not yet implemented
error_deferred	No program running; operation deferred
error_exception	An exception was raised in the DBMS_DEBUG or Probe packages on the server
error_communication	Some error other than a timeout occurred
error_timeout	Timout occurred

Table 67-7 illegal_init Exceptions

Exception	Description
illegal_init	DEBUG_ON was called prior to INITIALIZE

The following exceptions are raised by procedure SELF_CHECK:

Table 67-8 SELF_CHECK Procedure Exceptions

Exception	Description
pipe_creation_failure	Could not create a pipe
pipe_send_failure	Could not write data to the pipe
pipe_receive_failure	Could not read data from the pipe
pipe_datatype_mismatch	Datatype in the pipe was wrong
pipe_data_error	Data got garbled in the pipe

67.5 DBMS_DEBUG Operational Notes

There are two ways to ensure that debug information is generated: through a session switch, or through individual recompilation.

To set the session switch, enter the following statement:

```
ALTER SESSION SET PLSQL_DEBUG = true;
```

This instructs the compiler to generate debug information for the remainder of the session. It does not recompile any existing PL/SQL.

To generate debug information for existing PL/SQL code, use one of the following statements (the second recompiles a package or type body):

```
ALTER [PROCEDURE | FUNCTION | PACKAGE | TRIGGER | TYPE] <name> COMPILE DEBUG;
ALTER [PACKAGE | TYPE] <name> COMPILE DEBUG BODY;
```

Figure 67-1 and Figure 67-2 illustrate the flow of operations in the session to be debugged and in the debugging session.

Figure 67-1 Target Session

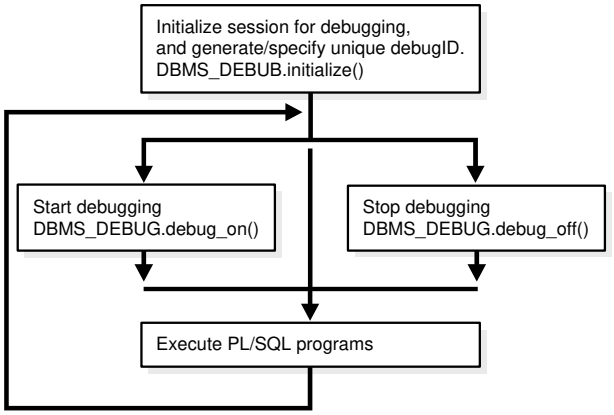

illustration: arpls001
release: 9
caption: Target Session
date: 1/29/01
platform: pc

Figure 67-2 Debug Session

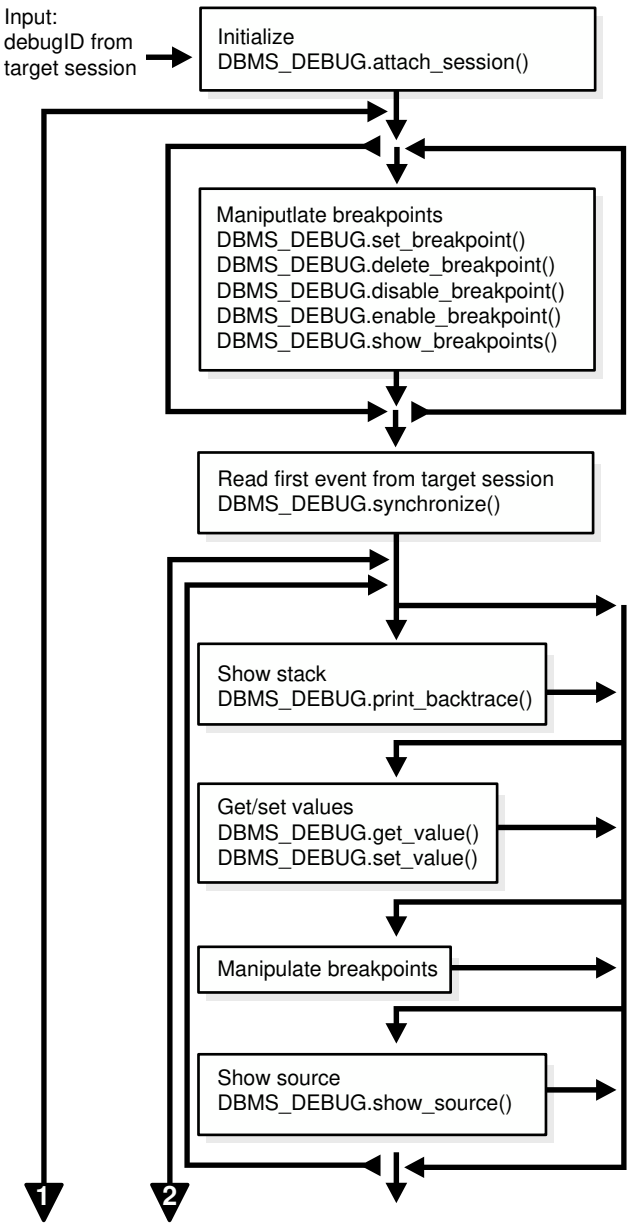

illustration: arpls003
release: 9
caption: Target Session
see arpls004
date: 1/30/01
platform: pc

Figure 67-3 Debug Session (Cont.)

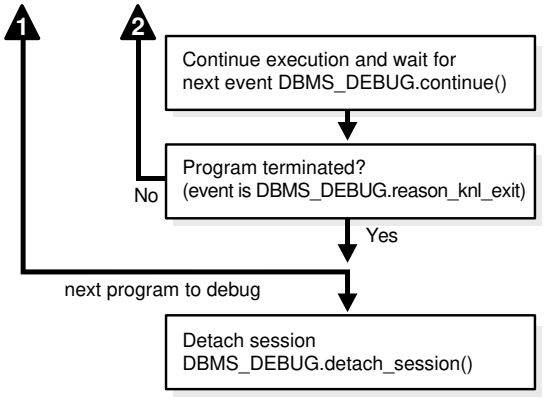

illustration: arpls004
release: 9
caption: Target Session
see arpls003
date: 1/30/01
platform: pc

Control of the Interpreter

The interpreter pauses execution at the following times:

1. At startup of the interpreter so any deferred breakpoints may be installed prior to execution.
2. At any line containing an enabled breakpoint.
3. At any line where an *interesting* event occurs. The set of interesting events is specified by the flags passed to `DBMS_DEBUG.CONTINUE` in the `breakflags` parameter.

Session Termination

There is no event for session termination. Therefore, it is the responsibility of the debug session to check and make sure that the target session has not ended. A call to `DBMS_DEBUG.SYNCHRONIZE` after the target session has ended causes the debug session to hang until it times out.

Deferred Operations

The diagram suggests that it is possible to set breakpoints prior to having a target session. This is true. In this case, Probe caches the breakpoint request and transmits it to the target session at first synchronization. However, if a breakpoint request is deferred in this fashion, then:

- `SET_BREAKPOINT` does not set the breakpoint number (it can be obtained later from `SHOW_BREAKPOINTS` if necessary).
- `SET_BREAKPOINT` does not validate the breakpoint request. If the requested source line does not exist, then an error silently occurs at synchronization, and no breakpoint is set.

Diagnostic Output

To debug Probe, there are *diagnostics* parameters to some of the calls in `DBMS_DEBUG`. These parameters specify whether to place diagnostic output in the RDBMS tracefile. If output to the RDBMS tracefile is disabled, these parameters have no effect.

Common Debug Session Sections

- Common Section
- Target Session
- Debug Session Section

Common Section

The following subprograms may be called in either the target or the debug session:

- PROBE_VERSION Procedure
- SELF_CHECK Procedure
- SET_TIMEOUT Function

Target Session

The following subprograms may be called only in the target session:

- INITIALIZE Function
- DEBUG_ON Procedure
- SET_TIMEOUT_BEHAVIOUR Procedure
- GET_TIMEOUT_BEHAVIOUR Function

Debug Session Section

The following subprograms should be run in the debug session only:

- ATTACH_SESSION Procedure
- SYNCHRONIZE Function
- SHOW_FRAME_SOURCE Procedure

- SHOW_SOURCE Procedures
- GET_MORE_SOURCE Procedure
- PRINT_BACKTRACE Procedure
- CONTINUE Function
- SET_BREAKPOINT Function
- DELETE_BREAKPOINT Function
- SET_OER_BREAKPOINT Function
- DELETE_OER_BREAKPOINT Function
- ENABLE_BREAKPOINT Function
- DISABLE_BREAKPOINT Function
- SHOW_BREAKPOINTS Procedures
- SET_VALUE Functionn
- GET_VALUE Function
- TARGET_PROGRAM_RUNNING Procedure
- DETACH_SESSION Procedure
- GET_RUNTIME_INFO Function
- PRINT_INSTANTIATIONS Procedure
- PING Procedure
- GET_LINE_MAP Function
- GET_RUNTIME_INFO Function
- GET_INDEXES Function
- EXECUTE Procedure

OER Breakpoints

Exceptions that are declared in PL/SQL programs are known as user-defined exceptions. In addition, there are Oracle Errors (OERs) that are returned from the Oracle kernel. To tie the two mechanisms together, PL/SQL provides the `exception_init` pragma that turns a user-defined exception into an OER, so that a PL/SQL handler may be used for it, and so that the PL/SQL engine can return OERs to the Oracle kernel. As of the current release, the only information available about an OER is its number. If two user-defined exceptions are exception_init'd to the same OER, they are indistinguishable.

Namespaces

Program units on the server reside in different namespaces. When setting a breakpoint, specify the desired namespace.

1. `Namespace_cursor` contains cursors (anonymous blocks).
2. `Namespace_pgkspec_or_toplevel` contains:
   - Package specifications.
   - Procedures and functions that are not nested inside other packages, procedures, or functions.
   - Object types.

3. `Namespace_pkg_body` contains package bodies and type bodies.
4. `Namespace_trigger` contains triggers.

Libunit Types

These values are used to disambiguate among objects in a given namespace. These constants are used in `PROGRAM_INFO` when Probe is giving a stack backtrace.

- `LibunitType_cursor`
- `LibunitType_procedure`
- `LibunitType_function`
- `LibunitType_package`
- `LibunitType_package_body`
- `LibunitType_trigger`
- `LibunitType_Unknown`

Breakflags

These are values to use for the `breakflags` parameter to `CONTINUE`, in order to tell Probe what events are of interest to the client. These flags may be combined.

Value	Description
`break_next_line`	Break at next source line (step over calls)
`break_any_call`	Break at next source line (step into calls)
`break_any_return`	Break after returning from current entrypoint (skip over any entrypoints called from the current routine)
`break_return`	Break the next time an entrypoint gets ready to return. (This includes entrypoints called from the current one. If interpreter is running `Proc1`, which calls `Proc2`, then `break_return` stops at the end of `Proc2`.)
`break_exception`	Break when an exception is raised
`break_handler`	Break when an exception handler is executed
`abort_execution`	Stop execution and force an 'exit' event as soon as `DBMS_DEBUG.CONTINUE` is called.

Information Flags

These are flags which may be passed as the `info_requested` parameter to `SYNCHRONIZE`, `CONTINUE`, and `GET_RUNTIME_INFO`.

Flag	Description
`info_getStackDepth`	Get the current depth of the stack
`info_getBreakpoint`	Get the breakpoint number
`info_getLineinfo`	Get program unit information

Reasons for Suspension

After `CONTINUE` is run, the program either runs to completion or breaks on some line.

Reason	Description
reason_none	-
reason_interpreter_starting	Interpreter is starting
reason_breakpoint	Hit a breakpoint
reason_enter	Procedure entry
reason_return	Procedure is about to return
reason_finish	Procedure is finished
reason_line	Reached a new line
reason_interrupt	An interrupt occurred
reason_exception	An exception was raised
reason_exit	Interpreter is exiting (old form)
reason_knl_exit	Kernel is exiting
reason_handler	Start exception-handler
reason_timeout	A timeout occurred
reason_instantiate	Instantiation block
reason_abort	Interpreter is aborting

67.6 DBMS_DEBUG Data Structures

The DBMS_DEBUG package defines RECORD types and TABLE types.

RECORD Types

- BREAKPOINT_INFO Record Type
- PROGRAM_INFO Record Type
- RUNTIME_INFO Record Type

TABLE Types

- BACKTRACE_TABLE Table Type
- BREAKPOINT_TABLE Table Type
- INDEX_TABLE Table Type
- VC2_TABLE Table Type

67.6.1 BREAKPOINT_INFO Record Type

This type gives information about a breakpoint, such as its current status and the program unit in which it was placed.

Syntax

```
TYPE breakpoint_info IS RECORD (
   name        VARCHAR2(30),
   owner       VARCHAR2(30),
   dblink      VARCHAR2(30),
```

```
  line#       BINARY_INTEGER,
  libunittype BINARY_INTEGER,
  status      BINARY_INTEGER);
```

Fields

Table 67-9 BREAKPOINT_INFO Fields

Field	Description
`name`	Name of the program unit
`owner`	Owner of the program unit
`dblink`	Database link, if remote
`line#`	Line number
`libunittype`	`NULL`, unless this is a nested procedure or function
`status`	See Constants for values of `breakpoint_status_*`

67.6.2 PROGRAM_INFO Record Type

The `PROGRAM_INFO` record type of the `DBMS_DEBG` package specifies a program location. It is a line number in a program unit.

This is used for stack backtraces and for setting and examining breakpoints. The read-only fields are currently ignored by Probe for breakpoint operations. They are set by Probe only for stack backtraces.

Syntax

```
TYPE program_info IS RECORD(
    -- The following fields are used when setting a breakpoint
    namespace        BINARY_INTEGER,
    name             VARCHAR2(30),
    owner            VARCHAR2(30),
    dblink           VARCHAR2(30),
    line#            BINARY_INTEGER,
    -- Read-only fields (set by Probe when doing a stack backtrace)
    libunittype      BINARY_INTEGER,
    entrypointname   VARCHAR2(30));
```

Fields

Table 67-10 PROGRAM_INFO Fields

Field	Description
`namespace`	See DBMS_DEBUG Operational Notes for more information about namespaces.
`name`	Name of the program unit
`owner`	Owner of the program unit
`dblink`	Database link, if remote
`line#`	Line number
`libunittype`	`A read-only field, NULL`, unless this is a nested procedure or function

Table 67-10 (Cont.) PROGRAM_INFO Fields

Field	Description
`entrypointname`	A read-only field, to disambiguate among objects that share the same namespace (for example, procedure and package specifications). See DBMS_DEBUG Operational Notes for more information about the libunit types.

67.6.3 RUNTIME_INFO Record Type

This type gives context information about the running program.

Syntax

```
TYPE runtime_info IS RECORD(
    line#            BINARY_INTEGER,
    terminated       binary_integer,
    breakpoint       binary_integer,
    stackdepth       BINARY_INTEGER,
    interpreterdepth BINARY_INTEGER,
    reason           BINARY_INTEGER,
    program          program_info);
```

Fields

Table 67-11 RUNTIME_INFO Fields

Field	Description
`line#`	Duplicate of `program.line#`
`terminated`	Whether the program has terminated
`breakpoint`	Breakpoint number
`stackdepth`	Number of frames on the stack
`interpreterdepth`	[A reserved field]
`reason`	Reason for suspension
`program`	Source location

67.6.4 BACKTRACE_TABLE Table Type

This type is used by `PRINT_BACKTRACE`.

Syntax

```
TYPE backtrace_table IS TABLE OF program_info INDEX BY BINARY_INTEGER;
```

67.6.5 BREAKPOINT_TABLE Table Type

This type is used by `SHOW_BREAKPOINTS`.

Syntax

```
TYPE breakpoint_table IS TABLE OF breakpoint_info INDEX BY BINARY_INTEGER;
```

67.6.6 INDEX_TABLE Table Type

This type is used by `GET_INDEXES` to return the available indexes for an indexed table.

Syntax

```
TYPE index_table IS table of BINARY_INTEGER INDEX BY BINARY_INTEGER;
```

67.6.7 VC2_TABLE Table Type

This type is used by `SHOW_SOURCE`.

Syntax

```
TYPE vc2_table IS TABLE OF VARCHAR2(90) INDEX BY BINARY_INTEGER;
```

67.7 Summary of DBMS_DEBUG Subprograms

This table lists the `DBMS_DEBUG` subprograms in alphabetical order and briefly describes them.

Table 67-12 DBMS_DEBUG Package Subprograms

Subprogram	Description
ATTACH_SESSION Procedure	Notifies the debug session about the target debugID
CONTINUE Function	Continues execution of the target program
DEBUG_OFF Procedure	Turns debug-mode off
DEBUG_ON Procedure	Turns debug-mode on
DELETE_BREAKPOINT Function	Deletes a breakpoint
DELETE_OER_BREAKPOINT Function	Deletes an OER breakpoint
DETACH_SESSION Procedure	Stops debugging the target program
DISABLE_BREAKPOINT Function	Disables a breakpoint
ENABLE_BREAKPOINT Function	Activates an existing breakpoint
EXECUTE Procedure	Executes SQL or PL/SQL in the target session
GET_INDEXES Function	Returns the set of indexes for an indexed table
GET_MORE_SOURCE Procedure	Provides additional source in the event of buffer overflow when using SHOW_SOURCE
GET_LINE_MAP Function	Returns information about line numbers in a program unit
GET_RUNTIME_INFO Function	Returns information about the current program
GET_TIMEOUT_BEHAVIOUR Function	Returns the current timeout behavior
GET_VALUE Function	Gets a value from the currently-running program
INITIALIZE Function	Sets debugID in target session
PING Procedure	Pings the target session to prevent it from timing out

Table 67-12 (Cont.) DBMS_DEBUG Package Subprograms

Subprogram	Description
PRINT_BACKTRACE Procedure	Prints a stack backtrace
PRINT_INSTANTIATIONS Procedure	Prints a stack backtrace
PROBE_VERSION Procedure	Returns the version number of `DBMS_DEBUG` on the server
SELF_CHECK Procedure	Performs an internal consistency check
SET_BREAKPOINT Function	Sets a breakpoint in a program unit
SET_OER_BREAKPOINT Function	Sets an OER breakpoint
SET_TIMEOUT Function	Sets the timeout value
SET_TIMEOUT_BEHAVIOUR Procedure	Tells Probe what to do with the target session when a timeout occurs
SET_VALUE Function	Sets a value in the currently-running program
SHOW_BREAKPOINTS Procedures	Returns a listing of the current breakpoints
SHOW_FRAME_SOURCE Procedure	Fetches the frame source
SHOW_SOURCE Procedures	Fetches program source
SYNCHRONIZE Function	Waits for program to start running
TARGET_PROGRAM_RUNNING Procedure	Returns `TRUE` if the target session is currently executing a stored procedure, or `FALSE` if it is not

67.7.1 ATTACH_SESSION Procedure

This procedure notifies the debug session about the target program.

Syntax

```
DBMS_DEBUG.ATTACH_SESSION (
   debug_session_id  IN VARCHAR2,
   diagnostics       IN BINARY_INTEGER := 0);
```

Parameters

Table 67-13 ATTACH_SESSION Procedure Parameters

Parameter	Description
`debug_session_id`	Debug ID from a call to `INITIALIZE` in target session
`diagnostics`	Generate diagnostic output if nonzero

67.7.2 CONTINUE Function

This function passes the given breakflags (a mask of the events that are of interest) to Probe in the target process. It tells Probe to continue execution of the target process, and it waits until the target process runs to completion or signals an event.

If `info_requested` is not `NULL`, then calls `GET_RUNTIME_INFO`.

Syntax

```
DBMS_DEBUG.CONTINUE (
   run_info       IN OUT runtime_info,
   breakflags     IN     BINARY_INTEGER,
   info_requested IN     BINARY_INTEGER := NULL)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-14 CONTINUE Function Parameters

Parameter	Description
`run_info`	Information about the state of the program
`breakflags`	Mask of events that are of interest (see the discussion about break flags under DBMS_DEBUG Operational Notes)
`info_requested`	Which information should be returned in `run_info` when the program stops (see the discussion of information flags under DBMS_DEBUG Operational Notes)

Return Values

Table 67-15 CONTINUE Function Return Values

Return	Description
`success`	
`error_timeout`	Timed out before the program started running
`error_communication`	Other communication error

67.7.3 DEBUG_OFF Procedure

This procedure notifies the target session that debugging should no longer take place in that session. It is not necessary to call this function before ending the session.

> **WARNING:**
>
> There must be a debug session waiting if immediate is `TRUE`.

Syntax

```
DBMS_DEBUG.DEBUG_OFF;
```

Usage Notes

The server does not handle this entrypoint specially. Therefore, it attempts to debug this entrypoint.

67.7.4 DEBUG_ON Procedure

This procedure marks the target session so that all PL/SQL is run in debug mode. This must be done before any debugging can take place.

Syntax

```
DBMS_DEBUG.DEBUG_ON (
   no_client_side_plsql_engine BOOLEAN := TRUE,
   immediate                   BOOLEAN := FALSE);
```

Parameters

Table 67-16 DEBUG_ON Procedure Parameters

Parameter	Description
`no_client_side_plsql_engine`	Should be left to its default value unless the debugging session is taking place from a client-side PL/SQL engine
`immediate`	If this is `TRUE`, then the interpreter immediately switches itself into debug-mode, instead of continuing in regular mode for the duration of the call.

67.7.5 DELETE_BREAKPOINT Function

This function deletes a breakpoint.

Syntax

```
DBMS_DEBUG.DELETE_BREAKPOINT (
   breakpoint IN BINARY_INTEGER)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-17 DELETE_BREAKPOINT Function Parameters

Parameter	Description
`breakpoint`	Breakpoint number from a previous call to `SET_BREAKPOINT`

Return Values

Table 67-18 DELETE_BREAKPOINT Function Return Values

Return	Description
`success`	

Table 67-18 (Cont.) DELETE_BREAKPOINT Function Return Values

Return	Description
`error_no_such_breakpt`	No such breakpoint exists
`error_idle_breakpt`	Cannot delete an unused breakpoint
`error_stale_breakpt`	The program unit was redefined since the breakpoint was set

67.7.6 DELETE_OER_BREAKPOINT Function

This function deletes an OER breakpoint.

Syntax

```
DBMS_DEBUG.DELETE_OER_BREAKPOINT (
   oer  IN PLS_INTEGER)
RETURN PLS_INTEGER;
```

Parameters

Table 67-19 DELETE_OER_BREAKPOINT Function Parameters

Parameter	Description
`oer`	The OER (positive 4-byte number) to delete

67.7.7 DETACH_SESSION Procedure

This procedure stops debugging the target program.

This procedure may be called at any time, but it does not notify the target session that the debug session is detaching itself, and it does not terminate execution of the target session. Therefore, care should be taken to ensure that the target session does not hang itself.

Syntax

```
DBMS_DEBUG.DETACH_SESSION;
```

67.7.8 DISABLE_BREAKPOINT Function

This function makes an existing breakpoint inactive but leaves it in place.

Syntax

```
DBMS_DEBUG.DISABLE_BREAKPOINT (
   breakpoint IN BINARY_INTEGER)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-20 DISABLE_BREAKPOINT Function Parameters

Parameter	Description
breakpoint	Breakpoint number from a previous call to SET_BREAKPOINT

Return Values

Table 67-21 DISABLE_BREAKPOINT Function Return Values

Returns	Description
success	
error_no_such_breakpt	No such breakpoint exists
error_idle_breakpt	Cannot disable an unused breakpoint

67.7.9 ENABLE_BREAKPOINT Function

This function is the reverse of disabling. This enables a previously disabled breakpoint.

Syntax

```
DBMS_DEBUG.ENABLE_BREAKPOINT (
   breakpoint IN BINARY_INTEGER)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-22 ENABLE_BREAKPOINT Function Parameters

Parameter	Description
breakpoint	Breakpoint number from a previous call to SET_BREAKPOINT

Return Values

Table 67-23 ENABLE_BREAKPOINT Function Return Values

Return	Description
success	Success
error_no_such_breakpt	No such breakpoint exists
error_idle_breakpt	Cannot enable an unused breakpoint

67.7.10 EXECUTE Procedure

This procedure executes SQL or PL/SQL code in the target session. The target session is assumed to be waiting at a breakpoint (or other event). The call to `DBMS_DEBUG.EXECUTE` occurs in the debug session, which then asks the target session to execute the code.

Syntax

```
DBMS_DEBUG.EXECUTE (
   what         IN VARCHAR2,
   frame#       IN BINARY_INTEGER,
   bind_results IN BINARY_INTEGER,
   results      IN OUT NOCOPY dbms_debug_vc2coll,
   errm         IN OUT NOCOPY VARCHAR2);
```

Parameters

Table 67-24 EXECUTE Procedure Parameters

Parameter	Description
`what`	SQL or PL/SQL source to execute
`frame#`	The context in which to execute the code. Only -1 (global context) is supported at this time.
`bind_results`	Whether the source wants to bind to `results` in order to return values from the target session: 0 = No 1 = Yes
`results`	Collection in which to place results, if `bind_results` is not 0
`errm`	Error message, if an error occurred; otherwise, `NULL`

Examples

Example 1

This example executes a SQL statement. It returns no results.

```
DECLARE
   coll sys.dbms_debug_vc2coll; -- results (unused)
   errm VARCHAR2(100);
BEGIN
   dbms_debug.execute('insert into emp(ename,empno,deptno) ' ||
                      'values(''LJE'', 1, 1)',
                      -1, 0, coll, errm);
END;
```

Example 2

This example executes a PL/SQL block, and it returns no results. The block is an autonomous transaction, which means that the value inserted into the table becomes visible in the debug session.

```
DECLARE
   coll sys.dbms_debug_vc2coll;
   errm VARCHAR2(100);
BEGIN
   dbms_debug.execute(
```

```
       'DECLARE PRAGMA autonomous_transaction; ' ||
       'BEGIN ' ||
       '   insert into emp(ename, empno, deptno) ' ||
       '   values(''LJE'', 1, 1); ' ||
       ' COMMIT; ' ||
       'END;',
       -1, 0, coll, errm);
END;
```

Example 3

This example executes a PL/SQL block, and it returns some results.

```
DECLARE
   coll sys.dbms_debug_vc2coll;
   errm VARCHAR2(100);
BEGIN
   dbms_debug.execute(
      'DECLARE ' ||
      '   pp SYS.dbms_debug_vc2coll := SYS.dbms_debug_vc2coll(); ' ||
      '   x  PLS_INTEGER; ' ||
      '   i  PLS_INTEGER := 1; ' ||
      'BEGIN ' ||
      '   SELECT COUNT(*) INTO x FROM emp; ' ||
      '   pp.EXTEND(x * 6); ' ||
      '   FOR c IN (SELECT * FROM emp) LOOP ' ||
      '      pp(i) := ''Ename: '' || c.ename; i := i+1; ' ||
      '      pp(i) := ''Empno: '' || c.empno; i := i+1; ' ||
      '      pp(i) := ''Job:   '' || c.job;   i := i+1; ' ||
      '      pp(i) := ''Mgr:   '' || c.mgr;   i := i+1; ' ||
      '      pp(i) := ''Sal:   '' || c.sal;   i := i+1; ' ||
      '      pp(i) := null;                   i := i+1; ' ||
      '   END LOOP; ' ||
      '   :1 := pp;' ||
      'END;',
       -1, 1, coll, errm);
   each := coll.FIRST;
   WHILE (each IS NOT NULL) LOOP
      dosomething(coll(each));
      each := coll.NEXT(each);
   END LOOP;
END;
```

67.7.11 GET_INDEXES Function

Given a name of a variable or parameter, this function returns the set of its indexes, if it is an indexed table. An error is returned if it is not an indexed table.

Syntax

```
DBMS_DEBUG.GET_INDEXES (
   varname   IN  VARCHAR2,
   frame#    IN  BINARY_INTEGER,
   handle    IN  program_info,
   entries   OUT index_table)
RETURN BINARY_INTEGER;
```

Parameters

Table 67-25 GET_INDEXES Function Parameters

Parameter	Description
varname	Name of the variable to get index information about
frame#	Number of frame in which the variable or parameter resides; NULL for a package variable
handle	Package description, if object is a package variable
entries	1-based table of the indexes: if non-NULL, then entries(1) contains the first index of the table, entries(2) contains the second index, and so on.

Return Values

Table 67-26 GET_INDEXES Function Return Values

Return	Description
error_no_such_object	One of the following: - The package does not exist - The package is not instantiated - The user does not have privileges to debug the package - The object does not exist in the package

67.7.12 GET_MORE_SOURCE Procedure

When the source does not fit in the buffer provided by the SHOW_SOURCE Procedure version which produced a formatted buffer, this procedure provides additional source.

Syntax

```
DBMS_DEBUG.GET_MORE_SOURCE (
   buffer          IN OUT VARCHAR2,
   buflen          IN BINARY_INTEGER,
   piece#          IN BINARY_INTEGER);
```

Parameters

Table 67-27 GET_MORE_SOURCE Procedure Parameters

Parameter	Description
buffer	The buffer
buflen	The length of the buffer
piece#	A value between 2 and the value returned in the parameter pieces from the call to the relevant version of the SHOW_SOURCE Procedures

Usage Notes

This procedure should be called only after the version of SHOW_SOURCE that returns a formatted buffer.

Related Topics

- SHOW_SOURCE Procedures
  The procedure gets the source code. There are two overloaded `SHOW_SOURCE` procedures.

67.7.13 GET_LINE_MAP Function

This function finds line and entrypoint information about a program so that a debugger can determine the source lines at which it is possible to place breakpoints.

Syntax

```
DBMS_DEBUG.GET_LINE_MAP (
   program                IN   program_info,
   maxline                OUT  BINARY_INTEGER,
   number_of_entry_points OUT  BINARY_INTEGER,
   linemap                OUT  RAW)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-28 GET_LINE_MAP Function Parameters

Parameter	Description
`program`	A top-level program unit (procedure / package / function / package body, and so on). Its `Namespace`, `Name`, and `Owner` fields must be initialized, the remaining fields are ignored.
`maxline`	The largest source code line number in 'program'
`number_of_entry_points`	The number of subprograms in 'program'
`linemap`	A bitmap representing the executable lines of 'program'. If line number N is executable, bit number N MOD 8 will be set to 1 at linemap position N / 8. The length of returned linemap is either `maxline` divided by 8 (plus one if `maxline` MOD 8 is not zero) or 32767 in the unlikely case of `maxline` being larger than 32767 * 8.

Return Values

Table 67-29 GET_LINE_MAP Function Return Values

Return	Description
`success`	A successful completion
`error_no_debug_info`	The program unit exists, but has no debug info
`error_bad_handle`	No such program unit exists

67.7.14 GET_RUNTIME_INFO Function

This function returns information about the current program. It is only needed if the `info_requested` parameter to `SYNCHRONIZE` or `CONTINUE` was set to `0`.

> **Note:**
>
> This is currently only used by client-side PL/SQL.

Syntax

```
DBMS_DEBUG.GET_RUNTIME_INFO (
   info_requested  IN  BINARY_INTEGER,
   run_info        OUT runtime_info)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-30 GET_RUNTIME_INFO Function Parameters

Parameter	Description
`info_requested`	Which information should be returned in `run_info` when the program stops (see DBMS_DEBUG Operational Notes for information about information flags)
`run_info`	Information about the state of the program

67.7.15 GET_TIMEOUT_BEHAVIOUR Function

This procedure returns the current timeout behavior. This call is made in the target session.

Syntax

```
DBMS_DEBUG.GET_TIMEOUT_BEHAVIOUR
 RETURN BINARY_INTEGER;
```

Parameters

Table 67-31 GET_TIMEOUT_BEHAVIOUR Function Parameters

Parameter	Description
`oer`	The OER (a 4-byte positive number)

Return Values

Table 67-32 GET_TIMEOUT_BEHAVIOUR Function Return Values

Return	Description
`success`	A successful completion

Information Flags

```
info_getOerInfo CONSTANT PLS_INTEGER:= 32;
```

Usage Notes

Less functionality is supported on OER breakpoints than on code breakpoints. In particular, note that:

- No "breakpoint number" is returned - the number of the OER is used instead. Thus it is impossible to set duplicate breakpoints on a given OER (it is a no-op).
- It is not possible to disable an OER breakpoint (although clients are free to simulate this by deleting it).
- OER breakpoints are deleted using `delete_oer_breakpoint`.

67.7.16 GET_VALUE Function

This function gets a value from the currently-running program. There are two overloaded `GET_VALUE` functions.

Syntax

```
DBMS_DEBUG.GET_VALUE (
   variable_name  IN  VARCHAR2,
   frame#         IN  BINARY_INTEGER,
   scalar_value   OUT VARCHAR2,
   format         IN  VARCHAR2 := NULL)
RETURN BINARY_INTEGER;
```

Parameters

Table 67-33 GET_VALUE Function Parameters

Parameter	Description
`variable_name`	Name of the variable or parameter
`frame#`	Frame in which it lives; 0 means the current procedure
`scalar_value`	Value
`format`	Optional date format to use, if meaningful

Return Values

Table 67-34 GET_VALUE Function Return Values

Return	Description
`success`	A successful completion
`error_bogus_frame`	Frame does not exist
`error_no_debug_info`	Entrypoint has no debug information
`error_no_such_object`	`variable_name` does not exist in `frame#`
`error_unknown_type`	The type information in the debug information is illegible
`error_nullvalue`	Value is `NULL`

Table 67-34 (Cont.) GET_VALUE Function Return Values

Return	Description
error_indexed_table	The object is a table, but no index was provided

This form of GET_VALUE is for fetching package variables. Instead of a frame#, it takes a handle, which describes the package containing the variable.

Syntax

```
DBMS_DEBUG.GET_VALUE (
   variable_name  IN  VARCHAR2,
   handle         IN  program_info,
   scalar_value   OUT VARCHAR2,
   format         IN  VARCHAR2 := NULL)
RETURN BINARY_INTEGER;
```

Parameters

Table 67-35 GET_VALUE Function Parameters

Parameter	Description
variable_name	Name of the variable or parameter
handle	Description of the package containing the variable
scalar_value	Value
format	Optional date format to use, if meaningful

Return Values

Table 67-36 GET_VALUE Function Return Values

Return	Description
error_no_such_object	One of the following: - Package does not exist - Package is not instantiated - User does not have privileges to debug the package - Object does not exist in the package
error_indexed_table	The object is a table, but no index was provided

Examples

This example illustrates how to get the value with a given package PACK in schema SCOTT, containing variable VAR:

```
DECLARE
   handle     dbms_debug.program_info;
   resultbuf  VARCHAR2(500);
   retval     BINARY_INTEGER;
BEGIN
   handle.Owner     := 'SCOTT';
   handle.Name      := 'PACK';
```

```
   handle.namespace := dbms_debug.namespace_pkgspec_or_toplevel;
   retval           := dbms_debug.get_value('VAR', handle, resultbuf, NULL);
END;
```

67.7.17 INITIALIZE Function

This function initializes the target session for debugging.

Syntax

```
DBMS_DEBUG.INITIALIZE (
   debug_session_id  IN VARCHAR2       := NULL,
   diagnostics       IN BINARY_INTEGER := 0)
  RETURN VARCHAR2;
```

Parameters

Table 67-37 INITIALIZE Function Parameters

Parameter	Description
debug_session_id	Name of session ID. If NULL, then a unique ID is generated.
diagnostics	Indicates whether to dump diagnostic output to the tracefile: 0 = (default) no diagnostics 1 = print diagnostics

Return Values

The newly-registered debug session ID (debugID)

Usage Notes

You cannot use DBMS_DEBUG and the JDWP-based debugging interface simultaneously. This call will either fail with an ORA-30677 error if the session is currently being debugged with the JDWP-based debugging interface or, if the call succeeds, any further use of the JDWP-based interface to debug this session will be disallowed.

Calls to DBMS_DEBUG will succeed only if either the caller or the specified debug role carries the DEBUG CONNECT SESSION privilege. Failing that, an ORA-1031 error will be raised. Other exceptions are also possible if a debug role is specified but the password does not match, or if the calling user has not been granted the role, or the role is application-enabled and this call does not originate from within the role-enabling package.

The CREATE ANY PROCEDURE privilege does not affect the visibility of routines through the debugger. A privilege DEBUG for each object has been introduced with a corresponding DEBUG ANY PROCEDURE variant. These are required in order to see routines owned by users other than the session's login user.

Authentication of the debug role and the check for DEBUG CONNECT SESSION privilege will be done in the context of the caller to this routine. If the caller is a definer's rights routine or has been called from one, only privileges granted to the defining user, the debug role, or PUBLIC will be used to check for DEBUG CONNECT SESSION. If this call is from within a definer's rights routine, the debug role, if specified, must be one that has been granted to that definer, but it need not also have been granted to the session login user or be enabled in the calling session at the time the call is made.

The checks made by the debugger after this call is made looking for the DEBUG privilege on individual procedures will be done in the context of the session's login user, the roles that were enabled at session level at the moment this call was made (even if those roles were not available within a definer's rights environment of the call), and the debug role.

67.7.18 PING Procedure

This procedure pings the target session to prevent it from timing out. Use this procedure when execution is suspended in the target session, for example at a breakpoint.

If the timeout_behaviour is set to retry_on_timeout then this procedure is not necessary.

Syntax

```
DBMS_DEBUG.PING;
```

Exceptions

Oracle will display the no_target_program exception if there is no target program or if the target session is not currently waiting for input from the debug session.

Usage Notes

Timeout options for the target session are registered with the target session by calling set_timeout_behaviour:

- retry_on_timeout - Retry. Timeout has no effect. This is like setting the timeout to an infinitely large value.
- continue_on_timeout - Continue execution, using same event flags.
- nodebug_on_timeout - Turn debug-mode OFF (in other words, call debug_off) and then continue execution. No more events will be generated by this target session unless it is re-initialized by calling debug_on.
- abort_on_timeout - Continue execution, using the abort_execution flag, which should cause the program to terminate immediately. The session remains in debug-mode.

```
retry_on_timeout CONSTANT BINARY_INTEGER:= 0;

continue_on_timeout CONSTANT BINARY_INTEGER:= 1;

nodebug_on_timeout CONSTANT BINARY_INTEGER:= 2;

abort_on_timeout CONSTANT BINARY_INTEGER:= 3;
```

67.7.19 PRINT_BACKTRACE Procedure

This procedure prints a backtrace listing of the current execution stack. This should only be called if a program is currently running.

There are two overloaded PRINT_BACKTRACE procedures.

Syntax

```
DBMS_DEBUG.PRINT_BACKTRACE (
  listing IN OUT VARCHAR2);

DBMS_DEBUG.PRINT_BACKTRACE (
  backtrace OUT backtrace_table);
```

Parameters

Table 67-38 PRINT_BACKTRACE Procedure Parameters

Parameter	Description
`listing`	A formatted character buffer with embedded newlines
`backtrace`	1-based indexed table of backtrace entries. The currently-running procedure is the last entry in the table (that is, the frame numbering is the same as that used by `GET_VALUE`). Entry 1 is the oldest procedure on the stack.

67.7.20 PRINT_INSTANTIATIONS Procedure

This procedure returns a list of the packages that have been instantiated in the current session.

Syntax

```
DBMS_DEBUG.PRINT_INSTANTIATIONS (
   pkgs   IN OUT NOCOPY backtrace_table,
   flags  IN BINARY_INTEGER);
```

Parameters

Table 67-39 PRINT_INSTANTIATIONS Procedure Parameters

Parameter	Description
`pkgs`	The instantiated packages
`flags`	Bitmask of options: • 1 - show specs • 2 - show bodies • 4 - show local instantiations • 8 - show remote instantiations (NYI) • 16 - do a fast job. The routine does not test whether debug information exists or whether the libunit is shrink-wrapped.

Exceptions

`no_target_program` - target session is not currently executing

Usage Notes

On return, `pkgs` contains a `program_info` for each instantiation. The valid fields are: `Namespace`, `Name`, `Owner`, and `LibunitType`.

In addition, Line# contains a bitmask of:

- 1 - the libunit contains debug info
- 2 - the libunit is shrink-wrapped

67.7.21 PROBE_VERSION Procedure

This procedure returns the version number of `DBMS_DEBUG` on the server.

Syntax

```
DBMS_DEBUG.PROBE_VERSION (
   major out BINARY_INTEGER,
   minor out BINARY_INTEGER);
```

Parameters

Table 67-40 PROBE_VERSION Procedure Parameters

Parameter	Description
`major`	Major version number
`minor`	Minor version number: increments as functionality is added

67.7.22 SELF_CHECK Procedure

This procedure performs an internal consistency check. `SELF_CHECK` also runs a communications test to ensure that the Probe processes are able to communicate.

If `SELF_CHECK` does not return successfully, then an incorrect version of `DBMS_DEBUG` was probably installed on this server. The solution is to install the correct version (`pbload.sql` loads `DBMS_DEBUG` and the other relevant packages).

Syntax

```
DBMS_DEBUG.SELF_CHECK (
   timeout IN binary_integer := 60);
```

Parameters

Table 67-41 SELF_CHECK Procedure Parameters

Parameter	Description
`timeout`	The timeout to use for the communication test. Default is 60 seconds.

Exceptions

Table 67-42 SELF_CHECK Procedure Exceptions

Exception	Description
`OER-6516`	Probe version is inconsistent
`pipe_creation_failure`	Could not create a pipe
`pipe_send_failure`	Could not write data to the pipe
`pipe_receive_failure`	Could not read data from the pipe
`pipe_datatype_mismatch`	Datatype in the pipe was wrong

Table 67-42 (Cont.) SELF_CHECK Procedure Exceptions

Exception	Description
pipe_data_error	Data got garbled in the pipe

All of these exceptions are fatal. They indicate a serious problem with Probe that prevents it from working correctly.

67.7.23 SET_BREAKPOINT Function

This function sets a breakpoint in a program unit, which persists for the current session.

Execution pauses if the target program reaches the breakpoint.

Syntax

```
DBMS_DEBUG.SET_BREAKPOINT (
   program     IN  program_info,
   line#       IN  BINARY_INTEGER,
   breakpoint# OUT BINARY_INTEGER,
   fuzzy       IN  BINARY_INTEGER := 0,
   iterations  IN  BINARY_INTEGER := 0)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-43 SET_BREAKPOINT Function Parameters

Parameter	Description
program	Information about the program unit in which the breakpoint is to be set. (In version 2.1 and later, the namespace, name, owner, and dblink may be set to NULL, in which case the breakpoint is placed in the currently-running program unit.)
line#	Line at which the breakpoint is to be set
breakpoint#	On successful completion, contains the unique breakpoint number by which to refer to the breakpoint
fuzzy	Only applicable if there is no executable code at the specified line: 0 means return error_illegal_line 1 means search forward for an adjacent line at which to place the breakpoint -1 means search backward for an adjacent line at which to place the breakpoint
iterations	Number of times to wait before signalling this breakpoint

Return Values

> **Note:**
>
> The fuzzy and iterations parameters are not yet implemented

Table 67-44 SET_BREAKPOINT Function Return Values

Return	Description
`success`	A successful completion
`error_illegal_line`	Cannot set a breakpoint at that line
`error_bad_handle`	No such program unit exists

67.7.24 SET_OER_BREAKPOINT Function

This function sets an OER breakpoint.

Syntax

```
DBMS_DEBUG.SET_OER_BREAKPOINT (
   oer  IN PLS_INTEGER)
RETURN PLS_INTEGER;
```

Parameters

Table 67-45 SET_OER_BREAKPOINT Function Parameters

Parameter	Description
`oer`	The OER (positive 4-byte number) to set

Return Values

Table 67-46 SET_OER_BREAKPOINT Function Return Values

Return	Description
`success`	A successful completion
`error_no_such_breakpt`	No such OER breakpoint exists

67.7.25 SET_TIMEOUT Function

This function sets the timeout value and returns the new timeout value.

Syntax

```
DBMS_DEBUG.SET_TIMEOUT (
   timeout BINARY_INTEGER)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-47 SET_TIMEOUT Function Parameters

Parameter	Description
`timeout`	The timeout to use for communication between the target and debug sessions

67.7.26 SET_TIMEOUT_BEHAVIOUR Procedure

This procedure tells Probe what to do with the target session when a timeout occurs. This call is made in the target session.

Syntax

```
DBMS_DEBUG.SET_TIMEOUT_BEHAVIOUR (
   behaviour IN PLS_INTEGER);
```

Parameters

Table 67-48 SET_TIMEOUT_BEHAVIOUR Procedure Parameters

Parameter	Description
`behaviour - One of the following:`	
`retry_on_timeout`	Retry. Timeout has no effect. This is like setting the timeout to an infinitely large value.
`continue_on_timeout`	Continue execution, using same event flags
`nodebug_on_timeout`	Turn debug-mode OFF (in other words, call `debug_off`) and continue execution. No more events will be generated by this target session unless it is re-initialized by calling `debug_on`.
`abort_on_timeout`	Continue execution, using the `abort_execution` flag, which should cause the program to terminate immediately. The session remains in debug-mode.

Exceptions

unimplemented - the requested behavior is not recognized

Usage Notes

The default behavior (if this procedure is not called) is `continue_on_timeout`, since it allows a debugger client to reestablish control (at the next event) but does not cause the target session to hang indefinitely.

67.7.27 SET_VALUE Function

This function sets a value in the currently-running program. There are two overloaded `SET_VALUE` functions.

Syntax

```
DBMS_DEBUG.SET_VALUE (
   frame#               IN binary_integer,
   assignment_statement IN varchar2)
  RETURN BINARY_INTEGER;

DBMS_DEBUG.SET_VALUE (
   handle               IN program_info,
   assignment_statement IN VARCHAR2)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-49 SET_VALUE Function Parameters

Parameter	Description
`frame#`	Frame in which the value is to be set; 0 means the currently executing frame.
`handle`	Description of the package containing the variable
`assignment_statement`	An assignment statement (which must be legal PL/SQL) to run in order to set the value. For example, 'x := 3;'. Only scalar values are supported in this release. The right side of the assignment statement must be a scalar.

Return Values

Table 67-50 SET_VALUE Function Return Values

Return	Description
`success`	-
`error_illegal_value`	Not possible to set it to that value
`error_illegal_null`	Cannot set to `NULL` because object type specifies it as 'not `NULL`'
`error_value_malformed`	Value is not a scalar
`error_name_incomplete`	The assignment statement does not resolve to a scalar. For example, 'x := 3;', if x is a record.
`error_no_such_object`	One of the following: - Package does not exist - Package is not instantiated - User does not have privileges to debug the package - Object does not exist in the package

Usage Notes

In some cases, the PL/SQL compiler uses temporaries to access package variables, and does not guarantee to update such temporaries. It is possible, although unlikely, that modification to a package variable using `SET_VALUE` might not take effect for a line or two.

Examples

To set the value of `SCOTT.PACK.var` to 6:

```
DECLARE
   handle  dbms_debug.program_info;
   retval  BINARY_INTEGER;
BEGIN
   handle.Owner     := 'SCOTT';
   handle.Name      := 'PACK';
   handle.namespace := dbms_debug.namespace_pkgspec_or_toplevel;
   retval           := dbms_debug.set_value(handle, 'var := 6;');
END;
```

67.7.28 SHOW_BREAKPOINTS Procedures

There are two overloaded procedures that return a listing of the current breakpoints. There are three overloaded `SHOW_BREAKPOINTS` procedures.

Syntax

```
DBMS_DEBUG.SHOW_BREAKPOINTS (
   listing    IN OUT VARCHAR2);

DBMS_DEBUG.SHOW_BREAKPOINTS (
   listing    OUT breakpoint_table);

DBMS_DEBUG.SHOW_BREAKPOINTS (
   code_breakpoints  OUT breakpoint_table,
   oer_breakpoints   OUT oer_table);
```

Parameters

Table 67-51 SHOW_BREAKPOINTS Procedure Parameters

Parameter	Description
`listing`	A formatted buffer (including newlines) of the breakpoints. Indexed table of breakpoint entries. The breakpoint number is indicated by the index into the table. Breakpoint numbers start at 1 and are reused when deleted.
`code_breakpoints`	The indexed table of breakpoint entries, indexed by breakpoint number
`oer_breakpoints`	The indexed table of OER breakpoints, indexed by OER

67.7.29 SHOW_FRAME_SOURCE Procedure

The procedure gets the source code. There are two overloaded `SHOW_SOURCE` procedures.

Syntax

```
DBMS_DEBUG.SHOW_FRAME_SOURCE (
   first_line  IN            BINARY_INTEGER,
   last_line   IN            BINARY_INTEGER,
   source      IN OUT NOCOPY vc2_table,
   frame_num   IN            BINARY_INTEGER);
```

Parameters

Table 67-52 SHOW_FRAME_SOURCE Procedure Parameters

Parameter	Description
`first_line`	Line number of first line to fetch (PL/SQL programs always start at line 1 and have no holes)
`last_line`	Line number of last line to fetch. No lines are fetched past the end of the program.
`source`	The resulting table, which may be indexed by line#
`frame_num`	1-based frame number

Usage Notes

- You use this function only when backtrace shows an anonymous unit is executing at a given frame position and you need to view the source in order to set a breakpoint.
- If frame number is top of the stack and it's an anonymous block then `SHOW_SOURCE` can also be used.
- If it's a stored PL/SQL package/function/procedure then use SQL as described in the Usage Notes to SHOW_SOURCE Procedures.

67.7.30 SHOW_SOURCE Procedures

The procedure gets the source code. There are two overloaded `SHOW_SOURCE` procedures.

Syntax

```
DBMS_DEBUG.SHOW_SOURCE (
   first_line  IN   BINARY_INTEGER,
   last_line   IN   BINARY_INTEGER,
   source      OUT  vc2_table);

DBMS_DEBUG.SHOW_SOURCE (
   first_line   IN     BINARY_INTEGER,
   last_line    IN     BINARY_INTEGER,
   window       IN     BINARY_INTEGER,
   print_arrow  IN     BINARY_INTEGER,
   buffer       IN OUT VARCHAR2,
   buflen       IN     BINARY_INTEGER,
   pieces       OUT    BINARY_INTEGER);
```

Parameters

Table 67-53 SHOW_SOURCE Procedure Parameters

Parameter	Description
`first_line`	Line number of first line to fetch (PL/SQL programs always start at line 1 and have no holes)
`last_line`	Line number of last line to fetch. No lines are fetched past the end of the program.
`source`	The resulting table, which may be indexed by line#
`window`	'Window' of lines (the number of lines around the current source line)
`print_arrow`	Nonzero means to print an arrow before the current line
`buffer`	Buffer in which to place the source listing
`buflen`	Length of buffer
`pieces`	Set to nonzero if not all the source could be placed into the given buffer

Return Values

An indexed table of source-lines. The source lines are stored starting at `first_line`. If any error occurs, then the table is empty.

Usage Notes

The best way to get the source code (for a program that is being run) is to use SQL. For example:

```
DECLARE
    info DBMS_DEBUG.runtime_info;
BEGIN
   -- call DBMS_DEBUG.SYNCHRONIZE, CONTINUE,
   -- or GET_RUNTIME_INFO to fill in 'info'
   SELECT text INTO <buffer> FROM all_source
   WHERE owner = info.Program.Owner
     AND name  = info.Program.Name
     AND line  = info.Line#;
END;
```

However, this does not work for nonpersistent programs (for example, anonymous blocks and trigger invocation blocks). For nonpersistent programs, call `SHOW_SOURCE`. There are two flavors: one returns an indexed table of source lines, and the other returns a packed (and formatted) buffer.

The second overloading of `SHOW_SOURCE` returns the source in a formatted buffer, complete with line-numbers. It is faster than the indexed table version, but it does not guarantee to fetch all the source.

If the source does not fit in bufferlength (`buflen`), then additional pieces can be retrieved using the `GET_MORE_SOURCE` procedure (`pieces` returns the number of additional pieces that need to be retrieved).

67.7.31 SYNCHRONIZE Function

This function waits until the target program signals an event. If `info_requested` is not `NULL`, then it calls `GET_RUNTIME_INFO`.

Syntax

```
DBMS_DEBUG.SYNCHRONIZE (
   run_info       OUT  runtime_info,
   info_requested IN   BINARY_INTEGER := NULL)
  RETURN BINARY_INTEGER;
```

Parameters

Table 67-54 SYNCHRONIZE Function Parameters

Parameter	Description
`run_info`	Structure in which to write information about the program. By default, this includes information about what program is running and at which line execution has paused.
`info_requested`	Optional bit-field in which to request information other than the default (which is `info_getStackDepth` + `info_getLineInfo`). 0 means that no information is requested at all (see DBMS_DEBUG Operational Notes for more about information flags).

Return Values

Table 67-55 SYNCHRONIZE Function Return Values

Return	Description
`success`	A successful completion
`error_timeout`	Timed out before the program started execution
`error_communication`	Other communication error

67.7.32 TARGET_PROGRAM_RUNNING Procedure

This procedure returns `TRUE` if the target session is currently executing a stored procedure, or `FALSE` if it is not.

Syntax

```
DBMS_DEBUG.TARGET_PROGRAM_RUNNING
  RETURN BOOLEAN;
```

68
DBMS_DEBUG_JDWP

The DBMS_DEBUG_JDWP provides the interface to initiate and control the debugging of PL/SQL stored procedures and Java stored procedures over Java Debug Wire Protocol (JDWP).

This chapter contains the following topics:

- DBMS_DEBUG_JDWP Overview
- DBMS_DEBUG_JDWP Security Model
- Summary of DBMS_DEBUG_JDWP Subprograms

68.1 DBMS_DEBUG_JDWP Overview

Oracle supports the debugging of PL/SQL stored procedures and Java stored procedures over JDWP. Using the DBMS_DEBUG_JDWP package, you can:

- Retrieve the session ID of the current session and serial number
- Connect a database session to a debugger over JDWP
- Set the NLS parameters to change the formats in which PL/SQL program values are represented over JDWP

68.2 DBMS_DEBUG_JDWP Security Model

The debugging user invoking a DBMS_DEBUG_JDWP subprogram requires a privilege to connect the target database session to the debugger.

Table 68-1 System Privileges

System Privilege	Description
DEBUG CONNECT SESSION	Allows a user to connect his current session to a debugger
DEBUG CONNECT ANY	Allows a user to connect a session by any logon user to a debugger

In addition, the user can also be granted the following user privilege to debug another user's session.

Table 68-2 User Privileges

User Privilege	Description
DEBUG CONNECT ON USER <*user*>	Allows a user to connect any of the specified user's logon sessions to a debugger to debug another user's session or his own

68.3 Summary DBMS_DEBUG_JDWP Subprograms

This table lists the `DBMS_DEBUG_JDWP` subprograms and briefly describes them.

Table 68-3 DBMS_DEBUG_JDWP Package Subprograms

Subprogram	Description
CONNECT_TCP Procedure	Connects the specified session to the debugger waiting at `host:port`
CURRENT_SESSION_ID Function	Gets the current session's session ID
CURRENT_SESSION_SERIAL Function	Gets the current session's session serial number
DISCONNECT Procedure	Disconnects the specified session from any debugger with which it is connected
GET_NLS_PARAMETER Function	Gets the value of the specified NLS parameter affecting the format in which `NUMBER`, `DATE`, `TIME (WITH TIME ZONE)` and `TIMESTAMP (WITH TIME ZONE)` runtime values of PL/SQL programs are converted to strings as they are presented through JDWP
PROCESS_CONNECT_STRING Procedure	Connects a session to a debugger without having to directly modify an application's code
SET_NLS_PARAMETER Procedure	Sets the value of the specified NLS parameter affecting the format in which `NUMBER`, `DATE`, `TIME (WITH TIME ZONE)` and `TIMESTAMP (WITH TIME ZONE)` runtime values of PL/SQL programs are converted to strings as they are presented through JDWP

68.3.1 CONNECT_TCP Procedure

This procedure connects the specified session to the debugger waiting at `host:port`.

Syntax

```
DBMS_DEBUG_JDWP.CONNECT_TCP(
   host                 IN  VARCHAR2,
   port                 IN  VARCHAR2,
   session_id           IN  PLS_INTEGER := NULL,
   session_serial       IN  PLS_INTEGER := NULL,
   debug_role           IN  VARCHAR2  := NULL,
   debug_role_pwd       IN  VARCHAR2  := NULL,
   option_flags         IN  PLS_INTEGER := 0,
   extensions_cmd_set   IN  PLS_INTEGER := 128);
```

Parameters

Table 68-4 CONNECT_TCP Parameters

Parameter	Description
`host`	The host name the debugger is waiting at
`port`	The port number the debugger is waiting at

Table 68-4 (Cont.) CONNECT_TCP Parameters

Parameter	Description
session_id	Session ID
session_serial	Session number
debug_role	Debug role
debug_role_pwd	Debug password
option_flags	Values: • 1 : Does not suspend the program until the next client/server request begins. This can be used to hide the startup sequence from end users, who may only want to see their own code • 2 : Forces the connection even if the session appears to be connected to a debugger. This should best only be specified after some human-interaction confirmation step has occurred; i.e., if an attempt without this option raised ORA-30677, then if the user confirms, retry with this bit set. These may be added together to select multiple option choices.
extensions_cmd_set	The ID of the Oracle JDWP extension command set

Exceptions

Table 68-5 CONNECT_TCP Exceptions

Exception	Description
ORA-00022	Invalid session ID
ORA-01031	Insufficient privilege
ORA-30677	Session is already connected to a debugger
ORA-30681	Improper value for argument EXTENSIONS_CMD_SET
ORA-30682	Improper value for argument OPTION_FLAGS
ORA-30683	Failure establishing connection to debugger

Usage Notes

- To connect the current session to a debugger, you can pass NULL to both the session_id and session_serial parameters.
- To connect a different session, you need to find out its ID and serial. These are available in the v$plsql_debuggable_sessions view. The Instance Manager option of Oracle Enterprise Manager is one example of a user interface that displays these values to users. You can also find the values of these for your own session using the CURRENT_SESSION_ID Function and CURRENT_SESSION_SERIAL Function.
- The debug_role and debug_role_pwd arguments allow the user to name any role as the "debug role", which will be available to privilege checking when checking for permissions to

connect the session and when checking permissions available on objects within the debugged session. Both the role and its password are passed here as strings and not as identifiers, so double quotes should not be used but case matters. If the original role name wasn't double-quoted, it should be specified here in upper case.

- An `ORA-30677` indicates that the requested session is already being debugged. It is suggested in this case that the user be asked to confirm that (s)he desires to steal the session from the existing connection, and then either an explicit disconnect call or the use of the `connect_force_connect` option bit can be used to allow the connection to succeed on a second attempt. Note that using the `connect_force_connect` bit will avoid the session being allowed to run freely if it is currently suspended through the debugger - in other words, this bit lets you steal a session from one debugger to another without actually disturbing the state of the session.

68.3.2 CURRENT_SESSION_ID Function

This function gets the current session's session ID

Syntax

```
DBMS_DEBUG_JDWP.CURRENT_SESSION_ID
  RETURN PLS_INTEGER;
```

68.3.3 CURRENT_SESSION_SERIAL Function

This function gets the current session's session number.

Syntax

```
DBMS_DEBUG_JDWP.CURRENT_SESSION_SERIAL
  RETURN PLS_INTEGER;
```

68.3.4 DISCONNECT Procedure

This procedure disconnects the specified session from any debugger with which it is connected.

Syntax

```
DBMS_DEBUG_JDWP.DISCONNECT(
   session_id        IN  PLS_INTEGER  := NULL,
   session_serial    IN  PLS_INTEGER := NULL);
```

Parameters

Table 68-6 DISCONNECT Procedure Parameters

Parameter	Description
`session_id`	Session ID
`session_serial`	Session number

Usage Notes

- If the session to disconnect is the current session, the session will be allowed to run freely after disconnecting the debugger. Otherwise, the session will be terminated.

- The same rights are required for this call as for connect, except when disconnecting the current session and the effective user at the time of the call is the same as the login user of the session where no privilege is required.

68.3.5 GET_NLS_PARAMETER Function

This function gets the value of the specified NLS parameter affecting the format in which the `NUMBER`, `DATE`, `TIME (WITH TIME ZONE)` and `TIMESTAMP (WITH TIME ZONE)` runtime values of PL/SQL programs are converted to strings, as they are presented through JDWP.

These values are private to the current session, but further are private to the debugger mechanisms, separate from the values used to convert values within the debugged program itself.

Syntax

```
DBMS_DEBUG_JDWP.GET_NLS_PARAMETER(
   name   IN  VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 68-7 GET_NLS_PARAMETER Function Parameters

Parameter	Description
`name`	Specified NLS parameter

Usage Notes

- When any variable value is read or assigned through JDWP, or when either `GET_NLS_PARAMETER Function` or `SET_NLS_PARAMETER Procedure` is first invoked in a session, the debugger mechanisms make a private copy of the then-current `NLS_LANGUAGE`, `NLS_TERRITORY`, `NLS_CALENDAR`, `NLS_DATE_LANGUAGE`, `NLS_NUMERIC_CHARACTERS`, `NLS_TIMESTAMP_FORMAT`, `NLS_TIMESTAMP_TZ_FORMAT`, `NLS_TIME_FOMAT` and `NLS_TIME_TZ_FORMAT` values. These private copies may be read using this `GET_NLS_PARAMETER Function`call and changed using the following call to the SET_NLS_PARAMETER Procedure.
- Once the debugger's private copy of the NLS parameters is established, changes made to the NLS parameters in the current session using the `ALTER SESSION` statement will have no effect on the formatting of values as seen through JDWP. To modify the NLS parameters used for JDWP, one must use the `SET_NLS_PARAMETER Procedure` . By the same token, changes made to the debugger's private copy of the NLS parameters using`SET_NLS_PARAMETER Procedure` will have no effect on the debugged program itself.
- Date values are always formatted for JDWP use using the `NLS_TIMESTAMP_FORMAT`. The default format for `DATE (NLS_DATE_FORMAT)` used in a session most often does not show the time information that is in fact present in the value, and for debugging purposes it seems beneficial to always display that information.

68.3.6 PROCESS_CONNECT_STRING Procedure

This procedure connects a session to a debugger in two ways, so that you do not have to directly modify the application code.

The two ways are:

- Using the ORA_DEBUG_JDWP environment variable, when running an OCI program
- Setting a web browser cookie called OWA_DEBUG_<dad>, when running an application through the PL/SQL Web Gateway

Syntax

```
DBMS_DEBUG_JDWP.PROCESS_CONNECT_STRING (
   connect_string          IN  VARCHAR2,
   connect_string_type     IN  PLS_INTEGER);
```

Parameters

Table 68-8 PROCESS_CONNECT_STRING Procedure Parameters

Parameter	Description
connect_string	The ORA_DEBUG_JDWP environment variable or OWA_DEBUG_<dad> cookie value that contains the JDWP connection information such as the host and port number of the debugger to connect to
connect_string_type	Can have the following two values: • 1 if the connect string value is retrieved from the ORA_DEBUG_JDWP environment • 2 if the value is from the OWA_DEBUG_<dad> cookie

Exceptions

Table 68-9 PROCESS_CONNECT_STRING Procedure Exceptions

Exception	Description
ORA-00022	Invalid session ID
ORA-01031	Insufficient privilege
ORA-30677	Session is already connected to a debugger
ORA-30681	Improper value for argument EXTENSIONS_CMD_SET
ORA-30682	Improper value for argument OPTION_FLAGS
ORA-30683	Failure establishing connection to debugger
ORA-30689	Improper value for environment variable ORA_DEBUG_JDWP

68.3.7 SET_NLS_PARAMETER Procedure

This function sets the value of the specified NLS parameter affecting the format in which NUMBER, DATE, TIME (WITH TIME ZONE) and TIMESTAMP (WITH TIME ZONE) runtime values of PL/SQL programs are converted to strings as they are presented through JDWP.

These values are private to the current session, but further are private to the debugger mechanisms, separate from the values used to convert values within the debugged program itself.

Syntax

```
DBMS_DEBUG_JDWP.SET_NLS_PARAMETER(
   name      IN  VARCHAR2,
   value     IN  VARCHAR2);
```

Parameters

Table 68-10 SET_NLS_PARAMETER Procedure Parameters

Parameter	Description
name	Specified NLS parameter
value	Value of specified NLS parameter

Usage Notes

- When any variable value is read or assigned through JDWP, or when either GET_NLS_PARAMETER Function or SET_NLS_PARAMETER Procedure is first invoked in a session, the debugger mechanisms make a private copy of the then-current NLS_LANGUAGE, NLS_TERRITORY, NLS_CALENDAR, NLS_DATE_LANGUAGE, NLS_NUMERIC_CHARACTERS, NLS_TIMESTAMP_FORMAT, NLS_TIMESTAMP_TZ_FORMAT, NLS_TIME_FOMAT and NLS_TIME_TZ_FORMAT values. These private copies may be read by calling the GET_NLS_PARAMETER Function and changed using the following call to the SET_NLS_PARAMETER Procedure.
- Once the debugger's private copy of the NLS parameters is established, changes made to the NLS parameters in the current session using the ALTER SESSION statement will have no effect on the formatting of values as seen through JDWP. To modify the NLS parameters used for JDWP, one must use the SET_NLS_PARAMETER Procedure. By the same token, changes made to the debugger's private copy of the NLS parameters usingSET_NLS_PARAMETER Procedure will have no effect on the debugged program itself.
- Date values are always formatted for JDWP use using the NLS_TIMESTAMP_FORMAT. The default format for DATE (NLS_DATE_FORMAT) used in a session most often does not show the time information that is in fact present in the value, and for debugging purposes it seems beneficial to always display that information.

69
DBMS_DEBUG_JDWP_CUSTOM

The DBMS_DEBUG_JDWP_CUSTOM package provides database users a means to perform custom handling of a debug connection request of a database session to a debugger using the Java Debug Wire Protocol (JDWP).

This chapter contains the following topics:

- DBMS_DEBUG_JDWP_CUSTOM Overview
- DBMS_DEBUG_JDWP_CUSTOM Security Model
- Summary of DBMS_DEBUG_JDWP_CUSTOM Subprograms

69.1 DBMS_DEBUG_JDWP_CUSTOM Overview

DBMS_DEBUG_JDWP_CUSTOM package is invoked when you attempt to connect your database session to a debugger, using the Java Debug Wire Protocol (JDWP). Before connecting to the database, you must set the ORA_DEBUG_JDWP environment variable in your client-side OCI application. Defining this package in your schema enables you to perform custom actions before connecting your database session to the debugger.

69.2 DBMS_DEBUG_JDWP_CUSTOM Security Model

This default implementation of the DBMS_DEBUG_JDWP_CUSTOM package does not perform any additional security checks and only invokes the DBMS_DEBUG_JDWP package for making the debug connection. When the DBMS_DEBUG_JDWP package is invoked, it checks for the necessary DEBUG CONNECT system or user privilege before making the connection.

A database user who wants to perform additional custom security checks, must perform the following steps to override this default implementation:

- Define the DBMS_DEBUG_JDWP_CUSTOM package in the user's own schema.
- Implement the check before invoking the DBMS_DEBUG_JDWP package in the user's local copy of the package.

69.3 Summary DBMS_DEBUG_JDWP_CUSTOM Subprograms

This table summarizes the DBMS_DEBUG_JDWP_CUSTOM subprograms.

Table 69-1 DBMS_DEBUG_JDWP_CUSTOM Subprograms

Subprogram	Description
CONNECT_DEBUGGER Procedure	This procedure is invoked to handle the debug connection request of the current database session to the debugger.

69.3.1 CONNECT_DEBUGGER Procedure

This procedure is invoked to handle the debug connection request of the current database session to the debugger.

Syntax

```
PROCEDURE CONNECT_DEBUGGER
(
   HOST                VARCHAR2,
   PORT                VARCHAR2,
   DEBUG_ROLE          VARCHAR2 := NULL,
   DEBUG_ROLE_PWD      VARCHAR2 := NULL,
   OPTION_FLAGS        PLS_INTEGER := 0,
   EXTENSIONS_CMD_SET  PLS_INTEGER := 128
)
```

Parameters

The preceding parameters are for the default implementation of the procedure for handling the debug connection request. A user, who wants to customize the handling of the request, must override this default implementation of the `DBMS_DEBUG_JDWP_CUSTOM` package procedure by defining the package (specification and body) with a procedure of the same name in the user's own schema. The user may customize the number and names of the arguments, but must have the same name for the package and the procedure, namely, `DBMS_DEBUG_JDWP_CUSTOM` and `CONNECT_DEBUGGER` respectively. The user's customized version of the package may contain overloaded versions of the `CONNECT_DEBUGGER` procedure with different arguments. All of the arguments to the custom package procedure must either be of VARCHAR2 type or of types that PL/SQL can implicitly convert from VARCHAR2.

For example, if a user wants to disallow debugging outside business hours, while applying optional NLS settings to the program values displayed during debugging, then the user can define a custom implementation of the package in the user's schema in the following way:

```
CREATE OR REPLACE PACKAGE DBMS_DEBUG_JDWP_CUSTOM AUTHID CURRENT_USER IS
  PROCEDURE CONNECT_DEBUGGER(HOST VARCHAR2,
                             PORT VARCHAR2,
                             NLS_LANGUAGE VARCHAR2 DEFAULT NULL,
                             NLS_TERRITORY VARCHAR2 DEFAULT NULL);
END;
/

CREATE OR REPLACE PACKAGE BODY DBMS_DEBUG_JDWP_CUSTOM IS

  PROCEDURE CONNECT_DEBUGGER(HOST VARCHAR2,
                             PORT VARCHAR2,
                             NLS_LANGUAGE VARCHAR2 DEFAULT NULL,
                             NLS_TERRITORY VARCHAR2 DEFAULT NULL) IS
  BEGIN
    SELECT EXTRACT(HOUR FROM LOCALTIMESTAMP) INTO HOUR FROM DUAL;
    IF (HOUR < 9 OR HOUR > 5) THEN
      RAISE_APPLICATION_ERROR(-20000,
        'Debug connection disallowed outside business hours');
```

```
    ELSE
      IF (NLS_LANGUAGE IS NOT NULL) THEN
        DBMS_DEBUG_JDWP.SET_NLS_PARAMETER('NLS_LANGUAGE', NLS_LANGUAGE);
      END IF;
      IF (NLS_TERRITORY IS NOT NULL) THEN
        DBMS_DEBUG_JDWP.SET_NLS_PARAMETER('NLS_TERRITORY', NLS_TERRITORY);
      END IF;
      DBMS_DEBUG_JDWP.CONNECT_TCP(HOST, PORT);
    END IF;
  END;

END;
/
```

Before executing the client application for passing arguments to the procedure parameters, the user must set the `ORA_DEBUG_JDWP` environment variable as follows :

```
> setenv ORA_DEBUG_JDWP "host=123.45.67.89;port=4000;nls_language=american"
> myapp -user HR/<password>
...
```

70
DBMS_DESCRIBE

You can use the `DBMS_DESCRIBE` package to get information about a PL/SQL object. When you specify an object name, `DBMS_DESCRIBE` returns a set of indexed tables with the results. Full name translation is performed and security checking is also checked on the final object.

This chapter contains the following topics:

- Overview
- Security Model
- Types
- Exceptions
- Examples
- Summary of DBMS_DESCRIBE Subprograms

70.1 DBMS_DESCRIBE Overview

This package provides the same functionality as the Oracle Call Interface `OCIDescribeAny` call.

> **See Also:**
>
> *Oracle Call Interface Developer's Guide*

70.2 DBMS_DESCRIBE Security Model

This package is available to `PUBLIC` and performs its own security checking based on the schema object being described.

70.3 DBMS_DESCRIBE Types

The `DBMS_DESCRIBE` package declares two PL/SQL table types, which are used to hold data returned by `DESCRIBE_PROCEDURE` in its `OUT` parameters.

The types are:

```
TYPE VARCHAR2_TABLE IS TABLE OF VARCHAR2(30)
    INDEX BY BINARY_INTEGER;

TYPE NUMBER_TABLE IS TABLE OF NUMBER
    INDEX BY BINARY_INTEGER;
```

70.4 DBMS_DESCRIBE Exceptions

`DBMS_DESCRIBE` can raise application errors in the range -20000 to -20004.

Table 70-1 DBMS_DESCRIBE Errors

Error	Description
ORA-20000	ORU 10035: cannot describe a package ('X') only a procedure within a package.
ORA-20001	ORU-10032: procedure 'X' within package 'Y' does not exist.
ORA-20002	ORU-10033: object 'X' is remote, cannot describe; expanded name 'Y'.
ORA-20003	ORU-10036: object 'X' is invalid and cannot be described.
ORA-20004	Syntax error attempting to parse 'X'.

70.5 DBMS_DESCRIBE Examples

One use of the DESCRIBE_PROCEDURE procedure is as an external service interface.

For example, consider a client that provides an OBJECT_NAME of SCOTT.ACCOUNT_UPDATE, where ACCOUNT_UPDATE is an overloaded function with specification:

```
TABLE account (accnt_no NUMBER, person_id NUMBER,
               balance NUMBER(7,2))
TABLE person  (person_id number(4), person_nm varchar2(10))

CREATE OR REPLACE PACKAGE ACCOUNT_PKG is   FUNCTION ACCOUNT_UPDATE (accnt_no    NUMBER,
                           person      person%rowtype,
                           amounts     DBMS_DESCRIBE.NUMBER_TABLE,
                           trans_date  DATE)
                           return       account.balance%type;

   FUNCTION ACCOUNT_UPDATE (accnt_no     NUMBER,
                           person       person%rowtype,
                           amounts      DBMS_DESCRIBE.NUMBER_TABLE,
                           trans_no     NUMBER)
                           return       account.balance%type;
END;
```

This procedure might look similar to the following output:

```
overload position  argument level  datatype length prec scale rad
-------- --------- -------- ------ -------- ------ ---- ----- ---
       1         0                0        2     22    7     2  10
       1         1   ACCNT_NO     0        2      0    0     0   0
       1         2   PERSON       0      250      0    0     0   0
       1         1   PERSON_ID    1        2     22    4     0  10
       1         2   PERSON_NM    1        1     10    0     0   0
       1         3   AMOUNTS      0      251      0    0     0   0
       1         1                1        2     22    0     0   0
       1         4   TRANS_DATE   0       12      0    0     0   0
       2         0                0        2     22    7     2  10
       2         1   ACCNT_NO     0        2     22    0     0   0
       2         2   PERSON       0        2     22    4     0  10
       2         3   AMOUNTS      0      251     22    4     0  10
       2         1                1        2      0    0     0   0
       2         4   TRANS_NO     0        2      0    0     0   0
```

The following PL/SQL procedure has as its parameters all of the PL/SQL datatypes:

```
CREATE OR REPLACE PROCEDURE p1 (
        pvc2    IN     VARCHAR2,
```

```
        pvc     OUT    VARCHAR,
        pstr    IN OUT STRING,
        plong   IN     LONG,
        prowid  IN     ROWID,
        pchara  IN     CHARACTER,
        pchar   IN     CHAR,
        praw    IN     RAW,
        plraw   IN     LONG RAW,
        pbinint IN     BINARY_INTEGER,
        pplsint IN     PLS_INTEGER,
        pbool   IN     BOOLEAN,
        pnat    IN     NATURAL,
        ppos    IN     POSITIVE,
        pposn   IN     POSITIVEN,
        pnatn   IN     NATURALN,
        pnum    IN     NUMBER,
        pintgr  IN     INTEGER,
        pint    IN     INT,
        psmall  IN     SMALLINT,
        pdec    IN     DECIMAL,
        preal   IN     REAL,
        pfloat  IN     FLOAT,
        pnumer  IN     NUMERIC,
        pdp     IN     DOUBLE PRECISION,
        pdate   IN     DATE,
        pmls    IN     MLSLABEL) AS

BEGIN
    NULL;
END;
```

If you describe this procedure using the following:

```
CREATE OR REPLACE PACKAGE describe_it AS

    PROCEDURE desc_proc (name VARCHAR2);

END describe_it;

CREATE OR REPLACE PACKAGE BODY describe_it AS

  PROCEDURE prt_value(val VARCHAR2, isize INTEGER) IS
    n INTEGER;
  BEGIN
    n := isize - LENGTHB(val);
    IF n < 0 THEN
      n := 0;
    END IF;
    DBMS_OUTPUT.PUT(val);
    FOR i in 1..n LOOP
      DBMS_OUTPUT.PUT(' ');
    END LOOP;
  END prt_value;

  PROCEDURE desc_proc (name VARCHAR2) IS

      overload     DBMS_DESCRIBE.NUMBER_TABLE;
      position     DBMS_DESCRIBE.NUMBER_TABLE;
      c_level      DBMS_DESCRIBE.NUMBER_TABLE;
      arg_name     DBMS_DESCRIBE.VARCHAR2_TABLE;
      dty          DBMS_DESCRIBE.NUMBER_TABLE;
      def_val      DBMS_DESCRIBE.NUMBER_TABLE;
```

```
        p_mode       DBMS_DESCRIBE.NUMBER_TABLE;
        length       DBMS_DESCRIBE.NUMBER_TABLE;
        precision    DBMS_DESCRIBE.NUMBER_TABLE;
        scale        DBMS_DESCRIBE.NUMBER_TABLE;
        radix        DBMS_DESCRIBE.NUMBER_TABLE;
        spare        DBMS_DESCRIBE.NUMBER_TABLE;
        idx          INTEGER := 0;

    BEGIN
        DBMS_DESCRIBE.DESCRIBE_PROCEDURE(
                name,
                null,
                null,
                overload,
                position,
                c_level,
                arg_name,
                dty,
                def_val,
                p_mode,
                length,
                precision,
                scale,
                radix,
                spare);

        DBMS_OUTPUT.PUT_LINE('Position    Name        DTY  Mode');
        LOOP
            idx := idx + 1;
            prt_value(TO_CHAR(position(idx)), 12);
            prt_value(arg_name(idx), 12);
            prt_value(TO_CHAR(dty(idx)), 5);
            prt_value(TO_CHAR(p_mode(idx)), 5);
            DBMS_OUTPUT.NEW_LINE;
        END LOOP;
    EXCEPTION
       WHEN NO_DATA_FOUND THEN
          DBMS_OUTPUT.NEW_LINE;
          DBMS_OUTPUT.NEW_LINE;

    END desc_proc;
END describe_it;
```

Then the results list all the numeric codes for the PL/SQL datatypes:

```
Position  Name     Datatype_Code  Mode
1         PVC2       1              0
2         PVC        1              1
3         PSTR       1              2
4         PLONG      8              0
5         PROWID     11             0
6         PCHARA     96             0
7         PCHAR      96             0
8         PRAW       23             0
9         PLRAW      24             0
10        PBININT    3              0
11        PPLSINT    3              0
12        PBOOL      252            0
13        PNAT       3              0
14        PPOS       3              0
15        PPOSN      3              0
16        PNATN      3              0
```

```
17        PNUM      2              0
18        PINTGR    2              0
19        PINT      2              0
20        PSMALL    2              0
21        PDEC      2              0
22        PREAL     2              0
23        PFLOAT    2              0
24        PNUMER    2              0
25        PDP       2              0
26        PDATE     12             0
27        PMLS      106            0
```

70.6 Summary of DBMS_DESCRIBE Subprograms

The `DBMS_DESCRIBE` package includes the `DESCRIBE_PROCEDURE` procedure.

Table 70-2 DBMS_DESCRIBE Package Subprograms

Subprogram	Description
DESCRIBE_PROCEDURE Procedure	Provides a brief description of a PL/SQL stored procedure

70.6.1 DESCRIBE_PROCEDURE Procedure

The procedure `DESCRIBE_PROCEDURE` provides a brief description of a PL/SQL stored procedure.

It takes the name of a stored procedure and returns information about each parameter of that procedure.

Syntax

```
DBMS_DESCRIBE.DESCRIBE_PROCEDURE(
   object_name                    IN  VARCHAR2,
   reserved1                      IN  VARCHAR2,
   reserved2                      IN  VARCHAR2,
   overload                       OUT NUMBER_TABLE,
   position                       OUT NUMBER_TABLE,
   level                          OUT NUMBER_TABLE,
   argument_name                  OUT VARCHAR2_TABLE,
   datatype                       OUT NUMBER_TABLE,
   default_value                  OUT NUMBER_TABLE,
   in_out                         OUT NUMBER_TABLE,
   length                         OUT NUMBER_TABLE,
   precision                      OUT NUMBER_TABLE,
   scale                          OUT NUMBER_TABLE,
   radix                          OUT NUMBER_TABLE,
   spare                          OUT NUMBER_TABLE
   include_string_constraints     OUT BOOLEAN DEFAULT FALSE);
```

Parameters

Table 70-3 DBMS_DESCRIBE.DESCRIBE_PROCEDURE Parameters

Parameter	Description
object_name	Name of the procedure being described. The syntax for this parameter follows the rules used for identifiers in SQL. The name can be a synonym. This parameter is required and may not be null. The total length of the name cannot exceed 197 bytes. An incorrectly specified OBJECT_NAME can result in one of the following exceptions: ORA-20000 - A package was specified. You can only specify a stored procedure, stored function, packaged procedure, or packaged function. ORA-20001 - The procedure or function that you specified does not exist within the given package. ORA-20002 - The object that you specified is a remote object. This procedure cannot currently describe remote objects. ORA-20003 - The object that you specified is invalid and cannot be described. ORA-20004 - The object was specified with a syntax error.
reserved1 reserved2	Reserved for future use -- must be set to NULL or the empty string.
overload	A unique number assigned to the procedure's signature. If a procedure is overloaded, then this field holds a different value for each version of the procedure.
position	Position of the argument in the parameter list. Position 0 returns the values for the return type of a function.
level	If the argument is a composite type, such as record, then this parameter returns the level of the datatype. See the *Oracle Call Interface Developer's Guide* for a description of the ODESSP call for an example.
argument_name	Name of the argument associated with the procedure that you are describing.

Table 70-3 (Cont.) DBMS_DESCRIBE.DESCRIBE_PROCEDURE Parameters

Parameter	Description
datatype	Oracle datatype of the argument being described. The datatypes and their numeric type codes are: 0 placeholder for procedures with no arguments 1 VARCHAR, VARCHAR, STRING 2 NUMBER, INTEGER, SMALLINT, REAL, FLOAT, DECIMAL 3 BINARY_INTEGER, PLS_INTEGER, POSITIVE, NATURAL 8 LONG 11 ROWID 12 DATE 23 RAW 24 LONG RAW 58 OPAQUE TYPE 96 CHAR (ANSI FIXED CHAR), CHARACTER 106 MLSLABEL 121 OBJECT 122 NESTED TABLE 123 VARRAY 178 TIME 179 TIME WITH TIME ZONE 180 TIMESTAMP 181 TIMESTAMP WITH TIME ZONE 231 TIMESTAMP WITH LOCAL TIME ZONE 250 PL/SQL RECORD 251 PL/SQL TABLE 252 PL/SQL BOOLEAN
default_value	1 if the argument being described has a default value; otherwise, the value is 0.
in_out	Describes the mode of the parameter: 0 IN 1 OUT 2 IN OUT
length	For %rowtype formal arguments, the length constraint is returned, otherwise 0 is returned.If the include_string_constraints parameter is set to TRUE, the argument's formal length constraint is passed back if it is of the appropriate type. Those are the string types: 1;8;23;24;96
precision	If the argument being described is of datatype 2 (NUMBER), then this parameter is the precision of that number.
scale	If the argument being described is of datatype 2 (NUMBER), then this parameter is the scale of that number.
radix	If the argument being described is of datatype 2 (NUMBER), then this parameter is the radix of that number.
spare	Reserved for future functionality.
include_string_constraints	The default is FALSE. If the parameter is set to TRUE, the arguments' formal type constraints is passed back if it is of the appropriate type.Those are the string types: 1;8;23;24;96

Return Values

All values from DESCRIBE_PROCEDURE are returned in its OUT parameters. The datatypes for these are PL/SQL tables, to accommodate a variable number of parameters.

71
DBMS_DG

The `DBMS_DG` package allows applications to notify the primary database or the fast-start failover target database in an Oracle Data Guard broker environment to initiate a fast-start failover when the application encounters a condition that warrants a failover.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This chapter contains the following topics:

- Using DBMS_DG
- Security Model
- Summary of the DBMS_DG Subprogram

> **See Also:**
>
> *Oracle Data Guard Broker*

71.1 Using DBMS_DG

There are conditions detectable by applications running outside of the Oracle database that may warrant the Oracle Data Guard broker to perform a fast-start failover. Because the range of possible conditions is virtually unlimited, it is left to the applications to determine which conditions warrant a fast-start failover.

When such conditions occur, the application calls the `DBMS_DG.INITIATE_FS_FAILOVER` procedure to alert either the primary or fast-start failover target standby database that the application wants a fast-start failover to occur immediately. The database on which the procedure was called then notifies the observer, which immediately initiates a fast-start failover as long as the standby database is in a valid fast-start failover state ("observed" and either "synchronized" or "within lag") to accept a failover.If the configuration is not in a valid fast-start failover state, the `INITIATE_FS_FAILOVER` subprogram returns an ORA error message (it will not signal an exception) to inform the calling application that a fast-start failover could not be performed.

> **Note:**
>
> If you are working in a multitenant container database (CDB), then functions within DBMS_DG are only executed at the root level. Ensure you are connected at the root level, not at the individual pluggable database (PDB) level.

71.2 DBMS_DG Security Model

The DBMS_DG package runs with invoker's rights and requires the SYSDBA privilege.

71.3 Summary of the DBMS_DG Subprogram

The DBMS_DG package contains one subprogram, the INITIATE_FS_FAILOVER procedure.

Table 71-1 DBMS_DG Package Subprogram

Subprogram	Description
INITIATE_FS_FAILOVER Procedure	Enables an application to notify either the primary or fast-start failover target standby database that a fast-start failover is necessary when the application encounters conditions that warrant a failover. This procedure can only be called while connected to a primary database or a fast-start failover standby database.

71.3.1 INITIATE_FS_FAILOVER Procedure

Use this procedure to specify a condition string that, when encountered by an application, allows the application to request that a fast-start failover be invoked.

Syntax

```
DBMS_DG.INITIATE_FS_FAILOVER (
     condstr          IN VARCHAR2)
RETURN BINARY_INTEGER;
```

Parameters

Table 71-2 INITIATE_FS_FAILOVER Procedure Parameters

Parameter	Description
condstr	Specifies the condition string for which a fast-start failover should be requested. If no condition string argument is supplied, the default string of "Application Failover Requested" will be logged in the broker log file and in the database alert log of the database on which the procedure was called.

Usage Notes

- This procedure returns a binary integer.
- Query the V$FS_FAILOVER_STATS view to see the time of the last fast-start failover and the reason it was performed.

- This procedure can only be called while connected to a primary database or a fast-start failover standby database.

Errors

Table 71-3 INITIATE_FS_FAILOVER Procedure Errors

Error	Description
`ORA-00000: normal, successful completion`	The request to initiate a fast-start failover has been posted to the observer.
`ORA-16646: fast-start failover is disabled`	Either a broker configuration does not exist or fast-start failover has not been enabled.
`ORA-16666: unable to initiate fast-start failover on a bystander standby database`	`DBMS_DG.INITIATE_FS_FAILOVER` was invoked on a bystander standby database. That is, it was not invoked on the primary or on the fast-start failover target standby database.
`ORA-16817: unsynchronized fast-start failover configuration`	`DBMS_DG.INITIATE_FS_FAILOVER` was invoked in a maximum available fast-start failover configuration when the configuration was not synchronized.
`ORA-16819: fast-start failover observer not started`	`DBMS_DG.INITIATE_FS_FAILOVER` was invoked but an observer had not yet been started.
`ORA-16820: fast-start failover observer is no longer observing this database`	`DBMS_DG.INITIATE_FS_FAILOVER` was invoked but the configuration detects that the observer may not be running.
`ORA-16829: lagging fast-start failover configuration`	`DBMS_DG.INITIATE_FS_FAILOVER` was invoked in a maximum performance fast-start failover configuration when the configuration was not in the user-specified redo lag limit.

Example

In this example, the program attempts to initiate a fast-start failover when fast-start failover is disabled. To use this example, connect as user `SYS` with `SYDDBA` privileges.

```
set serveroutput on

declare
status integer;

begin
status := dbms_dg.initiate_fs_failover(''Failover Requested'');

dbms_output.put_line(''Fast-Start Failover is disabled: Expected status = ORA-16646'');
dbms_output.put_line(''                 Actual Status = ORA-'' || status);

end;
/
exit;
```

72
DBMS_DICTIONARY_CHECK

DBMS_DICTIONARY_CHECK is a read-only and lightweight PL/SQL package procedure that helps you identify Oracle Database dictionary inconsistencies.

72.1 Overview of Oracle Database Dictionary Check

DBMS_DICTIONARY_CHECK is a read-only and lightweight PL/SQL package procedure that helps you identify Oracle Database dictionary inconsistencies that are manifested in unexpected entries in the Oracle Database dictionary tables or invalid references between dictionary tables. Oracle Database dictionary inconsistencies can cause process failures and, in some cases, instance crash. Such inconsistencies may be exposed to internal ORA-00600 errors. DBMS_DICTIONARY_CHECK assists you in identifying such inconsistencies and in some cases provides guided remediation to resolve the problem and avoid such database failures.

Unexpected entries in the dictionary tables or invalid references between dictionary tables, for example, include the following:

- A lob segment not in OBJ$
- An entry in SOURCE$ not in OBJ$
- Invalid data between OBJ$-PARTOBJ$ and TABPART$
- A segment with no owner
- A materialized segment with no entry in seg$
- A segment with no object entry
- A recycle bin object not in the recyclebin$
- Check if Control Seq is near the limit

72.2 Using DBMS_DICTIONARY_CHECK

To run all the checks or only the critical checks defined by DBMS_DICTIONARY_CHECK, connect to the SYS schema, and then run the following commands as SYS user:

Full check

```
SQL> set serveroutput on size unlimited
SQL> execute dbms_dictionary_check.full
```

```
SQL> set serveroutput on size unlimited
SQL> EXECUTE dbms_dictionary_check.full(repair=>TRUE)
```

While running a full check, optionally, you can use the repair option to resolve inconsistencies. Valid values: TRUE|FALSE. Default: FALSE.

Critical check

```
SQL> set serveroutput on size unlimited
SQL> execute dbms_dictionary_check.critical
```

Optionally, turn on the spool to redirect the output to a server-side flat file. By default, when you query the `SYS` schema, the `DBMS_DICTIONARY_CHECK` package creates a trace file named, `DICTCHECK.trc`.

For example: `/<path>/diag/rdbms/<db_name>/<oracle_sid>/trace/<oracle_sid>_<ora>_<pid>_DICTCHECK.trc`.

The execution reports the result as:

- **CRITICAL**: Requires an immediate fix.
- **FAIL**: Requires resolution on priority.
- **WARN**: Good to resolve.
- **PASS**: No issues.

> **Note:**
>
> In all cases, any output reporting "problems" must be triaged by Oracle Support to confirm if any action is required.

Example 72-1 Full check run

```
SQL> set serveroutput on size unlimited
SQL> execute dbms_dictionary_check.full
dbms_dictionary_check on 07-MAR-2023
03:17:48
-------------------------------------------

Catalog Version 21.0.0.0.0
(2300000000)
db_name:
ORCL

Is CDB?:
NO

Trace File: /oracle/log/diag/rdbms/orcl/orcl/trace/
orcl_ora_2574906_DICTCHECK.trc

                                  Catalog
Fixed
Procedure Name                    Version    Vs Release    Timestamp
Result
------------------------------ ... ---------- -- ---------- --------------
------
.- OIDOnObjCol                 ... 2300000000 <=  *All Rel* 03/07 03:17:48
PASS
```

```
.- LobNotInObj                 ... 2300000000 <=  *All Rel* 03/07 03:17:48
PASS
.- SourceNotInObj              ... 2300000000 <=  *All Rel* 03/07 03:17:48
PASS
.- OversizedFiles              ... 2300000000 <=  *All Rel* 03/07 03:17:48
PASS
.- PoorDefaultStorage          ... 2300000000 <=  *All Rel* 03/07 03:17:48
PASS
.- PoorStorage                 ... 2300000000 <=  *All Rel* 03/07 03:17:48
PASS
.- TabPartCountMismatch        ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- TabComPartObj               ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- Mview                       ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- ValidDir                    ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- DuplicateDataobj            ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- ObjSyn                      ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- ObjSeq                      ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- UndoSeg                     ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- IndexSeg                    ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- IndexPartitionSeg           ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- IndexSubPartitionSeg        ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- TableSeg                    ... 2300000000 <=  *All Rel* 03/07 03:17:49
FAIL

HCKE-0019: Orphaned TAB$ (no SEG$) (Doc ID
1360889.1)
ORPHAN TAB$: OBJ#=83241 DOBJ#=83241 TS=5 RFILE/BLOCK=5/11 TABLE=SYS.ORPHANSEG
BOBJ#=

.- TablePartitionSeg           ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- TableSubPartitionSeg        ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- PartCol                     ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- ValidSeg                    ... 2300000000 <=  *All Rel* 03/07 03:17:49
FAIL

HCKE-0023: Orphaned SEG$ Entry (Doc ID
1360934.1)
ORPHAN SEG$: SegType=LOB TS=5 RFILE/
BLOCK=5/26

.- IndPartObj                  ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
```

```
.- DuplicateBlockUse            ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- FetUet                       ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- Uet0Check                    ... 2300000000 <=  *All Rel* 03/07 03:17:49
PASS
.- SeglessUET                   ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ValidInd                     ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ValidTab                     ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- IcolDepCnt                   ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ObjIndDobj                   ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- TrgAfterUpgrade              ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ObjType0                     ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ValidOwner                   ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- StmtAuditOnCommit            ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- PublicObjects                ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- SegFreelist                  ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ValidDepends                 ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- CheckDual                    ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ObjectNames                  ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ChkIotTs                     ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- NoSegmentIndex               ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- NextObject                   ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- DroppedROTS                  ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- FilBlkZero                   ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- DbmsSchemaCopy               ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- IdnseqObj                    ... 2300000000 >  1201000000 03/07 03:17:50
PASS
.- IdnseqSeq                    ... 2300000000 >  1201000000 03/07 03:17:50
PASS
.- ObjError                     ... 2300000000 >  1102000000 03/07 03:17:50
PASS
.- ObjNotLob                    ... 2300000000 <=  *All Rel* 03/07 03:17:50
FAIL

HCKE-0049: OBJ$ LOB entry has no LOB$ or LOBFRAG$ entry (Doc ID
```

```
2125104.1)
OBJ$ LOB has no LOB$ entry: Obj=83243 Owner: SYS LOB Name:
LOBC1

.- MaxControlfSeq              ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- SegNotInDeferredStg         ... 2300000000 >  1102000000 03/07 03:17:50
PASS
.- SystemNotRfile1             ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- DictOwnNonDefaultSYSTEM     ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ValidateTrigger             ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- ObjNotTrigger               ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
.- InvalidTSMaxSCN             ... 2300000000 >  1202000000 03/07 03:17:50
CRITICAL

HCKE-0054: TS$ has Tablespace with invalid Maximum SCN (Doc ID
1360208.1)
TS$ has Tablespace with invalid Maximum SCN: TS#=5 Tablespace=HCHECK
Online$=1

.- OBJRecycleBin               ... 2300000000 <=  *All Rel* 03/07 03:17:50
PASS
---------------------------------------

07-MAR-2023 03:17:50  Elapsed: 2
secs
---------------------------------------

Found 4 potential problem(s) and 0
warning(s)
Found 1 CRITICAL problem(s) needing
attention
Contact Oracle Support with the output and trace
file
to check if the above needs attention or
not
BEGIN dbms_dictionary_check.full; END;

*
ERROR at line 1:
ORA-20000: dbms_dictionary_check found 1 critical issue(s).  Trace file:
/oracle/log/diag/rdbms/orcl/orcl/trace/orcl_ora_2574906_DICTCHECK.trc

SQL>
```

Example 72-2 Full check run with repair option

```
SQL> set serveroutput on size unlimited
SQL> EXECUTE dbms_dictionary_check.full(repair=>TRUE)
dbms_dictionary_check on 04-OCT-2023 01:35:37
---------------------------------------------
```

```
Catalog Version 23.0.0.0.0 (2300000000)
db_name: orcl
Is CDB?: NO
Trace File: /oracle/log/diag/rdbms/orcl/orcl/trace/
orcl_ora_3831454_DICTCHECK.trc
                                    Catalog       Fixed
Procedure Name                      Version    Vs Release    Timestamp
Result
------------------------------ ... ---------- -- ---------- --------------
------
.- OIDOnObjCol                 ... 2300000000 <=  *All Rel* 10/04 01:35:37
PASS
.- LobNotInObj                 ... 2300000000 <=  *All Rel* 10/04 01:35:37
PASS
.- SourceNotInObj              ... 2300000000 <=  *All Rel* 10/04 01:35:38
FAIL

HCKE-0003: SOURCE$ for OBJ# not in OBJ$ (Doc ID 1360233.1)
SOURCE$ has 10 rows for 1 OBJ# values not in OBJ$


INCONSISTENCY REPAIRED - Check the trace file for repair details:
SourceNotInObj_Repair: DELETED 10 objects from SOURCE$ not found in OBJ$

.- OversizedFiles              ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- PoorDefaultStorage          ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- PoorStorage                 ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- TabPartCountMismatch        ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- TabComPartObj               ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- Mview                       ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- ValidDir                    ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- DuplicateDataobj            ... 2300000000 <=  *All Rel* 10/04 01:35:38
PASS
.- ObjSyn                      ... 2300000000 <=  *All Rel* 10/04 01:35:39
PASS
.- ObjSeq                      ... 2300000000 <=  *All Rel* 10/04 01:35:39
PASS
.- ValidateSeg                 ... 2300000000 <=  *All Rel* 10/04 01:35:39
FAIL

HCKE-0023: Orphaned SEG$ Entry (Doc ID 1360934.1)
ORPHAN SEG$: SegType=DATA TS=5 RFILE/BLOCK=5/11 HWMINCR(DOBJ#)=73271

.- TableSeg                    ... 2300000000 <=  *All Rel* 10/04 01:35:39
PASS
.- TablePartitionSeg           ... 2300000000 <=  *All Rel* 10/04 01:35:39
PASS
.- TableSubPartitionSeg        ... 2300000000 <=  *All Rel* 10/04 01:35:39
PASS
```

```
ORPHAN SEG$: SegType=DATA TS=5 RFILE/BLOCK=5/11
^ Segment entry repaired - Converted to TEMPORARY
INCONSISTENCY REPAIRED - Check the trace file for repair details:
ValidateSeg repaired 1 Orphan Seg$ entries

.- UndoSeg                    ... 2300000000 <=  *All Rel* 10/04 01:35:39
PASS
...
---------------------------------------
04-OCT-2023 01:35:40  Elapsed: 3 secs
---------------------------------------
Found 3 potential problem(s) and 0 warning(s)
Repaired 11 item(s)
Contact Oracle Support with the output and trace file
to check if the above needs attention or not
BEGIN dbms_dictionary_check.full(repair=>TRUE); END;

*
ERROR at line 1:
ORA-20001: dbms_dictionary_check found 3 problem(s) and repaired 11 item(s).
Trace file:
/oracle/log/diag/rdbms/orcl/orcl/trace/orcl_ora_3831454_DICTCHECK.trc
```

Example 72-3 Critical check run

```
SQL> set serveroutput on size unlimited
SQL> execute dbms_dictionary_check.critical
dbms_dictionary_check on 07-MAR-2023
03:12:23
-------------------------------------------

Catalog Version 21.0.0.0.0
(2100000000)
db_name:
ORCL

Is CDB?:
NO

Trace File: /oracle/log/diag/rdbms/orcl/orcl/trace/
orcl_ora_2574058_DICTCHECK.trc

                                  Catalog
Fixed
Procedure Name                    Version    Vs Release    Timestamp
Result
------------------------------ ... ---------- -- ---------- --------------
------
.- UndoSeg                    ... 2300000000 <=  *All Rel* 03/07 03:12:23
PASS
.- MaxControlfSeq             ... 2300000000 <=  *All Rel* 03/07 03:12:23
PASS
.- InvalidTSMaxSCN            ... 2300000000 >  1202000000 03/07 03:12:23
CRITICAL
```

```
HCKE-0054: TS$ has Tablespace with invalid Maximum SCN (Doc ID
1360208.1)
TS$ has Tablespace with invalid Maximum SCN: TS#=5 Tablespace=HCHECK
Online$=1

------------------------------------

07-MAR-2023 03:12:23  Elapsed: 0
secs
------------------------------------

Found 1 potential problem(s) and 0
warning(s)
Found 1 CRITICAL problem(s) needing
attention
Contact Oracle Support with the output and trace
file
to check if the above needs attention or
not
BEGIN dbms_dictionary_check.critical; END;

*
ERROR at line 1:
ORA-20000: dbms_dictionary_check found 1 critical issue(s).  Trace file:
/oracle/log/diag/rdbms/orcl/orcl/trace/orcl_ora_2574058_DICTCHECK.trc

SQL>
```

72.3 Summary of DBMS_DICTIONARY_CHECK Subprograms

`DBMS_DICTIONARY_CHECK` package includes the following procedures:

- **LobNotInObj**: Checks if a LOB segment is not in `OBJ$` (My Oracle Support Note 1360208.1)
- **OIDOnObjCol**: Checks if an object type column is not in `OID$` (My Oracle Support Note 1360268.1)
- **SourceNotInObj**: Checks if an entry in `SOURCE$` is not in `OBJ$` (My Oracle Support Note 1360233.1)
  While running the `SourceNotInObj` procedure, optionally, you can use the `repair` option to resolve inconsistencies. Valid values: `TRUE|FALSE`. Default: `FALSE`.
- **IndIndparMismatch**: Checks for index name mismatch between partitions (My Oracle Support Note 1360285.1)
- **InvCorrAudit**: Checks for invalid `AUDIT$` entries (My Oracle Support Note 1360489.1)
- **OversizedFiles**: Checks for oversized database files (My Oracle Support Note 1360490.1)
- **PoorDefaultStorage**: Checks tablespace default storage clauses (My Oracle Support Note 1360493.1)
- **PoorStorage**: Checks objects storage clause (My Oracle Support Note 1360496.1)
- **PartSubPartMismatch**: Checks valid partition methods (My Oracle Support Note 1360504.1)

- **TabPartCountMismatch**: Checks invalid data between `OBJ$-PARTOBJ$` and `TABPART$` (My Oracle Support Note 1360514.1)
- **TabComPartObj**: Checks that the composite partition has a valid entry in `OBJ$` (My Oracle Support Note 1360515.1)
- **Mview**: Check invalid entries for materialized view (My Oracle Support Note 1360517.1)
- **ValidDir**: Checks that the directory object has valid entries (My Oracle Support Note 1360518.1)
- **DuplicateDataobj**: Checks for duplicate segment `data_object_id` (My Oracle Support Note 1360519.1)
- **ObjSyn**: Checks that a synonym has a valid entry in `OBJ$` (My Oracle Support Note 1360520.1)
- **ObjSeq**: Checks that a sequence has a valid entry in `OBJ$` (My Oracle Support Note 1360524.1)
- **UndoSeg**: Checks that undo segment has a valid entry in `SEG$` (My Oracle Support Note 1360527.1)
- **IndexSeg**: Checks that an index segment has a valid entry in `SEG$` (My Oracle Support Note 1360531.1)
- **IndexPartitionSeg**: Checks that an index partition has a valid entry in `SEG$` (My Oracle Support Note 1360535.1)
- **IndexSubPartitionSeg**: Checks that an index sub-partition has a valid entry in `SEG$` (My Oracle Support Note 1360536.1)
- **TableSeg**: Checks that a table has a valid entry in `SEG$` (My Oracle Support Note 1360889.1)
- **TablePartitionSeg**: Checks that a table partition has a valid entry in `SEG$` (My Oracle Support Note 1360890.1)
- **TableSubPartitionSeg**: Checks that a table sub-partition has a valid entry in `SEG$` (My Oracle Support Note 1360891.1)
- **PartCol**: Checks for valid entry of column partition (My Oracle Support Note 1360892.1)
- **ValidateSeg**: Checks that a segment in `SEG$` has an entry in its parent (My Oracle Support Note 1360934.1)
  While running the `ValidateSeg` procedure, optionally, you can use the `repair` option to resolve inconsistencies. Valid values: `TRUE|FALSE`. Default: `FALSE`.
- **IndPartObj**: Checks that an index partition has an entry in `OBJ$` (My Oracle Support Note 1360935.1)
- **DuplicateBlockUse**: Checks for a segment header block is used by only one segment (My Oracle Support Note 1360880.1)
- **FetUet**: Checks valid free/used space in a dictionary-managed tablespace (My Oracle Support Note 1360882.1)
- **Uet0Check**: Checks valid first extent in a dictionary-managed tablespace (My Oracle Support Note 1360883.1)
- **ExtentlessSeg**: Checks `SEG$`/`UET$` mismatch in a dictionary-managed tablespace (My Oracle Support Note 1360944.1)
- **SeglessUET**: Checks `UET$`/`SEG$` mismatch in a dictionary-managed tablespace (My Oracle Support Note 1360944.1)

- **ValidInd**: Checks that an index in `OBJ$` has a corresponding entry in the index dictionary (My Oracle Support Note 1360528.1)
- **ValidTab**: Checks that a table in `OBJ$` has a corresponding entry in `TAB$` (My Oracle Support Note 1360538.1)
- **IcolDepCnt**: Checks valid entries in `ICOLDEP$` (My Oracle Support Note 1360938.1)
- **WarnIcolDep**: Checks that an index does not have an ADT (object column) (My Oracle Support Note 1360939.1)
- **ObjIndDobj**: Checks index data_object_id mismatch between `OBJ$` and `IND$` (My Oracle Support Note 1360968.1)
- **TrgAfterUpgrade**: Checks valid entries in triggers (My Oracle Support Note 1361014.1)
- **ObjType0**: Checks that `OBJ$` has a valid type greater than 0 (My Oracle Support Note 1361015.1)
- **ValidOwner**: Checks that an entry in `OBJ$` has a valid user ID (My Oracle Support Note 1361020.1)
- **StmtAuditOnCommit**: Checks valid entries for `STMT_AUDIT_OPTION_MAP` (My Oracle Support Note 1361021.1)
- **PublicObjects**: Checks that objects are not owned by `PUBLIC` (My Oracle Support Note 1361022.1)
- **SegFreelist**: Checks that a LOB segment has a valid free list group (My Oracle Support Note 1361023.1)
- **ValidDepends**: Checks for valid dependency timestamps (My Oracle Support Note 1361045.1)
- **CheckDual**: Checks valid entries in DUAL (My Oracle Support Note 1361046.1)
- **ObjectNames**: Checks if an object has the same name as its schema owner (My Oracle Support Note 2363142.1)
- **CboHiLo**: Checks for valid entries in histograms (My Oracle Support Note 1361047.1)
- **ChkIotTs**: Checks that an IOT object does not have a segment (My Oracle Support Note 1361048.1)
- **NoSegmentIndex**: Checks for `NOSEGMENT` indexes (My Oracle Support Note 1361049.1)
- **NextObject**: Checks for valid `data_object_id` (My Oracle Support Note 2124772.1)
- **DroppedROTS**: Checks for valid entries in a read-only tablespace (My Oracle Support Note 2124774.1)
- **FilBlkZero**: Checks for zero data block address (My Oracle Support Note 2124783.1)
- **DbmsSchemaCopy**: Checks for invalid execution of `DBMS_SCHEMA_COPY` (My Oracle Support Note 2124795.1)
- **IdnseqObj**: Checks that Identity column has a valid object (My Oracle Support Note 2124805.1)
- **IdnseqSeq**: Checks that a sequence has a valid object (My Oracle Support Note 2124787.1)
- **ObjError**: Checks that an object error is valid (My Oracle Support Note 2124788.1)
- **ObjNotLob**: Checks that a LOB object has an entry in `LOB$` (My Oracle Support Note 2125104.1)

- **MaxControlfSeq**: Checks if `Control Seq` is near the limit (My Oracle Support Note 2128446.1)
- **SegNotInDeferredStg**: Checks for an invalid deferred segment (My Oracle Support Note 2298947.1)
- **SystemNotRfile1**: Checks that the system tablespace has a relative file number 1 (My Oracle Support Note 2364065.1)
- **DictOwnNonDefaultSYSTEM**: Checks that the users SYS and SYSTEM have default tablespace SYSTEM (My Oracle Support Note 2377270.1)
- **ValidateTrigger**: Checks that triggers have valid entries in their parents (My Oracle Support Note 2384373.1)
- **ObjNotTrigger**: Checks if an object trigger is not in `TRIGGER$` (My Oracle Support Note 2384392.1)
- **WarningTSMaxSCN**: Checks exposed SCN entries in tablespaces
- **InvalidTSMaxSCN**: Checks for invalid SCN entries in tablespaces (My Oracle Support Note 1360208.1)
- **OBJRecycleBin**: Checks that recycle bin objects in `OBJ$` exist in `RECYCLEBIN$` (My Oracle Support Note 2902943.1)
- **LobSeg**: Checks that a LOB segment has a valid entry in `SEG$` (My Oracle Support Note 2948392.1 and 2948408.1)
- **ObjLogicalConstraints:** Checks logical constraints in OBJ$ (My Oracle Support Note 2977609.1 and Note 2977591.1)
- **SysSequences:** Checks inconsistencies with critical sequences owned by SYS (My Oracle Support Note 2992123.1, Note 2992124.1 and Note 2992107.1)
- **ValidateObjStub:** Check orphaned stub entries in OBJ$ (My Oracle Support Note 3037043.1)
- **Critical**: Executes only critical checks
- **Full**: Executes all checks

Example 72-4 ValidateSeg with repair option

```
SQL> EXECUTE dbms_dictionary_check.ValidateSeg(repair=>TRUE)
.- ValidateSeg                 ... 2300000000 <=  *All Rel* 10/04 01:30:21
FAIL

HCKE-0023: Orphaned SEG$ Entry (Doc ID 1360934.1)
ORPHAN SEG$: SegType=DATA TS=5 RFILE/BLOCK=5/11 HWMINCR(DOBJ#)=73270

.- TableSeg                    ... 2300000000 <=  *All Rel* 10/04 01:30:21
PASS
.- TablePartitionSeg           ... 2300000000 <=  *All Rel* 10/04 01:30:21
PASS
.- TableSubPartitionSeg        ... 2300000000 <=  *All Rel* 10/04 01:30:21
PASS
ORPHAN SEG$: SegType=DATA TS=5 RFILE/BLOCK=5/11
^ Segment entry repaired - Converted to TEMPORARY
INCONSISTENCY REPAIRED - Check the trace file for repair details:
ValidateSeg repaired 1 Orphan Seg$ entries
```

```
PL/SQL procedure successfully completed.
```

Example 72-5 SourceNotInObj with repair option

```
SQL> EXECUTE dbms_dictionary_check.SourceNotInObj(repair=>TRUE)
dbms_dictionary_check on 04-OCT-2023 01:30:20
-----------------------------------------------
Catalog Version 23.0.0.0.0 (2300000000)
db_name: orcl
Is CDB?: NO
Trace File: /oracle/log/diag/rdbms/orcl/orcl/trace/
orcl_ora_3831239_DICTCHECK.trc

                                    Catalog       Fixed
Procedure Name                      Version    Vs Release    Timestamp
Result
------------------------------ ... ---------- -- ---------- --------------
------
.- SourceNotInObj              ... 2300000000 <=  *All Rel* 10/04 01:30:20
FAIL

HCKE-0003: SOURCE$ for OBJ# not in OBJ$ (Doc ID 1360233.1)
SOURCE$ has 10 rows for 1 OBJ# values not in OBJ$


INCONSISTENCY REPAIRED - Check the trace file for repair details:
SourceNotInObj_Repair: DELETED 10 objects from SOURCE$ not found in OBJ$

PL/SQL procedure successfully completed.
```

73
DBMS_DIMENSION

DBMS_DIMENSION enables you to verify dimension relationships and provides an alternative to the Enterprise Manager Dimension Wizard for displaying a dimension definition.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for detailed conceptual and usage information about the DBMS_DIMENSION package

This chapter contains the following topics:

- Security Model
- Summary of DBMS_DIMENSION Subprograms

73.1 DBMS_DIMENSION Security Model

Security on this package can be controlled by granting EXECUTE to selected users or roles.

A user can validate or describe all the dimensions in his own schema. To validate or describe a dimension in another schema, you must have either an object privilege on the dimension or one of the following system privileges: CREATE ANY DIMENSION, ALTER ANY DIMENSION, and DROP ANY DIMENSION.

73.2 Summary of DBMS_DIMENSION Subprograms

This table lists the DBMS_DIMENSION subprograms and briefly describes them.

Table 73-1 DBMS_DIMENSION Package Subprograms

Subprogram	Description
DESCRIBE_DIMENSION Procedure	Prints out the definition of the input dimension, including dimension owner and name, levels, hierarchies, and attributes
VALIDATE_DIMENSION Procedure	Verifies that the relationships specified in a dimension are correct

73.2.1 DESCRIBE_DIMENSION Procedure

This procedure displays the definition of the dimension, including dimension name, levels, hierarchies, and attributes. It displays the output using the DBMS_OUTPUT package.

Syntax

```
DBMS_DIMENSION.DESCRIBE_DIMENSION (
   dimension  IN VARCHAR2);
```

Parameters

Table 73-2 DESCRIBE_DIMENSION Procedure Parameter

Parameter	Description
dimension	The owner and name of the dimension in the format of owner.name.

73.2.2 VALIDATE_DIMENSION Procedure

This procedure verifies that the relationships specified in a dimension are valid. The rowid for any row that is found to be invalid will be stored in the table DIMENSION_EXCEPTIONS in the user's schema.

Syntax

```
DBMS_DIMENSION.VALIDATE_DIMENSION (
   dimension               IN VARCHAR2,
   incremental             IN BOOLEAN := TRUE,
   check_nulls             IN BOOLEAN := FALSE,
   statement_id            IN VARCHAR2 := NULL );
```

Parameters

Table 73-3 VALIDATE_DIMENSION Procedure Parameters

Parameter	Description
dimension	The owner and name of the dimension in the format of owner.name.
incremental	If TRUE, check only the new rows for tables of this dimension. If FALSE, check all the rows.
check_nulls	If TRUE, then all level columns are verified to be non-null. If FALSE, this check is omitted. Specify FALSE when non-NULLness is guaranteed by other means, such as NOT NULL constraints.
statement_id	A client-supplied unique identifier to associate output rows with specific invocations of the procedure.

74
DBMS_DISTRIBUTED_TRUST_ADMIN

DBMS_DISTRIBUTED_TRUST_ADMIN procedures maintain the Trusted Servers List. Use these procedures to define whether a server is trusted. If a database is not trusted, Oracle refuses current user database links from the database.

This chapter contains the following topics:

- Overview
- Security Model
- Examples
- Summary of DBMS_DISTRIBUTED_TRUST_ADMIN Subprograms

74.1 DBMS_DISTRIBUTED_TRUST_ADMIN Overview

Oracle uses local Trusted Servers Lists, along with enterprise domain membership lists stored in the enterprise LDAP directory service, to determine if another database is trusted. The LDAP directory service entries are managed with the Enterprise Security Manager Tool in Oracle Enterprise Manager.

Oracle considers another database to be "trusted" if it meets the following criteria:

1. It is in the same enterprise domain in the directory service as the local database.
2. The enterprise domain is marked as trusted in the directory service.
3. It is not listed as untrusted in the local Trusted Servers List. Current user database links will only be accepted from another database if both databases involved trust each other.

You can list a database server locally in the Trusted Servers List regardless of what is listed in the directory service. However, if you list a database that is not in the same domain as the local database, or if that domain is untrusted, the entry will have no effect.

This functionality is part of the Enterprise User Security feature of the Oracle Advanced Security Option.

74.2 DBMS_DISTRIBUTED_TRUST_ADMIN Security Model

To execute DBMS_DISTRIBUTED_TRUST_ADMIN, the EXECUTE_CATALOG_ROLE role must be granted to the DBA. To select from the view TRUSTED_SERVERS, the SELECT_CATALOG_ROLE role must be granted to the DBA.

It is important to know whether all servers are trusted or not trusted. Trusting a particular server with the ALLOW_SERVER procedure does not have any effect if the database already trusts all databases, or if that database is already trusted. Similarly, denying a particular server with the DENY_SERVER procedure does not have any effect if the database already does not trust any database or if that database is already untrusted.

The procedures DENY_ALL and ALLOW_ALL delete all entries (in other words, server names) that are explicitly allowed or denied using the ALLOW_SERVER procedure or DENY_SERVER procedure respectively.

74.3 DBMS_DISTRIBUTED_TRUST_ADMIN Examples

If you have not yet used the package `DBMS_DISTRIBUTED_TRUST_ADMIN` to change the trust listing, by default you trust all databases in the same enterprise domain if that domain it listed as trusted in the directory service:

```
SELECT * FROM TRUSTED_SERVERS;
TRUST
NAME
--------- --------------------
Trusted
All
```

Because all servers are currently trusted, you can execute the DENY_SERVER Procedure and specify that a particular server is not trusted:

```
EXECUTE DBMS_DISTRIBUTED_TRUST_ADMIN.DENY_SERVER
                           ('SALES.US.AMERICAS.ACME_AUTO.COM');
PL/SQL procedure successfully completed.

SELECT * FROM TRUSTED_SERVERS;
TRUST
NAME
--------- ----------------------------------------------
Untrusted SALES.US.AMERICAS.ACME_AUTO.COM
```

By executing the DENY_ALL Procedure, you can choose to not trust any database server:

```
EXECUTE DBMS_DISTRIBUTED_TRUST_ADMIN.DENY_ALL;

PL/SQL procedure successfully completed.

SELECT * FROM TRUSTED_SERVERS;

TRUST
NAME
--------- ----------------------------------------------
Untrusted
All
```

The ALLOW_SERVER Procedure can be used to specify that one particular database is to be trusted:

```
EXECUTE DBMS_DISTRIBUTED_TRUST_ADMIN.ALLOW_SERVER
                            ('SALES.US.AMERICAS.ACME_AUTO.COM');
PL/SQL procedure successfully completed.

SELECT * FROM TRUSTED_SERVERS;
TRUST
NAME
--------- ----------------------------------------------
Trusted   SALES.US.AMERICAS.ACME_AUTO.COM
```

74.4 Summary of DBMS_DISTRIBUTED_TRUST_ADMIN Subprograms

This table lists and briefly describes the `DBMS_DISTRIBUTED_TRUST_ADMIN` subprograms.

Table 74-1 ***DBMS_DISTRIBUTED_TRUST_ADMIN Package Subprograms***

Subprogram	Description
ALLOW_ALL Procedure	Empties the list and inserts a row indicating that all servers should be trusted
ALLOW_SERVER Procedure	Enables a specific server to be allowed access even though `deny` all is indicated in the list
DENY_ALL Procedure	Empties the list and inserts a row indicating that all servers should be untrusted
DENY_SERVER Procedure	Enables a specific server to be denied access even though allow all is indicated in the list

74.4.1 ALLOW_ALL Procedure

This procedure empties the Trusted Servers List and specifies that all servers that are members of a trusted domain in an enterprise directory service and that are in the same domain are allowed access.

The view `TRUSTED_SERVERS` will show "`TRUSTED ALL`" indicating that the database trusts all servers that are currently trusted by the enterprise directory service.

Syntax

```
DBMS_DISTRIBUTED_TRUST_ADMIN.ALLOW_ALL;
```

Usage Notes

`ALLOW_ALL` only applies to servers listed as trusted in the enterprise directory service and in the same enterprise domain.

74.4.2 ALLOW_SERVER Procedure

This procedure ensures that the specified server is considered trusted (even if you have previously specified "`deny all`").

Syntax

```
DBMS_DISTRIBUTED_TRUST_ADMIN.ALLOW_SERVER (
   server IN VARCHAR2);
```

Parameters

Table 74-2 ALLOW_SERVER Procedure Parameters

Parameter	Description
`server`	Unique, fully-qualified name of the server to be trusted.

Usage Notes

If the Trusted Servers List contains the entry "`deny all`", then this procedure adds a specification indicating that a specific database (for example, `DBx`) is to be trusted.

If the Trusted Servers List contains the entry "`allow all`", and if there is no "`deny DBx`" entry in the list, then executing this procedure causes no change.

If the Trusted Servers List contains the entry "`allow all`", and if there is a "`deny DBx`" entry in the list, then that entry is deleted.

74.4.3 DENY_ALL Procedure

This procedure empties the Trusted Servers List and specifies that all servers are denied access.

The view `TRUSTED_SERVERS` will show "`UNTRUSTED ALL`" indicating that no servers are currently trusted.

Syntax

```
DBMS_DISTRIBUTED_TRUST_ADMIN.DENY_ALL;
```

74.4.4 DENY_SERVER Procedure

This procedure ensures that the specified server is considered untrusted (even if you have previously specified `allow all`).

Syntax

```
DBMS_DISTRIBUTED_TRUST_ADMIN.DENY_SERVER (
   server IN VARCHAR2);
```

Parameters

Table 74-3 DENY_SERVER Procedure Parameters

Parameter	Description
`server`	Unique, fully-qualified name of the server to be untrusted.

Usage Notes

If the Trusted Servers List contains the entry `allow all`, then this procedure adds an entry indicating that the specified database (for example, `DBx`) is not to be trusted.

If the Trusted Servers List contains the entry "`deny all`", and if there is no "`allow DBx`" entry in the list, then this procedure causes no change.

If the Trusted Servers List contains the entry "`deny all`", and if there is an "`allow DBx`" entry, then this procedure causes that entry to be deleted.

75
DBMS_DNFS

The `DBMS_DNFS` package provides an interface to assists in creating a database using files in the backup set.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_DNFS Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide*

75.1 DBMS_DNFS Security Model

This package has to be executed by users with `SYSDBA` privileges.

75.2 Summary of DBMS_DNFS Subprograms

The `DBMS_DNFS` package includes the `CLONEDB_RENAMEFILE`, `RESTORE_DATAFILE_PERMISSIONS`, and `UNMOUNTVOLUME` procedures.

Table 75-1 DBMS_DNFS Package Subprograms

Subprogram	Description
CLONEDB_RENAMEFILE Procedure	Renames datafiles that were pointing to the backup set to the actual file name in cloned database.
RESTORE_DATAFILE_PERMISSIONS Procedure	This procedure will restore the file permissions for the data files belonging to the PDB to read write. PDB data files are converted to read only when one or more clones are created using this source PDB.
UNMOUNTVOLUME Procedure	This procedure cleans up the cached Direct NFS file handles in SGA.

75.2.1 CLONEDB_RENAMEFILE Procedure

This procedure is used to rename datafiles that were pointing to the backup set to the actual file name in cloned database.

The `srcfile` is the file name that represents the data file in the backup image copy or a read-only storage snapshot. The `destfile` destination file path must point to a NFS volume where cloneDB datafiles will be created. When the procedure is run successfully, the control file record is updated with the new datafile name.

Syntax

```
DBMS_DNFS.CLONEDB_RENAMEFILE (
   srcfile        IN         VARCHAR2,
   destfile       IN         VARCHAR2);
```

Parameters

Table 75-2 *CLONEDB_RENAMEFILE Procedure Parameters*

Parameter	Description
`srcfile`	Source datafile name in the control file
`destfile`	New datafile name

75.2.2 RESTORE_DATAFILE_PERMISSIONS Procedure

This procedure restores the file permissions for the data files belonging to the PDB to `read write`. PDB data files are converted to `read only` when one or more clones are created using this source PDB. This procedure must be invoked after the last clone is dropped to restore the data file permissions for the source PDB files.

Syntax

```
DBMS_DNFS.RESTORE_DATAFILE_PERMISSIONS (
   pdb_name       IN         VARCHAR2  DEFAULT NULL);
```

Parameters

Table 75-3 *RESTORE_DATAFILE_PERMISSIONS Procedure Parameters*

Parameter	Description
`pdb_name`	Name of pluggable database whose file permissions need to be restored.

75.2.3 UNMOUNTVOLUME Procedure

This procedure cleans up the cached Direct NFS file handles in SGA. This procedure must be used when any NFS mount changes are being made via the operating system and the OS unmount and mount commands are invoked.

The server is the host name or IP address of the NFS server. The volume is the exported path from the NFS server. This procedure will fail if the database has any open files in the volume that is being unmounted.

Syntax

```
DBMS_DNFS.UNMOUNTVOLUME (
   server         IN         VARCHAR2,
   volume         IN         VARCHAR2);
```

Parameters

Table 75-4 *UNMOUNTVOLUME Procedure Parameters*

Parameter	Description
`server`	NFS server host name or IP address.
`volume`	Exported path from the NFS server.

76
DBMS_DST

The DBMS_DST package provides an interface to apply the Daylight Saving Time (DST) patch to the Timestamp with Time Zone datatype.

This chapter contains the following topics:

- Overview
- Security Model
- Views
- Summary of DBMS_DST Subprograms

> **See Also:**
>
> - *Oracle Database Globalization Support Guide*
> - *Oracle Database Reference*

76.1 DBMS_DST Overview

The transition period during which Daylight Saving Time comes into effect, or stops being in effect, has the potential for problems, such as data loss, when handling timestamps with time zone data. The `DBMS_DST` package enables working with these transitions in the context of a set of rules.

76.2 DBMS_DST Security Model

The `DBMS_DST` package is an invoker's rights package.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information about using Invoker Rights or Definer Rights

The execute privilege on the package is granted to the `EXECUTE_CATALOG_ROLE` role. This role is normally granted to selected users to allow `EXECUTE` privileges for packages and procedures in the data dictionary.

The user that invokes the package must have the following privileges:

- `CREATE ANY TABLE`
- `ALTER ANY TABLE`

- DROP ANY TABLE
- SELECT ANY TABLE
- LOCK ANY TABLE
- ALTER ANY INDEX
- ALTER ANY TRIGGER
- UPDATE ANY TABLE
- EXECUTE ANY TYPE

76.3 DBMS_DST Views

The DBMS_DST package uses views to display table information.

These views are shown in the following table. They are further described in the *Oracle Database Reference*:

Table 76-1 Views used by DBMS_DST

View	Description
DBA_TSTZ_TABLES	Displays information about all tables in the database, which have columns defined on TIMESTAMP WITH TIME ZONE datatypes or object types containing attributes of TIMESTAMP WITH TIME ZONE datatypes. Its columns are the same as those in ALL_TSTZ_TABLES
USER_TSTZ_TABLES	Displays information about the tables owned by the current user, which have columns defined on TIMESTAMP WITH TIME ZONE datatypes or object types containing attributes of TIMESTAMP WITH TIME ZONE datatypes. Its columns (except for OWNER) are the same as those in ALL_TSTZ_TABLES.
ALL_TSTZ_TABLES	Displays information about the tables accessible to the current user, which have columns defined on TIMESTAMP WITH TIME ZONE datatypes or object types containing attributes of TIMESTAMP WITH TIME ZONE datatypes

76.4 Summary of DBMS_DST Subprograms

This table lists and describes the DBMS_DST package subprograms.

Table 76-2 DBMS_DST Package Subprograms

Subprogram	Description
BEGIN_PREPARE Procedure	Starts a prepare window
BEGIN_UPGRADE Procedure	Starts an upgrade window
CREATE_AFFECTED_TABLE Procedure	Creates a table that has the schema shown in the comments for the FIND_AFFECTED_TABLES Procedure
CREATE_ERROR_TABLE Procedure	Creates a log error table

Table 76-2 (Cont.) DBMS_DST Package Subprograms

Subprogram	Description
CREATE_TRIGGER_TABLE Procedure	Creates a table that is used to record active triggers disabled before performing upgrade on the table, having not been enabled due to fatal failure during the upgrading process
END_PREPARE Procedure	Ends a prepare window
END_UPGRADE Procedure	Ends an upgrade window
FIND_AFFECTED_TABLES Procedure	Finds all the tables that have affected TSTZ data due to the new timezone version
UPGRADE_DATABASE Procedure	Upgrades all tables in the database that have one or more columns defined on the TSTZ type, or an ADT containing the TSTZ type
UPGRADE_SCHEMA Procedure	Upgrades tables in a specified list of schemas that has one or more columns defined on the TSTZ type, or an ADT containing the TSTZ type
UPGRADE_TABLE Procedure	Upgrades a specified list of tables that has one or more columns defined on the TSTZ type or an ADT containing the TSTZ type

76.4.1 BEGIN_PREPARE Procedure

This procedure starts a prepare window. Once a prepare window is started successfully, the database property '`DST_UPGRADE_STATE`' is set to '`PREPARE`', and the database property '`SECONDARY_TT_VERSION`' is set to a new timezone version.

The prepare window lets a DBA investigate data affected by the upgrade, and so judge when it is optimal to perform the upgrade. The prepare window can overlap normal database operation.

Syntax

```
DBMS_DST.BEGIN_PREPARE (
   new_version               IN  BINARY_INTEGER);
```

Parameters

Table 76-3 BEGIN_PREPARE Procedure Parameters

Parameter	Description
`new_version`	New timezone version to which the database is to be prepared to upgrade

76.4.2 BEGIN_UPGRADE Procedure

This procedure starts an upgrade window.

Syntax

```
DBMS_DST.BEGIN_UPGRADE (
   new_version               IN  BINARY_INTEGER,
   error_on_overlap_time     IN  BOOLEAN := FALSE,
   error_on_nonexisting_time IN  BOOLEAN := FALSE);
```

Parameters

Table 76-4 BEGIN_UPGRADE Procedure Parameters

Parameter	Description
`new_version`	New timezone version to which the database is to be upgraded
`error_on_overlap_time`	Boolean flag indicating whether to report errors on the 'overlap' time semantic conversion error. The default is `FALSE`. For more information about boundary cases, see *Oracle Database SQL Language Reference*.
`error_on_nonexisting_time`	Boolean flag indicating whether to report errors on the 'non-existing' time semantic conversion error. The default value is `FALSE`.

76.4.3 CREATE_AFFECTED_TABLE Procedure

This procedure creates a table that has the schema shown in the comments for the FIND_AFFECTED_TABLES Procedure.

Syntax

```
DBMS_DST.CREATE_AFFECTED_TABLE (
   table_name      IN  VARCHAR2);
```

Parameters

Table 76-5 CREATE_AFFECTED_TABLE Procedure Parameters

Parameter	Description
`table_name`	Name of the table created

Usage Notes

This procedures takes a `table_name` without schema qualification, creating a table within the current user schema.

Related Topics

- FIND_AFFECTED_TABLES Procedure
  This procedure finds all the tables which have affected TSTZ data due to the new timezone version.

76.4.4 CREATE_ERROR_TABLE Procedure

This procedure creates a log error table.

The table has the following schema:

```
CREATE TABLE dst$error_table(
   table_owner     VARCHAR2(30),
   table_name      VARCHAR2(30),
   column_name     VARCHAR2(4000),
   rid             ROWID,
   error_number    NUMBER)
```

Syntax

```
DBMS_DST.CREATE_ERROR_TABLE (
   table_name      IN  VARCHAR2);
```

Parameters

Table 76-6 CREATE_ERROR_TABLE Procedure Parameters

Parameter	Description
table_name	Name of the table created

Usage Notes

- This procedures takes a `table_name` without schema qualification, creating a table within the current user schema.
- The error number is found when upgrading time zone file and timestamp with time zone data. For more information about error handling when upgrading time zone file and timestamp with time zone data, see Oracle Database Globalization Support Guide

76.4.5 CREATE_TRIGGER_TABLE Procedure

This procedure creates a table to record active triggers that are disabled before performing upgrade on the table, having not been enabled due to fatal failure during the upgrading process.

The table that has the following schema.

```
CREATE TABLE dst_trigger_table (
trigger_owner  VARCHAR2(30),
trigger_name   VARCHAR2(30));
```

Syntax

```
DBMS_DST.CREATE_TRIGGER_TABLE (
   table_name     IN    VARCHAR2);
```

Parameters

Table 76-7 CREATE_TRIGGER_TABLE Procedure Parameters

Parameter	Description
table_name	Name of table to be created

Usage Notes

This procedures takes a `table_name` without schema qualification, creating a table within the current user schema.

76.4.6 END_PREPARE Procedure

This procedure ends a prepare window.

Syntax

```
DBMS_DST.END_PREPARE;
```

76.4.7 END_UPGRADE Procedure

This procedure ends an upgrade window. An upgraded window is ended if all the affected user tables have been upgraded. Otherwise, the `OUT` parameter `num_of_failures` indicates how many tables have not been converted.

Syntax

```
DBMS_DST.END_UPGRADE (
  num_of_failures    OUT  BINARY_INTEGER);
```

Parameters

Table 76-8 END_UPGRADE Procedure Parameters

Parameter	Description
`num_of_failures`	Number of tables that fail to complete

76.4.8 FIND_AFFECTED_TABLES Procedure

This procedure finds all the tables which have affected TSTZ data due to the new timezone version.

This procedure can only be invoked during a prepare window. The tables which have affected TSTZ data are recorded into a table indicated by parameter `affected_tables`. If semantic errors must be logged, they are recorded into a table indicated by parameter `log_errors_table`.

Syntax

```
DBMS_DST.FIND_AFFECTED_TABLES (
   affected_tables             IN  VARCHAR2 =: 'sys.dst$affected_tables',
   log_errors                  IN  BOOLEAN := FALSE,
   log_errors_table            IN  VARCHAR2 =: 'sys.dst$error_table',
   parallel                    IN  BOOLEAN := FALSE);
```

Parameters

Table 76-9 FIND_AFFECTED_TABLES Procedure Parameters

Parameter	Description

affected_tables	Name of table with the following schema:
	`CREATE TABLE dst$affected_tables (`
	`table_owner VARCHAR2(30),`
	`table_name VARCHAR2(30),`
	`column_name VARCHAR2(4000),`
	`row_count NUMBER,`
	`error_count NUMBER)`
	The table can be created with the CREATE_AFFECTED_TABLE Procedure.

log_errors	Boolean flag indicating whether to log errors during upgrade. If FALSE, no error is logged into the log_errors_table after aborting conversion of the current table. If TRUE, the error is logged to the log_errors_table. The default is FALSE.

log_errors_table	Table name with the following schema:
	`CREATE TABLE dst$error_table (`
	`table_owner VARCHAR2(30),`
	`table_name VARCHAR2(30),`
	`column_name VARCHAR2(4000),`
	`rid ROWID,`
	`error_number NUMBER)`
	The table can be created with the CREATE_ERROR_TABLE Procedure. The rid column records the rowids of the offending rows, and the error_number column records the corresponding error number.

parallel	Boolean flag indicating whether to find the affected tables using parallel queries or serial queries. The default is FALSE.

76.4.9 UPGRADE_DATABASE Procedure

This procedure upgrades all tables in the database, which have one or more columns defined on the TSTZ type or an ADT containing the TSTZ type.

This procedure can only be invoked after an upgrade window has been started. Each table is upgraded in an atomic transaction. Note that, a base table and its materialized view log table are upgraded in an atomic transaction.

Syntax

```
DBMS_DST.UPGRADE_DATABASE  (
   num_of_failures             OUT BINARY_INTEGER,
   upgrade_data                IN  BOOLEAN := TRUE,
   parallel                    IN  BOOLEAN := FALSE,
   continue_after_errors       IN  BOOLEAN := TRUE,
   log_errors                  IN  BOOLEAN := FALSE,
   log_errors_table            IN  VARCHAR2 =: 'sys.dst$error_table' ,
   error_on_overlap_time       IN  BOOLEAN := FALSE,
```

```
  error_on_nonexisting_time   IN  BOOLEAN := FALSE,
  log_triggers_table          IN  VARCHAR2 := 'sys.dst$trigger_table');
```

Parameters

Table 76-10 UPGRADE_DATABASE Procedure Parameters

Parameter	Description
num_of_failures	Number of tables that fail to complete
upgrade_data	Boolean flag indicating whether to convert TSTZ data using the new Time Zone patch File (TRUE), or to leave it unconverted (FALSE).The default is TRUE.
parallel	Boolean flag indicating whether to convert tables using PDML (Parallel DML) or Serial DML.The default is FALSE.
continue_after_errors	Boolean flag indicating whether to continue after upgrade fails on the current table. The default is TRUE.
log_errors	Boolean flag indicating whether to log errors during upgrade. If FALSE, no error is logged into the log_errors_table after aborting conversion of the current table. If TRUE, errors are logged to the log_errors_table. The default is FALSE.

log_errors_table	Table name with the following schema:
	`CREATE TABLE dst$error_table (`
	`table_owner VARCHAR2(30),`
	`table_name VARCHAR2(30),`
	`column_name VARCHAR2(4000),`
	`rid ROWID,`
	`error_number NUMBER)`
	The table can be created with the CREATE_ERROR_TABLE Procedure. The rid column records the rowids of the offending rows, and the error_number column records the corresponding error number.

error_on_overlap_time	Boolean flag indicating whether to report errors on the 'overlap' time semantic conversion error. The default is TRUE.
error_on_nonexisting_time	Boolean flag indicating whether to report errors on the 'non-existing' time semantic conversion error. The default is TRUE.
log_triggers_table	Table to log triggers which are disabled before upgrade, having not been enabled due to a fatal failure when performing an upgrade

76.4.10 UPGRADE_SCHEMA Procedure

This procedure upgrades tables in a specified list of schemas that have one or more columns defined on the TSTZ type, or an ADT containing the TSTZ type.

This procedure can be invoked only after an upgrade window has been started. Each table is upgraded in an atomic transaction. Note that a base table and its materialized view log table are upgraded in an atomic transaction.

Syntax

```
DBMS_DST.UPGRADE_SCHEMA (
   num_of_failures             OUT BINARY_INTEGER,
   schema_list                 IN  VARCHAR2,
   upgrade_data                IN  BOOLEAN := TRUE,
   parallel                    IN  BOOLEAN := FALSE,
   continue_after_errors       IN  BOOLEAN := TRUE,
   log_errors                  IN  BOOLEAN := FALSE,
   log_errors_table            IN  VARCHAR2 =: 'sys.dst$error_table' ,
   error_on_overlap_time       IN  BOOLEAN := FALSE,
   error_on_nonexisting_time   IN  BOOLEAN := FALSE,
   log_triggers_table          IN  VARCHAR2 := 'sys.dst$trigger_table');
```

Parameters

Table 76-11 UPGRADE_SCHEMA Procedure Parameters

Parameter	Description
num_of_failures	Number of tables that fail to complete
schema_list	Schema name list (comma separated strings)
upgrade_data	Boolean flag indicating whether to convert TSTZ data using the new Time Zone patch File (TRUE) or to leave unconverted (FALSE).The default is TRUE.
parallel	Boolean flag indicating whether to convert tables using PDML (Parallel DML) or Serial DML.The default is FALSE.
continue_after_errors	Boolean flag indicating whether to continue after upgrade fails on the current table.The default is TRUE.
log_errors	Boolean flag indicating whether to log errors during upgrade. If FALSE, no error is logged into the log_errors_table after aborting conversion of the current table. If TRUE, the error is logged to the log_errors_table. The default is FALSE.

log_errors_table	Table name with the following schema:
	`CREATE TABLE dst$error_table (`
	`table_owner VARCHAR2(30),`
	`table_name VARCHAR2(30),`
	`column_name VARCHAR2(4000),`
	`rid ROWID,`
	`error_number NUMBER)`
	The table can be created with the CREATE_ERROR_TABLE Procedure. The rid column records the rowids of the offending rows, and the error_number column records the corresponding error number.

error_on_overlap_time	Boolean flag indicating whether to report errors on the 'overlap' time semantic conversion error. The default is TRUE.
error_on_nonexisting_time	Boolean flag indicating whether to report errors on the 'non-existing' time semantic conversion error. The default is TRUE.

Table 76-11 (Cont.) UPGRADE_SCHEMA Procedure Parameters

Parameter	Description
log_triggers_table	Table to log triggers that are disabled before upgrade, having not been enabled due to a fatal failure when performing an upgrade

76.4.11 UPGRADE_TABLE Procedure

This procedure upgrades a specified list of tables that have one or more columns defined on the TSTZ type, or an ADT containing the TSTZ type.

Syntax

```
DBMS_DST.UPGRADE_TABLE  (
   num_of_failures             OUT BINARY_INTEGER,
   table_list                  IN  VARCHAR2,
   upgrade_data                IN  BOOLEAN := TRUE,
   parallel                    IN  BOOLEAN := FALSE,
   continue_after_errors       IN  BOOLEAN := TRUE,
   log_errors                  IN  BOOLEAN := FALSE,
   log_errors_table            IN  VARCHAR2 =: 'sys.dst$error_table' ,
   error_on_overlap_time       IN  BOOLEAN := FALSE,
   error_on_nonexisting_time   IN  BOOLEAN := FALSE,
   log_triggers_table          IN  VARCHAR2 := 'sys.dst$trigger_table',
   atomic_upgrade              IN  BOOLEAN := FALSE);
```

Parameters

Table 76-12 UPGRADE_TABLE Procedure Parameters

Parameter	Description
num_of_failures	Number of tables that fail to complete
table_list	Table name list (comma separated strings)
upgrade_data	Boolean flag indicating whether to convert TSTZ data using the new Time Zone patch File (TRUE), or to leave unconverted (FALSE). The default is TRUE.
parallel	Boolean flag indicating whether to convert tables using PDML (Parallel DML), or Serial DML. The default is FALSE.
continue_after_errors	Boolean flag indicating whether to continue after upgrade fails on the current table. The default is TRUE.
log_errors	Boolean flag indicating whether to log errors during upgrade. If FALSE, no error is logged into the log_errors_table after aborting conversion of the current table. If TRUE, the error is logged to the log_errors_table. The default is FALSE.

Table 76-12 (Cont.) UPGRADE_TABLE Procedure Parameters

Parameter	Description

log_errors_table	Table name with the following schema:
	`CREATE TABLE dst$error_table (`
	`table_owner VARCHAR2(30),`
	`table_name VARCHAR2(30),`
	`column_name VARCHAR2(4000),`
	`rid ROWID,`
	`error_number NUMBER)`
	The table can be created with the CREATE_ERROR_TABLE Procedure. The `rid` parameter records the rowids of the offending rows and the corresponding error number.

error_on_overlap_time	Boolean flag indicating whether to report errors on the 'overlap' time semantic conversion error. The default is `TRUE`.
error_on_nonexisting_time	Boolean flag indicating whether to report errors on the 'non-existing' time semantic conversion error. The default is `TRUE`.
log_triggers_table	Table to log triggers that are disabled before upgrade, having not been enabled due to a fatal failure when performing an upgrade
atomic_upgrade	Boolean flag indicating whether to convert the listed tables atomically (in a single transaction). If `FALSE`, each table is converted in its own transaction. The default is `FALSE`.

Usage Notes

This procedure can only be invoked after an upgrade window has been started. The table list has to satisfy the following partial ordering:

1. If a base table has a materialized view log table, the log table must be the next item in the list.
2. If the container table for a materialized view appears in the list, the materialized view's 'non-upgraded' base tables and log tables must appear in the table list and before the container table.

A base table and its materialized view log table need to be upgraded in an atomic transaction by specifying `atomic_upgrade` to `TRUE`.

77
DBMS_EDITIONS_UTILITIES

The `DBMS_EDITIONS_UTILITIES` package provides helper functions for edition-related operations.

The chapter contains the following topics:

- Overview
- Security Model
- Exceptions
- Summary of DBMS_EDITIONS_UTILITIES Subprograms

77.1 DBMS_EDITIONS_UTILITIES Overview

The `DBMS_EDITIONS_UTILITIES` package implements an interface which provides helper functions for edition-related operations.

77.2 DBMS_EDITIONS_UTILITIES Security Model

This package is owned by `SYS` with execute access granted to `PUBLIC`. It runs with invoker's rights, that is, with the security profile of the caller.

77.3 DBMS_EDITIONS_UTILITIES Exceptions

The table in this topic lists exception messages created by DBMS_EDITIONS_UTILITIES.

Table 77-1 DBMS_EDITIONS_UTILITIES Error Messages

Error	Description
`ORA-38817`	Insufficient privileges
`ORA-942`	Missing table

77.4 Summary of DBMS_EDITIONS_UTILITIES Subprograms

This table lists and describes the `DBMS_EDITIONS_UTILITIES` package subprograms.

Table 77-2 DBMS_EDITIONS_UTILITIES Package Subprograms

Subprogram	Description
CLEAN_UNUSABLE_EDITIONS Procedure	Drops covered objects in unusable editions, and drops empty unusable editions if possible.

Table 77-2 (Cont.) DBMS_EDITIONS_UTILITIES Package Subprograms

Subprogram	Description
SET_EDITIONING_VIEWS_READ_ONLY Procedure	Given the schema name and table name, this procedure sets the corresponding editioning views in all editions to `READ ONLY` or `READ/WRITE`
SET_NULL_COLUMN_VALUES_TO_EXPR Procedure	For use only during an edition-based redefinition (EBR) exercise

77.4.1 CLEAN_UNUSABLE_EDITIONS Procedure

Drops all covered objects in any unusable editions and drops empty unusable editions.

The `CLEAN_UNUSABLE_EDITIONS` procedure executes immediately and in its entirety.

Syntax

```
DBMS_EDITIONS_UTILITIES.CLEAN_UNUSABLE_EDITIONS ( );
```

Usage Notes

The `COMPATIBLE` parameter must be set to 12.2.0 or higher for this procedure to execute.

> **See Also:**
>
> - *Oracle Database SQL Language Reference* for information about `DROP EDITION`
> - *Oracle Database Development Guide* for more information about using edition-based redefinition

77.4.2 SET_EDITIONING_VIEWS_READ_ONLY Procedure

Given the schema name and table name, this procedure sets the corresponding editioning views in all editions to `READ ONLY` or `READ/WRITE`.

Syntax

```
DBMS_EDITIONS_UTILITIES.SET_EDITIONING_VIEWS_READ_ONLY (
   table_name IN VARCHAR2,
   owner      IN VARCHAR2 DEFAULT NULL,
   read_only  IN BOOLEAN  DEFAULT TRUE);
```

Parameters

Table 77-3 SET_EDITIONING_VIEWS_READ_ONLY Procedure Parameters

Parameter	Description
`table_name`	Base table of the editioning views
`owner`	Base table schema. The default (or `NULL`) is the current schema.

Table 77-3 (Cont.) SET_EDITIONING_VIEWS_READ_ONLY Procedure Parameters

Parameter	Description
`read_only`	`TRUE` to set the views to read-only; `FALSE` (or `NULL`) sets the views to `READ/WRITE`. Default is `TRUE`.

Usage Notes

The user must have the following privileges:

- Owner of the table, or have the `ALTER ANY TABLE` system privileges
- `USE` object privilege on all the editions for which the views are defined

77.4.3 SET_NULL_COLUMN_VALUES_TO_EXPR Procedure

This procedure replaces `NULL` values in a replacement column with the value of an expression.

The expression evaluation cost is deferred to future updates and queries. The procedure is intended for use only during an edition-based redefinition (EBR) exercise.

> **See Also:**
>
> - *Oracle Database Development Guide* regarding transforming pre- to post-upgrade representation

Syntax

```
DBMS_EDITIONS_UTILITIES.SET_NULL_COLUMN_VALUES_TO_EXPR;
   table_name     IN  VARCHAR2,
   column_name    IN  VARCHAR2,
   expression     IN  VARCHAR2);
```

Parameters

Table 77-4 SET_NULL_COLUMN_VALUES_TO_EXPR Procedure Parameters

Parameter	Description
`table_name`	A potentially schema-qualified table name
`column_name`	Name of the column to be updated
`expression`	An expression composed of columns in the same table, constants, and SQL functions

78
DBMS_EPG

The DBMS_EPG package implements the embedded PL/SQL gateway that enables a Web browser to invoke a PL/SQL stored procedure through an HTTP listener.

This chapter contains the following topics:

- Overview
- Security Model
- Exceptions
- Data Structures
  - VARCHAR2_TABLE Table Type
- Subprogram Groups
- Summary of DBMS_EPG Subprograms

78.1 DBMS_EPG Overview

The DBMS_EPG package is a platform on which PL/SQL users develop and deploy PL/SQL Web applications. The embedded PL/SQL gateway is an embedded version of the gateway that runs in the XML database HTTP server in the Oracle database. It provides the core features of mod_plsql in the database but does not require the Oracle HTTP server powered by Apache.

In order to make a PL/SQL application accessible from a browser by way of HTTP, a Database Access Descriptor (DAD) must be created and mapped to a virtual path. A DAD is a set of configuration values used for database access and the virtual path mapping makes the application accessible under a virtual path of the XML DB HTTP Server. A DAD is represented as a servlet in XML DB HTTP Server.

78.2 DBMS_EPG Security Model

The XDBADMIN role is required to invoke the configuration interface. It may invoked by the database user "XDB".

The authorization interface can be invoked by any user.

78.3 DBMS_EPG Exceptions

The table in this topic lists the exceptions raised by the DBMS_EPG package.

Table 78-1 DBMS_EPG Exceptions

Exception	Error Code	Description
DAD_NOT_FOUND	20000	Database Access Descriptor (DAD) %s not found. Ensure that the name of the DAD is correct and that it exists.

78.4 DBMS_EPG Data Structures

The `DBMS_EPG` package defines a `TABLE` type.

VARCHAR2_TABLE Table Type

This type is used by the procedures `GET_ALL_GLOBAL_ATTRIBUTES`, `GET_ALL_DAD_ATTRIBUTES`, `GET_ALL_DAD_MAPPINGS`, and `GET_DAD_LIST` to return lists of attribute names, attribute values, virtual paths, and database access descriptors (DAD).

```
TYPE VARCHAR2_TABLE IS TABLE OF VARCHAR2(4000) INDEX BY BINARY_INTEGER;
```

78.5 DBMS_EPG Subprogram Groups

The DBMS_EPG consists of two interfaces: configuration subprograms and authorization subprograms.

- Configuration Subprograms
- Authorization Subprograms

78.5.1 DBMS_EPG Configuration Subprograms

The Configuration subprogram group contain the subprogram interfaces to examine and modify the global and database access descriptor (DAD) specific settings of the embedded PL/SQL gateway.

Table 78-2 Configuration Subprogram Group

Subprogram	Description
CREATE_DAD Procedure	Creates a new DAD
DELETE_DAD_ATTRIBUTE Procedure	Deletes a DAD attribute
DELETE_GLOBAL_ATTRIBUTE Procedure	Deletes a global attribute
DROP_DAD Procedure	Drops a DAD
GET_ALL_DAD_ATTRIBUTES Procedure	Retrieves all the attributes of a DAD.
GET_ALL_DAD_MAPPINGS Procedure	Retrieves all virtual paths to which the specified DAD is mapped.
GET_ALL_GLOBAL_ATTRIBUTES Procedure	Retrieves all global attributes and values
GET_DAD_ATTRIBUTE Function	Retrieves the value of a DAD attribute
GET_DAD_LIST Procedure	Retrieves a list of all DADs for an Embedded Gateway instance.
GET_GLOBAL_ATTRIBUTE Function	Retrieves the value of a global attribute
MAP_DAD Procedure	Maps a DAD to the specified virtual path.
SET_DAD_ATTRIBUTE Procedure	Sets the value for a DAD
SET_GLOBAL_ATTRIBUTE Procedure	Sets the value of a global attribute
UNMAP_DAD Procedure	Unmaps a DAD from the specified virtual path

78.5.2 DBMS_EPG Authorization Subprograms

The Authorization subprogram group contains the subprogram interfaces to authorize and deauthorize the use of a database user's privileges by the embedded PL/SQL gateway through a specific database access descriptor (DAD)

Table 78-3 Authorization Subprogram Group

Subprogram	Description
AUTHORIZE_DAD Procedure	Authorizes a DAD to invoke procedures and access document tables with a database user's privileges
DEAUTHORIZE_DAD Procedure	Deauthorizes a DAD with regard to invoking procedures and accessing document tables with a database user's privileges

78.6 Summary of DBMS_EPG Subprograms

This table lists the `DBMS_ALERT` subprograms and briefly describes them.

Table 78-4 DBMS_EPG Package Subprograms

Subprogram	Description
AUTHORIZE_DAD Procedure	authorizes a DAD to invoke procedures and access document tables with a database user's privileges
CREATE_DAD Procedure	Creates a new DAD
DEAUTHORIZE_DAD Procedure	Deauthorizes a DAD with regard to invoking procedures and accessing document tables with a database user's privileges
DELETE_DAD_ATTRIBUTE Procedure	Deletes a DAD attribute
DELETE_GLOBAL_ATTRIBUTE Procedure	Deletes a global attribute
DROP_DAD Procedure	Drops a DAD
GET_ALL_DAD_ATTRIBUTES Procedure	Retrieves all the attributes of a DAD.
GET_ALL_DAD_MAPPINGS Procedure	Retrieves all virtual paths to which the specified DAD is mapped.
GET_ALL_GLOBAL_ATTRIBUTES Procedure	Retrieves all global attributes and values
GET_DAD_ATTRIBUTE Function	Retrieves the value of a DAD attribute
GET_DAD_LIST Procedure	Retrieves a list of all DADs for an Embedded Gateway instance.
GET_GLOBAL_ATTRIBUTE Function	Retrieves the value of a global attribute
MAP_DAD Procedure	Maps a DAD to the specified virtual path.
SET_DAD_ATTRIBUTE Procedure	Sets the value for a DAD
SET_GLOBAL_ATTRIBUTE Procedure	Sets the value of a global attribute
UNMAP_DAD Procedure	Unmaps a DAD from the specified virtual path

78.6.1 AUTHORIZE_DAD Procedure

This procedure authorizes a DAD to invoke procedures and access document tables with a database user's privileges. The invoker can always authorize the use of her/his own privileges.

> **See Also:**
>
> Authorization Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.AUTHORIZE_DAD (
   dad_name  IN  VARCHAR2,
   path      IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 78-5 AUTHORIZE_DAD Procedure Parameters

Parameter	Description
dad_name	The name of the DAD to create
user	The user whose privileges to deauthorize. If use, the invoker is assumed.

Usage Notes

- To authorize the use of another user's privileges, the invoker must have the `ALTER USER` system privilege.
- The DAD must exist but its "database-username" DAD attribute does not have to be set to user to authorize.
- Multiple users can authorize the same DAD and it is up to the DAD's "database-username" setting to decide which user's privileges to use.

Exceptions

Raises an error if the DAD or user does not exist, or the invoker does not have the needed system privilege.

Examples

```
DBMS_EPG.AUTHORIZE_DAD('HR');
```

78.6.2 CREATE_DAD Procedure

This procedure creates a new DAD.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.CREATE_DAD (
   dad_name  IN  VARCHAR2,
   path      IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 78-6 CREATE_DAD Procedure Parameters

Parameter	Description
dad_name	The name of the DAD to create
path	The virtual path to which to map the DAD

78.6.3 DEAUTHORIZE_DAD Procedure

This procedure deauthorizes a DAD with regard to invoking procedures and accessing document tables with a database user's privileges. The invoker can always deauthorize the use of his own privileges.

> **See Also:**
>
> Authorization Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.DEAUTHORIZE_DAD (
   dad_name  IN  VARCHAR2,
   path      IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 78-7 DEAUTHORIZE_DAD Procedure Parameters

Parameter	Description
dad_name	The name of the DAD for which to deauthorize use
user	The user whose privileges to deauthorize. If use, the invoker is assumed.

Usage Notes

To deauthorize the use of another user's privileges, the invoker must have the `ALTER USER` system privilege.

Exceptions

Raises an error if the DAD or user does not exist, or the invoker does not have the needed system privilege.

Examples

```
DBMS_EPG.DEAUTHORIZE_DAD('HR');
```

78.6.4 DELETE_DAD_ATTRIBUTE Procedure

This procedure deletes a DAD attribute.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.DELETE_DAD_ATTRIBUTE (
   dad_name      IN  VARCHAR2,
   attr_name     IN  VARCHAR2);
```

Parameters

Table 78-8 DELETE_DAD_ATTRIBUTE Procedure Parameters

Parameter	Description
`dad_name`	The name of the DAD for which to delete a DAD attribute
`attr_name`	The name of the DAD attribute to delete

Exceptions

Raises an error if DAD does not exist

78.6.5 DELETE_GLOBAL_ATTRIBUTE Procedure

This procedure deletes a global attribute.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.DELETE_GLOBAL_ATTRIBUTE (
   attr_name      IN  VARCHAR2);
```

Parameters

Table 78-9 DELETE_GLOBAL_ATTRIBUTE Procedure Parameters

Parameter	Description
attr_name	The global attribute to delete

78.6.6 DROP_DAD Procedure

This procedure drops a DAD. All the virtual-path mappings of the DAD will be dropped also.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.DROP_DAD (
   dadname  IN  VARCHAR2);
```

Parameters

Table 78-10 DROP_DAD Procedure Parameters

Parameter	Description
dad_name	The DAD to drop

Exceptions

Raises an error if the DAD does not exist.

78.6.7 GET_ALL_DAD_ATTRIBUTES Procedure

This procedure retrieves all the attributes of a DAD. The outputs are 2 correlated index-by tables of the name/value pairs.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.GET_ALL_DAD_ATTRIBUTES (
   dad_name      IN          VARCHAR2,
   attr_names    OUT NOCOPY  VARCHAR2_TABLE,
   attr_values   OUT NOCOPY  VARCHAR2_TABLE);
```

Parameters

Table 78-11 GET_ALL_DAD_ATTRIBUTES Procedure Parameters

Parameter	Description
dad_names	The name of the DAD
attr_names	The attribute names
attr_values	The attribute values

Exceptions

Raises an error if DAD does not exist.

Usage Notes

If the DAD has no attributes set, then `attr_names` and `attr_values` will be set to empty arrays.

78.6.8 GET_ALL_DAD_MAPPINGS Procedure

This procedure retrieves all the virtual paths to which the specified DAD is mapped.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.GET_ALL_DAD_MAPPINGS (
   dad_name      IN          VARCHAR2,
   paths         OUT NOCOPY  VARCHAR2_TABLE);
```

Parameters

Table 78-12 GET_ALL_DAD_MAPPINGS Procedure Parameters

Parameter	Description
dad_names	The name of the DAD
paths	The virtual paths to which h the DAD is mapped

Exceptions

Raises an error if DAD does not exist.

Usage Notes

If the DAD is not mapped to any virtual path, `paths` will be set to empty arrays.

78.6.9 GET_ALL_GLOBAL_ATTRIBUTES Procedure

This procedure retrieves all global attributes and values. The outputs are 2 correlated index-by tables of the name/value pairs.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.GET_ALL_GLOBAL_ATTRIBUTES (
   attr_names     OUT   NOCOPY  VARCHAR2_TABLE,
   attr_values    OUT   NOCOPY  VARCHAR2_TABLE);
```

Parameters

Table 78-13 GET_ALL_GLOBAL_ATTRIBUTES Procedure Parameters

Parameter	Description
`attr_names`	The global attribute names
`attr_values`	The values of the global attributes

Usage Notes

If the gateway instance has no global attributes set, then `attr_names` and `attr_values` will be set to empty arrays.

78.6.10 GET_DAD_ATTRIBUTE Function

This procedure retrieves the value of a DAD attribute.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.GET_DAD_ATTRIBUTE (
   dad_name     IN  VARCHAR2,
   attr_name    IN  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 78-14 GET_DAD_ATTRIBUTE Function Parameters

Parameter	Description
dad_name	The name of the DAD for which to delete an attribute
attr_name	The name of the attribute to delete

Return values

Returns the DAD attribute value. Returns NULL if attribute is unknown or has not been set.

Exceptions

Raises an error if DAD does not exist.

78.6.11 GET_DAD_LIST Procedure

This procedure retrieves a list of all DADs for an Embedded Gateway instance.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.GET_DAD_LIST (
   dad_names     OUT NOCOPY  VARCHAR2_TABLE);
```

Parameters

Table 78-15 GET_DAD_LIST Procedure Parameters

Parameter	Description
dad_names	The list of all DADs

Usage Notes

If no DADs exist then dad_names will be set to an empty array.

78.6.12 GET_GLOBAL_ATTRIBUTE Function

This function retrieves the value of a global attribute.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.GET_GLOBAL_ATTRIBUTE (
   attr_name  IN  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 78-16 GET_GLOBAL_ATTRIBUTE Procedure Parameters

Parameter	Description
attr_name	The global attribute to retrieve

Return Values

Returns the global attribute value. Returns NULL if attribute has not been set or is not a valid attribute.

78.6.13 MAP_DAD Procedure

This procedure maps a DAD to the specified virtual path. If the virtual path exists already, the old virtual-path mapping will be overridden.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.MAP_DAD (
   dad_name  IN  VARCHAR2,
   path      IN  VARCHAR2);
```

Parameters

Table 78-17 MAP_DAD Procedure Parameters

Parameter	Description
dad_name	The name of the DAD to map
path	The virtual path to map

Exceptions

Raises and error if the DAD does not exist.

78.6.14 SET_DAD_ATTRIBUTE Procedure

This procedure sets the value for a DAD.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.SET_DAD_ATTRIBUTE (
   dad_name    IN  VARCHAR2,
   attr_name   IN  VARCHAR2,   attr_value  IN  VARCHAR2);
```

Parameters

Table 78-18 SET_DAD_ATTRIBUTE Procedure Parameters

Parameter	Description
dad_name	The name of the DAD for which to set the attribute
attr_name	The name of the attribute to set
attr_value	The attribute value to set

Table 78-19 Mapping Between mod_plsql and Embedded PL/SQL Gateway DAD Attributes

mod_plsql DAD Attribute	Embedded PL/SQL Gateway DAD Attribute	Allows Multiple Occurr-ences	Legal Values
PlsqlAfterProcedure	after-procedure	No	String
PlsqlAlwaysDescribeProcedure	always-describe-procedure	No	Enumeration of On, Off
PlsqlAuthenticationMode	authentication-mode	No	Enumeration of Basic, SingleSignOn, GlobalOwa, CustomOwa, PerPackageOwa
PlsqlBeforeProcedure	before-procedure	No	String
PlsqlBindBucketLengths	bind-bucket-lengths	Yes	Unsigned integer
PlsqlBindBucketWidths	bind-bucket-widths	Yes	Unsigned integer
PlsqlCGIEnvironmentList	cgi-environment-list	Yes	String
PlsqlCompatibilityMode	compatibility-mode	No	Unsigned integer

Table 78-19 (Cont.) Mapping Between mod_plsql and Embedded PL/SQL Gateway DAD Attributes

mod_plsql DAD Attribute	Embedded PL/SQL Gateway DAD Attribute	Allows Multiple Occurr-ences	Legal Values
PlsqlDatabaseUsername	database-username	No	String
PlsqlDefaultPage	default-page	No	String
PlsqlDocumentPath	document-path	No	String
PlsqlDocumentProcedure	document-procedure	No	String
PlsqlDocumentTablename	document-table-name	No	String
PlsqlErrorStyle	error-style	No	Enumeration of ApacheStyle, ModplsqlStyle, DebugStyle
PlsqlExclusionList	exclusion-list	Yes	String
PlsqlFetchBufferSize	fetch-buffer-size	No	Unsigned integer
PlsqlInfoLogging	info-logging	No	Enumeration of InfoDebug
PlsqlOWADebugEnable	owa-debug-enable	No	Enumeration of On, Off
PlsqlMaxRequestsPerSession	max-requests-per-session	No	Unsigned integer
PlsqlNLSLanguage	nls-language	No	String
PlsqlPathAlias	path-alias	No	String
PlsqlPathAliasProcedure	path-alias-procedure	No	String
PlsqlRequestValidationFunction	request-validation-function	No	String
PlsqlSessionCookieName	session-cookie-name	No	String
PlsqlSessionStateManagement	session-state-management	No	Enumeration of StatelessWithResetPackageState, StatelessWithFastRestPackageState, StatelessWithPreservePackageState
PlsqlTransferMode	transfer-mode	No	Enumeration of Char, Raw
PlsqlUploadAsLongRaw	upload-as-long-raw	No	String

Exceptions

Raises an error if DAD does not exist or the attribute is unknown.

Usage Notes

- If `attr_name` attribute has been set before, then the old value will be overwritten with the new `attr_value` argument.

- The embedded gateway assumes default values when the attributes are not set. The default values of the DAD attributes should be sufficient for most users of the embedded gateway. mod_plsql users should note the following
  - The `PlsqlDatabasePassword` attribute is not needed.
  - The `PlsqlDatabaseConnectString` attribute is not needed because the embedded gateway does not support logon to external databases.

Examples

```
DBMS_EPG.SET_DAD_ATTRIBUTE('HR', 'default-page', 'HRApp.home');
```

78.6.15 SET_GLOBAL_ATTRIBUTE Procedure

This procedure sets the value of a global attribute.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.SET_GLOBAL_ATTRIBUTE (
   attr_name    IN VARCHAR2,
   attr_value   IN VARCHAR2);
```

Parameters

Table 78-20 SET_GLOBAL_ATTRIBUTE Procedure Parameters

Parameter	Description
`attr_name`	The global attribute to set
`attr_value`	The attribute value to set

Table 78-21 Mapping Between mod_plsql and Embedded PL/SQL Gateway Global Attributes

mod_plsql Global Attribute	Embedded PL/SQL Gateway Global Attribute	Allows Multiple Occurr-ences	Legal Values
`PlsqlLogLevel`	`log-level`	No	Unsigned integer
`PlsqlMaxParameters`	`max-parameters`	No	Unsigned integer

Usage Notes

- The attribute name is case sensitive. The value may or may not be case-sensitive depending on the attribute.
- If `attr_name` attribute has been set before, then the old value will be overwritten with the new `attr_value` argument.

Exceptions

Raises an error if the attribute is unknown.

Examples

```
dbms_epg.set_global_attribute('max-parameters', '100');
```

78.6.16 UNMAP_DAD Procedure

This procedure unmaps a DAD from the specified virtual path. If path is `NULL`, the procedure removes all virtual-path mappings for the DAD but keeps the DAD.

> **See Also:**
>
> Configuration Subprograms for other subprograms in this group

Syntax

```
DBMS_EPG.UNMAP_DAD (
   dad_name IN VARCHAR2,
   path     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 78-22 UNMAP_DAD Procedure Parameters

Parameter	Description
`dad_name`	The name of the DAD to unmap
`path`	The virtual path to unmap

Usage Notes

Raises and error if the DAD does not exist.

79
DBMS_ERRLOG

The DBMS_ERRLOG package provides a procedure that enables you to create an error logging table so that DML operations can continue after encountering errors rather than abort and roll back. This enables you to save time and system resources.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_ERRLOG Subprograms

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information regarding how to use DBMS_ERRLOG and *Oracle Database SQL Language Reference* for error_logging_clause syntax

79.1 DBMS_ERRLOG Security Model

Security on this package can be controlled by granting EXECUTE on this package to selected users or roles. The EXECUTE privilege is granted publicly. However, to create an error logging table, you need SELECT access on the base table or view, the CREATE TABLE privilege, as well as tablespace quota for the target tablespace.

79.2 Summary of DBMS_ERRLOG Subprograms

The DBMS_ERRORLOG package includes the CREATE_ERROR_LOG procedure subprogram.

Table 79-1 DBMS_ERRLOG Package Subprograms

Subprogram	Description
CREATE_ERROR_LOG Procedure	Creates the error logging table used in DML error logging

79.2.1 CREATE_ERROR_LOG Procedure

This procedure creates the error logging table needed to use the DML error logging capability.

LONG, CLOB, BLOB, BFILE, and ADT datatypes are not supported in the columns.

Syntax

```
DBMS_ERRLOG.CREATE_ERROR_LOG (
   dml_table_name            IN VARCHAR2,
   err_log_table_name        IN VARCHAR2 := NULL,
```

```
   err_log_table_owner       IN VARCHAR2 := NULL,
   err_log_table_space       IN VARCHAR2 := NULL,
   skip_unsupported          IN BOOLEAN := FALSE);
```

Parameters

Table 79-2 CREATE_ERROR_LOG Procedure Parameters

Parameter	Description
dml_table_name	The name of the DML table to base the error logging table on. The name can be fully qualified (for example, emp, scott.emp, "EMP", "SCOTT"."EMP"). If a name component is enclosed in double quotes, it will not be upper cased.
err_log_table_name	The name of the error logging table you will create. The default is the first 25 characters in the name of the DML table prefixed with 'ERR$_'. Examples are the following: dml_table_name: 'EMP', err_log_table_name: 'ERR$_EMP' dml_table_name: '"Emp2"', err_log_table_name: 'ERR$_Emp2'
err_log_table_owner	The name of the owner of the error logging table. You can specify the owner in dml_table_name. Otherwise, the schema of the current connected user is used.
err_log_table_space	The tablespace the error logging table will be created in. If not specified, the default tablespace for the user owning the DML error logging table will be used.
skip_unsupported	When set to TRUE, column types that are not supported by error logging will be skipped over and not added to the error logging table. When set to FALSE, an unsupported column type will cause the procedure to terminate. The default is FALSE.

Examples

First, create an error log table for the channels table in the SH schema, using the default name generation.

Then, see all columns of the table channels:

```
SQL> DESC channels
Name                                  Null?       Type
----------------------------          -------     -----
CHANNEL_ID                            NOT NULL    CHAR(1)
CHANNEL_DESC                          NOT NULL    VARCHAR2(20)
CHANNEL_CLASS                                     VARCHAR2(20)
```

Finally, see all columns of the generated error log table. Note the mandatory control columns that are created by the package:

```
SQL> DESC ERR$_CHANNELS
Name                              Null?    Type
-----------------                 ----     ----
ORA_ERR_NUMBER$                            NUMBER
ORA_ERR_MESG$                              VARCHAR2(2000)
ORA_ERR_ROWID$                             ROWID
ORA_ERR_OPTYP$                             VARCHAR2(2)
ORA_ERR_TAG$                               VARCHAR2(2000)
CHANNEL_ID                                 VARCHAR2(4000)
```

```
CHANNEL_DESC                             VARCHAR2(4000)
CHANNEL_CLASS                            VARCHAR2(4000)
```

See *Oracle Database Administrator's Guide* for more information regarding control columns.

80
DBMS_FGA

The `DBMS_FGA` package provides fine-grained security functions.

This chapter contains the following topics:

- Security Model
- Operational Notes
- Summary of DBMS_FGA Subprograms

80.1 DBMS_FGA Security Model

You must have the `AUDIT_ADMIN` role or the `EXECUTE` privilege on the `DBMS_FGA` package to create audit policies. `DBMS_FGA` is an invoker rights package.

> **Note:**
>
> Starting in Oracle Database 23ai, fine-grained audit policies that are created with the `DBMS_FGA` package will generate audit records in the unified audit trail, viewable with the `UNIFIED_AUDIT_TRAIL` data dictionary view.

To analyze and audit data, you must have the `AUDIT_VIEWER` role. Because the audit function can potentially capture all user environment and application context values, policy administration should be executable by privileged users only. The policy event handler module is executed with the module owner's privilege.

80.2 DBMS_FGA Operational Notes

This package is available for only cost-based optimization. The rule-based optimizer may generate unnecessary audit records since audit monitoring can occur before row filtering.

For both the rule-based optimizer and the cost-based optimizer, you can query the `SQL_TEXT` and `SQL_BINDS` columns of the `UNIFIED_AUDIT_TRAIL` view to analyze the SQL text and corresponding bind variables that are issued.

80.3 Summary of DBMS_FGA Subprograms

This table describes the `DBMS_FGA` subprograms.

Table 80-1 ***DBMS_FGA Package Subprograms***

Subprogram	Description
ADD_POLICY Procedure	Creates an audit policy using the supplied predicate as the audit condition

Table 80-1 (Cont.) *DBMS_FGA Package Subprograms*

Subprogram	Description
DISABLE_POLICY Procedure	Disables an audit policy
DROP_POLICY Procedure	Drops an audit policy
ENABLE_POLICY Procedure	Enables an audit policy

80.3.1 ADD_POLICY Procedure

This procedure creates an audit policy using the supplied predicate as the audit condition.

Syntax

```
DBMS_FGA.ADD_POLICY(
   object_schema      IN  VARCHAR2 DEFAULT NULL,
   object_name        IN  VARCHAR2,
   policy_name        IN  VARCHAR2,
   audit_condition    IN  VARCHAR2 DEFAULT NULL,
   audit_column       IN  VARCHAR2 DEFAULT NULL,
   handler_schema     IN  VARCHAR2 DEFAULT NULL,
   handler_module     IN  VARCHAR2 DEFAULT NULL,
   enable             IN  BOOLEAN DEFAULT TRUE,
   statement_types    IN  VARCHAR2 DEFAULT SELECT,
   audit_trail        IN  BINARY_INTEGER DEFAULT NULL,
   audit_column_opts  IN  BINARY_INTEGER DEFAULT ANY_COLUMNS,
   policy_owner       IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 80-2 ADD_POLICY Procedure Parameters

Parameter	Description
object_schema	Schema of the object to be audited. If NULL, the current schema is assumed.
object_name	Name of the object to be audited
policy_name	Unique name of the policy. Do not enter special characters such as spaces or commas. If you want to use special characters for the policy name, then enclose the name in quotation marks.
audit_condition	A condition in a row that indicates a monitoring condition. NULL is allowed and acts as TRUE.
audit_column	Columns to be checked for access. These can include OLS hidden columns or object type columns. The default, NULL, causes audit if any column is accessed or affected.
handler_schema	Schema that contains the event handler. The default, NULL, causes the current schema to be used.
handler_module	Function name of the event handler; includes the package name if necessary. This function is invoked only after the first row that matches the audit condition in the query is processed. If the procedure fails with an exception, the user SQL statement will fail as well.
enable	Enables the policy if TRUE, which is the default

Table 80-2 (Cont.) ADD_POLICY Procedure Parameters

Parameter	Description
statement_types	SQL statement types to which this policy is applicable: INSERT, UPDATE, DELETE, or SELECT only
audit_trail	Do not set this parameter; it is desupported. All audit records are written to the unified audit trail, viewable by querying the UNIFIED_AUDIT_TRAIL data dictionary view.
audit_column_opts	Establishes whether a statement is audited when the query references *any* column specified in the audit_column parameter or only when *all* such columns are referenced
policy_owner	User who owns the fine-grained auditing policy. However, this setting is not a user-supplied argument. The Oracle Data Pump client uses this setting internally to recreate the fine-grained audit policies appropriately.

Usage Notes

- A table or view can have a maximum of 256 fine-grained audit policies applied to it.
- If object_schema is not specified, the current schema is assumed.
- An FGA policy should not be applied to out-of-line columns such as LOB columns.
- Each audit policy is applied to the query individually. However, at most one audit record may be generated for each policy, no matter how many rows being returned satisfy that policy's audit_condition. In other words, whenever any number of rows being returned satisfy an audit condition defined on the table, a single audit record will be generated for each such policy.
- If a table with an FGA policy defined on it receives a Fast Path insert or a vectored update, the hint is automatically disabled before any such operations. Disabling the hint allows auditing to occur according to the policy's terms. (One example of a Fast Path insert is the statement INSERT-WITH-APPEND-hint.)
- The audit_condition must be a boolean expression that can be evaluated using the values in the row being inserted, updated, or deleted. The expression can also use functions, such as the USER or SYS_CONTEXT functions.

  The expression must not combine conditions using operators such as AND and OR. audit_condition can be NULL (or omitted), which is interpreted as TRUE, but it cannot contain the following elements:

  – Subqueries or sequences
  – The following attributes of the USERENV namespace when accessed using the SYS_CONTEXT function:
    * CURRENT_SQL
    * CURRENT_SQL_LENGTH
    * CURRENT_BIND
  – Any use of the pseudo columns LEVEL, PRIOR, or ROWNUM.

  Specifying an audit condition of 1=1 to force auditing of all specified statements ("statement_types") affecting the specified column ("audit_column") is no longer needed

to achieve this purpose. A NULL value for audit_condition causes audit to happen even if no rows are processed, so that all actions on a table with this policy are audited.

- The audit_condition is evaluated using the privileges of the user who creates the policy.
- For the audit_condition setting, do not include functions, which execute the auditable statement on the same base table, in the audit_condition setting. For example, suppose you create a function that executes an INSERT statement on the HR.EMPLOYEES table. The policy audit_condition contains this function and it is for INSERT statements (as set by the statement_types parameter). When the policy is used, the function executes recursively until the system has run out of memory. This can raise the error ORA-1000: maximum open cursors exceeded or ORA-00036: maximum number of recursive SQL levels (50) exceeded.
- Do not issue the DBMS_FGA.ENABLE_POLICY or DBMS_FGA.DISABLE_POLICY statement from a policy function in a condition.
- The audit function (handler_module) is an alerting mechanism for the administrator. The required interface for such a function is as follows:

  ```
  PROCEDURE fname ( object_schema VARCHAR2, object_name VARCHAR2, policy_name
  VARCHAR2 )  AS ...
  ```

  where *fname* is the name of the procedure, *object_schema* is the name of the schema of the table audited, *object_name* is the name of the table to be audited, and *policy_name* is the name of the policy being enforced. The audit function will be executed with the function owner's privilege.
- Because traditional auditing is desupported, omit the audit_trail parameter because the audit records are written to the unified audit trail, viewable by querying the UNIFIED_AUDIT_TRAIL data dictionary view.
- Be aware that sensitive data, such as credit card information, can be recorded in clear text.
- You can change the operating system destination by using the following statement:

  ```
  ALTER SYSTEM SET AUDIT_FILE_DEST =  new_directory DEFERRED
  ```

  Starting with Oracle Database 23ai, the AUDIT_FILE_DEST parameter is deprecated.
- The audit_column_opts parameter establishes whether a statement is audited
  - when the query references *any* column specified in the audit_column parameter (audit_column_opts = DBMS_FGA.ANY_COLUMNS), or
  - only when *all* such columns are referenced (audit_column_opts = DBMS_FGA.ALL_COLUMNS).

  The default is DBMS_FGA.ANY_COLUMNS.

  The ALL_AUDIT_POLICIES view also shows audit_column_opts.
- When audit_column_opts is set to DBMS_FGA.ALL_COLUMNS, a SQL statement is audited only when all the columns mentioned in audit_column have been explicitly referenced in the statement. And these columns must be referenced in the same SQL-statement or in the sub-select.

  All these columns must refer to a single table/view or alias.

  If a SQL statement selects the columns from different table aliases, the statement will not be audited.

- For `SQL_TEXT` and `SQL_BIND` element values (`CLOB` type columns), the dynamic view shows only the first 4000 characters. The underlying XML file may have more than 4000 characters for such `SQL_TEXT` and `SQL_BIND` values.
- Error handling is the same as when `AUDIT_TRAIL=OS`. If any error occurs in writing an audit record, the audited operation fails and an alert message is logged.
- The policy event handler module will be executed with the module owner's privilege.
- Do not create recursive fine-grained audit handlers. For example, suppose you create a handler that executes an `INSERT` statement on the `HR.EMPLOYEES` table. The policy that is associated with this handler is for `INSERT` statements (as set by the statement_types parameter). When the policy is used, the handler executes recursively until the system has run out of memory. This can raise the error `ORA-1000: maximum open cursors exceeded` or `ORA-00036: maximum number of recursive SQL levels (50) exceeded`. See also *Oracle Database Security Guide* with regard to creating a fine-grained audit policy.
- The fine-grained audit handler module should not have explicit `COMMIT`, `ROLLBACK`, and `DDL` statements mentioned in it.

> **See Also:**
>
> *Oracle Database Security Guide* for an example of creating an email alert handler for a fine-grained audit policy

Examples

```
DBMS_FGA.ADD_POLICY (
   object_schema      =>  'scott',
   object_name        =>  'emp',
   policy_name        =>  'mypolicy1',
   audit_condition    =>  'sal < 100',
   audit_column       =>  'comm,sal',
   handler_schema     =>   NULL,
   handler_module     =>   NULL,
   enable             =>   TRUE,
   statement_types    =>  'INSERT, UPDATE',
   audit_column_opts  =>   DBMS_FGA.ANY_COLUMNS,
   policy_owner       =>  'sec_admin);
```

80.3.2 DISABLE_POLICY Procedure

This procedure disables an audit policy.

Syntax

```
DBMS_FGA.DISABLE_POLICY(
   object_schema  IN  VARCHAR2,
   object_name    IN  VARCHAR2,
   policy_name    IN  VARCHAR2);
```

Parameters

Table 80-3 DISABLE_POLICY Procedure Parameters

Parameter	Description
object_schema	Schema of the object to be audited. If NULL, the current schema is assumed.
object_name	Name of the object to be audited
policy_name	Unique name of the policy

The default value for object_schema is NULL. If NULL, the current schema is assumed.

Examples

```
DBMS_FGA.DISABLE_POLICY (
object_schema    =>  'scott',
object_name      =>  'emp',
policy_name      =>  'mypolicy1');
```

80.3.3 DROP_POLICY Procedure

This procedure drops an audit policy.

Syntax

```
DBMS_FGA.DROP_POLICY(
   object_schema  IN  VARCHAR2,
   object_name    IN  VARCHAR2,
   policy_name    IN  VARCHAR2);
```

Parameters

Table 80-4 DROP_POLICY Procedure Parameters

Parameter	Description
object_schema	Schema of the object to be audited. If NULL, the current schema is assumed.
object_name	Name of the object to be audited
policy_name	Unique name of the policy

Usage Notes

The DBMS_FGA procedures cause current DML transactions, if any, to commit before the operation unless they are inside a DDL event trigger. With DDL transactions, the DBMS_FGA procedures are part of the DDL transaction. The default value for object_schema is NULL. If NULL, the current schema is assumed.

> **Note:**
>
> Oracle Database automatically drops the audit policy if you remove the object specified in the `object_name` parameter of the `DBMS_FGA.ADD_POLICY` procedure, or if you drop the user who created the audit policy.

Examples

```
DBMS_FGA.DROP_POLICY (
object_schema   =>  'scott',
object_name     =>  'emp',
policy_name     =>  'mypolicy1');
```

80.3.4 ENABLE_POLICY Procedure

This procedure enables an audit policy.

Syntax

```
DBMS_FGA.ENABLE_POLICY(
   object_schema  IN  VARCHAR2,
   object_name    IN  VARCHAR2,
   policy_name    IN  VARCHAR2,
   enable         IN  BOOLEAN);
```

Parameters

Table 80-5 ENABLE_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema of the object to be audited. If `NULL`, the current schema is assumed.
`object_name`	Name of the object to be audited
`policy_name`	Unique name of the policy
`enable`	Defaults to `TRUE` to enable the policy

Examples

```
DBMS_FGA.ENABLE_POLICY (
object_schema    =>  'scott',
object_name      =>  'emp',
policy_name      =>  'mypolicy1',
enable           =>   TRUE);
```

81
DBMS_FILE_TRANSFER

The DBMS_FILE_TRANSFER package provides procedures to copy a binary file within a database or to transfer a binary file between databases.

This chapter contains the following topics:

- Overview
- Security Model
- Operating Notes
- Summary of DBMS_FILE_TRANSFER Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide* for instructions about using file transfer

81.1 DBMS_FILE_TRANSFER Overview

The DBMS_FILE_TRANSFER package provides procedures to copy a binary file within a database or to transfer a binary file between databases.

The destination database converts each block when it receives a file from a platform with different endianness. Datafiles can be imported after they are moved to the destination database as part of a transportable operation without RMAN conversion. Both GET and PUT operations will converted the file across platform difference at the destination. However, COPY is a local operation and therefore no conversion is required.

81.2 DBMS_FILE_TRANSFER Security Model

The DBMS_FILE_TRANSFER package must be created under SYS (CONNECT INTERNAL). Operations provided by this package are performed under the current calling user, not the package owner (SYS).

To use this interface the following users must have the following privileges:

- The current user at the local database must have READ privilege on the directory object specified in the source_directory_object parameter.
- The connected user at the destination database must have WRITE privilege to the directory object specified in the destination_directory_object parameter.

81.3 DBMS_FILE_TRANSFER Operational Notes

DBMS_FILE_TRANSFER supports online backup. You should therefore be careful in copying or transferring a file that is being modified by the database because this can result in an

inconsistent file, and require recovery. To guarantee consistency, bring files offline when the database is in use.

If you want to use `DBMS_FILE_TRANSFER` for performing backups, note that you are implementing self-managed backups, and should therefore put the files in hot backup mode.

81.4 Summary of DBMS_FILE_TRANSFER Subprograms

This table lists the `DBMS_FILE_TRANSFER` subprograms and briefly describes them.

Table 81-1 DBMS_FILE_TRANSFER Package Subprograms

Subprogram	Description
COPY_FILE Procedure	Reads a file from a source directory and creates a copy of it in a destination directory. The source and destination directories can both be in a local file system, or both be in an Automatic Storage Management (ASM) disk group, or between local file system and ASM with copying in either direction.
GET_FILE Procedure	Contacts a remote database to read a remote file and then creates a copy of the file in the local file system or ASM
PUT_FILE Procedure	Reads a local file or ASM and contacts a remote database to create a copy of the file in the remote file system

81.4.1 COPY_FILE Procedure

This procedure reads a file from a source directory and creates a copy of it in a destination directory. The source and destination directories can both be in a local file system, or both be in an Automatic Storage Management (ASM) disk group, or between local file system and ASM with copying in either direction.

You can copy any type of file to and from a local file system. However, you can copy only database files (such as datafiles, tempfiles, controlfiles, and so on) to and from an ASM disk group.

The destination file is not closed until the procedure completes successfully.

Syntax

```
DBMS_FILE_TRANSFER.COPY_FILE(
   source_directory_object       IN  VARCHAR2,
   source_file_name              IN  VARCHAR2,
   destination_directory_object  IN  VARCHAR2,
   destination_file_name         IN  VARCHAR2);
```

Parameters

Table 81-2 *COPY_FILE Procedure Parameters*

Parameter	Description
source_directory_object	The directory object that designates the source directory. The directory object must already exist. (You create directory objects with the `CREATE DIRECTORY` command).

Table 81-2 (Cont.) *COPY_FILE Procedure Parameters*

Parameter	Description
source_file_name	The name of the file to copy. This file must exist in the source directory.
destination_directory_object	The directory object that designates the destination directory. The directory object must already exist. If the destination is ASM, the directory object must designate either a disk group name (for example, +diskgroup1) or a directory created for alias names. In the case of a directory, the full path to the directory must be specified (for example: +diskgroup1/dbs/control).
destination_file_name	The name to assign to the file in the destination directory. A file with the same name must not exist in the destination directory. If the destination is ASM: • The file is given a fully qualified ASM filename and created in the appropriate directory (depending on the database name and file type) • The file type tag assigned to the file is COPY_FILE • The value of the destination_file_name argument becomes the file's alias name in the designated destination directory The file name can be followed by an ASM template name in parentheses. The file is then given the attributes specified by the template.

Usage Notes

To run this procedure successfully, the current user must have the following privileges:

- READ privilege on the directory object specified in the source_directory_object parameter
- WRITE privilege on directory object specified in the destination_directory_object parameter

This procedure converts directory object parameters to uppercase unless they are surrounded by double quotation marks, but this procedure does not convert file names to uppercase.

Also, the copied file must meet the following requirements:

- The size of the copied file must be a multiple of 512 bytes.
- The size of the copied file must be less than or equal to two terabytes.

The source_file_name parameter must specify a file that is in the directory specified by the source_directory_object parameter before running the procedure, and the destination_file_name parameter must specify the new name of the file in the new location specified in the destination_directory_object parameter. Relative paths and symbolic links are not allowed in the directory objects for the source_directory_object and destination_directory_object parameters.

Transferring the file is not transactional. To monitor the progress of a long file copy, query the V$SESSION_LONGOPS dynamic performance view.

> **See Also:**
>
> *Oracle Automatic Storage Management Administrator's Guide*for instructions about using file transfer

Examples

```
SQL> create directory DGROUP as '+diskgroup1/dbs/backup';

Directory created.

SQL>  BEGIN
   2    DBMS_FILE_TRANSFER.COPY_FILE('SOURCEDIR','t_xdbtmp.f', 'DGROUP',
                                     't_xdbtmp.f');
   3  END;
   4  /

PL/SQL procedure successfully completed.

SQL> EXIT
$ASMCMD
ASMCMD> ls
DISKGROUP1/
ASMCMD> cd diskgroup1/dbs/backup
ASMCMD> ls
t_xdbtmp.f => +DISKGROUP1/ORCL/TEMPFILE/COPY_FILE.267.546546525
```

81.4.2 GET_FILE Procedure

This procedure contacts a remote database to read a remote file and then creates a copy of the file in the local file system or ASM. The file that is copied is the source file, and the new file that results from the copy is the destination file. The destination file is not closed until the procedure completes successfully.

Syntax

```
DBMS_FILE_TRANSFER.GET_FILE
   source_directory_object      IN  VARCHAR2,
   source_file_name             IN  VARCHAR2,
   source_database              IN  VARCHAR2,
   destination_directory_object IN  VARCHAR2,
   destination_file_name        IN  VARCHAR2);
```

Parameters

Table 81-3 GET_FILE Procedure Parameters

Parameter	Description
source_directory_object	The directory object from which the file is copied at the source site. This directory object must exist at the source site.
source_file_name	The name of the file that is copied in the remote file system. This file must exist in the remote file system in the directory associated with the source directory object.

Table 81-3 (Cont.) GET_FILE Procedure Parameters

Parameter	Description
source_database	The name of a database link to the remote database where the file is located.
destination_directory_object	The directory object into which the file is placed at the destination site. This directory object must exist in the local file system.
destination_file_name	The name of the file copied to the local file system. A file with the same name must not exist in the destination directory in the local file system.

Usage Notes

To run this procedure successfully, the following users must have the following privileges:

- The connected user at the source database must have read privilege on the directory object specified in the source_directory_object parameter.
- The current user at the local database must have write privilege on the directory object specified in the destination_directory_object parameter.

This procedure converts directory object parameters to uppercase unless they are surrounded by double quotation marks, but this procedure does not convert file names to uppercase.

Also, the copied file must meet the following requirements:

- The size of the copied file must be a multiple of 512 bytes.
- The size of the copied file must be less than or equal to two terabytes.

Transferring the file is not transactional. To monitor the progress of a long file transfer, query the V$SESSION_LONGOPS dynamic performance view.

Examples

```
CREATE OR REPLACE DIRECTORY df AS '+datafile' ;
GRANT WRITE ON DIRECTORY df TO "user";
CREATE DIRECTORY DSK_FILES AS ''^t_work^'';
GRANT WRITE ON DIRECTORY dsk_files TO "user";

-- asumes that dbs2 link has been created and we are connected to the instance.
-- dbs2 could be a loopback or point to another instance.

BEGIN
-- asm file to an os file
-- get an asm file from dbs1.asm/a1 to dbs2.^t_work^/oa5.dat
  DBMS_FILE_TRANSFER.GET_FILE ( 'df' , 'a1' , 'dbs1', 'dsk_files' , 'oa5.dat' );

-- os file to an os file
-- get an os file from dbs1.^t_work^/a2.dat to dbs2.^t_work^/a2back.dat
  DBMS_FILE_TRANSFER.GET_FILE ( 'dsk_files' , 'a2.dat' , 'dbs1', 'dsk_files' ,
'a2back.dat' );

END ;
/
```

81.4.3 PUT_FILE Procedure

This procedure reads a local file or ASM and contacts a remote database to create a copy of the file in the remote file system.

The file that is copied is the source file, and the new file that results from the copy is the destination file. The destination file is not closed until the procedure completes successfully.

Syntax

```
DBMS_FILE_TRANSFER.PUT_FILE(
   source_directory_object       IN  VARCHAR2,
   source_file_name              IN  VARCHAR2,
   destination_directory_object  IN  VARCHAR2,
   destination_file_name         IN  VARCHAR2,
   destination_database          IN  VARCHAR2);
```

Parameters

Table 81-4 PUT_FILE Procedure Parameters

Parameter	Description
source_directory_object	The directory object from which the file is copied at the local source site. This directory object must exist at the source site.
source_file_name	The name of the file that is copied from the local file system. This file must exist in the local file system in the directory associated with the source directory object.
destination_directory_object	The directory object into which the file is placed at the destination site. This directory object must exist in the remote file system.
destination_file_name	The name of the file placed in the remote file system. A file with the same name must not exist in the destination directory in the remote file system.
destination_database	The name of a database link to the remote database to which the file is copied.

Usage Notes

To run this procedure successfully, the following users must have the following privileges:

- The current user at the local database must have read privilege on the directory object specified in the source_directory_object parameter.
- The connected user at the destination database must have write privilege to the directory object specified in the destination_directory_object parameter.

This procedure converts directory object parameters to uppercase unless they are surrounded by double quotation marks, but this procedure does not convert file names to uppercase.

Also, the copied file must meet the following requirements:

- The size of the copied file must be a multiple of 512 bytes.
- The size of the copied file must be less than or equal to two terabytes.

Transferring the file is not transactional. To monitor the progress of a long file transfer, query the `V$SESSION_LONGOPS` dynamic performance view.

Examples

```
CREATE OR REPLACE DIRECTORY df AS '+datafile' ;
GRANT WRITE ON DIRECTORY df TO "user";
CREATE OR REPLACE DIRECTORY ft1 AS '+datafile/ft1' ;
GRANT READ,WRITE ON DIRECTORY ft1 TO "user";
CREATE OR REPLACE DIRECTORY ft1_1 AS '+datafile/ft1/ft1_1' ;

CONNECT user;
Enter password: password

-- - put a1.dat to a4.dat (using dbs2 dblink)
-- - level 2 sub dir to parent dir
-- - user has read privs on ft1_1 at dbs1 and write on df in dbs2
BEGIN
 DBMS_FILE_TRANSFER.PUT_FILE ( 'ft1_1' , 'a2.dat' , 'df' , 'a4.dat' ,
                               'dbs2' ) ;
END ;
```

82
DBMS_FLASHBACK

Using `DBMS_FLASHBACK`, you can flash back to a version of the database at a specified time or a specified system change number (SCN).

This chapter contains the following topics:

- Overview
- Security Model
- Types
- Exceptions
- Operational Notes
- Examples
- Summary of DBMS_FLASHBACK Subprograms

> **See Also:**
>
> For detailed information about `DBMS_FLASHBACK`:
>
> - *Oracle Database Development Guide*
> - *Oracle Database SQL Language Reference*.

82.1 DBMS_FLASHBACK Overview

`DBMS_FLASHBACK` provides an interface for the user to view the database at a particular time in the past, with the additional capacity provided by transaction back out features that allow for selective removal of the effects of individual transactions. This is different from a flashback database which moves the database back in time.

When `DBMS_FLASHBACK` is enabled, the user session uses the Flashback version of the database, and applications can execute against the Flashback version of the database.

You may want to use `DBMS_FLASHBACK` for the following reasons:

- Self-service repair: If you accidentally delete rows from a table, you can recover the deleted rows.
- Packaged applications such as email and voicemail: You can use Flashback to restore deleted email by re-inserting the deleted message into the current message box.
- Decision support system (DSS) and online analytical processing (OLAP) applications: You can perform data analysis or data modeling to track seasonal demand.

82.2 DBMS_FLASHBACK Security Model

To use the `DBMS_FLASHBACK` package, you must have the `EXECUTE` privilege on it.

82.3 DBMS_FLASHBACK Types

The following table describes the types used by DBMS_FLASHBACK.

Table 82-1 DBMS_FLASHBACK

Type	Description
TXNAME_ARRAY	Creates a VARRAY for holding Transaction Names or Identifiers (XIDs)

82.4 DBMS_FLASHBACK Exceptions

DBMS_FLASHBACK creates the following error messages.

Table 82-2 DBMS_FLASHBACK Error Messages

Error	Description
ORA-08180	Time specified is too old
ORA-08181	Invalid system change number specified
ORA-08182	User cannot begin read-only or serializable transactions in Flashback mode
ORA-08183	User cannot enable Flashback within an uncommitted transaction
ORA-08184	User cannot enable Flashback within another Flashback session
ORA-08185	SYS cannot enable Flashback mode

82.5 DBMS_FLASHBACK Operational Notes

DBMS_FLASHBACK is automatically turned off when the session ends, either by disconnection or by starting another connection.

PL/SQL cursors opened in Flashback mode return rows as of the flashback time or SCN. Different concurrent sessions (connections) in the database can perform Flashback to different wall-clock times or SCNs. DML and DDL operations and distributed operations are not allowed while a session is running in Flashback mode. You can use PL/SQL cursors opened before disabling Flashback to perform DML.

Under Automatic Undo Management (AUM) mode, you can use retention control to control how far back in time to go for the version of the database you need. If you need to perform a Flashback over a 24-hour period, the DBA must set the undo_retention parameter to 24 hours. This way, the system retains enough undo information to regenerate the older versions of the data.

You can set the RETENTION GUARANTEE clause for the undo tablespace to ensure that unexpired undo is not discarded. UNDO_RETENTION is not in itself a guarantee because, if the system is under space pressure, unexpired undo may be overwritten with freshly generated undo. In such cases, RETENTION GUARANTEE prevents this. For more information, see the *Oracle Database Administrator's Guide*.

In a Flashback-enabled session, `SYSDATE` is not affected; it continues to provide the current time.

`DBMS_FLASHBACK` can be used within logon triggers to enable Flashback without changing the application code.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information on setting the minimum undo retention period.

82.6 DBMS_FLASHBACK Examples

The following example illustrates how Flashback can be used when the deletion of a senior employee triggers the deletion of all the personnel reporting to him. Using the Flashback feature, you can recover and re-insert the missing employees.

```
DROP TABLE employee;
DROP TABLE keep_scn;

REM -- Keep_scn is a temporary table to store scns that we are interested in

CREATE TABLE keep_scn (scn number);
SET ECHO ON
CREATE TABLE employee (
   employee_no   number(5) PRIMARY KEY,
   employee_name varchar2(20),
   employee_mgr  number(5)
      CONSTRAINT mgr_fkey REFERENCES EMPLOYEE ON DELETE CASCADE,
   salary        number,
   hiredate      date
);

REM -- Populate the company with employees
INSERT INTO employee VALUES (1, 'John Doe', null, 1000000, '5-jul-81');
INSERT INTO employee VALUES (10, 'Joe Johnson', 1, 500000, '12-aug-84');
INSERT INTO employee VALUES (20, 'Susie Tiger', 10, 250000, '13-dec-90');
INSERT INTO employee VALUES (100, 'Scott Tiger', 20, 200000, '3-feb-86');
INSERT INTO employee VALUES (200, 'Charles Smith', 100, 150000, '22-mar-88');
INSERT INTO employee VALUES (210, 'Jane Johnson', 100, 100000, '11-apr-87');
INSERT INTO employee VALUES (220, 'Nancy Doe', 100, 100000, '18-sep-93');
INSERT INTO employee VALUES (300, 'Gary Smith', 210, 75000, '4-nov-96');
INSERT INTO employee VALUES (310, 'Bob Smith', 210, 65000, '3-may-95');
COMMIT;

REM -- Show the entire org
SELECT lpad(' ', 2*(level-1)) || employee_name Name
FROM employee
CONNECT BY PRIOR employee_no = employee_mgr
START WITH employee_no = 1
ORDER BY LEVEL;

REM -- Sleep for a short time (approximately 10 to 20  seconds) to avoid
REM -- querying close to table creation

EXECUTE DBMS_LOCK.SLEEP(10);
```

```
REM -- Store this snapshot for later access through Flashback
DECLARE
I NUMBER;
BEGIN
I := DBMS_FLASHBACK.GET_SYSTEM_CHANGE_NUMBER;
INSERT INTO keep_scn VALUES (I);
COMMIT;
END;
/

REM -- Scott decides to retire but the transaction is done incorrectly
DELETE FROM EMPLOYEE WHERE employee_name = 'Scott Tiger';
COMMIT;

REM -- notice that all of scott's employees are gone
SELECT lpad(' ', 2*(level-1)) || employee_name Name
FROM EMPLOYEE
CONNECT BY PRIOR employee_no = employee_mgr
START WITH employee_no = 1
ORDER BY LEVEL;

REM -- Flashback to see Scott's organization
DECLARE
   restore_scn number;
BEGIN
   SELECT  scn INTO restore_scn FROM keep_scn;
   DBMS_FLASHBACK.ENABLE_AT_SYSTEM_CHANGE_NUMBER (restore_scn);
END;
/

REM -- Show Scott's org.
SELECT lpad(' ', 2*(level-1)) || employee_name Name
FROM employee
CONNECT BY PRIOR employee_no = employee_mgr
START WITH employee_no =
   (SELECT employee_no FROM employee WHERE employee_name = 'Scott Tiger')
ORDER BY LEVEL;

REM -- Restore scott's organization.
DECLARE
   scotts_emp NUMBER;
   scotts_mgr NUMBER;
   CURSOR c1 IS
      SELECT employee_no, employee_name, employee_mgr, salary, hiredate
      FROM employee
      CONNECT BY PRIOR employee_no = employee_mgr
      START WITH employee_no =
         (SELECT employee_no FROM employee WHERE employee_name = 'Scott Tiger');
   c1_rec c1 % ROWTYPE;
BEGIN
   SELECT employee_no, employee_mgr INTO scotts_emp, scotts_mgr FROM employee
   WHERE employee_name = 'Scott Tiger';
   /* Open c1 in flashback mode */
   OPEN c1;
   /* Disable Flashback */
   DBMS_FLASHBACK.DISABLE;
 LOOP
   FETCH c1 INTO c1_rec;
   EXIT WHEN c1%NOTFOUND;
   /*
     Note that all the DML operations inside the loop are performed
     with Flashback disabled
```

```
   */
   IF (c1_rec.employee_mgr = scotts_emp) then
      INSERT INTO employee VALUES (c1_rec.employee_no,
         c1_rec.employee_name,
         scotts_mgr,
         c1_rec.salary,
         c1_rec.hiredate);
   ELSE
   IF (c1_rec.employee_no != scotts_emp) THEN
   INSERT INTO employee VALUES (c1_rec.employee_no,
         c1_rec.employee_name,
         c1_rec.employee_mgr,
         c1_rec.salary,
         c1_rec.hiredate);
      END IF;
    END IF;
 END LOOP;
END;
/

REM -- Show the restored organization.
select lpad(' ', 2*(level-1)) || employee_name Name
FROM employee
CONNECT BY PRIOR employee_no = employee_mgr
START WITH employee_no = 1
ORDER BY LEVEL;
```

82.7 Summary of DBMS_FLASHBACK Subprograms

This table lists the `DBMS_FLASHBACK` subprograms and briefly describes them.

Table 82-3 DBMS_FLASHBACK Package Subprograms

Subprogram	Description
DISABLE Procedure	Disables the Flashback mode for the entire session
ENABLE_AT_SYSTEM_CHANGE_NUMBER Procedure	Enables Flashback for the entire session. Takes an SCN as an Oracle number and sets the session snapshot to the specified number. Inside the Flashback mode, all queries return data consistent as of the specified wall-clock time or SCN
ENABLE_AT_TIME Procedure	Enables Flashback for the entire session. The snapshot time is set to the SCN that most closely matches the time specified in `query_time`
GET_SYSTEM_CHANGE_NUMBER Function	Returns the current SCN as an Oracle number. You can use the SCN to store specific snapshots
TRANSACTION_BACKOUT Procedures	Provides the mechanism to back out a transaction

82.7.1 DISABLE Procedure

This procedure disables the Flashback mode for the entire session.

Syntax

```
DBMS_FLASHBACK.DISABLE;
```

Examples

The following example queries the salary of an employee, Joe, on August 30, 2000:

```
EXECUTE dbms_flashback.enable_at_time('30-AUG-2000');
SELECT salary FROM emp where name = 'Joe'
EXECUTE dbms_flashback.disable;
```

82.7.2 ENABLE_AT_SYSTEM_CHANGE_NUMBER Procedure

This procedure takes an SCN as an input parameter and sets the session snapshot to the specified number.

In the Flashback mode, all queries return data consistent as of the specified wall-clock time or SCN. It enables Flashback for the entire session.

Syntax

```
DBMS_FLASHBACK.ENABLE_AT_SYSTEM_CHANGE_NUMBER (
   query_scn IN NUMBER);
```

Parameters

Table 82-4 ENABLE_AT_SYSTEM_CHANGE_NUMBER Procedure Parameters

Parameter	Description
query_scn	The system change number (SCN), a version number for the database that is incremented on every transaction commit.

82.7.3 ENABLE_AT_TIME Procedure

This procedure enables Flashback for the entire session.

The snapshot time is set to the SCN that most closely matches the time specified in query_time. It enables Flashback for the entire session.

Syntax

```
DBMS_FLASHBACK.ENABLE_AT_TIME (
   query_time   IN TIMESTAMP);
```

Parameters

Table 82-5 ENABLE_AT_TIME Procedure Parameters

Parameter	Description

query_time	This is an input parameter of type TIMESTAMP. A time stamp can be specified in the following ways:
	• Using the TIMESTAMP constructor
	`EXECUTE DBMS_FLASHBACK.ENABLE_AT_TIME(TIMESTAMP '2001-01-09 12:31:00').`
	Use the Globalization Support (NLS) format and supply a string. The format depends on the Globalization Support settings.
	• Using the TO_TIMESTAMP function:
	`EXECUTE DBMS_FLASHBACK.ENABLE_AT_TIME(TO_TIMESTAMP('12-02-2001 14:35:00', 'DD-MM-YYYY HH24:MI:SS'))`
	You provide the format you want to use. This example shows the TO_TIMESTAMP function for February 12, 2001, 2:35 PM.
	• If the time is omitted from query time, it defaults to the beginning of the day, that is, 12:00 A.M.
	• Note that if the query time contains a time zone, the time zone information is truncated.

82.7.4 GET_SYSTEM_CHANGE_NUMBER Function

This function returns the current SCN as an Oracle number datatype. You can obtain the current change number and store it for later use. This helps you retain specific snapshots.

Syntax

```
DBMS_FLASHBACK.GET_SYSTEM_CHANGE_NUMBER
 RETURN NUMBER;
```

82.7.5 TRANSACTION_BACKOUT Procedures

This procedure provides a mechanism to back out a set of transactions. The user can call these procedures with either transaction names or transaction identifiers (XIDS).

The procedure analyzes the transactional dependencies, perform DMLs and generates an extensive report on the operation performed by the subprogram. This procedure does not commit the DMLs performed as part of transaction back out. However it holds all the required locks on rows and tables in the right form, so that no other dependencies can enter the system. To make the changes permanent you must explicitly commit the transaction.

A report is generated in the system tables DBA_FLASHBACK_TRANSACTION_STATE and DBA_FLASHBACK_TRANSACTION_REPORT.

Syntax

```
DBMS_FLASHBACK.TRANSACTION_BACKOUT
   numtxns            NUMBER,
   xids               XID_ARRAY,
   options            NUMBER default NOCASCADE,
```

```
   timeHint           TIMESTAMP default MINTIME);

DBMS_FLASHBACK.TRANSACTION_BACKOUT
   numtxns             NUMBER,
   xids                XID_ARRAY,
   options             NUMBER default NOCASCADE,
   scnHint             TIMESTAMP default 0   );

DBMS_FLASHBACK.TRANSACTION_BACKOUT
   numtxns           NUMBER,
   txnnames          TXNAME_ARRAY,
   options           NUMBER default NOCASCADE,
   timehint          TIMESTAMP MINTIME );

DBMS_FLASHBACK.TRANSACTION_BACKOUT
   numtxns           NUMBER,
   txnNames         TXNAME_ARRAY,
   options          NUMBER default NOCASCADE,
   scnHint          NUMBER 0);
```

Parameters

Table 82-6 TRANSACTION_BACKOUT Procedure Parameters

Parameter	Description
numtxns	Number of transactions passed as input
xids	List of transaction IDs in the form of an array
txnnames	List of transaction names in the form of an array
options	Back out dependent transactions: • NOCASCADE - No dependency is expected. If a dependency is found, this raises an error, with the first dependent transaction provided in the report. • NOCASCADE_FORCE - The user forcibly backs out the given transactions without considering the dependent transactions. The RDBMS executes the UNDO SQL for the given transactions in reverse order of their commit times. If no constraints break, and the result is satisfactory, the user can either COMMIT the changes or else ROLL BACK. • NONCONFLICT_ONLY - This option lets the user back out the changes to the nonconflicting rows of the given transactions. Note that a transaction dependency can happen due to a row conflict through either WAW or primary/unique key constraints. If the user chooses to back out only the nonconflicting rows, this does not cause any problem with database consistency, although transaction atomicity is lost. As this is a recovery operation, the user can correct the data. • CASCADE - This completely removes the given transactions including their dependents in a post order fashion (reverse order of commit times).
timehint	Time hint on the start of the transaction
scnhint	SCN hint on the start of the transaction

Usage Notes

> **Note:**
>
> For information about restrictions in using `TRANSACTION_BACKOUT`, see "Using Flashback Transaction" in the *Oracle Database Development Guide*.

- If transaction name is used, a time hint must be provided. The time hint should be a time before the start of all the given transactions to back out.
- If the SCN hint is provided, it must be before the start of the earliest transaction in the specified input set, or this raises an error and terminates. If it is not provided and the transaction has committed within undo retention, the database system is able to determine the start time.

83
DBMS_FLASHBACK_ARCHIVE

The `DBMS_FLASHBACK_ARCHIVE` package contains procedures for performing various flashback archive tasks.

These include:

- Disassociation and reassociation of a Flashback Archive enabled table from/with its underlying Flashback Time Travel
- Tamper-proofing the tables of an application
- Importing of user history

> **Caution:**
>
> Importing user-generated history can lead to inaccurate, or unreliable results. This procedure should only be used after consulting with Oracle Support.

- Enabling and disabling of session-level support for valid-time

> **See Also:**
>
> *Oracle Database Development Guide* for more information about Using Flashback Time Travel

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_FLASHBACK_ARCHIVE Subprograms

83.1 DBMS_FLASHBACK_ARCHIVE Overview

Flashback Time Travel provides strict protection on the internal history tables that it creates and maintains for users.

> **Caution:**
>
> The `DBMS_FLASHBACK_ARCHIVE` package does not support migration between databases. It allows you to move Flashback Time Travel history within the same database only, and not outside it. This package does not allow you to move Flashback Time Travel table data. It allows you to import customer history only and not Flashback Time Travel history.

The read-only semantics prohibit users, including a DBA, from doing updates, deletes, and inserts on the Flashback Time Travel internal history tables. Users are also prevented from issuing any DDL statements on these tables. This strict security enforcement helps meet the requirements of applications in regulatory / compliance environments. Flashback Time Travel supports most common DDL statements, including those that alter the table definition or incur data movement. However, some DDL statements are not supported on Flashback Archive enabled tables. Since most application schemas are modified during application software upgrades, the ability to perform DDL operations on tracked tables is critical.

To support schema evolution during application upgrades and other table maintenance tasks that require use of DDL statements not supported by Flashback Time Travel, the `DBMS_FLASHBACK_ARCHIVE` package provides a set of simple-to-use PL/SQL procedures:

- To disassociate a Flashback Archive enabled base table from the underlying Flashback Time Travel
- To reassociate a temporarily disassociated base table with its underlying Flashback Time Travel

After a user has disassociated the base table from its Flashback Archive, it's possible to issue any DDL statements on the base table or the history tables in the Flashback Archive. Having finished with the schema changes, the user can then reassociate the base table with its Flashback Archive so that Flashback Time Travel protection is in operation and automatic tracking and archiving is resumed.

83.2 DBMS_FLASHBACK_ARCHIVE Security Model

Users need the `FLASHBACK_ARCHIVE_ADMINISTER` privilege to import user-generated history, to set context level, and to tamper-proof tables. After a table is disassociated, users can perform DDL and DML statements on the table if they have the necessary privileges. Enabling and disabling session-level Valid Time Temporal flashback needs no additional privileges.

83.3 DBMS_FLASHBACK_ARCHIVE Constants

The `DBMS_FLASHBACK_ARCHIVE` package uses the constants shown in the following table.

Table 83-1 DBMS_FLASHBACK_ARCHIVE Constants

Constant	Type	Value	Description
NODROP	BINARY_INTEGER	1	Do not drop temporary history table
NOCOMMIT	BINARY_INTEGER	2	Do not commit transaction
NODELETE	BINARY_INTEGER	4	Do not delete data in history table

83.4 Summary of DBMS_FLASHBACK_ARCHIVE Subprograms

This table lists the DBMS_FLASHBACK_ARCHIVE subprograms and briefly describes them.

Table 83-2 DBMS_FLASHBACK_ARCHIVE Package Subprograms

Subprogram	Description
ADD_TABLE_TO_APPLICATION Procedure	Takes an application name and adds a table to the application as a security table
CREATE_TEMP_HISTORY_TABLE Procedure	Creates a table called TEMP_HISTORY with the correct definition in schema
DISABLE_APPLICATION Procedure	Takes an application name and marks a table in it as a security table
DISABLE_ASOF_VALID_TIME Procedure	Disables session level valid-time flashback
DISASSOCIATE_FBA Procedure	Disassociates the given table from the Flashback Time Travel
DROP_APPLICATION Procedure	Takes an application name and removes it from the list of applications
ENABLE_APPLICATION Procedure	Takes an application name and enables Flashback Time Travel on all the security tables for this application
ENABLE_AT_VALID_TIME Procedure	Enables session level valid time flashback
EXTEND_MAPPINGS Procedure	Extends time mappings to times in the past
GET_CURRENT_LIFESPAN_DIGEST Function	Generates the current lifespan digest for the specified row in a user table
GET_SYS_CONTEXT Function	Gets the context previously selected by the SET_CONTEXT_LEVEL Procedure
IMPORT_HISTORY Procedure	Imports history from a table called TEMP_HISTORY in the given schema.
LOCK_DOWN_APPLICATION Procedure	Takes an application name and makes all the security tables read-only. The group called SYSTEM cannot be locked
PURGE_CONTEXT Procedure	Purges the context to be saved selected by the SET_CONTEXT_LEVEL Procedure
REASSOCIATE_FBA Procedure	Reassociates the given table with the Flashback Time Travel
REGISTER_APPLICATION Procedure	Takes an application name and optionally a Time Travel, and registers an application for database hardening
REMOVE_TABLE_FROM_APPLICATION Procedure	Takes an application name and marks a table in it as no longer being a security table

Table 83-2 (Cont.) DBMS_FLASHBACK_ARCHIVE Package Subprograms

Subprogram	Description
SET_CONTEXT_LEVEL Procedure	Defines how much of the user context is to be saved
VERIFY_BLOCKCHAIN_LIFESPAN Procedure	Verifies the content of the current, historical, or all lifespans of user table rows

83.4.1 ADD_TABLE_TO_APPLICATION Procedure

This procedure takes an application name and adds a table to the application as a security table. If the application is enabled for Flashback Time Travel, then this table will also be enabled for Flashback Time Travel.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.ADD_TABLE_TO_APPLICATION (
   application_name           IN   VARCHAR2,
   table_name                 IN   VARCHAR2,
   schema_name                IN   VARCHAR2 := NULL);
```

Parameters

Table 83-3 ADD_TABLE_TO_APPLICATION Procedure Parameters

Parameter	Description
application_name	Name of the application for which a table has been added as a security table
table_name	Name of the table to add as a security table for the given application
schema_name	Name of the schema containing the desired table. If no schema name is specified, the current schema is used.

83.4.2 CREATE_TEMP_HISTORY_TABLE Procedure

This procedure creates a table called `TEMP_HISTORY` with the correct definition in schema.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.CREATE_TEMP_HISTORY_TABLE (
   owner_name1      IN   VARCHAR2,
   table_name1      IN   VARCHAR2);
```

Parameters

Table 83-4 CREATE_TEMP_HISTORY_TABLE Procedure Parameters

Parameter	Description
owner_name1	Schema of the Flashback Time Travel-enabled table
table_name1	Name of the Flashback Time Travel-enabled table

83.4.3 DISABLE_APPLICATION Procedure

This procedure takes an application name and disables Flashback Time Travel on all of its security tables.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.DISABLE_APPLICATION (
   application_name           IN   VARCHAR2);
```

Parameters

Table 83-5 DISABLE_APPLICATION Procedure Parameters

Parameter	Description
application_name	Name of the application whose security tables will be disabled for Flashback Time Travel

83.4.4 DISABLE_ASOF_VALID_TIME Procedure

This procedure disables session level valid-time flashback.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.DISABLE_ASOF_VALID_TIME;
```

83.4.5 DISASSOCIATE_FBA Procedure

This procedure disassociates the given table from the Flashback Time Travel.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.DISASSOCIATE_FBA (
   owner_name     IN    VARCHAR2,
   table_name     IN    VARCHAR2);
```

Parameters

Table 83-6 DISASSOCIATE_FBA Procedure Parameters

Parameter	Description
owner_name	Schema of the Flashback Time Travel enabled base table
table_name	Name of the Flashback Time Travel enabled base table

Exceptions

Table 83-7 DISASSOCIATE_FBA Procedure Exceptions

Exception	Description
ORA-55602	User table is not enabled for Flashback Time Travel
ORA-55634	Cannot acquire the lock on the table for disassociation

83.4.6 DROP_APPLICATION Procedure

This procedure takes an application name and removes it from the list of applications. As part of this procedure, Flashback Time Travel will be disabled on all security-enabled tables and all history data will be lost. The group called `SYSTEM` cannot be dropped.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.DROP_APPLICATION (
   application_name           IN   VARCHAR2);
```

Parameters

Table 83-8 DROP_APPLICATION Procedure Parameters

Parameter	Description
`application_name`	Name of the application for which a table has been added as a security table

83.4.7 ENABLE_APPLICATION Procedure

This procedure takes an application name and enables Flashback Time Travel on all the security tables for this application. Once an application is enabled, every change to an Flashback Time Travel enabled table will be tracked.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.ENABLE_APPLICATION (
   application_name           IN   VARCHAR2);
```

Parameters

Table 83-9 ENABLE_APPLICATION Procedure Parameters

Parameter	Description
`application_name`	Name of the application for which to enable Flashback Time Travel on all its security tables

83.4.8 ENABLE_AT_VALID_TIME Procedure

This procedure enables session level valid time flashback.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.ENABLE_AT_VALID_TIME (
   level          IN    VARCHAR2,
   query_time     IN    TIMESTAMP DEFAULT SYSTIMESTAMP);
```

Parameters

Table 83-10 ENABLE_AT_VALID_TIME Procedure Parameters

Parameter	Description
`level`	Options: • `All` - Sets the visibility of temporal data to the full table, which is the default temporal table visibility • `CURRENT` - Sets the visibility of temporal data to currently valid data within the valid time period at the session level • `ASOF` - Sets the visibility of temporal data to data valid as of the given time as defined by the timestamp
`query_time`	Used only if level is `ASOF`. Data which is valid at this `query_time` will only be shown.

83.4.9 EXTEND_MAPPINGS Procedure

This procedure extends time mappings to times in the past.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.EXTEND_MAPPINGS;
```

83.4.10 GET_CURRENT_LIFESPAN_DIGEST Function

This function generates the current lifespan digest for the specified row in a user table.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.GET_CURRENT_LIFESPAN_DIGEST (
            rid          IN   VARCHAR2,
            start_scn    IN   RAW,
            xid          IN   RAW,
            op           IN   VARCHAR2,
            objno        IN   NUMBER)
RETURN RAW;
```

Parameters

Table 83-11 GET_CURRENT_LIFESPAN_DIGEST Function Parameters

Parameter	Description
`rid`	The ROWID of the user table row.
`start_scn`	The latest start SCN of the row.
`xid`	The latest transaction ID of the row.
`op`	The latest operation that is performed on the row.
`objno`	The table's object number.

The cryptographic hash is generated over all column values in the ascending order of the column IDs. Each column value is prefixed by column metadata. For more information, refer to

the Appendix B.1 Column Content in Blockchain Tables section of the *Database Administrator's Guide*.

- Each row of the user table transitions through a series of "lifespans" created by transactions that commit changes to the row. The most recent lifespan is referred to as the "current lifespan," and all other previous lifespans are referred to as "historical lifespans."
- The historical lifespans are recorded into a Flashback Archive (FBA) internal table: `SYS_FBA_HIST_<objno>`.
- The current lifespans are always present in the user table. However, the relevant metadata of current lifespans are recorded in another FBA internal table: `SYS_FBA_TCRV_<objno>`.
- This API can be used to compute the current lifespan digest of the specific row of a user table using its associated metadata from the FBA internal table: `SYS_FBA_TCRV_<objno>`. A user can record this digest for detecting corruptions, if any, in the row data.

83.4.11 GET_SYS_CONTEXT Function

This function gets the context previously selected by the SET_CONTEXT_LEVEL Procedure.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.GET_SYS_CONTEXT (
   xid            IN   RAW,
   namespace      IN   VARCHAR2,
   parameter      IN   VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 83-12 GET_SYS_CONTEXT Function Parameters

Parameter	Description
`xid`	Transaction identifier is an opaque handle to a transaction obtained from the versions query
`namespace`	Namespace
`parameter`	If undefined, the subprogram returns `NULL`

Related Topics

- SET_CONTEXT_LEVEL Procedure
  This procedure defines how much of the user context is to be saved.

83.4.12 IMPORT_HISTORY Procedure

This procedure is called after invoking the CREATE_TEMP_HISTORY_TABLE procedure, and after the `TEMP_HISTORY` table is populated with user-generated history data

> **Caution:**
>
> Importing user-generated history can lead to inaccurate, or unreliable results. This procedure should only be used after consulting with Oracle Support.
>
> The `DBMS_FLASHBACK_ARCHIVE` package does not support migration between databases. It allows you to move Flashback Time Travel history within the same database only, and not outside it. This package does not allow you to move Flashback Time Travel table data. It allows you to import customer history only and not Flashback Time Travel history.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.IMPORT_HISTORY (
   owner_name1          IN   VARCHAR2,
   table_name1          IN   VARCHAR2
   temp_history_name    IN   VARCHAR2 DEFAULT 'TEMP_HISTORY',
   options              IN   BINARY_INTEGER DEFAULT 0);
```

Parameters

Table 83-13 IMPORT_HISTORY Procedure Parameters

Parameter	Description
`owner_name1`	Schema of the Flashback Time Travel-enabled table
`table_name1`	Name of the Flashback Time Travel-enabled table
`temp_history_name`	Optional temporary history table from which we import history data
`options`	One (or a combination) of constants (`NODROP`, `NOCOMMIT`, and `NODELETE`) to specify if we want to drop, commit changes of, or truncate the temporary history table

Usage Notes

The database function `TIMESTAMP_TO_SCN` can be used to convert times to SCN when populating the temporary history table.

Related Topics

- CREATE_TEMP_HISTORY_TABLE Procedure
  This procedure creates a table called `TEMP_HISTORY` with the correct definition in schema.

83.4.13 LOCK_DOWN_APPLICATION Procedure

This procedure takes an application name and makes all the security tables read-only. The group called `SYSTEM` cannot be locked.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.LOCK_DOWN_APPLICATION (
   application_name           IN   VARCHAR2);
```

Parameters

Table 83-14 LOCK_DOWN_APPLICATION Procedure Parameters

Parameter	Description
`application_name`	Name of the application for which a table has been added as a security table

83.4.14 PURGE_CONTEXT Procedure

This procedure purges the context to be saved selected by the SET_CONTEXT_LEVEL Procedure.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.PURGE_CONTEXT;
```

Related Topics

- SET_CONTEXT_LEVEL Procedure
  This procedure defines how much of the user context is to be saved.

83.4.15 REASSOCIATE_FBA Procedure

This procedure reassociates the given table with the Flashback Time Travel.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.REASSOCIATE_FBA (
   owner_name      VARCHAR2,
   table_name      VARCHAR2);
```

Parameters

Table 83-15 REASSOCIATE_FBA Procedure Parameters

Parameter	Description
`owner_name`	Schema of the Flashback Time Travel enabled base table
`table_name`	Name of the Flashback Time Travel enabled base table

Exceptions

Table 83-16 REASSOCIATE_FBA Procedure Exceptions

Parameter	Description
ORA-55602	User table is not enabled for Flashback Time Travel
ORA-55636	Table definition validation failed

Usage Notes

- The procedure will signal an error if the base table and the history table do not have identical data definitions. For example, when columns are added or table is split, the resulting base table and history table need to have the same schema.
- The Flashback Time Travel internal history table schema has some row versions metadata columns. The procedure will signal an error if any metadata column is dropped by users.

83.4.16 REGISTER_APPLICATION Procedure

This procedure takes an application name and optionally a Flashback Archive, and registers an application for database hardening.

When database hardening is enabled, then all the security tables for that application are enabled for Flashback Archive using the given Flashback Archive. If no Flashback Archive is specified, the default Flashback Archive is used.

> **See Also:**
>
> Using Flashback Time Travel in *Oracle Database Development Guide* regarding database hardening

Syntax

```
DBMS_FLASHBACK_ARCHIVE.REGISTER_APPLICATION (
   application_name           IN   VARCHAR2,
   flashback_archive_name     IN    VARCHAR2 := NULL);
```

Parameters

Table 83-17 REGISTER_APPLICATION Procedure Parameters

Parameter	Description
application_name	Name of the application which is being registered. The application SYSTEM is already registered when the package is created and is populated with list of tables needed for database hardening.
flashback_archive_name	Name of the Flashback Archive in which the historical data for the security tables for given application is stored. If no Flashback Archive is specified, the default Flashback Archive is used.

83.4.17 REMOVE_TABLE_FROM_APPLICATION Procedure

This procedure takes an application name and marks a table in it as no longer being a security table.

If the application is already enabled for Flashback Archive, Flashback Archive will be disabled for this table.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.REMOVE_TABLE_TO_APPLICATION (
   application_name           IN   VARCHAR2,
   table_name                 IN   VARCHAR2,
   schema_name                IN   VARCHAR2 := NULL);
```

Parameters

Table 83-18 REMOVE_TABLE_FROM_APPLICATION Procedure Parameters

Parameter	Description
application_name	Name of the application for which a table is being removed from the list of security tables
table_name	Name of the table to mark as being no longer a security table for the given application
schema_name	Name of the schema containing the desired table. If no schema name is specified, the current schema is used.

83.4.18 SET_CONTEXT_LEVEL Procedure

This procedure defines how much of the user context is to be saved.

Syntax

```
DBMS_FLASHBACK_ARCHIVE.SET_CONTEXT_LEVEL (
   level          VARCHAR2);
```

Parameters

Table 83-19 SET_CONTEXT_LEVEL Procedure Parameters

Parameter	Description
level	Depending on how much of the user context needs to be saved: • ALL - the entire SYS_CONTEXT • TYPICAL - the user ID, global user ID and the hostname • NONE - nothing

83.4.19 VERIFY_BLOCKCHAIN_LIFESPAN Procedure

This procedure verifies the contents of the current lifespans, historical lifespans, or all lifespans of the user table rows that are protected using blockchain Flashback Archive (FBA). The

verification is done by comparing the row's cryptographic hash with the `ORAFBA_CURRENT_LIFESPAN_DIGEST$` column value recorded in its previous historical lifespan.

Syntax

```
PROCEDURE verify_blockchain_lifespan(
          schema_name                    IN  VARCHAR2,
          table_name                     IN  VARCHAR2,
          number_of_lifespan_verified   OUT NUMBER,
          type_of_lifespan               IN  NUMBER DEFAULT LIFESPAN_ALL);
```

Parameters

Table 83-20 VERIFY_BLOCKCHAIN_LIFESPAN Procedure Parameters

Parameter	Description
`schema_name`	The user schema name.
`table_name`	The user table name.
`number_of_lifespan_verified`	The number of lifespan rows verified.
`type_of_lifespan`	The type of lifespan to be verified. The following are the constants for the `type_of_lifespan` parameter. • `LIFESPAN_CURRENT CONSTANT NUMBER := 1;` • `LIFESPAN_HISTORICAL CONSTANT NUMBER := 2;` • `LIFESPAN_ALL CONSTANT NUMBER := 3;`

84
DBMS_FLASHBACK_ARCHIVE_MIGRATE

Using the DBMS_FLASHBACK_ARCHIVE_MIGRATE package, you can export and import the Flashback Archive base tables along with their history to another database using the Transportable Tablespaces.

Restrictions

You can use this migration package only on a non-CDB environment.

This chapter contains the following topics:

- DBMS_FLASHBACK_ARCHIVE_MIGRATE Overview
- DBMS_FLASHBACK_ARCHIVE_MIGRATE Security Model
- DBMS_FLASHBACK_ARCHIVE_MIGRATE Operational Notes
- DBMS_FLASHBACK_ARCHIVE_MIGRATE Examples
- Summary of DBMS_FLASHBACK Subprograms

> **See Also:**
>
> For detailed information about DBMS_FLASHBACK_ARCHIVE_MIGRATE:
>
> - *Oracle Database Development Guide*
> - *Oracle Database SQL Language Reference*.

84.1 DBMS_FLASHBACK_ARCHIVE_MIGRATE Overview

Using this package one can export and import the Flashback Archive base tables along with their history to another database using the Transportable Tablespaces.

84.2 DBMS_FLASHBACK_ARCHIVE_MIGRATE Security Model

The DBMS_FLASHBACK_ARCHIVE_MIGRATE package must be compiled on both the source and target databases as SYS. The source file is located at ?/rdbms/admin/crefbamig.sql, using which the package can be created or compiled.

The export and import procedures must be executed as SYS. Since the package uses DBMS_DATAPUMP, DBMS_LOCK, DBMS_SYSTEM, DBMS_SQL. and DBMS_SCHEDULER PL/SQL packages, their Security Models are also applicable. Along with the user's FDA tables that are being exported the export function also exports the following SYS tables' data:

- SMON_SCN_TIME
- SYS_FBA_FA
- SYS_FBA_TSFA

- SYS_FBA_TRACKEDTABLES
- SYS_FBA_PARTITIONS
- SYS_FBA_USERS

84.3 DBMS_FLASHBACK_ARCHIVE_MIGRATE Operational Notes

This section describes how to use the `DBMS_FLASHBACK_ARCHIVE_MIGRATE` package.

Prerequisites

Database version >= 11.2

If database version is `11.2`, following conditions should be met:

- `shared_pool_size` >= 500M
- `streams_pool_size` >= 40M

or

- `sga_target` >= 2G

Constants

None

1. On the source database run `flashback_archive.export` to export the given Flashback Archive enabled tables, their history tables and other relevant metadata information using the transportable tablespaces.
2. Copy the transportable tablespace export dump and relevant datafiles from the source to the target database.
3. On the target database run `flashback_archive.import` to finish the import of the exported Flashback Archive enabled table and their history.

Extraneous objects

The tablespaces which host the Flashback Archive enabled tables and Flashback archive itself may have totally unrelated objects. During the export all those objects are also exported as the process uses the transportable tablespaces feature. Such extraneous objects are listed in the table `SYS_FBA_EXTRANEOUSOBJS`.

Upon successful import, the extraneous objects are left as is. DBA should drop the extraneous objects.

Self Containment of Tablespaces

Since transportable tablespaces feature is used for export and import, all the objects in those tablespaces must be self contained. Otherwise, the function aborts.

Reduction in History Granularity

After the import, it is possible to see the reduction in the granularity of the history.

For example, The multiple history generated on the source database in a small duaration of time may be coalesced into one history.

Logs

Export and import logs are available in the OS directory given as the input data_pump_dir argument.

The format of export logs is fda_mig_expdp_YYYY_MMDD_HH24MI.log

The format of import logs is fda_mig_impdp_YYYY_MMDD_HH24MI.log

Operational Notes

- After export, involved tablespaces are kept in READONLY mode so that the DBA can copy the involved datafiles.
- Keeping the involved tablespaces in READONLY mode also ensures the exported tables are in immutable state during exporting and copying of the datafiles.
- The procedures provided cannot be used for exporting and importing to change the table blocksizes. Because, internally these procedures use transportable tablespace export and import functionalities.
- After successful import, the tablespaces are again kept in READONLY mode. After the DBA verify the imported data, tablespaces can be made READ-WRITE.
- After successful import, the small tablespace that was given as input to the export operation can be dropped.

84.4 DBMS_FLASHBACK_ARCHIVE_MIGRATE Examples

This section illustrates exporting and importing Flashback Archive tables along with their history.

Step 1: Export the Flashback Archive enables tables along with their history

In the example, three Flashback Archive enabled tables USR1_TAB1, USR1_TAB2, and USR1_TAB3 are exported. These tables are spread across the FA1_TBSP, FA1_TBSP_1, USR1_TBSP, and USR1_TBSP_1 tablespaces.

The TTS_TBSP is a small tablespace that holds export related metadata.

/some/dir/with/enough/disk/space is the OS directory where data pump files and logs are created.

```
set serveroutput on;
declare
  l_fda_tabs dbms_sql.varchar2_table;
begin
  l_fda_tabs(1) := upper('usr1_tab1');
  l_fda_tabs(2) := upper('usr1_tab2');
  l_fda_tabs(3) := upper('usr1_tab3');
  flashback_archive_migrate.export
  ( schema_owner  => 'USR1'
  , fda_tables    => l_fda_tabs
  , tts           => 'TTS_TBSP'
  , data_pump_dir => '/some/dir/with/enough/disk/space'
  );
end;
/
```

On completion, the export output lists the datapump dump file and the datafiles that are to be copied to the target database to complete the import.

`Step 2`: Copy the datapump dump file and datafiles listed in the `Step 1` to the target database.

`Step 3`: Import the Flashback Archive tables along with their history

In the example, three Flashback Archive enabled tables `USR1_TAB1`, `USR1_TAB2`, and `USR1_TAB3` are imported from `FA1_TBSP`, `FA1_TBSP_1`, `USR1_TBSP`, and `USR1_TBSP_1` tablespaces spread over the data files.

For the example given, following are the data files captured in the export (`Step 1`) log file:

- `/u01/app/oracle/oradata/TGT_DB/tts_tbsp.dbf`
- `/u01/app/oracle/oradata/TGT_DB/usr1_tbsp.dbf`
- `/u01/app/oracle/oradata/TGT_DB/usr1_tbsp1.dbf`
- `/u01/app/oracle/oradata/TGT_DB/FA1_tbsp.dbf`
- `/u01/app/oracle/oradata/TGT_DB/FA1_tbsp_1.dbf`

The `TTS_TBSP` is a small tablespace that holds the export related metadata whose datafile is `/ade/nkedlaya_fda2/rdbms/dbs/tts_tbsp.dbf`.

`/some/dir/with/enough/disk/space` is the `OS` directory where data pump files are created.

```
set serveroutput on;
   declare
     l_data_files dbms_sql.varchar2_table;
   begin
     l_data_files(1) := '/u01/app/oracle/oradata/TGT_DB/tts_tbsp.dbf';
     l_data_files(2) := '/u01/app/oracle/oradata/TGT_DB/usr1_tbsp.dbf';
     l_data_files(3) := '/u01/app/oracle/oradata/TGT_DB/usr1_tbsp1.dbf';
     l_data_files(4) := '/u01/app/oracle/oradata/TGT_DB/FA1_tbsp.dbf';
     l_data_files(5) := '/u01/app/oracle/oradata/TGT_DB/FA1_tbsp_1.dbf';

     flashback_archive_migrate.import
     ( schema_owner   => 'USR1'
     , tts            => 'TTS_TBSP'
     , data_pump_dir  => '/data/pump/dir/where/export/dump/file/resides'
     , tts_data_files => l_data_files
     );
   end;
   /
```

Upon successful completion of the import, Flashback Archive is enabled on the imported tables and the prior history is available along with the new history that will be generated.

84.5 Summary of DBMS_FLASHBACK Subprograms

This table lists the `DBMS_FLASHBACK_ARCHIVE_MIGRATE` subprograms and briefly describes them.

Table 84-1 DBMS_FLASHBACK_ARCHIVE_MIGRATE Package Subprograms

Subprogram	Description
EXPORT Procedure	This procedure exports the given Flashback Archive enabled base tables, their history, and related tablespaces.

Table 84-1 (Cont.) DBMS_FLASHBACK_ARCHIVE_MIGRATE Package Subprograms

Subprogram	Description
EXPORT_ANALYZE Procedure	This procedure analyzes the given Flashback Archive enabled base tables, their history, and related tablespaces for self containment using transportable tablesapce checks.
IMPORT Procedure	This procedure imports the Flashback Archive enabled base tables that were exported, their history, and related tablespaces.

84.5.1 EXPORT Procedure

This procedure exports the given Flashback Archive enabled base tables, their history, and related tablespaces.

Syntax

```
DBMS_FLASHBACK_ARCHIVE_MIGRATE.EXPORT (
   schema_owner        IN   VARCHAR2,
   fda_tables          IN   DBMS_SQL.VARCHAR2,
   tts                 IN   VARCHAR2,
   data_pump_dir       IN   VARCHAR2,
   ignore_errors       IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 84-2 EXPORT Procedure Parameters

Parameter	Description
schema_owner	Flashback Archive enabled tables' owner
fda_tables	Array of Flashback Archive enabled tables
tts	A small tablespace that can be used to hold the export related metadata
data_pump_dir	The directory path in the operating system where the export dump is placed
ignore_errors	Ignore any transportable tablespaces errors. The default value is FALSE.

84.5.2 EXPORT_ANALYZE Procedure

This procedure analyzes the given Flashback Archive enabled base tables, their history, and related tablespaces for self containment using Transportable tablesapce checks.

Syntax

```
DBMS_FLASHBACK_ARCHIVE_MIGRATE.EXPORT_ANALYZE (
   schema_owner        IN   VARCHAR2,
   fda_tables          IN   DBMS_SQL.VARCHAR2,
   tts                 IN   VARCHAR2);
```

Parameters

Table 84-3 EXPORT_ANALYZE Procedure Parameters

Parameter	Description
schema_owner	Flashback Archive enabled tables' owner
fda_tables	Array of Flashback Archive enabled tables
tts	A small tablespace that can be used to hold the export related metadata

84.5.3 IMPORT Procedure

This procedure Imports the Flashback Archive enabled base tables that were exported, their history, and related tablespaces.

Syntax

```
DBMS_FLASHBACK_ARCHIVE_MIGRATE.IMPORT (
   schema_owner        IN   VARCHAR2,
   tts                 IN   VARCHAR2,
   data_pump_dir       IN   VARCHAR2,
   tts_data_files      IN   DBMS_SQL.VARCHAR2_TABLE);
```

Parameters

Table 84-4 IMPORT Procedure Parameters

Parameter	Description
schema_owner	Flashback Archive enabled tables' owner
tts	A small tablespace that can be used to hold the export related metadata
data_pump_dir	The directory path in the operating system where the export dump is placed
tts_data_files	The data files list of TTS exported tablespaces.

85
DBMS_FREQUENT_ITEMSET

The DBMS_FREQUENT_ITEMSET package enables frequent itemset counting. The two functions are identical except in the input cursor format difference.

This chapter contains the following topics:

- Summary of DBMS_FREQUENT_ITEMSET Subprograms

85.1 Summary of DBMS_FREQUENT_ITEMSET Subprograms

The `DBMS_FREQUENT_ITEMSET` package includes the `FI_HORIZONTAL` function and `FI_TRANSACTIONAL` function subprograms.

Table 85-1 DBMS_FREQUENT_ITEMSET Package Subprograms

Subprogram	Description
FI_HORIZONTAL Function	Counts all frequent itemsets given a cursor for input data which is in 'HORIZONTAL' row format, support threshold, minimum itemset length, maximum itemset length, items to be included, items to be excluded
FI_TRANSACTIONAL Function	Counts all frequent itemsets given a cursor for input data which is in 'TRANSACTIONAL' row format, support threshold, minimum itemset length, maximum itemset length, items to be included, items to be excluded

85.1.1 FI_HORIZONTAL Function

The purpose of this table function is to count all frequent itemsets given a cursor for input data which is in 'HORIZONTAL' row format, support threshold, minimum itemset length, maximum itemset length, items to be included, items to be excluded. The result will be a table of rows in form of itemset, support, length, total transactions counted.

In 'HORIZONTAL' row format, each row contains all of the item ids for a single transaction. Since all of the items come together, no transaction id is necessary.

The benefit of this table function is that if an application already has data in horizontal format, the database can skip the step of transforming rows that are in transactional format into horizontal format.

Syntax

```
DBMS_FREQUENT_ITEMSET.FI_HORIZONTAL(
   tranx_cursor         IN    SYSREFCURSOR,
   support_threshold    IN    NUMBER,
   itemset_length_min   IN    NUMBER,
   itemset_length_max   IN    NUMBER,
   including_items      IN    SYS_REFCURSOR DEFAULT NULL,
   excluding_items      IN    SYS_REFCURSOR DEFAULT NULL)
  RETURN TABLE OF ROW (
     itemset [Nested Table of Item Type DERIVED FROM tranx_cursor],
```

```
    support         NUMBER,
    length          NUMBER,
    total_tranx     NUMBER);
```

Parameters

Table 85-2 FI_HORIZONTAL Function Parameters

Parameter	Description
tranx_cursor	The cursor parameter that the user will supply when calling the function. There is no limits on the number of returning columns.Each column of cursor represents an item. All columns of the cursor must be of the same datatype. The item id must be number or character type (for example, VARCHAR2(n)).
support_threshold	A fraction number of total transaction count. An itemset is termed "frequent" if [the number of transactions it occurs in] divided by [the total number of transactions] exceed the fraction. The parameter must be a NUMBER.
itemset_length_min	The minimum length for interested frequent itemset. The parameter must be a NUMBER between 1 and 20, inclusive.
itemset_length_max	The maximum length for interested frequent itemset. This parameter must be a NUMBER between 1 and 20, inclusive, and must not be less than itemset_length_min.
including_items	A cursor from which a list of items can be fetched. At least one item from the list must appear in frequent itemsets that are returned. The default is NULL.
excluding_items	A cursor from which a list of items can be fetched. No item from the list can appear in frequent itemsets that are returned.The default is NULL.

Return Values

Table 85-3 FI_HORIZONTAL Return Values

Parameter	Description
support	The number of transactions in which a frequent itemset occurs. This will be returned as a NUMBER.
itemset	A collection of items which is computed as frequent itemset. This will be returned as a nested table of item type which is the item column type of the input cursor.
length	Number of items in a frequent itemset. This will be returned as a NUMBER.
total_tranx	The total transaction count. This will be returned as a NUMBER.

Example

Suppose you have a table horiz_table_in.

```
horiz_table_in(iid1 VARCHAR2(30), iid2 VARCHAR2(30), iid3 VARCHAR2(30), iid4
VARCHAR2(30), iid5 VARCHAR2(30));
```

and the data in horiz_table_in looks as follows:

```
('apple', 'banana', NULL, NULL, NULL)
('apple', 'milk', 'banana', NULL, NULL)
('orange', NULL, NULL, NULL, NULL)
```

Suppose you want to find out what combinations of items is frequent with a given support threshold of 30%, requiring itemset containing at least one of ('apple','banana','orange'), but excluding any of ('milk') in any itemset. You use the following query:

```
CREATE TYPE fi_varchar_nt AS TABLE OF VARCHAR2(30);
SELECT CAST(itemset as FI_VARCHAR_NT)itemset, support, length, total_tranx
   FROM table(DBMS_FREQUENT_ITEMSET.FI_HORIZONTAL(
                CURSOR(SELECT iid1, iid2, iid3, iid4, iid5
                         FROM horiz_table_in),
                0.3,
                2,
                5,
                CURSOR(SELECT * FROM table(FI_VARCHAR_NT
                                           ('apple','banana','orange'))),
                CURSOR(SELECT * FROM table(FI_VARCHAR_NT('milk'))))));
```

85.1.2 FI_TRANSACTIONAL Function

This procedure counts all frequent itemsets given a cursor for input data which is in 'TRANSACTIONAL' row format, support threshold, minimum itemset length, maximum itemset length, items to be included, items to be excluded. The result will be a table of rows in form of itemset, support, length, total number of transactions.

In 'TRANSACTIONAL' row format, each transaction is spread across multiple rows. All the rows of a given transaction have the same transaction id, and each row has a different item id. Combining all of the item ids which share a given transaction id results in a single transaction.

Syntax

```
DBMS_FREQUENT_ITEMSET.FI_TRANSACTIONAL (
   tranx_cursor         IN    SYSREFCURSOR,
   support_threshold    IN    NUMBER,
   itemset_length_min   IN    NUMBER,
   itemset_length_max   IN    NUMBER,
   including_items      IN    SYS_REFCURSOR DEFAULT NULL,
   excluding_items      IN    SYS_REFCURSOR DEFAULT NULL)
  RETURN TABLE OF ROW (
     itemset [Nested Table of Item Type DERIVED FROM tranx_cursor],
     support        NUMBER,
     length         NUMBER,
     total_tranx    NUMBER);
```

Parameters

Table 85-4 *FI_TRANSACTIONAL Function Parameters*

Parameter	Description
tranx_cursor	The cursor parameter that the user will supply when calling the function. It should return two columns in its returning row, the first column being the transaction id, the second column being the item id. The item id must be number or character type (for example, VARCHAR2(n)).

Table 85-4 (Cont.) *FI_TRANSACTIONAL Function Parameters*

Parameter	Description
support_threshold	A fraction number of total transaction count. An itemset is termed "frequent" if [the number of transactions it occurs in] divided by [the total number of transactions] exceed the fraction. The parameter must be a NUMBER.
itemset_length_min	The minimum length for interested frequent itemset. The parameter must be a NUMBER between 1 and 20, inclusive.
itemset_length_max	The maximum length for interested frequent itemset. This parameter must be a NUMBER between 1 and 20, inclusive, and must not be less than itemset_length_min.
including_items	A cursor from which a list of items can be fetched. At least one item from the list must appear in frequent itemsets that will be returned. The default is NULL.
excluding_items	A cursor from which a list of items can be fetched. No item from the list can appear in frequent itemsets that will returned. The default is NULL.

Return Values

Table 85-5 FI_TRANSACTIONAL Return Values

Parameter	Description
support	The number of transactions in which a frequent itemset occurs. This will be returned as a NUMBER.
itemset	A collection of items which is computed as frequent itemset. This will be returned as a nested table of item type which is the item column type of the input cursor.
length	Number of items in a frequent itemset. This will be returned as a NUMBER.
total_tranx	The total transaction count. This will be returned as a NUMBER, and will be the same for all returned rows, similar to a reporting aggregate.

Usage Notes

Applications must predefine a nested table type of the input item type and cast the output itemset into this predefined nested table type before further processing, such as loading into a table.

Examples

Suppose that the input table tranx_table_in looks as follows:

```
(1, 'apple')
(1, 'banana')
(2, 'apple')
(2, 'milk')
(2, 'banana')
(3, 'orange')
```

and the user is trying to find itemsets that satisfy a support-threshold of 60% and have the itemset-length greater than 1 (namely, (apple, banana)).

The output of this function would contain the following output row:

```
itemset=('apple','banana'), support=2, length=2, total_tranx=3
```

You need to create a nested table of item type before you submit a query to perform the frequent itemset counting. In this example, since item is of `VARCHAR2(30)`, you must create a nested table of `VARCHAR2(30)`:

```
CREATE TYPE fi_varchar_nt AS TABLE OF VARCHAR2(30);
SELECT CAST(itemset as FI_VARCHAR_NT) itemset, support, length, total_tranx
   FROM table(DBMS_FREQUENT_ITEMSET.FI_TRANSACTIONAL(
                cursor(SELECT tid, iid FROM tranx_table_in),
                0.6,
                2,
                5,
                NULL,
                NULL));
```

Here is another example to illustrate how to include certain items and exclude certain items in the counting.

```
SELECT CAST(itemset as FI_VARCHAR_NT)itemset, support, length, total_tranx
   FROM table(DBMS_FREQUENT_ITEMSET.FI_TRANSACTIONAL(
                CURSOR(SELECT tid, iid FROM tranx_table_in),
                0.6,
                2,
                5,
               CURSOR(SELECT * FROM table(FI_VARCHAR_NT
                                ('apple','banana','orange'))),
               CURSOR(SELECT * FROM table(FI_VARCHAR_NT('milk')))));
```

Using the including/excluding items parameter, you are able to further optimize the execution by ignoring itemsets that are not expected by application.

You can also use transactional output through collection unnesting:

```
  SELECT
    bt.setid, nt.*
  FROM
  (SELECT cast(Itemset as FI_VARCHAR_NT) itemset, rownum setid
   FROM table(
     DBMS_FREQUENT_ITEMSET.FI_TRANSACTIONAL(
       CURSOR(SELECT tid, iid FROM tranx_table_in), 0.6, 2, 5,
       NULL, NULL))) bt,
  table(bt.itemset) nt;
```

If you want to use an insert statement to load frequent itemsets into a nested table, it is better to use the `NESTED_TABLE_FAST_INSERT` hint for performance:

```
  CREATE TABLE fq_nt (coll FI_VARCHAR_NT) NESTED TABLE coll STORE AS
    coll_nest;
  INSERT /*+ NESTED_TABLE_FAST_INSERT */  INTO fq_nt
   SELECT cast(itemset as FI_VARCHAR_NT)
   FROM table(DBMS_FREQUENT_ITEMSET.FI_TRANSACTIONAL(
          cursor(SELECT tid, iid FROM tranx_table_in), 0.6, 2, 5,
          NULL, NULL));
```

Note that if you want to use the package inside a PL/SQL cursor, you must cast the return type of the table function:

```
  CREATE TYPE fi_res AS OBJECT (
    itemset      FI_VARCHAR_NT,
    support      NUMBER,
    length       NUMBER,
    total_tranx  NUMBER
  );
  /
  CREATE TYPE fi_coll AS TABLE OF fi_res;
  /

  DECLARE
    cursor freqC is
      SELECT Itemset
      FROM table(
        CAST(DBMS_FREQUENT_ITEMSET.FI_TRANSACTIONAL(
          cursor(SELECT tid, iid FROM tranx_table_in), 0.6, 2, 5,
          NULL, NULL) AS fi_coll));
    coll_nt  FI_VARCHAR_NT;
    num_rows int;
    num_itms int;
  BEGIN
    num_rows := 0;
    num_itms := 0;
    OPEN freqC;
    LOOP
      FETCH freqC INTO coll_nt;
      EXIT WHEN freqC%NOTFOUND;
      num_rows := num_rows + 1;
      num_itms := num_itms + coll_nt.count;
    END LOOP;
  CLOSE freqC;
  DBMS_OUTPUT.PUT_LINE('Totally ' || num_rows || ' rows ' || num_itms || '
items were produced.');
END;
/
```

86
DBMS_FS

The `DBMS_FS` package for performing operations on an Oracle file system (make, mount, unmount and destroy operations) in an Oracle database.

This chapter contains the following topics:

- DBMS_FS Overview
- DBMS_FS Security Model
- Summary of DBMS_FS Subprograms

86.1 DBMS_FS Overview

The `DBMS_FS` package contains Oracle file system (OFS) procedures that you can use to create, mount, unmount, and destroy an Oracle file system.

Starting 19c release, the file systems are supported by PDB. Oracle Database supports maximum 5 file systems per PDB and 1000 file systems in total.

The `DBMS_FS` package enables applications to access database objects from a universal client such as an NFS server. This feature interfaces with Oracle SecureFiles to provide the file system access.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for a detailed description of managing an NFS server in Oracle Database

86.2 DBMS_FS Security Model

You must have the `SYSDBA` administrative privilege to use the `DBMS_FS` package.

The operations that you perform using the `DBMS_FS` package are equivalent to the file system operations that are performed in an operating system by the root user. Access to the individual file system that is created and mounted by this package is enforced using Access Control Lists (ACLs) and the permissions on the mounted directories to the operating system user.

86.3 Summary of DBMS_FS Subprograms

The following table lists the `DBMS_FS` subprograms and briefly describes them.

Table 86-1 DBMS_FS Subprograms

Subprogram	Description
DESTROY_ORACLE_FS Procedure	Destroys an Oracle file system, using the `fstype` and of name `fsname`.
FS_EXISTS Procedure	Validates if the specified file system exists. Run this procedure before executing any mount, unmount or destroy operation.
LIST_FILES Procedure	Lists all the files within a specified directory in a file system.
MAKE_ORACLE_FS Procedure	Creates a file system of type specified by `fstype` and of name `fsname`.
MOUNT_ORACLE_FS Procedure	Mounts an Oracle file system on the specified mount point.
UNMOUNT_ORACLE_FS Procedure	Unmounts an Oracle file system on the specified mount point.

86.3.1 DESTROY_ORACLE_FS Procedure

This procedure destroys an Oracle file system and then frees the resources that were associated with it. Run the `dbms_fs.destroy_oracle_fs()` procedure to destroy file systems that are no longer in use.

Syntax

```
DBMS_FS.DESTROY_ORACLE_FS (
   fstype       IN VARCHAR2,
   fsname       IN VARCHAR2);
```

Parameters

Table 86-2 DBMS_FS Parameters

Parameter	Description
`fstype`	File system type. Oracle File System (OFS) and Database File System (DBFS) are supported.
`fsname`	Name of the file system

Usage Notes

- You can find information about the currently mounted file systems by querying the `V$OFSMOUNT` dynamic view.
- For more information about the file system types, see the `fstype` description in MAKE_ORACLE_FS Procedure.
- Before you run the `DBMS_FS.DESTROY_ORACLE_FS` procedure, you must unmount the file system by using the `DBMS_FS.UNMOUNT_ORACLE_FS` procedure. If you run this procedure on a file system that is still in use, it results in an error.
- After you run `DBMS_FS.DESTROY_ORACLE_FS`, Oracle Database destroys the file system and frees the associated resources.

Example

The following sample code shows how to destroy a DBFS file system, `dbfs_fs1`.

```
BEGIN
 DBMS_FS.DESTROY_ORACLE_FS (
  fstype           => 'dbfs',
  fsname           => 'dbfs_fs1');
END;
/
```

86.3.2 FS_EXISTS Procedure

Use the `dbms_fs.fs_exists()` procedure to validate if the given file system exists before executing any mount, unmount or destroy operation.

Syntax

```
DBMS_FS.FS_EXISTS (
   fsname      IN VARCHAR2);
```

Parameters

Table 86-3 DBMS_FS Parameters

Parameter	Description
`fsname`	Name of the file system.

Example

The following sample code shows how to check if a file system exists with the specified file system name, `data`.

```
SQL> declare
     fsname varchar2(64) := 'data';
     fsexists integer;

     begin
      fs_exists := dbms_fs.fs_exists(fsname);
      dbms_output.put_line('fs exists: ' || fsexists);
     end;
     /
SQL> select dbms_fs.fs_exists('data') from dual;
```

86.3.3 LIST_FILES Procedure

Run the `dbms_fs.list_files()` procedure to list all the files within a specified directory in a file system.

Syntax

```
DBMS_FS.LIST_FILES (
   fsname      IN VARCHAR2,
   dirpath     IN VARCHAR2);
```

Parameters

Table 86-4 DBMS_FS Parameters

Parameter	Description
`fsname`	Name of the file system.
`dirpath`	Path to an existing directory in the file system. Enter a path that is relative to the mount path. For example, if you enter `'/'` as the directory path, it is relative to the mount path.

Example

The following sample code shows how to list all the files within a specified directory in a file system.

```
SQL> select * from DBMS_FS.LIST_FILES ('data', '/') from dual;
```

Where, the path to the directory (`'/'`) is relative to the mount path.

86.3.4 MAKE_ORACLE_FS Procedure

This procedure creates a new file system of type `DBFS` or `OFS`, on top of an existing Oracle tablespace or other database object.

Syntax

```
DBMS_FS.MAKE_ORACLE_FS (
   fstype      IN VARCHAR2,
   fsname      IN VARCHAR2,
   fsoptions   IN VARCHAR2);
```

Parameters

Table 86-5 DBMS_FS Parameters

Parameter	Description
`fstype`	File system type. Enter one of the following values: • `ofs`—to create an Oracle File System • `dbfs`—to create a Database File System
`fsname`	Name of the file system. Enter a string no longer than 256 characters, using alphanumeric characters.
`fsoptions`	Specify an existing tablespace to use for the Oracle file system, using the following format: `"tablespace=`*`tablespace_name`*`"`

Usage Notes

- If you want to create a database file system (DBFS), then you must run the `dbfs_create_filesystem.sql` script, which in turn calls the `dbfs_create_filesystem_advanced.sql` script. By default, this script is in the `$ORACLE_HOME/rdbms/admin` directory. When you run this script, provide the name of an

existing tablespace and a name for the file system that will be stored in the database. The size of the file system will be the same as the table size. For example, to create a file system in the `dbfs_ts` tablespace, in the file system `dbfs_tab`:

```
@/$ORACLE_HOME/rdbms/admin/dbfs_create_filesystem.sql dbfs_ts dbfs_tab
```

After you run this script, you can use the other procedures in the `DBMS_FS` package to mount, unmount, and destroy the file system.

> **Note:**
>
> Staring Oracle Database `19.3.1.0` release, the `DBMS_FS.MAKE_ORACLE_FS` is used to create a `DBFS` filesystem; hence no auxiliary SQL script is needed to create a `DBFS` filesystem.

- Running the `DBMS_FS.MAKE_ORACLE_FS` procedure on the database instance is equivalent to running the `mkfs` command by root in an operating system.
- The tablespace that you specified in the `fsoptions` parameter must already exist before you execute the `DBMS_FS.MAKE_ORACLE_FS` procedure. To find existing tablespaces, query the `DBA_TABLESPACES` data dictionary view.
- The size of the file system is the same size as this tablespace.

Example

The following example shows how to create a DBFS file system named `dbfs_fs1` in the tablespace `dbfs_fs1_tbspc`.

```
BEGIN
 DBMS_FS.MAKE_ORACLE_FS (
  fstype           => 'dbfs',
  fsname           => 'dbfs_fs1',
  fsoptions        => 'TABLESPACE=dbfs_fs1_tbspc');
END;
/
```

86.3.5 MOUNT_ORACLE_FS Procedure

Use the `dbms_fs.mount_oracle_fs()` procedure to mount an Oracle file system or OFS managed file system on the specified mount point.

Before you begin, complete the following checks to ensure that:

- You have created the file system using the `dbms_fs.make_oracle_fs()` procedure.
- The mount point that you specify exists in the local node.
- The Oracle user has access permissions.
- The mount path is empty, which means that the directory specified by the mount point does not have any files.

Oracle supports an extensive list of mount options that you can use to have better control on resources and achieve good performance. When you use the `persist` mount option, file systems are automatically remounted every time instance restarts. Since OFSD is a non-fatal background process, it gets automatically restarted after the death of a process. File systems that are mounted at the time of the death of the OFSD process are automatically remounted

after starting a new OFSD process. This ensures continuous availability of the mounted file systems even in the case of an error.

To improve the throughput, OFSD has its own local cache for writes and reads. The write cache uses 8 (1 MB) buffers per file connection, and it can use up to 256 MB per file system. You can modify this value through the mount option, `wcache_size`. A read-ahead algorithm is implemented for read operation and it uses 2 (1 MB) per file connection and it can use up to 256 MB per file system. Use the mount option, `rcache_size`, to modify this value.

The read-write cache is maintained per node, so this provides local cache consistency. In an RAC environment, consistency is guaranteed only after the `flush()` or `close()` operation is performed on the file. When two different processes on two different RAC nodes modify a single file and write to the same offset, then the first process that performs the `close()` operation on the file will have its data written into the file.

Syntax

```
DBMS_FS.MOUNT_ORACLE_FS (
  fstype           IN VARCHAR2,
  fsname           IN VARCHAR2,
  mount_point      IN VARCHAR2,
  mount_options    IN VARCHAR2);
```

Parameters

Table 86-6 MOUNT_ORACLE_FS Procedure Parameters

Parameter	Description
`fstype`	File system type. Oracle File System (OFS) and Database File System (DBFS) are supported.
`fsname`	Name of the file system. Enter a string no longer than 256 characters, using alpha numeric characters.
`mount_point`	Local directory where the file system should be mounted. This directory must already exist. Enter an absolute path. The maximum number of mount points that you can create is 5 mount points per PDB and 1000 mount points per instance.
`mount_options`	Comma-separated mount options, listed in Table 86-7.

Usage Notes

Table 86-7 Supported Mount Options for the MOUNT_ORACLE_FS procedure

Mount Option	Usage Description
`db_access`	Optimizes the read and write operations performed by database process when you access files managed through OFS or DBFS.
`default_permissions`	Enables permission check and restrict access based on file mode. This option is useful with the `allow_other` mount option.
`allow_other`	Allows other users apart from the operating system user that did the mount can access the files. This will be used in conjunction with permission checks in determining the file access. This option requires setting the `user_allow_other` parameter in the `/etc/fuse.conf` configuration file on Linux.
`max_read`	Maximum size of the read operation. No maximum size is set by default.

Table 86-7 (Cont.) Supported Mount Options for the MOUNT_ORACLE_FS procedure

Mount Option	Usage Description
max_write	Maximum write size in a single request. The default is 128K.
direct_io	Indicates to the operating system kernel not use file system cache.
nopersist	Does not store the mount options for use in next instance start up.
persist	Stores the mount entry persistently so that on subsequent instance start up it will be automatically mounted again. This option is supported for both ofs and dbfs.
ro	Mounts the file system in read-only mode. Files cannot be modified.
rw	Mounts the file system as read-write. This is the default.
nosuid	Specifies that the file system cannot contain set userid files.
suid	Specifies that the file system can contain set userid files. This is the default.
ofs_cache_attr_time	Specifies the OFS cache attribute timeout in seconds. Default value is 5 seconds. The permissible range is 0 to UB4MAXVAL. Oracle recommends that you specify low values when you use RAC. If you specify a larger value, such as 1000, the attributes from the files may not match if the same file system is mounted in another path. This is due to the aggressive caching that is done in the OFS layer.
exec	Allows executing a file under the mount path. This option is enabled by default.
noexec	Prohibits executing a file under the mount path. Use this option to disallow any file to be executed in the file system.
atime	Maintains information about the time when the file was accessed. This option is enabled by default.
noatime	Does not maintain information about the time when the was file accessed. When you enable this option, it increases throughput as the access time is not updated for every read operation.
max_readahead	Sets the maximum size, in bytes, for the read-ahead operations. The default value depends on the versions of the kernel and FUSE in your operating system.
sync_read	Permits only synchronous reads. If you use this option, it disables read-ahead caching for OFS.
async_read	Permits asynchronous reads. All read and write operations in the file system as performed asynchronously. This is the default option for FUSE versions 7.6 or later.
dirsync	Makes all directory operations synchronous. FUSE utilizes this option.
big_writes	Permits writes larger than 4KB. This is the default option for FUSE versions 7.6 or later.
no_big_writes	Does not allow writes larger than 4KB. This is not a recommended option for OFS.
xattr_enabled	Enables the use of extended attributes. Specify this option only if you want to use extended attributes.

Table 86-7 (Cont.) Supported Mount Options for the MOUNT_ORACLE_FS procedure

Mount Option	Usage Description
`kernel_cache_mode=writeback`	Makes the Linux kernel cache behave like a writeback cache when you enter `writeback` as the value for this option. A writeback cache provides better performance when there are many small input-output requirements. FUSE utilizes this option. Permitted values are: `writeback` and `write-through`. The default value is `write-through`.
`kernel_cache_attr_time=N`	Sets the Linux kernel attribute cache timeout to N seconds. The default value is infinity. Oracle recommends that you set low values for RAC setup.
`wcache_size=N`	Sets the OFS write cache size to N bytes. The default value is 256 MB.
`rcache_size`	Sets the OFS read cache size to N bytes. The default value is 256 MB.
`statfs_ctime`	Sets the time, in seconds, for which `statfs()` gets cached in OFS. The default value is 300 seconds or 5 minutes. Oracle recommends that you use the default value to avoid repeated `statfs()` calls from FUSE.
`no_rbt_cache`	Disallows the use of a red-black tree for caching directory entries. By default, a red-black tree is used to ensure accuracy when there are concurrent updates while listing directories.
`posix_locks`	Implements the POSIX file locks in OFS. This is not supported by OFS. If you want to use the file-locking mechanism, use the `noposix_locks` option.
`noposix_locks`	Implements POSIX file locks in Linux kernel. This is the default option.

Note:

The following options are exclusive options and cannot be used together:

- `nopersist/persist`
- `ro/rw`
- `nosuid/suid`
- `exec/noexec`
- `atime/noatime`
- `big_writes/no_big_writes`
- `sync_read/async_read`
- `posix_locks/noposix_locks`

Usage Notes

- This procedure makes the files system visible in the local database instance.

- For more information about the file system types, see the `fstype` description in MAKE_ORACLE_FS Procedure.
- You can find information about currently mounted file systems by querying the `V$OFSMOUNT` dynamic view.
- Run the `DBMS_FS.MOUNT_ORACLE_FS` procedure on a file system that has already been created with `DBMS_FS.MAKE_ORACLE_FS` in the local computer node where the Oracle database instance is running. You cannot run this procedure on file systems that were created outside of Oracle Database.
- You cannot update the mount options after mounting the file system. If you want to change the mount options later, you'll have to unmount the file system, and then remount it.

Example 1: Mounts a DBFS file system

Mounts a DBFS file system at `/oracle/dbfs/testfs`.

```
BEGIN
 DBMS_FS.MOUNT_ORACLE_FS (
  fstype           => 'dbfs',
  fsname           => 'dbfs_fs1',
  mount_point      => '/oracle/dbfs/testfs',
  mount_options    => 'default_permissions, allow_other, db_access');
END;
```

Example 2: Persist mount a DBFS file system

Persist mounts a DBFS file system at `/oracle/dbfs/testfs`.

```
BEGIN
 DBMS_FS.MOUNT_ORACLE_FS (
  fstype           => 'dbfs',
  fsname           => 'dbfs_fs1',
  mount_point      => '/oracle/dbfs/testfs',
  mount_options    => 'default_permissions, allow_other, persist, db_access');
END;
```

86.3.6 UNMOUNT_ORACLE_FS Procedure

This procedure unmounts an Oracle file system on the specified mount point.

File systems that are mounted in a PDB are automatically unmounted when the PDB is closed. You can also use `dbms_fs.unmount_oracle_fs()` to explicitly unmount an Oracle file system that you have mounted through OFS.

Syntax

```
DBMS_FS.UNMOUNT_ORACLE_FS (
  fsname             IN VARCHAR2,
  mount_point        IN VARCHAR2,
  unmount_options    IN VARCHAR2);
```

Table 86-8 UNMOUNT_ORACLE_FS Procedure Parameters

Parameter	Description
`fsname`	Name of the file system.
`mount_point`	Local directory where the file system had been mounted. Enter an absolute path.

Table 86-8 (Cont.) UNMOUNT_ORACLE_FS Procedure Parameters

Parameter	Description
unmount_options	Optionally, enter force or clean. Enter force to unmount the file system forcibly. This setting prevents new requests from being sent to the file system. All pending requests on the file system are either completed or canceled. If you omit this setting, then attempts to unmount a busy file system cause an EBUSY error. Enter clean if you don't want the file system to be remounted when the OFSD restarts.

Usage Notes

- Before you unmount the file system, ensure that all applications that use this file system are shut down. Also ensure that no processes reference the mounting file system.
- You can find information about the currently mounted file systems by querying the V$OFSMOUNT dynamic view.
- For more information about the file system types, see the fstype description in MAKE_ORACLE_FS Procedure.
- After unmounting the file system, the write permissions are removed from the mount point to prevent applications from writing to the underlying file system instead of the OFS supported file systems.
- When an Oracle instance is shut down in normal immediate mode, then all the mounted file systems are automatically unmounted.
- If a file system is mounted with the MOUNT_ORACLE_FS procedure with the persist option, it will be automatically mounted again when the database instance starts or the PDB is plugged. If this file system is unmounted by executing DBMS_FS.UNMOUNT_ORACLE_FS, it will remain unmounted even if the persist option was used to mount it.
- If you perform a SHUTDOWN ABORT, then the file system may still show as mounted but it may not be accessible. In this case, you can unmount the system manually by calling the unmount command at the operating system level or the fusermount procedure on Linux systems.
- Do not use fusermount -u to unmount a running file system as it causes inconsistency in Oracle views, such as v$ofsmount.
- You can export the local mount point of an Oracle file system to point to the remote system, and then NFS mount the file system from the remote system by using the operating system mount command. The DBMS_FS.MOUNT_ORACLE_FS procedure is similar to mount commands that are used for other local file systems.
- For better security, Oracle recommends that you use access control lists (ACLs) and Kerberos to control access to sensitive data.
- Do not attempt to unmount the file system from the operating system level. Doing so can leave the Oracle Database-created file system internal tables in an inconsistent state.

Example

The following sample code unmounts a DBFS mounted file system at /oracle/dbfs/testfs.

```
BEGIN
 DBMS_FS.UNMOUNT_ORACLE_FS (
  fsname           => 'dbfs_fs1',
  mount_point      => '/oracle/dbfs/testfs',
  mount_options    => 'force');
END;
```

87
DBMS_GOLDENGATE_ADM

The `DBMS_GOLDENGATE_ADM` package provides subprograms to configure and manage Oracle GoldenGate conflict detection and resolution.

This chapter contains the following topics:

- Using DBMS_GOLDENGATE_ADM
- Summary of DBMS_GOLDENGATE_ADM Subprograms

87.1 Using DBMS_GOLDENGATE_ADM

This section contains topics which relate to using the `DBMS_GOLDENGATE_ADM` package.

- DBMS_GOLDENGATE_ADM Overview
- DBMS_GOLDENGATE_ADM Security Model

87.1.1 DBMS_GOLDENGATE_ADM Overview

The `DBMS_GOLDENGATE_ADM` package provides interfaces to configure automatic conflict detection and resolution in an Oracle GoldenGate configuration that replicates tables between Oracle databases.

When more than one replica of a table allows changes to the table, a conflict can occur when a change is made to the same row in two different databases at nearly the same time. Oracle GoldenGate replicates changes using row logical change records (LCRs). It detects a conflict by comparing the old values in the row LCR with the current values of the corresponding table row identified by the key columns. If any column value does not match, then there is a conflict. After a conflict is detected, Oracle GoldenGate can resolve the conflict by overwriting values in the row with some values from the row LCR, ignoring the values in the row LCR, or computing a delta to update the row values.

XStream inbound servers and outbound servers can be used in an XStream configuration in a multitenant container database (CDB). A CDB is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs).

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product.

> **See Also:**
>
> - *Oracle Database XStream Guide*
> - *Oracle Database Concepts* for more information about CDBs and PDBs

87.1.2 DBMS_GOLDENGATE_ADM Security Model

Security on this package can be controlled either by granting `EXECUTE` on this package to selected users or roles or by granting `EXECUTE_CATALOG_ROLE` to selected users or roles.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

An Oracle GoldenGate administrator must be configured at each Oracle database in the table's replication environment, and Oracle GoldenGate must be configured to replicate the table at each Oracle database. You can configure an Oracle GoldenGate administrator using the `GRANT_ADMIN_PRIVILEGE` procedure in the `DBMS_GOLDENGATE_ADM` package.

> **See Also:**
>
> The Oracle GoldenGate documentation for more information about Oracle GoldenGate replication and configuring an Oracle GoldenGate administrator

87.2 Summary of DBMS_GOLDENGATE_ADM Subprograms

Table 87-1 DBMS_GOLDENGATE_ADM Package Subprograms

Subprogram	Description
ADD_AUTO_CDR Procedure	Configures Oracle GoldenGate automatic conflict detection and resolution for a table
ADD_AUTO_CDR_COLUMN_GROUP Procedure	Adds a column group and configures Oracle GoldenGate automatic conflict detection and resolution for the column group
ADD_AUTO_CDR_DELTA_RES Procedure	Configures Oracle GoldenGate automatic conflict detection and delta resolution for the column
ALTER_AUTO_CDR Procedure	Alters the Oracle GoldenGate automatic conflict detection and resolution for a table
ALTER_AUTO_CDR_COLUMN_GROUP Procedure	Alters a column group for Oracle GoldenGate automatic conflict detection and resolution
DELETE_PROCREP_EXCLUSION_OBJ Procedure	Deletes a database object from the exclusion list for Oracle GoldenGate procedural replication
GG_PROCEDURE_REPLICATION_ON Function	Returns 1 if Oracle GoldenGate procedural replication is enabled and returns 0 if it is disabled
INSERT_PROCREP_EXCLUSION_OBJ Procedure	Inserts a database object into the exclusion list for Oracle GoldenGate procedural replication
PURGE_TOMBSTONES Procedure	Purges rows that were deleted before the specified timestamp from the tombstone table
REMOVE_AUTO_CDR Procedure	Removes Oracle GoldenGate automatic conflict detection and resolution for a table
REMOVE_AUTO_CDR_COLUMN_GROUP Procedure	Removes a column group that was configured for Oracle GoldenGate automatic conflict detection and resolution

Table 87-1 (Cont.) DBMS_GOLDENGATE_ADM Package Subprograms

Subprogram	Description
REMOVE_AUTO_CDR_DELTA_RES Procedure	Removes Oracle GoldenGate automatic conflict detection and delta resolution for the column

> **Note:**
>
> All procedures commit unless specified otherwise.

87.2.1 ADD_AUTO_CDR Procedure

This procedure configures Oracle GoldenGate automatic conflict detection and resolution for a table.

The conflict detection and resolution configured by this procedure is based on the timestamp of the changes. The procedure adds one or more hidden columns of `TIMESTAMP` type to the table, and each hidden column is counted against the limit of 1,000 columns for each table.

The procedure automatically places the columns in the table into a default column group and into an unconditional supplemental log group, excluding nonscalar columns. To create column groups that include a subset of the columns in the table, use the `ADD_AUTO_CDR_COLUMN_GROUP` procedure in this package.

This procedure is overloaded. One version of this procedure contains the `additional_options` parameter, and the other does not.

Syntax

```
DBMS_GOLDENGATE_ADM.ADD_AUTO_CDR(
   schema_name             IN VARCHAR2,
   table_name              IN VARCHAR2,
   resolution_granularity  IN VARCHAR2 DEFAULT 'ROW',
   existing_data_timestamp IN TIMESTAMP WITH TIME ZONE DEFAULT 'SYSTIMESTAMP',
   tombstone_deletes       IN BOOLEAN DEFAULT TRUE,
   fetchcols               IN BOOLEAN DEFAULT TRUE,
   record_conflicts        IN BOOLEAN DEFAULT TRUE,
   use_custom_handlers     IN BINARY_INTEGER DEFAULT 0);


DBMS_GOLDENGATE_ADM.ADD_AUTO_CDR(
   schema_name             IN VARCHAR2,
   table_name              IN VARCHAR2,
   resolution_granularity  IN VARCHAR2 DEFAULT 'ROW',
   existing_data_timestamp IN TIMESTAMP WITH TIME ZONE DEFAULT 'SYSTIMESTAMP',
   tombstone_deletes       IN BOOLEAN DEFAULT TRUE,
   fetchcols               IN BOOLEAN DEFAULT TRUE,
   record_conflicts        IN BOOLEAN DEFAULT TRUE,
   use_custom_handlers     IN BINARY_INTEGER DEFAULT 0,
   additional_options      IN BINARY_INTEGER);
```

Parameters

Table 87-2 ADD_AUTO_CDR Procedure Parameters

Parameter	Description
schema_name	The name of the table's schema.
table_name	The name of the table.
resolution_granularity	ROW, the default, adds one hidden TIMESTAMP column for the row and one hidden TIMESTAMP column for each LOB column. COLUMN adds one hidden TIMESTAMP column for each column in the table.
existing_data_timestamp	Timestamp to assign to existing rows. If NULL, then the current system timestamp is used. If a time is specified, and the operating system time zone is not a valid Oracle time zone, then Oracle uses UTC as the default value.
tombstone_deletes	TRUE, the default, tracks deleted rows in a tombstone table. Tracking deleted rows might be required to detect and resolve some conflicts, but tracking deleted rows requires additional database resources. FALSE does not track deleted rows in a tombstone table.
fetchcols	TRUE, the default, fetches the value of LOBs during conflict detection and resolution. Fetching LOBs can be an expensive operation. FALSE does not fetch the value of LOBs during conflict detection and resolution.
record_conflicts	TRUE, the default, records the conflict in the DBA_APPLY_ERROR and DBA_APPLY_ERROR_MESSAGES views. FALSE does not record the conflict.
use_custom_handlers	0, the default, indicates that automatic conflict handlers are used. 1 indicates that automatic conflict handlers are not used and that a custom error handler must be specified using the SET_DML_HANDLER procedure in the DBMS_APPLY_ADM package.

Table 87-2 (Cont.) ADD_AUTO_CDR Procedure Parameters

Parameter	Description
additional_options	0 indicates NONE. You can also use the DBMS_GOLDENGATE_ADM.DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_NONE constant. 1 indicates EARLIEST_TIMESTAMP_RESOLUTION. In earliest timestamp resolution, the earliest timestamp is used for conflict resolution instead of the latest timestamp. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_EARLIEST_TIMESTAMP constant. 2 indicates DELETE_ALWAYS_WINS. In delete always wins, when there is a conflict between a delete operation and another operation, the delete operation succeeds and the other operation is discarded. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_DELETE_ALWAYS_WINS constant. 4 indicates IGNORE_GLOBAL_SITE_PRIORITY. In ignore global site priority, the locally defined resolution is used for the table even if global site priority is set in the Oracle GoldenGate replicat file. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_IGNORE_GLOBAL_SITE_PRIORITY constant. 8 indicates ADD_KEY_VERSION. Use this option to add a version number with TIMESTAMP data type to handle delete conflicts. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_ADD_KEY_VERSION constant. Add the options together to use more than one option. For example, to use earliest timestamp resolution and delete always wins, use 3.

Example

The following example illustrates how to use bitmask:

```
SQL> EXECUTE dbms_goldengate_adm.add_auto_cdr('scott','table1',
   additional_options =>
   dbms_goldengate_adm.additional_options_delete_always_wins +
      dbms_goldengate_adm.additional_options_add_key_version,
    record_conflicts => TRUE);
```

87.2.2 ADD_AUTO_CDR_COLUMN_GROUP Procedure

This procedure adds a column group to a table that is configured for Oracle GoldenGate automatic conflict detection and resolution.

For a table that has been configured for timestamp conflict detection and resolution, this procedure adds a column group that includes a specified subset of columns in the table. Any columns in the table that are not part of a column group remain in the default column group for the table.

When you add a column group to a table, conflict detection and resolution is performed on the columns in the column group separately from the other columns in the table. Column groups

enable different databases to update different columns in the same row at nearly the same time without causing a conflict.

For example, a replicated table that contains employee information might have a salary column and a bonus column as well as other columns that identify the employee and a location column for the employees office number. Assume that one department in the company updates its database to change the employee's salary while another department updates its database to change the employee's location. If the salary and bonus columns are in a column group, these changes are applied to the replicated table in each database without requiring conflict resolution.

The procedure automatically places the columns in the column group into an unconditional supplemental log group, excluding nonscalar columns.

Before this procedure can be run on a table, the `DBMS_GOLDENGATE_ADM.ADD_AUTO_CDR` procedure must be run in the table with `ROW` specified for the `resolution_granularity` parameter.

Syntax

```
DBMS_GOLDENGATE_ADM.ADD_AUTO_CDR_COLUMN_GROUP(
   schema_name              IN VARCHAR2,
   table_name               IN VARCHAR2,
   column_list              IN VARCHAR2,
   column_group_name        IN VARCHAR2 DEFAULT NULL,
   existing_data_timestamp IN TIMESTAMP WITH TIME ZONE DEFAULT NULL);
```

Parameters

Table 87-3 ADD_AUTO_CDR_COLUMN_GROUP Procedure Parameters

Parameter	Description
`schema_name`	The name of the table's schema.
`table_name`	The name of the table.
`column_list`	Group of columns for which the conflict detection and resolution is configured. Specify the columns in a comma-separated list. The same column cannot be in more than one column group. Also, the same column cannot be in a column group and specified in a delta resolution.
`column_group_name`	The name of the column group. If `NULL`, the column group name is system generated.
`existing_data_timestamp`	The time value for the added `TIMESTAMP` columns for existing table data.

87.2.3 ADD_AUTO_CDR_DELTA_RES Procedure

This procedure configures Oracle GoldenGate automatic conflict detection and delta resolution for the column.

The resolution method does not depend on a timestamp or an extra resolution column. With delta conflict resolution, the conflict is resolved by adding the difference between the new and old values in the LCR to the value in the table. For example, if a bank balance is updated at two sites concurrently, then the converged value accounts for all debits and credits. This resolution method is generally used for financial data such as an account balance.

The procedure automatically places the column into an unconditional supplemental log group.

Before this procedure can be run on a table, the `DBMS_GOLDENGATE_ADM.ADD_AUTO_CDR` procedure must be run in the table with `ROW` specified for the `resolution_granularity` parameter.

Syntax

```
DBMS_GOLDENGATE_ADM.ADD_AUTO_CDR_DELTA_RES(
   schema_name IN VARCHAR2,
   table_name  IN VARCHAR2,
   column_name IN VARCHAR2);
```

Parameters

Table 87-4 ADD_AUTO_CDR_DELTA_RES Procedure Parameters

Parameter	Description
`schema_name`	The name of the table's schema.
`table_name`	The name of the table.
`column_name`	The name of the column. The specified column must be a `NUMBER` or `FLOAT` data type column. The same column cannot be in a column group.

87.2.4 ALTER_AUTO_CDR Procedure

This procedure alters the Oracle GoldenGate automatic conflict detection and resolution for a table.

This procedure is overloaded. One version of this procedure contains the `additional_options` parameter, and the other does not.

Syntax

```
DBMS_GOLDENGATE_ADM.ALTER_AUTO_CDR(
   schema_name             IN VARCHAR2,
   table_name              IN VARCHAR2,
   tombstone_deletes       IN BOOLEAN DEFAULT TRUE,
   fetchcols               IN BOOLEAN DEFAULT TRUE,
   record_conflicts        IN BOOLEAN DEFAULT TRUE,
   use_custom_handlers     IN BINARY_INTEGER DEFAULT 0);


DBMS_GOLDENGATE_ADM.ALTER_AUTO_CDR(
   schema_name             IN VARCHAR2,
   table_name              IN VARCHAR2,
   tombstone_deletes       IN BOOLEAN DEFAULT TRUE,
   fetchcols               IN BOOLEAN DEFAULT TRUE,
   record_conflicts        IN BOOLEAN DEFAULT TRUE,
   use_custom_handlers     IN BINARY_INTEGER DEFAULT 0,
   additional_options      IN BINARY_INTEGER);
```

Parameters

Table 87-5 ALTER_AUTO_CDR Procedure Parameters

Parameter	Description
schema_name	The name of the table's schema.
table_name	The name of the table.
tombstone_deletes	TRUE tracks deleted rows in a tombstone table. Tracking deleted rows might be required to detect and resolve some conflicts, but tracking deleted rows requires additional database resources. FALSE does not track deleted rows in a tombstone table. NULL retains the current setting for the parameter.
fetchcols	TRUE fetches the value of nonscalar columns during conflict detection and resolution. FALSE does not fetch the value of nonscalar columns during conflict detection and resolution. NULL retains the current setting for the parameter.
record_conflicts	TRUE records the conflict. FALSE does not record the conflict. NULL retains the current setting for the parameter.
use_custom_handlers	0, default, indicates that automatic conflict handlers are used. 1 indicates that automatic conflict handlers are not used and that a custom error handler must be specified using the SET_DML_HANDLER procedure in the DBMS_APPLY_ADM package. NULL, the default, retains the current setting for the parameter.

Table 87-5 (Cont.) ALTER_AUTO_CDR Procedure Parameters

Parameter	Description
additional_options	A bit mask that indicates which of six types of actions to follow: • 0 - no changes are made from any previous settings. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_NONE constant. • 1 - sets earliest timestamp, but clears all other options/bits. You can also use the DBMS_GOLDENGATE_ADM.EARLIEST_TIMESTAMP_RESOLUTION constant. • 2 - sets delete always wins, but clears all other options/bits. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_DELETE_ALWAYS_WINS constant. • 4 - indicates IGNORE_GLOBAL_SITE_PRIORITY. When you specify this option, then use local site resolution configuration for the table, even if the global site priority is set in the Oracle GoldenGate replicat file. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_IGNORE_GLOBAL_SITE_PRIORITY constant. • 8 - indicates ADD_KEY_VERSION. Use this option to add a version number with TIMESTAMP data type to handle delete conflicts. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_ADD_KEY_VERSION constant. • 16 - clears all. You can also use the DBMS_GOLDENGATE_ADM.ADDITIONAL_OPTIONS_CLEAR_OPTIONS constant. **Note:** • Any combination of 1,2, 4, and 8 are possible. • Any combination that include 1 or 2 will implicitly set 8. You can add the options together to use more than one option. For example: If the previously used option was DELETE_ALWAYS_WINS, then when you specify a value of 1 for additional_options, EARLIEST_TIMESTAMP_RESOLUTION and 16 means clear all options. As a result, there are no additional options set.

87.2.5 ALTER_AUTO_CDR_COLUMN_GROUP Procedure

This procedure alters a column group for Oracle GoldenGate automatic conflict detection and resolution.

Syntax

```
DBMS_GOLDENGATE_ADM.ALTER_AUTO_CDR_COLUMN_GROUP(
   schema_name        IN VARCHAR2,
   table_name         IN VARCHAR2,
   column_group_name  IN VARCHAR2,
   add_column_list    IN VARCHAR2,
   remove_column_list IN VARCHAR2);
```

Parameters

Table 87-6 ALTER_AUTO_CDR_COLUMN_GROUP Procedure Parameters

Parameter	Description
`schema_name`	The name of the table's schema.
`table_name`	The name of the table.
`column_group_name`	The name of the column group.
`add_column_list`	A comma-separated list of columns to add to the column group.
`remove_column_list`	A comma-separated list of columns to remove from the column group.

87.2.6 DELETE_PROCREP_EXCLUSION_OBJ Procedure

This procedure deletes a database object from the exclusion list for Oracle GoldenGate procedural replication.

When a database object is on the exclusion list for Oracle GoldenGate procedural replication, execution of subprogram in the package is not replicated if the subprogram operates on the excluded object. For example, if `hr.employees` is an excluded database object for the `DBMS_REDEFINITION` package, then an execution of the `DBMS_REDEFINITION.START_REDEF_TABLE` procedure on the `hr.employees` table is not replicated.

> **Caution:**
>
> Run the `DELETE_PROCREP_EXCLUSION_OBJ` procedure only under the direction of Oracle Support.

Syntax

```
DBMS_GOLDENGATE_ADM.DELETE_PROCREP_EXCLUSION_OBJ(
   package_owner     IN VARCHAR2 DEFAULT NULL,
   package_name      IN VARCHAR2 DEFAULT NULL,
   object_owner      IN VARCHAR2 DEFAULT NULL,
   object_name       IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 87-7 DELETE_PROCREP_EXCLUSION_OBJ Procedure Parameters

Parameter	Description
package_owner	The owner of the package.
package_name	The name of the package.
object_owner	The owner of the object.
object_name	The name of the object.

87.2.7 GG_PROCEDURE_REPLICATION_ON Function

This function returns 1 if Oracle GoldenGate procedural replication is enabled and returns 0 if it is disabled.

Syntax

```
DBMS_GOLDENGATE_ADM.GG_PROCEDURE_REPLICATION_ON
RETURN NUMBER;
```

87.2.8 INSERT_PROCREP_EXCLUSION_OBJ Procedure

This procedure inserts a database object into the exclusion list for Oracle GoldenGate procedural replication.

When a database object is on the exclusion list for Oracle GoldenGate procedural replication, execution of subprogram in the package is not replicated if the subprogram operates on the excluded object. For example, if `hr.employees` is an excluded database object for the `DBMS_REDEFINITION` package, then an execution of the `DBMS_REDEFINITION.START_REDEF_TABLE` procedure on the `hr.employees` table is not replicated.

> **Caution:**
>
> Run the `INSERT_PROCREP_EXCLUSION_OBJ` procedure only under the direction of Oracle Support.

Syntax

```
DBMS_GOLDENGATE_ADM.INSERT_PROCREP_EXCLUSION_OBJ(
   package_owner     IN VARCHAR2 DEFAULT NULL,
   package_name      IN VARCHAR2 DEFAULT NULL,
   object_owner      IN VARCHAR2 DEFAULT NULL,
   object_name       IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 87-8 INSERT_PROCREP_EXCLUSION_OBJ Procedure Parameters

Parameter	Description
package_owner	The owner of the package.
package_name	The name of the package.
object_owner	The owner of the object.
object_name	The name of the object.

87.2.9 PURGE_TOMBSTONES Procedure

This procedure purges rows that were deleted before the specified timestamp from the tombstone table.

Syntax

```
DBMS_GOLDENGATE_ADM.PURGE_TOMBSTONES(
   purge_timestamp IN TIMESTAMP WITH TIME ZONE);
```

Parameters

Table 87-9 PURGE_TOMBSTONES Procedure Parameters

Parameter	Description
purge_timestamp	The timestamp before which records are purged.

87.2.10 REMOVE_AUTO_CDR Procedure

This procedure removes Oracle GoldenGate automatic conflict detection and resolution for a table.

Syntax

```
DBMS_GOLDENGATE_ADM.REMOVE_AUTO_CDR(
   schema_name             IN VARCHAR2,
   table_name              IN VARCHAR2);
```

Parameters

Table 87-10 REMOVE_AUTO_CDR Procedure Parameters

Parameter	Description
schema_name	The name of the table's schema.
table_name	The name of the table.

87.2.11 REMOVE_AUTO_CDR_COLUMN_GROUP Procedure

This procedure removes a column group that was configured for Oracle GoldenGate automatic conflict detection and resolution.

Syntax

```
DBMS_GOLDENGATE_ADM.REMOVE_AUTO_CDR_COLUMN_GROUP(
   schema_name             IN VARCHAR2,
   table_name              IN VARCHAR2,
   column_group_name       IN VARCHAR2);
```

Parameters

Table 87-11 REMOVE_AUTO_CDR_COLUMN_GROUP Procedure Parameters

Parameter	Description
schema_name	The name of the table's schema.
table_name	The name of the table.
column_group_name	The name of the column group.

87.2.12 REMOVE_AUTO_CDR_DELTA_RES Procedure

This procedure removes Oracle GoldenGate automatic conflict detection and delta resolution for the column.

Syntax

```
DBMS_GOLDENGATE_ADM.REMOVE_AUTO_CDR_DELTA_RES(
   schema_name IN VARCHAR2,
   table_name  IN VARCHAR2,
   column_name IN VARCHAR2);
```

Parameters

Table 87-12 REMOVE_AUTO_CDR_DELTA_RES Procedure Parameters

Parameter	Description
schema_name	The name of the table's schema.
table_name	The name of the table.
column_name	The name of the column.

88
DBMS_GOLDENGATE_AUTH

The `DBMS_GOLDENGATE_AUTH` package provides subprograms for granting privileges to and revoking privileges from GoldenGate administrators.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_GOLDENGATE_AUTH Subprograms

> **See Also:**
>
> GRANT_ADMIN_PRIVILEGE Procedure in the DBMS_XSTREAM_AUTH package

88.1 DBMS_GOLDENGATE_AUTH Overview

This package provides subprograms for granting privileges to GoldenGate administrators and revoking privileges from GoldenGate administrators. A GoldenGate administrator manages an integrated GoldenGate and XStream Out configuration.

GoldenGate administrators can be used in a multitenant container database (CDB). A CDB is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs).

> **See Also:**
>
> - GRANT_ADMIN_PRIVILEGE Procedure in the DBMS_XSTREAM_AUTH package
> - *Oracle Database XStream Guide*
> - *Oracle Database Concepts* for more information about CDBs and PDBs

88.2 DBMS_GOLDENGATE_AUTH Security Model

Security on this package can be controlled by granting `EXECUTE` on this package to selected users or roles, or by granting `EXECUTE_CATALOG_ROLE` to selected users or roles.

The user executing the subprograms in the `DBMS_GOLDENGATE_AUTH` package must have `SYSDBA` administrative privilege, and the user must exercise the privilege using `AS SYSDBA` at connect time.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

To ensure that the user who runs the subprograms in this package has the necessary privileges, connect as an administrative user who can create users, grant privileges, and create tablespaces when using this package.

88.3 Summary of DBMS_GOLDENGATE_AUTH Subprograms

The DBMS_XSTREAM_AUTH package includes the GRANT_ADMIN_PRIVILEGE procedure and REVOKE_ADMIN_PRIVILEGE procedure subprograms.

Table 88-1 DBMS_GOLDENGATE_AUTH Package Subprograms

Subprogram	Description
GRANT_ADMIN_PRIVILEGE Procedure	Either grants the privileges needed by a user to be a GoldenGate administrator directly, or generates a script that grants these privileges
REVOKE_ADMIN_PRIVILEGE Procedure	Either revokes GoldenGate administrator privileges from a user directly, or generates a script that revokes these privileges

> **Note:**
>
> All subprograms commit unless specified otherwise.

88.3.1 GRANT_ADMIN_PRIVILEGE Procedure

This procedure grants the privileges needed by a user to be a GoldenGate administrator.

> **See Also:**
>
> GRANT_ADMIN_PRIVILEGE Procedure in the DBMS_XSTREAM_AUTH package

Syntax

```
DBMS_GOLDENGATE_AUTH.GRANT_ADMIN_PRIVILEGE(
   grantee                    IN  VARCHAR2,
   privilege_type             IN  VARCHAR2  DEFAULT '*',
   grant_select_privileges    IN  BOOLEAN   DEFAULT TRUE,
   do_grants                  IN  BOOLEAN   DEFAULT TRUE,
   file_name                  IN  VARCHAR2  DEFAULT NULL,
   directory_name             IN  VARCHAR2  DEFAULT NULL
   grant_optional_privileges  IN  VARCHAR2  DEFAULT NULL,
   container                  IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameters

Table 88-2 GRANT_ADMIN_PRIVILEGE Procedure Parameters

Parameter	Description
grantee	The user to whom privileges are granted
privilege_type	Specify one of the following values: • CAPTURE Specifying CAPTURE grants the minimum privileges required by the user to administer Oracle GoldenGate integrated extract. • APPLY Specifying APPLY grants the minimum privileges required by the user to administer Oracle GoldenGate integrated replicat. • * Specifying * grants the minimum privileges required by the user to administer Oracle GoldenGate integrated extract and Oracle GoldenGate integrated replicat.
grant_select_privileges	If TRUE, then the procedure grants a set of privileges, including SELECT_CATALOG_ROLE, to the user. This setting is recommended for GoldenGate administrators. If FALSE, then the procedure does not grant the set of privileges to the user. SELECT_CATALOG_ROLE enables the user to select from the data dictionary.
do_grants	If TRUE, then the procedure grants the privileges to the specified grantee directly, and adds the grantee to the DBA_GOLDENGATE_PRIVILEGES data dictionary view. If the user already has an entry in this data dictionary view, then the procedure does not make another entry, and no error is raised. If TRUE and any of the grant statements fails, then the procedure raises an error. If FALSE, then the procedure does not grant the privileges to the specified grantee directly, and does not add the grantee to the DBA_GOLDENGATE_PRIVILEGES data dictionary view. You specify FALSE when the procedure is generating a file that you will run later. If you specify FALSE and either the file_name or directory_name parameter is NULL, then the procedure raises an error. Note: It is recommended that do_grants be set to TRUE because many APIs check for the presence of a user in the DBA_GOLDENGATE_PRIVILEGES view, which will not be populated if do_grants is set to FALSE.
file_name	The name of the file generated by the procedure. The file contains all of the statements that grant the privileges. If a file with the specified file name exists in the specified directory name, then the grant statements are appended to the existing file. If NULL, then the procedure does not generate a file.

Table 88-2 (Cont.) GRANT_ADMIN_PRIVILEGE Procedure Parameters

Parameter	Description
directory_name	The directory into which the generated file is placed. The specified directory must be a directory object created using the SQL statement CREATE DIRECTORY. If you specify a directory, then the user who invokes the procedure must have the WRITE privilege on the directory object. If the file_name parameter is NULL, then this parameter is ignored, and the procedure does not generate a file. If NULL and the file_name parameter is non-NULL, then the procedure raises an error.
grant_optional_privileges	A comma-separated list of optional privileges to grant to the grantee. You can specify the following roles and privileges: • XDBADMIN • DV_XSTREAM_ADMIN • DV_GOLDENGATE_ADMIN • EXEMPT_ACCESS_POLICY • EXEMPT_REDACTION_POLICY
container	If CURRENT, then grants privileges to the grantee only in the container where the procedure is invoked. CURRENT can be specified while connected to the root or to a PDB. If ALL, then grants privileges to the grantee in all containers in the CDB and all PDBs created after the procedure is invoked. To specify ALL, the procedure must be invoked in the root by a common user. If a container name, then grants privileges to the grantee only in the specified container. To specify root, use CDB$ROOT while connected to the root. To specify a PDB, the procedure must be invoked in the root. **Note:** This parameter only applies to CDBs.

Usage Notes

The user who runs the procedure must be an administrative user who can grant privileges to other users.

Specifically, the procedure grants the following privileges to the specified user:

- The RESTRICTED SESSION system privilege
- EXECUTE on the following packages:
  - DBMS_APPLY_ADM
  - DBMS_AQ
  - DBMS_AQADM
  - DBMS_AQIN
  - DBMS_AQELM
  - DBMS_CAPTURE_ADM
  - DBMS_FLASHBACK
  - DBMS_LOCK
  - DBMS_PROPAGATION_ADM

  - DBMS_RULE_ADM
  - DBMS_TRANSFORM
  - DBMS_XSTREAM_ADM
- Privileges to enqueue messages into and dequeue messages from any queue
- Privileges to manage any queue
- Privileges to create, alter, and execute any of the following types of objects in the user's own schema and in other schemas:
  - Evaluation contexts
  - Rule sets
  - Rules

  In addition, the grantee can grant these privileges to other users.
- SELECT_CATALOG_ROLE
- SELECT or READ privilege on data dictionary views related to GoldenGate and Oracle Replication
- The ability to allow a remote GoldenGate administrator to perform administrative actions through a database link by connecting to the grantee

  This ability is enabled by running the GRANT_REMOTE_ADMIN_ACCESS procedure in this package.

> **Note:**
>
> This procedure grants only the privileges necessary to configure and administer a GoldenGate environment. You can grant additional privileges to the grantee if necessary.

> **See Also:**
>
> GRANT_ADMIN_PRIVILEGE Procedure in the DBMS_XSTREAM_AUTH package

88.3.2 REVOKE_ADMIN_PRIVILEGE Procedure

This procedure revokes GoldenGate administrator privileges from a user.

Syntax

```
DBMS_GOLDENGATE_AUTH.REVOKE_ADMIN_PRIVILEGE(
   grantee                     IN  VARCHAR2,
   privilege_type              IN  VARCHAR2  DEFAULT '*',
   revoke_select_privileges    IN  BOOLEAN   DEFAULT FALSE,
   do_revokes                  IN  BOOLEAN   DEFAULT TRUE,
   file_name                   IN  VARCHAR2  DEFAULT NULL,
   directory_name              IN  VARCHAR2  DEFAULT NULL
   revoke_optional_privileges  IN  VARCHAR2  DEFAULT NULL,
   container                   IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameters

Table 88-3 *REVOKE_ADMIN_PRIVILEGE Procedure Parameters*

Parameter	Description
grantee	The user from whom privileges are revoked
privilege_type	Specify one of the following values: • CAPTURE Specifying CAPTURE revokes the minimum privileges required by the user to administer Oracle GoldenGate integrated extract. • APPLY Specifying APPLY revokes the minimum privileges required by the user to administer Oracle GoldenGate integrated replicat. • * Specifying * revokes the minimum privileges required by the user to administer Oracle GoldenGate integrated extract and Oracle GoldenGate integrated replicat.
revoke_select_privileges	If TRUE, then the procedure revokes SELECT_CATALOG_ROLE from the user. If FALSE, then the procedure does not revoke SELECT_CATALOG_ROLE to the user. SELECT_CATALOG_ROLE enables the user to select from the data dictionary.
do_revokes	If TRUE, then the procedure revokes the privileges from the specified user directly, and removes the user from the DBA_XSTREAM_ADMINISTRATOR data dictionary view. If the user does not have a record in this data dictionary view, then the procedure does not remove a record from the view, and no error is raised. If TRUE and any of the revoke statements fails, then the procedure raises an error. A revoke statement fails if the user is not granted the privilege that is being revoked. If FALSE, then the procedure does not revoke the privileges from the specified user directly, and does not remove the user from the DBA_XSTREAM_ADMINISTRATOR data dictionary view. You specify FALSE when the procedure is generating a file that you will run later. If you specify FALSE and either the file_name or directory_name parameter is NULL, then the procedure does not raise an error.
file_name	The name of the file generated by this procedure. The file contains all of the statements that revoke the privileges. If a file with the specified file name exists in the specified directory name, then the revoke statements are appended to the existing file. If NULL, then the procedure does not generate a file.
directory_name	The directory into which the generated file is placed. The specified directory must be a directory object created using the SQL statement CREATE DIRECTORY. If you specify a directory, then the user who invokes the procedure must have the WRITE privilege on the directory object. If the file_name parameter is NULL, then this parameter is ignored, and the procedure does not generate a file. If NULL and the file_name parameter is non-NULL, then the procedure raises an error.
revoke_optional_privileges	A comma-separated list of optional privileges to revoke from the grantee, such as the DV_XSTREAM_ADMIN and DV_GOLDENGATE_ADMIN privileges

Table 88-3 (Cont.) *REVOKE_ADMIN_PRIVILEGE Procedure Parameters*

Parameter	Description
`container`	If `CURRENT`, then revokes privileges from the grantee only in the container where the procedure is invoked. `CURRENT` can be specified while connected to the root or to a PDB. If `ALL`, then revokes privileges from the grantee in all containers in the CDB. To specify `ALL`, the procedure must be invoked in the root. If a container name, then revokes privileges from the grantee only in the specified container. To specify root, use `CDB$ROOT` while connected to the root. To specify a PDB, the procedure must be invoked in the root. **Note:** This parameter only applies to CDBs.

Usage Notes

The user who runs this procedure must be an administrative user who can revoke privileges from other users. Specifically, this procedure revokes the privileges granted by running the `GRANT_ADMIN_PRIVILEGE` procedure in this package.

> **See Also:**
>
> "GRANT_ADMIN_PRIVILEGE Procedure"

89
DBMS_HADOOP

The DBMS_HADOOP package provides a PL/SQL procedure called CREATE_EXTDDL_FOR_HIVE(), that creates an Oracle external table for a given hive table.

Big Data SQL needs to be correctly set up for DBMS_HADOOP to work.

This chapter contains the following topics:

- DBMS_HADOOP Overview
- DBMS_HADOOP Security Model
- Summary of DBMS_HADOOP Subprograms

89.1 DBMS_HADOOP Overview

The DBMS_HADOOP package provides two procedures for creating an Oracle external table and for synchronizing the Oracle external table partitions.

These procedures are:

- CREATE_EXTDDL_FOR_HIVE() — creates an Oracle external table for a given hive table
- SYNCHRONIZE_PARTITIONS_FOR_HIVE() — helps to synchronize the Oracle external table partitions with those in the corresponding hive table

89.2 DBMS_HADOOP Security Model

Users must have ALTER privileges on the table for modifying external table parameters. In addition the ALTER privileges, users must have READ privileges for the directory object that contains the external data source and WRITE privileges for the directory objects containing bad, log, and discard files.

Similar privileges are applicable to the partitioned external tables.

89.3 Summary of DBMS_HADOOP Subprograms

DBMS_HADOOP includes the CREATE_EXTDDL_FOR_HIVE procedure and SYNC_PARTITIONS_FOR_HIVE procedure subprograms.

Table 89-1 DBMS_HADOOP Subprograms

Subprogram	Description
CREATE_EXTDDL_FOR_HIVE Procedure	Given a hive table name, creates a text of DDL that can be executed to create an external table corresponding to the hive table
CREATE_HYBRID_PARTNED_TABLE Procedure	Merges a given partitioned Hive table with an Oracle partitioned table, and replaces it with the merged table.

Table 89-1 (Cont.) DBMS_HADOOP Subprograms

Subprogram	Description
SYNCHRONIZE_PARTITIONS_FOR_HIVE Procedure	Synchronizes the existing partitioning definition of a hive table in Oracle catalog

89.3.1 CREATE_EXTDDL_FOR_HIVE Procedure

This procedure creates an Oracle external table for a given hive table.

Syntax

```
DBMS_HADOOP.CREATE_EXTDDL_FOR_HIVE (
   cluster_id         IN  VARCHAR2,
   db_name            IN  VARCHAR2 := NULL,
   hive_table_name    IN  VARCHAR2,
   hive_partition     IN  BOOLEAN,
   table_name         IN  VARCHAR2 := NULL,
   perform_DDL        IN  BOOLEAN := FALSE,
   text_of_DDL        OUT CLOB);
```

Parameters

Table 89-2 CREATE_EXTDDL_FOR_HIVE Procedure Parameters

Parameter	Description
cluster_id	Hadoop cluster ID
hive_db_name	Database where the hive table is located
hive_table_name	Name of the hive table
hive_partition	If this argument is TRUE and if the original hive table is partitioned, the corresponding Oracle external table will also be partitioned using the same partition key(s). If the original hive table is not partitioned, hive_partition=TRUE has no effect. If hive_partition=FALSE, the corresponding Oracle external table will not be partitioned even if the original hive table is partitioned.
table_name	Name of the Oracle external table
perform_DDL	If this argument is TRUE, the external table will be automatically created. Otherwise, only the textual representation of the DDL statement will be generated and returned in text_of_DDL .
text_of_DDL	If the argument perform_DDL is FALSE, only the textual representation of the DDL statement will be generated and returned in text_of_DDL .

89.3.2 CREATE_HYBRID_PARTNED_TABLE Procedure

This procedure merges the partitioned Hive table with an Oracle partitioned table, and replace it with the merged table.

Syntax

```
DBMS_HADOOP.CREATE_HYBRID_PARTNED_TABLE (
   cluster_id                    IN     VARCHAR2,
   db_name            IN    VARCHAR2,
   hive_table_name    IN     VARCHAR2,
   table_name                   IN     VARCHAR2,
   table_owner        IN     VARCHAR2,
   perform_ddl                IN     BOOLEAN := TRUE,
   text_of_ddl        OUT    CLOB);
```

Parameters

Table 89-3 CREATE_HYBRID_PARTNED_TABLE Procedure Parameters

Parameter	Description
cluster_id	Hadoop cluster ID
db_name	Database where the partitioned Hive table is located
hive_table_name	Name of the partitioned Hive table
table_name	Name of the partitioned Oracle table
table_owner	The owner of the partitioned Oracle table
perform_DDL	If the value of this parameter is TRUE, the external table is automatically created. If the value of this parameter is FALSE, only the textual representation of the DDL statement is generated and returned in text_of_DDL.
text_of_DDL	If the argument perform_DDL is FALSE, only the textual representation of the DDL statement is generated and returned in text_of_DDL.

89.3.3 SYNCHRONIZE_PARTITIONS_FOR_HIVE Procedure

This procedure synchronizes the Oracle external table partitions with those in the corresponding hive table.

Syntax

```
DBMS_HADOOP.SYNCHRONIZE_PARTITIONS_FOR_HIVE (
   table_name     IN  VARCHAR2,
   table_owner    IN  VARCHAR2);
```

Parameters

Table 89-4 SYNCHRONIZE_PARTITIONS_FOR_HIVE Procedure Parameters

Parameter	Description
`table_name`	Oracle external table
`table_owner`	Schema name

90
DBMS_HEAT_MAP

The `DBMS_HEAT_MAP` package provides an interface to externalize heatmaps at various levels of storage including block, extent, segment, object and tablespace. A second set of subprograms externalize the heatmaps materialized by the background for top N tablespaces.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_HEAT_MAP Subprograms

> **See Also:**
>
> - Heat Map in *Oracle Database VLDB and Partitioning Guide*
> - DBMS_ILM
> - DBMS_ILM_ADMIN

90.1 DBMS_HEAT_MAP Overview

To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification. You can also use Automatic Data Optimization (ADO) to automate the compression and movement of data between different tiers of storage within the database.

The Heat Map tracks modification times at the block level, and multiple access statistics at the segment level. Objects in the `SYSTEM` and `SYSAUX` tablespaces are not tracked. `DBMS_HEAT_MAP` gives you access to the Heat Map statistics at various levels - block, extent, segment, object, and tablespace.

90.2 DBMS_HEAT_MAP Security Model

The execution privilege is granted to `PUBLIC`. Procedures in this package run under the caller security. The user must have `ANALYZE` privilege on the object.

90.3 Summary of DBMS_HEAT_MAP Subprograms

This table lists and briefly describes the `DBMS_HEAT_MAP` package subprograms.

Table 90-1 DBMS_HEAT_MAP Package Subprograms

Subprogram	Description
BLOCK_HEAT_MAP Function	Returns last modification time for each block in a table segment

Table 90-1 (Cont.) DBMS_HEAT_MAP Package Subprograms

Subprogram	Description
EXTENT_HEAT_MAP Function	Returns the extent level Heat Map statistics for a table segment
OBJECT_HEAT_MAP Function	Returns the minimum, maximum and average access times for all the segments belonging to the object
SEGMENT_HEAT_MAP Procedure	Returns the heatmap attributes for the given segment
TABLESPACE_HEAT_MAP Function	Returns the minimum, maximum and average access times for all the segments in the tablespace

90.3.1 BLOCK_HEAT_MAP Function

This table function returns the last modification time for each block in a table segment. It returns no information for segment types that are not data.

Syntax

```
DBMS_HEAT_MAP.BLOCK_HEAT_MAP (
   owner             IN VARCHAR2,
   segment_name      IN VARCHAR2,
   partition_name    IN VARCHAR2 DEFAULT NULL,
   sort_columnid     IN NUMBER DEFAULT NULL,
   sort_order        IN VARCHAR2 DEFAULT NULL)
RETURN hm_bls_row PIPELINED;
```

Parameters

Table 90-2 *BLOCK_HEAT_MAP Function Parameters*

Parameter	Description
`owner`	Owner of the segment
`segment_name`	Table name of a non-partitioned table or (sub)partition of partitioned table. Returns no rows when table name is specified for a partitioned table.
`partition_name`	Defaults to `NULL`. For a partitioned table, specify the partition or subpartition segment name.
`sort_columnid`	ID of the column on which to sort the output. Valid values 1..9. Invalid values are ignored.
`sort_order`	Defaults to `NULL`. Possible values: `ASC`, `DESC`

Return Values

Table 90-3 *BLOCK_HEAT_MAP Function Return Values (Output Parameters)*

Parameter	Description
`owner`	Owner of the segment
`segment_name`	Segment name of the non-partitioned table
`partition_name`	Partition or subpartition name

Table 90-3 (Cont.) *BLOCK_HEAT_MAP Function Return Values (Output Parameters)*

Parameter	Description
tablespace_name	Tablespace containing the segment
file_id	Absolute file number of the block in the segment
relative_fno	Relative file number of the block in the segment
block_id	Block number of the block
write time	Last modification time of the block

90.3.2 EXTENT_HEAT_MAP Function

This table function returns the extent level Heat Map statistics for a table segment. It returns no information for segment types that are not data. Aggregates at extent level, including minimum modification time and maximum modification time, are included.

Syntax

```
DBMS_HEAT_MAP.EXTENT_HEAT_MAP (
   owner             IN VARCHAR2,
   segment_name      IN VARCHAR2,
   partition_name    IN VARCHAR2 DEFAULT NULL,
RETURN hm_els_row PIPELINED;
```

Parameters

Table 90-4 *EXTENT_HEAT_MAP Function Parameters*

Parameter	Description
owner	Owner of the segment
segment_name	Table name of a non-partitioned table or (sub)partition of partitioned table. Returns no rows when table name is specified for a partitioned table.
partition_name	Defaults to NULL. For a partitioned table, specify the partition or subpartition segment name.

Return Values

Table 90-5 *EXTENT_HEAT_MAP Function Return Values (Output Parameters)*

Parameter	Description
owner	Owner of the segment
segment_name	Segment name of the non-partitioned table
partition_name	Partition or subpartition name
tablespace_name	Tablespace containing the segment
file_id	Absolute file number of the block in the segment
relative_fno	Relative file number of the block in the segment
block_id	Block number of the block

Table 90-5 (Cont.) *EXTENT_HEAT_MAP Function Return Values (Output Parameters)*

Parameter	Description
blocks	Number of blocks in the extent
bytes	Number of bytes in the extent
min_writetime	Minimum of last modification time of the block
max_writetime	Maximum of last modification time of the block
avg_writetime	Average of last modification time of the block

90.3.3 OBJECT_HEAT_MAP Function

This table function returns the minimum, maximum and average access times for all the segments belonging to the object.

The object must be a table. The table function raises an error if called on object tables other than table.

Syntax

```
DBMS_HEAT_MAP.OBJECT_HEAT_MAP (
   object_owner      IN VARCHAR2,
   object_name       IN VARCHAR2)
 RETURN hm_object_table PIPELINED;
```

Parameters

Table 90-6 *OBJECT_HEAT_MAP Function Parameters*

Parameter	Description
object_owner	Tablespace containing the segment
object_name	Segment header relative file number

Return Values

Table 90-7 *OBJECT_HEAT_MAP Function Return Values (Output Parameters)*

Parameter	Description
segment_name	Name of the top level segment
partition_name	Name of the partition
tablespace_name	Name of the tablespace
segment_type	Type of segment as in DBA_SEGMENTS.SEGMENT_TYPE
segment_size	Segment size in bytes
min_writetime	Oldest write time for the segment
max_writetime	Latest write time for the segment
avg_writetime	Average write time for the segment
min_readtime	Oldest read time for the segment

Table 90-7 (Cont.) *OBJECT_HEAT_MAP Function Return Values (Output Parameters)*

Parameter	Description
max_readtime	Latest read time for the segment
avg_writetime	Average write time for the segment
min_lookuptime	Oldest index lookup time for the segment
max_lookuptime	Latest index lookup time for the segment
avg_lookuptime	Average index lookup time for the segment
min_ftstime	Oldest full table scan time for the segment
max_ftstime	Latest full table scan time for the segment
avg_ftstime	Average full table scan time for the segment

90.3.4 SEGMENT_HEAT_MAP Procedure

This procedure returns the heatmap attributes for the given segment.

Syntax

```
DBMS_HEAT_MAP.SEGMENT_HEAT_MAP (
   tablespace_id          IN  NUMBER,
   header_file            IN  NUMBER,
   header_block           IN  NUMBER,
   segment_objd           IN  NUMBER,
   min_writetime          OUT DATE,
   max_writetime          OUT DATE,
   avg_writetime          OUT DATE,
   min_readtime           OUT DATE,
   max_readtime           OUT DATE,
   avg_readtime           OUT DATE,
   min_lookuptime         OUT DATE,
   max_lookuptime         OUT DATE,
   avg_lookuptime         OUT DATE,
   min_ftstime            OUT DATE,
   max_ftstime            OUT DATE,
   avg_ftstime            OUT DATE);
```

Parameters

Table 90-8 *SEGMENT_HEAT_MAP Procedure Parameters*

Parameter	Description
tablespace_id	Tablespace containing the segment
header_file	Segment header relative file number
header_block	Segment header block number
segment_objd	DATAOBJ of the segment

Return Values

Table 90-9 *SEGMENT_HEAT_MAP Procedure Return Values (Output Parameters)*

Parameter	Description
min_writetime	Oldest write time for the segment
max_writetime	Latest write time for the segment
avg_writetime	Average write time for the segment
min_readtime	Oldest read time for the segment
max_readtime	Latest read time for the segment
avg_writetime	Average write time for the segment
min_lookuptime	Oldest index lookup time for the segment
max_lookuptime	Latest index lookup time for the segment
avg_lookuptime	Average index lookup time for the segment
min_ftstime	Oldest full table scan time for the segment
max_ftstime	Latest full table scan time for the segment
avg_ftstime	Average full table scan time for the segment

90.3.5 TABLESPACE_HEAT_MAP Function

This table function returns the minimum, maximum and average access times for all the segments in the tablespace.

Syntax

```
DBMS_HEAT_MAP.TABLESPACE_HEAT_MAP (
    tablespace_name      IN VARCHAR2)
  RETURN hm_tablespace_table PIPELINED;
```

Parameters

Table 90-10 *TABLESPACE_HEAT_MAP Procedure Parameters*

Parameter	Description
tablespace_name	Name of the tablespace

Return Values

Table 90-11 *TABLESPACE_HEAT_MAP Procedure Return Values (Output Parameters)*

Parameter	Description
segment_count	Total number of segments in the tablespace
allocated_bytes	Space used by the segments in the tablespace
min_writetime	Oldest write time for the segment
max_writetime	Latest write time for the segment

Table 90-11 (Cont.) *TABLESPACE_HEAT_MAP Procedure Return Values (Output Parameters)*

Parameter	Description
avg_writetime	Average write time for the segment
min_readtime	Oldest read time for the segment
max_readtime	Latest read time for the segment
avg_writetime	Average write time for the segment
min_lookuptime	Oldest index lookup time for the segment
max_lookuptime	Latest index lookup time for the segment
avg_lookuptime	Average index lookup time for the segment
min_ftstime	Oldest full table scan time for the segment
max_ftstime	Latest full table scan time for the segment
avg_ftstime	Average full table scan time for the segment

91
DBMS_HIERARCHY

`DBMS_HIERARCHY` contains subprograms for validating the data in tables used by hierarchies and analytic views.

This chapter contains the following topics:

- DBMS_HIERARCHY Overview
- DBMS_HIERARCHY Security Model
- Summary of DBMS_HIERARCHY Subprograms

91.1 DBMS_HIERARCHY Overview

The `DBMS_HIERARCHY` package contains functions for validating that the contents of a database table are suitable for use by an analytic view or a hierarchy, a function for verifying the success of the validation, and a procedure for creating a table for logging validation operations.

> **Note:**
>
> Names specified by parameters of the `DBMS_HIERARCHY` subprograms are case-sensitive.

For information about using analytic views, see *Oracle Database Data Warehousing Guide*.

91.2 DBMS_HIERARCHY Security Model

Summarizes security considerations for the validation of analytic view and hierarchy objects.

All procedures in this package validate that the current user has the necessary privileges on the specified objects and return an error if those privileges are not found.

> **Note:**
>
> To ensure that the user has enough tablespace to log validation operations, do one of the following:
>
> - `GRANT UNLIMITED TABLESPACE TO` *username;*
> - `ALTER USERNAME` *username* `QUOTA` *size* `ON` *tablespace_name;*

The following system privileges are required to use this package:

To validate objects in the user's own schema:

- `CREATE TABLE` privilege for `CREATE_VALIDATE_LOG_TABLE` or to have `VALIDATE_ANALTYIC_VIEW` or `VALIDATE_HIERARCHY` automatically create a table
- `SELECT` privilege on the tables or views used by the analytic view or hierarchy
- `INSERT` privilege on the tables used by the attribute dimensions of the hierarchy or the fact table used by the analytic view

To validate objects in different schemas:

- `CREATE ANY TABLE` privilege for `CREATE_VALIDATE_LOG_TABLE` or to have the `VALIDATE_ANALTYIC_VIEW` or `VALIDATE_HIERARCHY` automatically create a table
- `INSERT ANY TABLE` privilege on the tables used by the attribute dimensions of the hierarchy or the fact table used by the analytic view

91.3 Summary of DBMS_HIERARCHY Subprograms

This table lists the `DBMS_HIERARCHY` subprograms and briefly describes them.

Subprogram	Description
CREATE_VALIDATE_LOG_TABLE Procedure	Creates a table that you can use for logging messages generated by the `VALIDATE_HIERARCHY` and `VALIDATE_ANALYTIC_VIEW` functions.
CREATE_VIEW_FOR_FACT_ROWS Procedure	This procedure creates a view for fact rows.
CREATE_VIEW_FOR_STAR_ROWS Procedure	This procedure creates a view for star rows.
GET_MV_SQL_FOR_AV_CACHE Function	This function either returns the SQL for a cache that is defined in the AV or the SQL for a given set of measures and levels. These measures and levels do not need to be part of a cache in the AV.
GET_MV_SQL_FOR_STAR_CACHE Function	This function returns the SQL in text form for the star cache.
UPGRADE_VALIDATE_LOG_TABLE Procedure	This procedure takes a log table and upgrades it to the newest format. Returns an error if the table cannot be upgraded due to other errors.
VALIDATE_ANALYTIC_VIEW Function	Validates that the data in a table is suitable for use by an analytic view.
VALIDATE_CHECK_SUCCESS Function	Indicates whether a prior call to `VALIDATE_HIERARCHY` or `VALIDATE_ANALYTIC_VIEW` was successful or produced validation errors.
VALIDATE_HIERARCHY Function	Validates that the data in a table is suitable for use by a hierarchy.

91.3.1 CREATE_VALIDATE_LOG_TABLE Procedure

This procedure creates a table that you can use for logging messages generated by the `VALIDATE_ANALYTIC_VIEW` or `VALIDATE_HIERARCHY` function, which validate data used by an analytic view or hierarchy.

The table that this procedure creates has the following structure.

```
NAME                NULL?    DATATYPE
------------------- -------- --------
LOG_NUMBER          NOT NULL NUMBER
ACTION_ORDER        NOT NULL NUMBER
OBJECT_OWNER        NOT NULL VARCHAR2(128 BYTE)
OBJECT_NAME         NOT NULL VARCHAR2(128 BYTE)
ACTION              NOT NULL VARCHAR2(10 BYTE)
TIME                NOT NULL TIMESTAMP(6)
ERROR_NUMBER                 NUMBER
ERROR_MESSAGE                VARCHAR2(4000)
```

Syntax

```
DBMS_HIERARCHY.CREATE_VALIDATE_LOG_TABLE (
    table_name        IN  VARCHAR2,
    owner_name        IN  VARCHAR2         DEFAULT NULL
    IGNORE_IF_EXISTS  IN  PL/SQL BOOLEAN   DEFAULT FALSE);
```

Parameters

Parameter	Description
`table_name`	The name of the table to create.
`owner_name`	The name of the schema in which to create the table. If `owner_name` is `NULL`, then the table is created in the current user's schema.
`IGNORE_IF_EXISTS`	A Boolean that indicates whether to create the table if a table by the same name exists. If you specify a table, it must have the same structure as the table that this procedure creates.

Examples

Example 91-1 Creating a Validation Log Table

```
BEGIN
  DBMS_HIERARCHY.CREATE_VALIDATE_LOG_TABLE (
    'VAL_AV_HIERARCHY_LOG',
    'AV_USER',
    FALSE
  );
END;
/
```

91.3.2 CREATE_VIEW_FOR_FACT_ROWS Procedure

This procedure creates a view over the analytic view that exposes only the rows of the fact table. This allows standard SQL group functions to be used to aggregate those rows. The advantage to querying such a view rather than the fact table directly is that the analytic view may be able to perform query optimizations otherwise unavailable. When `include_meas` is `TRUE`, the `AV_AGGREGATE` group function may be used to expose measure columns defined in the analytic view.

Syntax

```
DBMS_HIERARCHY.CREATE_VIEW_FOR_FACT_ROWS (
      analytic_view_name         IN VARCHAR2,
      view__name                 IN VARCHAR2,
      dim_hier_seq               IN ID2_SEQUENCE  DEFAULT NULL,
      analytic_view_owner_name   IN VARCHAR2      DEFAULT
SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA'),
      view_owner_name            IN VARCHAR2      DEFAULT
SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA'),
      dim_qual_sep               IN VARCHAR2      DEFAULT '_',
      all_join_keys              IN BOOLEAN       DEFAULT TRUE,
      include_meas               IN BOOLEAN       DEFAULT FALSE,
      include_hier_attr          IN BOOLEAN       DEFAULT FALSE);
```

Parameters

Parameter	Description
`analytic_view_name`	The name of the analytic view to be created.
`view__name`	The name of the view.
`dim_hier_seq`	A sequence of `ID2` records. An ID2 record contains two `VARCHAR2(128)` for the two parallel components, providing (`dim, hier`) name pairs. Default value is `NULL` which indicates that all hierarchies should be considered.
`analytic_view_owner_name`	The name of the analytic view owner. The default value is `SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA')`.
`view_owner_name`	The owner of the view to be created. The default value is `SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA')`.
`dim_qual_sep`	The separator to be used on the columns of the created view. For example, • for `fact` columns - `"FACT" || dim_qual_sep || {name of fact column alias}` • for `measures` - `"MEASURE" || dim_qual_sep || {measure name alias}` The default value is `_`.
`all_join_keys`	If `TRUE`, adds all of the join key columns in the AV to the view.

Parameter	Description
include_meas	If TRUE, include all of the base measures and calculated measures defined in the AV as columns in the created view.
include_hier_attr	If TRUE, include all of the hierarchical attributes defined in the AV as columns in the created view.

91.3.3 CREATE_VIEW_FOR_STAR_ROWS Procedure

This procedure creates a view over the analytic view that exposes only rows for a particular dimension star table. Such a view is intended to be used in conjunction with a corresponding fact rows view created using CREATE_VIEW_FOR_FACT_ROWS, joined using their corresponding keys as defined in the analytic view.

Syntax

```
DBMS_HIERARCHY.CREATE_VIEW_FOR_STAR_ROWS (
      analytic_view_name         IN VARCHAR2,
      dimension_alias            IN VARCHAR2,
      view__name                 IN VARCHAR2,
      analytic_view_owner_name   IN VARCHAR2      DEFAULT
SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA'),
      view_owner_name            IN VARCHAR2      DEFAULT
SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA'),
      hier_qual_sep              IN VARCHAR2      DEFAULT '_',
      include_hier_attr          IN BOOLEAN       DEFAULT FALSE);
```

Parameters

Parameter	Description
analytic_view_name	The name of the analytic view to be created.
dimension_alias	The alias of the attribute dimension that you wish to create a star rows view for.
view__name	The name of the view.
analytic_view_owner_name	The name of the analytic view owner. The default value is SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA').
view_owner_name	The owner of the view to be created. The default value is SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA').
hier_qual_sep	The separator to be used. The separator is between the attribute dimension name and the built-in hierarchical attributes, such as MEMBER_NAME, LEVEL_NAME, DEPTH, and so on. The default value is _.
include_hier_attr	If TRUE, include all of the hierarchical attributes defined in the AV as columns in the created view.

91.3.4 GET_MV_SQL_FOR_AV_CACHE Function

This function has two signatures. The first version of this method returns the SQL for a cache that is defined in the AV. The second version of this method returns SQL for a given set of measures and levels. These measures and levels do not need to be part of a cache in the AV.

Syntax

```
DBMS_HIERARCHY.GET_MV_SQL_FOR_AV_CACHE (
    analytic_view_name          IN    VARCHAR2,
    cache_idx                   IN    NUMBER,
    analytic_view_owner_name    IN    VARCHAR2  DEFAULT SYS_CONTEXT('USERENV',
'CURRENT_SCHEMA'))
  RETURN CLOB;



DBMS_HIERARCHY.GET_MV_SQL_FOR_AV_CACHE (
    analytic_view_name          IN    VARCHAR2,
    lvl_seq                     IN    ID3_SEQUENCE  DEFAULT NULL,
    meas_seq                    IN    ID_SEQUENCE   DEFAULT NULL,
    analytic_view_owner_name    IN    VARCHAR2  DEFAULT SYS_CONTEXT('USERENV',
'CURRENT_SCHEMA'))
  RETURN CLOB;
```

Parameters

Table 91-1 GET_MV_SQL_FOR_AV_CACHE Function Parameters

Parameter	Description
analytic_view_name	The name of the analytic view to be created.
cache_idx	Zero-based index for the caches defined on the AV.
analytic_view_owner_name	The name of the analytic view owner. The default value is SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA').
lvl_seq	The list of levels that the cache should be based on.
meas_seq	The list of measures that the cache should be based on.

91.3.5 GET_MV_SQL_FOR_STAR_CACHE Function

This function returns the SQL in text form for the star cache.

Syntax

```
DBMS_HIERARCHY.GET_MV_SQL_FOR_STAR_CACHE (
    attr_dim_name           IN  VARCHAR2,
    attr_dim_owner_name     IN  VARCHAR2 DEFAULT SYS_CONTEXT('USERENV',
'CURRENT_SCHEMA'))
  RETURN CLOB;
```

Parameters

Table 91-2 GET_MV_SQL_FOR_STAR_CACHE Function Parameters

Parameter	Description
attr_dim_name	The name of the attribute dimension on which to create a cache.
attr_dim_owner_name	The schema that owns the attribute dimension. The default value is SYS_CONTEXT('USERENV', 'CURRENT_SCHEMA').

91.3.6 UPGRADE_VALIDATE_LOG_TABLE Procedure

This procedure takes a log table and upgrades it to the newest format. Returns an error if the table cannot be upgraded due to other errors.

Syntax

```
DBMS_HIERARCHY.UPGRADE_VALIDATE_LOG_TABLE (
    table_name        IN  VARCHAR2,
    owner_name        IN  VARCHAR2         DEFAULT NULL);
```

Parameters

Parameter	Description
table_name	The name of the table to upgrade.
owner_name	The name of the schema in which to find the table. If owner_name is NULL, then the table is created in the current user's schema.

91.3.7 VALIDATE_ANALYTIC_VIEW Function

This function validates that the data in a table or view conforms to the logical constraints inherent in the definition of an analytic view.

Syntax

```
DBMS_HIERARCHY.VALIDATE_ANALYTIC_VIEW (
      analytic_view_name          IN VARCHAR2   DEFAULT NULL,
      analytic_view_owner_name    IN VARCHAR2   DEFAULT NULL,
      log_table_name              IN VARCHAR2   DEFAULT NULL,
      log_table_owner_name        IN VARCHAR2   DEFAULT NULL)
      RETURN NUMBER;
```

Parameters

Parameter	Description
analytic_view_name	The name of the analytic view to validate.
analytic_view_owner_name	The name of the owner of the schema that contains the analytic view.
log_table_name	The name of the validation log table in which to put the results of the validation operation.

Parameter	Description
log_table_owner_name	The name of the owner of the schema in which the validation log table exists or in which to create the table.

Returns

The number of the entry in the validation log table for the validation results.

Usage Notes

If the log_table_name parameter is NULL, then the VALIDATE_ANALYTIC_VIEW function creates a validation log table. The name of the table it creates is DBMS_HIERARCHY_LOG.

When the validation operation begins, a row is inserted into the log table with the action of START. When the operation completes, a row is inserted into the log table with the action of END. When an error is detected, a row is inserted into the log table with the action of ERROR, and the associated error_number and error_message columns are populated. All rows inserted into the validation log table include a log number and the time of the insert.

The VALIDATE_ANALYTIC_VIEW function verifies that the following conditions are true for each attribute dimension the analytic view is dimensioned by:

- The key values found in the fact table for the attribute dimension must exist in the star schema dimension table for that attribute dimension.
- The referenced attribute values for the attribute dimension must be unique across all rows of the star schema dimension table for that dimension.

Also, for every hierarchy in the analytic view, the function verifies that the following conditions are true:

- The primary key of a level determines a unique value for each attribute of the level.
- For each row of the table or view used by the attribute dimension of the hierarchy, the value for every level key column (including alternate keys) of a NOT NULL level is non-NULL.
- For each row of the table or view, either all level key columns and alternate key columns of a SKIP WHEN NULL level must be NULL or they must all be non-NULL. This verifies that the alternate level key is determined by the level key.
- For each group of rows that have the same alternate key column values for a level, the key column values must have the same column values. This verifies that the level key is determined by the alternate level key, which is required for an alternate key.

Examples

Example 91-2 Validating an Analytic View

```
DECLARE
  log_num NUMBER;
  obj_name VARCHAR2(8) := 'SALES_AV';
BEGIN
   log_num := DBMS_HIERARCHY.VALIDATE_ANALYTIC_VIEW(obj_name);
END;
/
```

91.3.8 VALIDATE_CHECK_SUCCESS Function

This function indicates whether a prior call to `VALIDATE_HIERARCHY` or `VALIDATE_ANALYTIC_VIEW` was successful or produced validation errors.

Syntax

```
DBMS_HIERARCHY.VALIDATE_CHECK_SUCCESS (
    TOPOBJ_NAME           IN  VARCHAR2,
    TOPOBJ_OWNER          IN  VARCHAR2,
    LOG_NUMBER            IN  VARCHAR2
    LOG_TABLE_NAME        IN  VARCHAR2
    LOG_TABLE_OWNER_NAME  IN  VARCHAR2 )
    RETURN VARCHAR2;
```

Parameters

Parameter	Description
`TOPOBJ_NAME`	The name of the hierarchy or analytic view.
`TOPOBJ_OWNDER`	The owner of the hierarchy or analytic view.
`LOG_NUMBER`	The number of the log entry.
`LOG_TABLE_NAME`	The name of the log table.
`LOG_TABLE_OWNER_NAME`	The name of the schema in which the table exits.

Returns

A `VARCHAR2` that is SUCCESS if no errors occurred or ERROR if errors did occur.

Examples

Example 91-3 Using VALIDATE_CHECK_SUCCESS

This example finds out whether the prior call to `VALIDATE_ANALTYIC_VIEW` encountered errors.

```
DECLARE
  log_num NUMBER;
  succ VARCHAR2(7);
  obj_name VARCHAR2(8) := 'SALES_AV';
BEGIN
   log_num := dbms_hierarchy.validate_analytic_view(obj_name);
   succ := dbms_hierarchy.validate_check_success(
      topobj_name => obj_name, log_number => log_num);
  IF (succ != 'SUCCESS') THEN
     RAISE_APPLICATION_ERROR(
       num => -20000,
       msg => 'Validate failed!');
  END IF;
END;
/
```

91.3.9 VALIDATE_HIERARCHY Function

This function validates that the data in a table or view conforms to the logical constraints inherent in the definitions of an attribute dimension that uses the table or view and a hierarchy that uses the attribute dimension.

Syntax

```
DBMS_HIERARCHY.VALIDATE_HIERARCHY (
      hier_name                IN VARCHAR2,
      hier_owner_name          IN VARCHAR2    DEFAULT NULL,
      log_table_name           IN VARCHAR2    DEFAULT NULL,
      log_table_owner_name     IN VARCHAR2    DEFAULT NULL)
      RETURN NUMBER;
```

Parameters

Parameter	Description
hier_name	The name of the hierarchy to validate.
hier_owner_name	The name of the owner of the schema that contains the hierarchy.
log_table_name	The name of the validation log table in which to put the results of the validation operation.
log_table_owner_name	The name of the owner of the schema in which the validation log table exists or in which to create the table.

Returns

The number of the entry in the validation log table for the validation results.

Usage Notes

If the log_table_name parameter is NULL, then the VALIDATE_HIERARCHY function creates a validation log table. The name of the table it creates is DBMS_HIERARCHY_LOG.

When the validation operation begins, a row is inserted into the log table with the action of START. When the operation completes, a row is inserted into the log table with the action of END. When an error is detected, a row is inserted into the log table with the action of ERROR, and the associated error_number and error_message columns are populated. All rows inserted into the validation log table include a log number and the time of the insert.

The VALIDATE_HIERARCHY function verifies that the following conditions are true for the hierarchy:

- The primary key of a level determines a unique value for each attribute of the level.
- For each row of the table or view used by the attribute dimension of the hierarchy, the value for every level key column (including alternate keys) of a NOT NULL level is non-NULL.
- For each row of the table or view, either all level key columns and alternate key columns of a SKIP WHEN NULL level must be NULL or they must all be non-NULL. This verifies that the alternate level key is determined by the level key.
- For each group of rows that have the same alternate key column values for a level, the key column values must have the same column values. This verifies that the level key is determined by the alternate level key, which is required for an alternate key.

Examples

Example 91-4 Validating a Hierarchy and Specifying a Table Name

This example validates the `PRODUCT_HIER` hierarchy and specifies that the results be inserted in the table named `VAL_AV_HIERARCHY_LOG`. The owner of the hierarchy and of the schema that contains the table is AV_USER.

```
-- Create a log table.
BEGIN
  DBMS_HIERARCHY.CREATE_VALIDATE_LOG_TABLE (
    'VAL_AV_HIERARCHY_LOG',
    'AV_USER',
    FALSE
  );
END;
/
-- Validate the hierarchy.
DECLARE
  log_num NUMBER;
  obj_name VARCHAR2(12) := 'PRODUCT_HIER';
  table_name  VARCHAR2(28) := 'VAL_AV_HIERARCHY_LOG';
BEGIN
   log_num := DBMS_HIERARCHY.VALIDATE_HIERARCHY(obj_name, 'AV_USER',
table_name);
END;
/
```

Query the log table.

```
SELECT LOG_NUMBER, ACTION, OBJECT_NAME, ERROR_NUMBER, ERROR_MESSAGE
  FROM AV_USER.VAL_AV_HIERARCHY_LOG;
  WHERE OBJECT_NAME = 'PRODUCT_HIER';

LOG_NUMBER ACTION     OBJECT_NAME          ERROR_NUMBER ERROR_MESSAGE
---------- ---------- -------------------- ------------ -------------
         1 START      PRODUCT_HIER
         1 END        PRODUCT_HIER
```

Example 91-5 Validating a Hierarchy Without Specifying a Table Name

This example shows that if you do not specify a validation log table, then the `VALIDATE_HIERARCHY` function creates one named DBMS_HIERARCHY_LOG.

```
DECLARE
  log_num NUMBER;
  obj_name VARCHAR2(12) := 'PRODUCT_HIER';
BEGIN
   log_num := DBMS_HIERARCHY.VALIDATE_HIERARCHY(obj_name);
END;
```

Query the log table.

```
SELECT LOG_NUMBER, ACTION, OBJECT_NAME, ERROR_NUMBER, ERROR_MESSAGE
  FROM DBMS_HIERARCHY_LOG
  WHERE OBJECT_NAME = 'PRODUCT_HIER';

LOG_NUMBER ACTION     OBJECT_NAME         ERROR_NUMBER ERROR_MESSAGE
---------- ---------- ------------------- ------------ -------------
         1 START      PRODUCT_HIER
         1 END        PRODUCT_HIER
```

92
DBMS_HM

This package contains constants and procedure declarations for health check management. Health Monitor provides facilities to run a check store and retrieve the reports through `DBMS_HM` package

This chapter contains the following topics:

- Security Model
- Summary of DBMS_HM Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more information about "Health Monitor"

92.1 DBMS_HM Security Model

Users must have `EXECUTE` privilege to run the procedures of `DBMS_HM` package.

92.2 Summary of DBMS_HM Subprograms

This table lists the `DBMS_CONNECTION_POOL` subprograms and briefly describes them.

Table 92-1 DBMS_HM Package Subprograms

Subprogram	Description
GET_RUN_REPORT Function	Returns the report for the specified checker run
RUN_CHECK Procedure	Runs the specified checker with the given arguments

92.2.1 GET_RUN_REPORT Function

This function returns the report for the specified checker run.

Syntax

```
DBMS_HM.GET_RUN_REPORT (
    run_name      IN  VARCHAR2,
    type           IN  VARCHAR2 := 'TEXT',
    level          IN  VARCHAR2 := 'BASIC',)
  RETURN CLOB;
```

Parameters

Table 92-2 GET_RUN_REPORT Function Parameters

Parameter	Description
run_name	Name of the check's run
type	Report format type. Possible values are 'HTML', 'XML' and 'TEXT'. Default report type is 'TEXT'.
level	Details of report, possible value are 'BASIC' and 'DETAIL'. Caution: Currently only 'BASIC' level is supported.

92.2.2 RUN_CHECK Procedure

This procedure runs the specified checker with the specified arguments.

You can specify a name for the run, the inputs needed and maximum time-out for the run. The run report will be maintained persistently in the database.

Syntax

```
DBMS_HM.RUN_CHECK (
   check_name     IN  VARCHAR2,
   run_name       IN  VARCHAR2 := NULL,
   timeout        IN  NUMBER := NULL,
   input_params   IN  VARCHAR2 := NULL);
```

Parameters

Table 92-3 RUN_CHECK Procedure Parameters

Parameter	Description
check_name	Name of the check to be invoked. Check names and their parameters can be accessed from the V$HM_CHECK and V$HM_CHECK_PARAM views. Users can run all checks which are not internal in nature: SELECT name FROM V$HM_CHECK WHERE INTERNAL_CHECK = 'N' retrieves the list of checks that can be run manually by users.
run_name	Name with which external users can uniquely identify this check's run. If NULL value is passed, then HM creates a unique name and associates with this check's run.
timeout	Maximum amount of time (in units of seconds), this checker run is allowed to run. HM will interrupt the run, if it the specified time elapses for the run. If NULL value is passed, HM doesn't impose any timeout limits on the run.

Table 92-3 (Cont.) RUN_CHECK Procedure Parameters

Parameter	Description

input_params	Input string: which consists of name, value pairs de-limited by a special character ';'.
	Example ('Data Block Integrity Check' invocation may take following type of input parameters.
	`'BLC_DF_NUM=1;BLC_BL_NUM=23456'`
	Input parameters `BLC_DF_NUM` and `BLC_BL_NUM` have values '1' and '23456' respectively.
	Every check will have well defined set of inputs associated with it. These Input parameters, their types, default values and descriptions can be obtained using `V$HM_CHECK_PARAM` view.
	Example: The following query gets the list of parameters, their default values and descriptions for a 'Data Block Integrity Check'
	`SELECT a.* FROM v$hm_check_param a, v$hm_check b`
	`WHERE a.check_id = b.id`
	`AND b.name = 'Data Block Integrity Check';`

93
DBMS_HPROF

The `DBMS_HPROF` package provides an interface for profiling the execution of PL/SQL applications. It provides services for collecting the hierarchical profiler data, analyzing the raw profiler output and profiling information generation.

This chapter contains the following topic:

- Summary of DBMS_HPROF Subprograms

> **See Also:**
>
> *Oracle Database Development Guide* for more information about the "PL/SQL Hierarchical Profiler"

93.1 DBMS_HPROF Security Model

You must have the following privileges to use the `DBMS_HPROF` package:

- An `EXECUTE` privilege on the `DBMS_HPROF` package.
- A `WRITE` privilege on the directory that you specify when you call the `DBMS_HPROF.START_PROFILING` procedure.

93.2 Summary of DBMS_HPROF Subprograms

This table lists and briefly describes the `DBMS_HPROF` package subprograms.

Table 93-1 DBMS_HPROF Package Subprograms

Subprogram	Description
ANALYZE Function	Analyzes the raw profiler output and produces hierarchical profiler information in database tables.
CREATE_TABLES Procedure	Creates the hierarchical profiler database tables and data structures in the user's session.
START_PROFILING Procedure	Starts hierarchical profiler data collection in the user's session.
STOP_PROFILING Procedure	Stops profiler data collection in the user's session.

93.2.1 ANALYZE Function

This function analyzes the raw profiler output and produces hierarchical profiler information in database tables or generates out-of-the-box HTML reports.

Syntax

```
DBMS_HPROF.ANALYZE (
   trace_id          IN NUMBER,
   summary_mode      IN BOOLEAN DEFAULT FALSE,
   trace             IN VARCHAR2 DEFAULT NULL,
   skip              IN PLS_INTEGER DEFAULT 0,
   collect           IN PLS_INTEGER DEFAULT NULL,
   run_comment       IN VARCHAR2 DEFAULT NULL)
  RETURN NUMBER;

DBMS_HPROF.ANALYZE (
   trace_id          IN NUMBER,
   report_clob       OUT CLOB,
   trace             IN VARCHAR2 DEFAULT NULL,
   skip              IN PLS_INTEGER DEFAULT 0,
   collect           IN PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 93-2 ANALYZE Function Parameters

Parameter	Description
`trace_id`	The `trace_id` of the raw profiler data entry in the raw profiler data table (`dbmshp_trace_data`).
`summary_mode`	By default (that is, when `summary_mode` is `FALSE`), a detailed analysis is done. When `summary_mode` is `TRUE`, only top-level summary information is generated into the database table.
`report_clob`	The analyzed HTML report.
`trace`	Analyzes only the subtrees rooted at the specified trace entry. By default (when trace is `NULL`), the analysis/reporting is generated for the entire run. The trace entry must be specified in a special quoted qualified format. For example, `'"HR"."PKG"."FOO"'` or `'"."."__plsql_vm"'`. If multiple overloads exist for the specified name, all of them will be analyzed.
`skip`	Used only when `trace` is specified. Analyze only the subtrees rooted at the specified trace, but ignore the first `skip` invocations to trace. The default value for skip is `0`.
`collect`	Used only when `trace` is specified. Analyze `collect` number of invocations of traces (starting from `skip`+1'th invocation). By default, only 1 invocation is collected.
`run_comment`	User-provided comment for this run.

Return Values

A unique run identifier for this run of the analyzer. This can then be used to look up the results corresponding to this run from the hierarchical profiler tables.

Usage Notes

- Use the `DBMS_HPROF.CREATE_TABLES` subprogram to create the hierarchical profiler database tables and other data structures required for persistently storing the results of analyzing the raw profiler data.
- Calling the `DBMS_HPROF.CREATE_TABLES` with default value (`FALSE`) will raise error if table already exists.
- Use `DBMS_HPROF.CREATE_TABLES(TRUE)` to drop any previously created hierarchical profiler tables.
- Use the `DBMS_HPROF.CREATE_TABLES` to drop any previously created hierarchical profiler tables. By default, `force_it` is `FALSE`; therefore, to drop any previously created hierarchical profiler tables you must set the value of `force_it` to `TRUE`.
- If `trace_id` entry is `NULL`, error is raised.
- If `trace_id` entry in the raw profiler data table does not exist, error is raised.
- If raw data of the `trace_id` entry in the raw profiler data table is `NULL` or is `zero` size, error is raised.

Examples

The following snippet installs the hierarchical profiler tables in HR schema.

```
connect HR/<password>;
```

The following example analyzes and generates HTML CLOB report from a raw profiler data table.

```
DECLARE
        reportclob clob;
        trace_id number;
BEGIN
    -- create raw profiler data and analysis tables
    -- force_it =>TRUE will dropped the tables if table exists
    DBMS_HPROF.CREATE_TABLES(force_it =>TRUE);

    -- Start profiling
    -- Write raw profiler data in raw profiler data table
    trace_id := DBMS_HPROF.START_PROFILING;

    -- Run procedure to be profiled
    test;

    -- Stop profiling
    DBMS_HPROF.STOP_PROFILING;

    -- analyzes trace_id entry in raw profiler data table and produce
    -- analyzed HTML report in reportclob
    DBMS_HPROF.ANALYZE(trace_id , reportclob);
END;
/
```

93.2.2 CREATE_TABLES Procedure

Creates the hierarchical profiler database tables and data structures in the user's session.

Syntax

```
DBMS_HPROF.CREATE_TABLES (
   force_it         IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 93-3 CREATE_TABLES Procedure Parameters

Parameter	Description
`force_it`	If `FALSE` and `DBMS_HPROF` tables are present, then an `HPROF` error is raised. If `TRUE`, then the procedure creates tables. If the tables already exist, then they are dropped and new tables are created.

> **Note:**
>
> Users need not use the `dbmshptab.sql` script located in the `rdbms/admin` directory to create the hierarchical profiler database tables and data structures anymore.
>
> The `dbmshptab.sql` script is deprecated starting in Oracle Database 18*c*.

93.2.3 START_PROFILING Procedure

This procedure starts hierarchical profiler data collection in the user's session.

Syntax

```
DBMS_HPROF.START_PROFILING(
   max_depth     IN PLS_INTEGER DEFAULT NULL,
   sqlmonitor    IN BOOLEAN DEFAULT TRUE,
   run_comment   IN VARCHAR2 DEFAULT NULL)
RETURN NUMBER;
```

Parameters

Table 93-4 START_PROFILING Procedure Parameters

Parameter	Description
`max_depth`	By default (that is, when `max_depth` value is `NULL`) profile information is gathered for all functions irrespective of their call depth. When a non-`NULL` value is specified for `max_depth`, the profiler collects data only for functions up to a call depth level of `max_depth`.

Table 93-4 (Cont.) START_PROFILING Procedure Parameters

Parameter	Description
sqlmonitor	Generates a real-time monitoring report for a profiler run when the profiler run ends. The default value is TRUE.
run_comment	User provided comment for the profiler data collection run.

Return Values

Unique run identifier for this profiler run. This can then be used to look up the results corresponding to this run from the hierarchical profiler tables.

Usage Notes

- Even though the profiler does not individually track functions at depth greater than max_depth, the time spent in such functions is charged to the ancestor function at depth max_depth.
- Raw profiler data is generated in the raw profiler data table with an unique trace_id.
- The unique trace_id is used to manage the raw profiler output stored in the raw profiler data table.

93.2.4 STOP_PROFILING Procedure

This procedure stops profiler data collection in the user's session. This subprogram also has the side effect of flushing data collected so far in the session, and it signals the end of a run. When the STOP_PROFILING procedure returns CLOB, it contains the Real-Time Monitoring report for the profiler run.

Syntax

```
DBMS_HPROF.STOP_PROFILING;


DBMS_HPROF.STOP_PROFILING
  RETURN CLOB;
```

Examples

Profiling with raw profiler data table

```
DECLARE
    analyze_runid number;
    trace_id number;
BEGIN
    -- create raw profiler data and analysis tables
    -- call create_tables with force_it =>FALSE (default) when
    -- raw profiler data and analysis tables do not exist already
    DBMS_HPROF.CREATE_TABLES;
    -- Start profiling
    -- Write raw profiler data in raw profiler data table
    trace_id := DBMS_HPROF.START_PROFILING;
    -- Run the procedure to be profiled
    test;
    -- Stop profiling
```

```
    DBMS_HPROF.STOP_PROFILING;
    -- analyzes trace_id entry in raw profiler data table and writes
    -- hierarchical profiler information in hprof's analysis tables
    analyze_runid := DBMS_HPROF.ANALYZE(trace_id);
END;
/
```

94
DBMS_HS_PARALLEL

The DBMS_HS_PARALLEL PL/SQL package enables parallel processing for heterogeneous targets access. This package is designed to improve performance when retrieving data from a large foreign table.

This chapter discusses the following topics:

- Using DBMS_HS_PARALLEL
- Summary of DBMS_HS_PARALLEL Subprograms

94.1 Using the DBMS_HS_PARALLEL Subprogram

DBMS_HS_PARALLEL is compiled with the authorization ID of CURRENT_USER, which uses invoker's rights. In other words, all procedures in this package are executed with the privileges of the calling user.

94.2 Summary of DBMS_HS_PARALLEL Subprograms

This table lists the DBMS_HS_PARALLEL subprograms in alphabetic order and briefly describes them.

Table 94-1 DBMS_HS_PARALLEL Package Subprograms

Subprogram	Description
CREATE_OR_REPLACE_VIEW	Creates (or replaces) a read-only view to be referenced for retrieving the data from a remote table in parallel.
CREATE_TABLE_TEMPLATE	Writes out a CREATE TABLE template based on information gathered from the remote table. You can use the information to add any optimal Oracle CREATE TABLE clauses.
DROP_VIEW	Drops the view and internal objects created by the CREATE_OR_REPLACE_VIEW procedure. If the view has not already been created by the CREATE_OR_REPLACE_VIEW procedure, an error message is returned.
LOAD_TABLE	Loads the data from a remote table to a local Oracle table in parallel. If the local Oracle table does not already exist, it is created automatically.

94.2.1 CREATE_OR_REPLACE_VIEW Procedure

This procedure creates (or replaces) a read-only view to be referenced for retrieving the data from a remote table in parallel.

Syntax

```
CREATE_OR_REPLACE_VIEW (remote_table, database_link, oracle_view, parallel_degree)
```

Parameters

Table 94-2 CREATE_OR_REPLACE_VIEW Procedure Parameters

Parameters	Description
remote_table	The name of the remote database table. It is specified as [remote_schema_name.]remote_table_name.
database_link	The remote database link name. The call can only be applied to a heterogeneous services database link.
oracle_view	The name of the Oracle view. It is specified as [schema_name.]oracle_view_name. The default schema name is the current user. If the oracle_view parameter is not specified, the remote table name will be used as the view name.
parallel_degree	The number of parallel processes for the operation is computed based on the range-partition number if applicable, or the number of CPUs. The range of values is 2 to 16.

Usage Notes

- The specified Oracle view is created and future reference of this view utilizes internal database objects for parallel retrieval of remote non-Oracle table data. If the Oracle view already exists, the following Oracle error message is raised:

  ```
  ORA-00955: name is already used by an existing object
  ```

- This view is created as a read-only view. If you attempt to insert and update the view, the following Oracle error message is raised:

  ```
  ORA-01733: virtual column not allowed here
  ```

- If the remote table or the database link does not exist, one of the following Oracle error messages is raised:

  ```
  ORA-00942: table or view does not exist
  or
  ORA-02019: connection description for remote database not found
  ```

- You need the CREATE VIEW, CREATE TABLE, CREATE TYPE, CREATE PACKAGE, and CREATE FUNCTION privileges to execute the CREATE_OR_REPLACE_VIEW procedure.
- If you encounter either of the following Oracle error messages, increase the PROCESSES and SESSIONS parameter in the Oracle initialization parameter file:

  ```
  ORA-12801: error signaled in parallel query server P003
  or
  ORA-00018: maximum number of session exceeded
  ```

- Because the CREATE_OR_REPLACE_VIEW procedure creates some internal objects, use the DROP_VIEW procedure to drop the view and the internal objects. The SQL DROP VIEW statement only drops the view and not the internal objects.

94.2.2 CREATE_TABLE_TEMPLATE Procedure

This procedure writes out a `CREATE TABLE` template based on information gathered from the remote table. You can use the information to add any optimal Oracle `CREATE TABLE` clauses.

Syntax

```
CREATE_TABLE_TEMPLATE (remote_table, database_link, oracle_table,
create_table_template_string)
```

Parameters

Table 94-3 CREATE_TABLE_TEMPLATE Procedure Parameters

Parameter	Description
`remote_table`	The name of the remote database table. It is specified as `[remote_schema_name.]remote_table_name`.
`database_link`	The remote database link name. The call can only be applied to a heterogeneous services database link.
`oracle_table`	The name of the local Oracle table the data will be loaded into. It is specified as `[schema_name.]oracle_table_name`. The default schema name is the current user. If the `oracle_table` parameter is not specified, the remote table name will be used as the local Oracle name.
`create_table_template_string`	Contains the Oracle `CREATE TABLE` SQL template when the procedure is returned.

94.2.3 DROP_VIEW Procedure

This procedure drops the view and internal objects created by the `CREATE_OR_REPLACE_VIEW` procedure. If the view has not already been created by the `CREATE_OR_REPLACE_VIEW` procedure, an error message is returned.

Syntax

```
DROP_VIEW (oracle_view)
```

Parameters

Table 94-4 DROP_VIEW Procedure Parameter

Parameter	Description
`oracle_view`	The name of the Oracle view created by the `CREATE_OR_REPLACE_VIEW` procedure. If the view has not been created by the `CREATE_OR_REPLACE_VIEW` procedure, an error is returned.

94.2.4 LOAD_TABLE Procedure

This procedure loads the data from a remote table to a local Oracle table in parallel. If the local Oracle table does not already exist, it is created automatically.

Syntax

```
LOAD_TABLE (remote_table, database_link, oracle_table, truncate, parallel_degree,
row_count)
```

Parameters

Table 94-5 LOAD_TABLE Procedure Parameters

Parameter	Description
`remote_table`	The name of the remote database table. It is specified as `[remote_schema_name.]remote_table_name`
`database_link`	The remote database link name. The call can only be applied to a heterogeneous services database link.
`oracle_table`	The name of the local Oracle table the data will be loaded into. It is specified as `[schema_name.]oracle_table_name`. The default schema name is the current user. If the `oracle_table` parameter is not specified, the remote table name will be used as the local Oracle name.
`truncate`	Determines whether the Oracle table is truncated before the data is loaded. The value is either `TRUE` or `FALSE`. The default value is `TRUE` which means the Oracle table is truncated first. When set to `FALSE`, the Oracle table will not be truncated before the data is loaded.
`parallel_degree`	The number of parallel processes for the operation is computed based on the range-partition number if applicable, or the number of CPUs. The range of values is `2` to `16`.
`row_count`	Contains the number of rows just added with the load table operation.

Usage Notes

- This procedure only loads the remote table data into Oracle local table. It does not create a key, index, constraints or any other dependencies such as triggers. It is recommended that you create these dependencies after the table data is loaded as performance will improve greatly. You will need to decide whether to create the dependencies before or after the data is loaded based on your knowledge of the remote table data and dependencies.
- If the local table does not exist, the `LOAD_TABLE` procedure creates a simple (non-partitioned) local table based on the exact column matching of the remote table after which the data is inserted into the local table.
- If the remote table or the database link does not exist, an error message is returned.
- If the local table is incompatible with the remote table, an error message is returned.
- You need the `CREATE TABLE`, `CREATE TYPE`, `CREATE PACKAGE`, and `CREATE FUNCTION` privileges to execute the `LOAD_TABLE` procedure.

- If you encounter either of the following Oracle error messages, increase the `PROCESSES` and `SESSIONS` parameter in Oracle initialization parameter file:

  ```
  ORA-12801: error signaled in parallel query server P003
  or
  ORA-00018: maximum number of session exceeded
  ```

- One of the following is required for parallel processing:
  - The remote table is range partitioned.
  - Histogram information for a numeric column is available.
  - There is a numeric index or primary key.
- To drop the local table, use the `DROP TABLE` SQL statement.

95
DBMS_HS_PASSTHROUGH

The DBMS_HS_PASSTHROUGH PL/SQL package allows you to send a statement directly to a non-Oracle system without being interpreted by the Oracle server. This can be useful if the non-Oracle system allows operations in statements for which there is no equivalent in Oracle.

This chapter discusses the following topics:

- DBMS_HS_PASSTHROUGH Overview
- DBMS_HS_PASSTHROUGH Operational Notes
- Summary of DBMS_HS_PASSTHROUGH Subprograms

> **See Also:**
>
> *Oracle Database Heterogeneous Connectivity User's Guide* for more information about this package

95.1 DBMS_HS_PASSTHROUGH Overview

You can execute passthrough SQL statements directly at the non-Oracle system using the PL/SQL package DBMS_HS_PASSTHROUGH. Any statement executed with this package is executed in the same transaction as standard SQL statements.

> **See Also:**
>
> *Oracle Database Heterogeneous Connectivity User's Guide* for information about this package

95.2 DBMS_HS_PASSTHROUGH Operational Notes

The DBMS_HS_PASSTHROUGH package is a virtual package. It conceptually resides at the non-Oracle system. In reality, however, calls to this package are intercepted by Heterogeneous Services and mapped to one or more Heterogeneous Services calls. The driver, in turn, maps these Heterogeneous Services calls to the API of the non-Oracle system. The client application should invoke the procedures in the package through a database link in exactly the same way as it would invoke a non-Oracle system stored procedure. The special processing done by Heterogeneous Services is transparent to the user.

95.3 Summary of DBMS_HS_PASSTHROUGH Subprograms

This table lists the DBMS_HS_PASSTHROUGH subprograms and briefly describes them.

Table 95-1 ***DBMS_HS_PASSTHROUGH Package Subprograms***

Subprogram	Description
BIND_INOUT_VARIABLE Procedure	Binds `IN OUT` bind variables
BIND_INOUT_VARIABLE_RAW Procedure	Binds `IN OUT` bind variables of datatype `RAW`
BIND_OUT_VARIABLE Procedure	Binds an `OUT` variable with a PL/SQL program variable
BIND_OUT_VARIABLE_RAW Procedure	Binds an `OUT` variable of datatype `RAW` with a PL/SQL program variable
BIND_VARIABLE Procedure	Binds an `IN` variable positionally with a PL/SQL program variable
BIND_VARIABLE_RAW Procedure	Binds `IN` variables of type `RAW`
CLOSE_CURSOR Procedure	Closes the cursor and releases associated memory after the SQL statement has been run at the non-Oracle system
EXECUTE_IMMEDIATE Procedure	Runs a (non-`SELECT`) SQL statement immediately, without bind variables
EXECUTE_NON_QUERY Function	Runs a (non-`SELECT`) SQL statement
FETCH_ROW Function	Fetches rows from a query
GET_VALUE Procedure	Retrieves column value from `SELECT` statement, or retrieves `OUT` bind parameters
GET_VALUE_RAW Procedure	Similar to `GET_VALUE`, but for datatype `RAW`
OPEN_CURSOR Function	Opens a cursor for running a passthrough SQL statement at the non-Oracle system
PARSE Procedure	Parses SQL statement at non-Oracle system

95.3.1 BIND_INOUT_VARIABLE Procedure

This procedure binds `IN OUT` bind variables.

Syntax

```
DBMS_HS_PASSTHROUGH.BIND_OUT_VARIABLE (
   c      IN  BINARY_INTEGER NOT NULL,
   p      IN  BINARY_INTEGER NULL,
   v      OUT <dty>,
   n      IN  VARCHAR2);
```

`<dty>` is either `DATE`, `NUMBER`, or `VARCHAR2`.

> **See Also:**
>
> For binding `OUT` variables of datatype `RAW`, see BIND_OUT_VARIABLE_RAW Procedure.

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-2 BIND_INOUT_VARIABLE Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable in the SQL statement: Starts at 1.
v	This value is used for two purposes: - To provide the IN value before the SQL statement is run. - To determine the size of the out value.
n	(Optional) Name of the bind variable. For example, in SELECT * FROM emp WHERE ename=:ename, the position of the bind variable :ename is 1, the name is :ename. This parameter can be used if the non-Oracle system supports "named binds" instead of positional binds. Passing the position is still required.

Exceptions

Table 95-3 BIND_INOUT_VARIABLE Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.2 BIND_INOUT_VARIABLE_RAW Procedure

This procedure binds IN OUT bind variables of datatype RAW.

Syntax

```
DBMS_HS_PASSTHROUGH.BIND_INOUT_VARIABLE_RAW (
   c      IN     BINARY_INTEGER NOT NULL,
   p      IN     BINARY_INTEGER NOT NULL,
   v      IN OUT RAW,
   n      IN     VARCHAR2);
```

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-4 BIND_INOUT_VARIABLE_RAW Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable in the SQL statement: Starts at 1.
v	This value is used for two purposes: - To provide the IN value before the SQL statement is run. - To determine the size of the out value.
n	(Optional) Name the bind variable. For example, in SELECT * FROM emp WHERE ename=:ename, the position of the bind variable :ename is 1, the name is :ename. This parameter can be used if the non-Oracle system supports "named binds" instead of positional binds. Passing the position is still required.

Exceptions

Table 95-5 BIND_INOUT_VARIABLE_RAW Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.3 BIND_OUT_VARIABLE Procedure

This procedure binds an OUT variable with a PL/SQL program variable.

Syntax

```
DBMS_HS_PASSTHROUGH.BIND_OUT_VARIABLE (
   c      IN  BINARY_INTEGER NOT NULL,
   p      IN  BINARY_INTEGER NULL,
   v      OUT <dty>,
   n      IN  VARCHAR2);
```

<dty> is either DATE, NUMBER, or VARCHAR2.

> **See Also:**
>
> For binding OUT variables of datatype RAW, see BIND_OUT_VARIABLE_RAW Procedure.

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-6 BIND_OUT_VARIABLE Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable in the SQL statement: Starts at 1.
v	Variable in which the OUT bind variable stores its value. The package remembers only the "size" of the variable. After the SQL statement is run, you can use GET_VALUE to retrieve the value of the OUT parameter. The size of the retrieved value should not exceed the size of the parameter that was passed using BIND_OUT_VARIABLE.
n	(Optional) Name of the bind variable. For example, in SELECT * FROM emp WHERE ename=:ename, the position of the bind variable :ename is 1, the name is :ename. This parameter can be used if the non-Oracle system supports "named binds" instead of positional binds. Passing the position is still required.

Exceptions

Table 95-7 BIND_OUT_VARIABLE Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.4 BIND_OUT_VARIABLE_RAW Procedure

This procedure binds an OUT variable of datatype RAW with a PL/SQL program variable.

Syntax

```
DBMS_HS_PASSTHROUGH.BIND_OUT_VARIABLE_RAW (
   c     IN  BINARY_INTEGER NOT NULL,
   p     IN  BINARY_INTEGER NOT NULL,
   v     OUT RAW,
   n     IN  VARCHAR2);
```

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-8 BIND_OUT_VARIABLE_RAW Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable in the SQL statement: Starts at 1.
v	Variable in which the OUT bind variable stores its value. The package remembers only the "size" of the variable. After the SQL statement is run, you can use GET_VALUE to retrieve the value of the OUT parameter. The size of the retrieved value should not exceed the size of the parameter that was passed using BIND_OUT_VARIABLE_RAW.
n	(Optional) Name of the bind variable. For example, in SELECT * FROM emp WHERE ename=:ename, the position of the bind variable :ename is 1, the name is :ename. This parameter can be used if the non-Oracle system supports "named binds" instead of positional binds. Passing the position is still required.

Exceptions

Table 95-9 BIND_OUT_VARIABLE_RAW Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.5 BIND_VARIABLE Procedure

This procedure binds an IN variable positionally with a PL/SQL program variable.

Syntax

```
DBMS_HS_PASSTHROUGH.BIND_VARIABLE (
   c      IN BINARY_INTEGER NOT NULL,
   p      IN BINARY_INTEGER NOT NULL,
   v      IN <dty>,
   n      IN VARCHAR2);
```

<dty> is either DATE, NUMBER, or VARCHAR2.

> **See Also:**
>
> To bind RAW variables use BIND_VARIABLE_RAW Procedure.

Pragmas

```
Purity level defined: WNDS, RNDS
```

Parameters

Table 95-10 BIND_VARIABLE Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable in the SQL statement: Starts at 1.
v	Value that must be passed to the bind variable name.
n	(Optional) Name of the bind variable. For example, in SELECT * FROM emp WHERE ename=:ename, the position of the bind variable :ename is 1, the name is :ename. This parameter can be used if the non-Oracle system supports "named binds" instead of positional binds. Passing the position is still required.

Exceptions

Table 95-11 BIND_VARIABLE Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.6 BIND_VARIABLE_RAW Procedure

This procedure binds IN variables of type RAW.

Syntax

```
DBMS_HS_PASSTHROUGH.BIND_VARIABLE_RAW (
   c    IN BINARY_INTEGER NOT NULL,
   p    IN BINARY_INTEGER NOT NULL,
   v    IN RAW,
   n    IN VARCHAR2);
```

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-12 BIND_VARIABLE_RAW Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable in the SQL statement: Starts at 1.
v	Value that must be passed to the bind variable.
n	(Optional) Name of the bind variable. For example, in SELECT * FROM emp WHERE ename=:ename, the position of the bind variable :ename is 1, the name is :ename. This parameter can be used if the non-Oracle system supports "named binds" instead of positional binds. Passing the position is still required.

Exceptions

Table 95-13 BIND_VARIABLE_RAW Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.7 CLOSE_CURSOR Procedure

This function closes the cursor and releases associated memory after the SQL statement has been run at the non-Oracle system. If the cursor was not open, then the operation is a "no operation".

Syntax

```
DBMS_HS_PASSTHROUGH.CLOSE_CURSOR (
   c IN BINARY_INTEGER NOT NULL);
```

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-14 CLOSE_CURSOR Procedure Parameters

Parameter	Description
c	Cursor to be released.

Exceptions

Table 95-15 CLOSE_CURSOR Procedure Exceptions

Exception	Description
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.8 EXECUTE_IMMEDIATE Procedure

This function runs a SQL statement immediately. Any valid SQL command except SELECT can be run immediately.

The statement must not contain any bind variables. The statement is passed in as a VARCHAR2 in the argument. Internally the SQL statement is run using the PASSTHROUGH SQL protocol sequence of OPEN_CURSOR, PARSE, EXECUTE_NON_QUERY, CLOSE_CURSOR.

Syntax

```
DBMS_HS_PASSTHROUGH.EXECUTE_IMMEDIATE (
   s IN VARCHAR2 NOT NULL)
RETURN BINARY_INTEGER;
```

Parameters

Table 95-16 EXECUTE_IMMEDIATE Procedure Parameters

Parameter	Description
s	VARCHAR2 variable with the statement to be executed immediately.

Return Values

The number of rows affected by the execution of the SQL statement.

Exceptions

Table 95-17 EXECUTE_IMMEDIATE Procedure Exceptions

Exception	Description
ORA-28551	SQL statement is invalid.
ORA-28554	Max open cursors.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.9 EXECUTE_NON_QUERY Function

This function runs a SQL statement. The SQL statement cannot be a `SELECT` statement. A cursor has to be open and the SQL statement has to be parsed before the SQL statement can be run.

Syntax

```
DBMS_HS_PASSTHROUGH.EXECUTE_NON_QUERY (
   c IN BINARY_INTEGER NOT NULL)
  RETURN BINARY_INTEGER;
```

Parameters

Table 95-18 EXECUTE_NON_QUERY Function Parameters

Parameter	Description
`c`	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines `OPEN_CURSOR` and `PARSE` respectively.

Return Values

The number of rows affected by the SQL statement in the non-Oracle system

Exceptions

Table 95-19 EXECUTE_NON_QUERY Function Exceptions

Exception	Description
`ORA-28550`	The cursor passed is invalid.
`ORA-28552`	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
`ORA-28555`	A `NULL` value was passed for a `NOT NULL` parameter.

95.3.10 FETCH_ROW Function

This function fetches rows from a result set.

The result set is defined with a SQL `SELECT` statement. When there are no more rows to be fetched, the exception `NO_DATA_FOUND` is raised. Before the rows can be fetched, a cursor has to be opened, and the SQL statement has to be parsed.

Syntax

```
DBMS_HS_PASSTHROUGH.FETCH_ROW (
   c   IN BINARY_INTEGER NOT NULL,
   f   IN BOOLEAN)
  RETURN BINARY_INTEGER;
```

Pragmas

```
Purity level defined : WNDS
```

Parameters

Table 95-20 FETCH_ROW Function Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
first	(Optional) Reexecutes SELECT statement. Possible values: – TRUE: reexecute SELECT statement. – FALSE: fetch the next row, or if run for the first time, then execute and fetch rows (default).

Return Values

The returns the number of rows fetched. The function returns "0" if the last row was already fetched.

Exceptions

Table 95-21 FETCH_ROW Function Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.11 GET_VALUE Procedure

This procedure has two purposes: it retrieves the select list items of SELECT statements after a row has been fetched, and it retrieves the OUT bind values, after the SQL statement has been run.

Syntax

```
DBMS_HS_PASSTHROUGH.GET_VALUE (
   c      IN  BINARY_INTEGER NOT NULL,
   p    IN  BINARY_INTEGER NOT NULL,
   v    OUT <dty>);
```

<dty> is either DATE, NUMBER, or VARCHAR2.

> **See Also:**
>
> For retrieving values of datatype RAW, see GET_VALUE_RAW Procedure.

Pragmas

```
Purity level defined : WNDS
```

Parameters

Table 95-22 GET_VALUE Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.
p	Position of the bind variable or select list item in the SQL statement: Starts at 1.
v	Variable in which the OUT bind variable or select list item stores its value.

Exceptions

Table 95-23 GET_VALUE Procedure Exceptions

Exception	Description
ORA-1403	Returns NO_DATA_FOUND exception when running the GET_VALUE after the last row was fetched (that is, FETCH_ROW returned "0").
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.12 GET_VALUE_RAW Procedure

This procedure is similar to GET_VALUE, but for datatype RAW.

Syntax

```
DBMS_HS_PASSTHROUGH.GET_VALUE_RAW (
   c    IN  BINARY_INTEGER NOT NULL,
   p    IN  BINARY_INTEGER NOT NULL,
   v    OUT RAW);
```

Pragmas

```
Purity level defined : WNDS
```

Parameters

Table 95-24 GET_VALUE_RAW Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened and parsed, using the routines OPEN_CURSOR and PARSE respectively.

Table 95-24 (Cont.) GET_VALUE_RAW Procedure Parameters

Parameter	Description
p	Position of the bind variable or select list item in the SQL statement: Starts at 1.
v	Variable in which the OUT bind variable or select list item stores its value.

Exceptions

Table 95-25 GET_VALUE_RAW Procedure Exceptions

Exception	Description
ORA-1403	Returns NO_DATA_FOUND exception when running the GET_VALUE after the last row was fetched (that is, FETCH_ROW returned "0").
ORA-28550	The cursor passed is invalid.
ORA-28552	Procedure is not run in right order. (Did you first open the cursor and parse the SQL statement?)
ORA-28553	The position of the bind variable is out of range.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

95.3.13 OPEN_CURSOR Function

This function opens a cursor for running a passthrough SQL statement at the non-Oracle system. This function must be called for any type of SQL statement.

The function returns a cursor, which must be used in subsequent calls. This call allocates memory. To deallocate the associated memory, call the procedure CLOSE_CURSOR.

Syntax

```
DBMS_HS_PASSTHROUGH.OPEN_CURSOR
  RETURN BINARY_INTEGER;
```

Pragmas

```
Purity level defined : WNDS, RNDS
```

Return Values

The cursor to be used on subsequent procedure and function calls.

Exceptions

Table 95-26 OPEN_CURSOR Function Exceptions

Exception	Description
ORA-28554	Maximum number of open cursor has been exceeded. Increase Heterogeneous Services' OPEN_CURSORS initialization parameter.

95.3.14 PARSE Procedure

This procedure parses an SQL statement at a non-Oracle system.

Syntax

```
DBMS_HS_PASSTHROUGH.PARSE (
   c       IN  BINARY_INTEGER NOT NULL,
   stmt    IN  VARCHAR2 NOT NULL);
```

Pragmas

```
Purity level defined : WNDS, RNDS
```

Parameters

Table 95-27 PARSE Procedure Parameters

Parameter	Description
c	Cursor associated with the passthrough SQL statement. Cursor must be opened using function OPEN_CURSOR.
stmt	Statement to be parsed.

Exceptions

Table 95-28 PARSE Procedure Exceptions

Exception	Description
ORA-28550	The cursor passed is invalid.
ORA-28551	SQL statement is illegal.
ORA-28555	A NULL value was passed for a NOT NULL parameter.

96
DBMS_HYBRID_VECTOR

The `DBMS_HYBRID_VECTOR` package contains a JSON-based `SEARCH` API that lets you search by *keywords* and *vectors* against hybrid vector indexes. By integrating traditional keyword-based text search with vector-based similarity search, you can improve the overall search experience and provide users with more accurate information.

Related Topics

- *Oracle Database AI Vector Search User's Guide*

96.1 SEARCH

Use the `DBMS_HYBRID_VECTOR.SEARCH` PL/SQL function to run textual queries, vector similarity queries, or hybrid queries against hybrid vector indexes.

Purpose

To search by vectors and keywords. This function lets you perform the following tasks:

- Facilitate a combined (hybrid) query of textual documents and vectorized chunks:

  You can query a hybrid vector index in multiple vector and keyword search combinations called search modes, as described in Understand Hybrid Search.

- Fuse and reorder the search results:

  Because a hybrid query by keywords and vectors involves independent searches on text and vector indexes, this API merges (fuses) the two search results into a unified result set as `CLOB`, which is reordered by a combined score. The results are fused using a range of fusion set operators. The new relevance scores are evaluated using standard algorithms such as Reciprocal Rank Fusion (RRF) and Relative Score Fusion (RSF).

- Run a default query for a simplified search experience:

  The minimum input parameters required are the hybrid vector index name and a search text string. The same search text string is used to query against a vectorized chunk index and a text document index. The search text is transformed into a `CONTAINS` query for keyword search, and is vectorized for a `VECTOR_DISTANCE` query for semantic search. By default, the results are fused using `INTERSECT` and scored using RSF.

Syntax

```
DBMS_HYBRID_VECTOR.SEARCH(
   json(
     '{ "hybrid_index_name" :  "hybrid_vector_index_name",
         "search_text"      :  "query string for vector and keyword search",
         "search_scorer"    :  "RRF | RSF",
         "search_fusion"    :  "INTERSECT | UNION | TEXT_ONLY | VECTOR_ONLY |
MINUS_TEXT | MINUS_VECTOR",
         "vector":
          {
           "search_text"    :  "query string for vector search",
```

```
           "search_vector"  :  "[vector_embedding]",
           "search_mode"    :  "DOCUMENT | CHUNK",
           "aggregator"     :  "MAX | AVG | MEDIAN | AVGN | MEDN |
MAXAVGMED",
           "score_weight"   :  score_weight,
           "rank_penalty"   :  rank_penalty
          },
         "text":
          {
           "contains"       :  "query keyword for text search",
           "score_weight"   :  score_weight,
           "rank_penalty"   :  rank_penalty
          },
         "return":
          {
           "topN"           :   "topN_value",
           "values"         :   ["rowid | score | vector_score | text_score |
vector_rank | text_rank | chunk_text | chunk_id"],
           "format"         :   "JSON | XML"
         }
     }'
  )
)
```

This API accepts a JSON specification for all query parameters.

hybrid_index_name

Specify the name of the hybrid vector index to use.

For information on how to create a hybrid vector index if not already created, see Manage Hybrid Vector Indexes.

search_text

Specify the search text (one or more query input), depending on the search construct to use.

This API supports the following two constructs of search:

- With common search text for hybrid search (default setting).

  The same text is used for a keyword query on document text index (by converting the `search_text` into a `CONTAINS ACCUM` operator syntax) and a semantic query on vectorized chunk index (by vectorizing or embedding the `search_text` for a `VECTOR_DISTANCE` search).

  Example of hybrid search with common search text:

  ```
  DBMS_HYBRID_VECTOR.SEARCH(
      json('{ "hybrid_index_name" : "my_hybrid_idx",
              "search_text"       : "C, Python"
            }'))
  from dual;
  ```

- With separate search texts using `vector` and `text` sub-elements:
  - Text-only search: The search text is converted into a `CONTAINS` query operator syntax for pure keyword search.

Example of text-only search:

```
DBMS_HYBRID_VECTOR.SEARCH(
    json('{ "hybrid_index_name" : "my_hybrid_idx",
            "text":
                   { "contains" : "C and Python" }
          }'))
from dual;
```

– Vector-only search: The search text is converted into a vector embedding, and is used in `VECTOR_DISTANCE`.

Example of vector-only search:

Here you use the `search_text` parameter to specify query input for vector search.

```
DBMS_HYBRID_VECTOR.SEARCH(
    json('{ "hybrid_index_name" : "my_hybrid_idx",
            "vector":
                    { "search_text" : "C, Python }
          }'))
from dual;
```

Here you use the `search_vector` parameter to directly pass a vector embedding as query input for your vector search.

```
dbms_hybrid_vector.search(
    json('{ "hybrid_index_name" : "my_hybrid_idx",
         "vector":
                  {
                   "search_vector" : vector_serialize(
                                          vector_embedding(doc_model
                                                  using
                                                  "C, Python,
Database"
                                                  as data)
                                        RETURNING CLOB),
                  },}'))
from dual;
```

– Hybrid search: Conducts two separate keyword and semantic queries, where keyword scores and semantic scores are combined. Specify two different search texts: One search text or a vector embedding to run a `VECTOR_DISTANCE` query for semantic search. A second search text to run a `CONTAINS` query for keyword search.

Example of hybrid search with separate search texts:

```
DBMS_HYBRID_VECTOR.SEARCH(
    json('{ "hybrid_index_name" : "my_hybrid_idx",
            "vector":
                    { "search_text" : "leadership experience" },
             "text":
                    { "contains" : "C and Python" }
          }'))
from dual;
```

> **Note:**
>
> You cannot use both the constructs in a query, that is, a common `search_text` field along with separate search texts (using the `vector` or `text` sub-elements).

search_scorer

Specify the method to evaluate the combined "fusion" search scores from separate text and vector searches. The values can be:

- `RSF` to use the relative score fusion algorithm (default setting)
- `RRF` to use the reciprocal rank fusion algorithm

For a deeper understanding of how these algorithms work in hybrid search modes, see Understand Hybrid Search.

Example of search_scorer (hybrid search with common search text):

```
DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "search_text"       : "C, Python",
         "search_scorer"     : "rsf"
      }'))
from dual;
```

Example of search_scorer (hybrid search with separate search texts):

```
DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "search_scorer"     : "rsf",
         "vector":
          { "search_text"    : "leadership experience" },
         "text":
          { "contains"       : "C and Python" }
      }'))
from dual;
```

search_fusion

Specify the method to combine the search results for hybrid search modes. Note that vector-only and text-only searches do not combine any results.

Parameter	Description
`INTERSECT` (default)	Combines all distinct rows selected by both text and vector searches. Score condition: `text_score > 0 AND vector_score > 0`
`UNION`	Combines all distinct rows selected by either text or vector search. Score condition: `text_score > 0 OR vector_score > 0`
`TEXT_ONLY`	Combines all distinct rows selected by text search. Score condition: `text_score > 0`

Parameter	Description
VECTOR_ONLY	Combines all distinct rows selected by vector-only search. Score condition: vector_score > 0
MINUS_TEXT	Combines all distinct rows selected by vector search but not text search. Score condition: text_score = 0
MINUS_VECTOR	Combines all distinct rows selected by text search but not vector search. Score condition: vector_score = 0

Example of search_fusion:

```
DBMS_HYBRID_VECTOR.SEARCH(
    json('{ "hybrid_index_name" : "my_hybrid_idx",
            "search_fusion"     : "INTERSECT",
            "vector":
                    { "search_text" : "leadership experience" },
             "text":
                    { "contains" : "C and Python" }
         }'))
from dual;
```

vector

Specify query parameters for a VECTOR_DISTANCE query on the vector index part of the hybrid vector index:

- search_text: Search text string to use for vector search.
- search_vector: Vector embedding, if available. Note that the vector embedding to be passed here must be generated using the same embedding model used for vector similarity search by the specified hybrid vector index.
- search_mode: Mode of the vector search at a document or chunk level:

Parameter	Description
DOCUMENT (default)	Returns the top-matched row IDs from your base table corresponding to the list of best documents identified.
CHUNK	Returns the top-matched chunk IDs and chunk texts from the files stored in your base table, regardless of whether the chunks come from the same document or different documents. The content from these chunks can be used as input for LLMs to formulate responses.

- aggregator: Aggregate function to apply for aggregating the vector scores for each document in DOCUMENT SEARCH_MODE.

Parameter	Description
MAX (default)	Standard database aggregate function that selects the top chunk score as the result score. With this function, TopN and minScore are ignored.
AVG	Standard database aggregate function that sums the chunk scores and divides by the count. With this function, TopN and minScore are ignored.
MEDIAN	Standard database aggregate function that computes the middle value or an interpolated value of the sorted scores. With this function, TopN and minScore are ignored.

Parameter	Description
AVGN(topN,minScore)	This is similar to the standard database AVG aggregate, but it sums the topN chunk scores (>= minScore) divided by topN.
MEDN(topN,minScore)	This is the similar to the standard database MEDIAN aggregate, but it computes the median value of the sorted topN scores (>= minScore).
MAXAVGMED(topN,minScore,maxWgt,avgWgt,medWgt)	This function computes a weighted sum of the MAX, AVGN, and MEDN values (with topN and minScore parameters in effect). The max, average, and median weights are passed as arg3, arg4, and arg5. MAXAVGMED = (MAX * maxWgt + AVGN * avgWgt + MEDN * medWgt) / (maxWgt + avgWgt + medWgt) If weight of zero is passed, then that corresponding value is not used.

- score_weight: Relative weight (degree of importance or preference) to assign to the semantic VECTOR_DISTANCE query. This value is used when combining the results of RSF ranking.

  Value: Any positive integer greater than 0 (zero)

  Default: 10 (implies 10 times more importance to vector query than text query)

- rank_penalty: Penalty (degree of reduction in the relevance score) to assign to vector query. This can help in balancing the relevance score by reducing the importance of unnecessary or repetitive words in a document. This value is used when combining the results of RRF ranking.

  Value: Any positive integer greater than 0 (zero)

  Default: 1 (implies a minimal reduction of relevance score by 1 point)

Example of vector-only search (chunk result)::

```
DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "vector":
          {
             "search_text"   : "leadership experience",
             "search_mode"   : "CHUNK"
          },
      }'))
from dual;
```

Example of vector-only search (document result):

```
DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "vector":
          {
             "search_text"   : "leadership experience",
             "search_mode"   : "DOCUMENT",
             "aggregator"    : "MAX"
          },
      }'))
from dual;
```

text

Specify query parameters for a `CONTAINS` query on the Oracle Text index part of the hybrid vector index:

Parameter	Description
`contains`	Oracle Text index operator to specify query expressions for full-text search. See Oracle Text CONTAINS Query Operators.
`score_weight`	Relative weight (degree of importance or preference) to assign to the text `CONTAINS` query. This value is used when combining the results of RSF ranking. Value: Any positive integer greater than `0` (zero) Default: `1` (implies neutral weight)
`rank_penalty`	Penalty (degree of reduction in the relevance score) to assign to text query. This can help in balancing the relevance score by reducing the importance of unnecessary or repetitive words in a document. This value is used when combining the results of RRF ranking. Value: Any positive integer greater than `0` (zero) Default: `5` (implies a reduction of relevance score by 5 points)

Example of text-only search with all text sub-elements:

```
DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "text":
          {
             "contains"      : "C and Python",
             "score_weight"  : 1,
             "rank_penalty"  : 5
          },
      }'))
from dual;
```

return

Specify which fields to return in the result set:

Parameter	Description
`topN`	Maximum number of best-matched results to be returned Value: Any integer greater than `0` (zero) Default: `20`
`values`	Return attributes for the search results: • `rowid`: Row ID value for each source document • `score`: Combined text-and-vector search score • `vector_score`: Vector-only search score • `text_score`: Text-only search score • `vector_rank`: Vector-only search rank • `text_rank`: Text-only search rank • `chunk_text`: Human-readable content from each chunk • `chunk_id`: ID of each chunk text Default: All the above return attributes are shown by default.

Parameter	Description
format	Format of the results as JSON (default) or XML.

Example of return sub-elements:

```
DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "search_text"       : "C, Python",
         "return":
          {
             "values"        : [ "rowid", "score" ],
             "topN"          : 3,
             "format"        : "JSON"
          }
      }'))
from dual;
```

Complete Example With All Query Parameters

The following example shows a hybrid search query that performs separate text and vector searches against `my_hybrid_idx`. This query specifies the `search_text` for vector search using the `vector_distance` function as `prioritize teamwork and leadership experience` and the keyword for text search using the `contains` Oracle Text operator as `C and Python`. The search mode is `DOCUMENT` to return the search results as topN documents.

```
select json_Serialize(
  DBMS_HYBRID_VECTOR.SEARCH(
    json(
      '{ "hybrid_index_name" : "my_hybrid_idx",
         "search_scorer"     : "rsf",
         "search_fusion"     : "INTERSECT",
         "vector":
          {
             "search_text"   : "prioritize teamwork and leadership
experience",
             "search_mode"   : "DOCUMENT",
             "score_weight"  : 10,
             "rank_penalty"  : 1,
             "aggregator"    : "MAX"
          },
         "text":
          {
             "contains"      : "C and Python",
             "score_weight"  : 1,
             "rank_penalty"  : 5
          },
         "return":
          {
             "format"        : "JSON",
             "topN"          : 3,
             "values"        : [ "rowid", "score", "vector_score",
                                 "text_score", "vector_rank",
```

```
                                    "text_rank", "chunk_text", "chunk_id" ]
          }
      }'
    )
  ) pretty)
from dual;
```

The top 3 rows are ordered by relevance, with higher scores indicating a better match. All the return attributes are shown by default:

```
[
  {
    "rowid"         : "AAAR9jAABAAAQeaAAA",
    "score"         : 58.64,
    "vector_score"  : 61,
    "text_score"    : 35,
    "vector_rank"   : 1,
    "text_rank"     : 2,
    "chunk_text"    : "Candidate 1: C Master. Optimizes low-level system
(i.e. Database)
                      performance with C. Strong leadership skills in
guiding teams to
                      deliver complex projects.",
    "chunk_id"      : "1"
  },
  {
    "rowid"         : "AAAR9jAABAAAQeaAAB",
    "score"         : 56.86,
    "vector_score"  : 55.75,
    "text_score"    : 68,
    "vector_rank"   : 3,
    "text_rank"     : 1,
    "chunk_text"    : "Candidate 3: Full-Stack Developer. Skilled in
Database, C, HTML,
                      JavaScript, and Python with experience in building
responsive web
                      applications. Thrives in collaborative team
environments.",
    "chunk_id"      : "1"
  },
  {
    "rowid"         : "AAAR9jAABAAAQeaAAD",
    "score"         : 51.67,
    "vector_score"  : 56.64,
    "text_score"    : 2,
    "vector_rank"   : 2,
    "text_rank"     : 3,
    "chunk_text"    : "Candidate 2: Database Administrator (DBA). Maintains
and secures
                      enterprise database (Oracle, MySql, SQL Server).
Passionate about
                      data integrity and optimization. Strong mentor for
junior DBA(s).",
    "chunk_id"      : "1"
```

```
    }
]
```

End-to-end example:

To see how to create a hybrid vector index and explore all types of queries against the index, see Query Hybrid Vector Indexes End-to-End Example.

97
DBMS_ILM

The `DBMS_ILM` package provides an interface for implementing Information Lifecycle Management (ILM) strategies using Automatic Data Optimization (ADO) policies.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Exceptions
- Summary of DBMS_ILM Subprograms

> **See Also:**
>
> - *Oracle Database VLDB and Partitioning Guide* for information about managing Automatic Data Optimization (ADO) with this package
> - DBMS_ILM_ADMIN
> - DBMS_HEAT_MAP

97.1 DBMS_ILM Overview

To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification. You can also use Automatic Data Optimization (ADO) to automate the compression and movement of data between different tiers of storage within the database. The `DBMS_ILM` package supports immediate evaluation or execution of ADO related tasks. T

he package supports the following two ways for scheduling ADO actions.

- A database user schedules immediate ADO policy execution on a set of objects.
- A database user views the results of evaluation of ADO policies on a set of objects. The user then adds or deletes objects to this set and reviews the results of ADO policy evaluation again. The user repeats this step to determine the set of objects for ADO execution. The user can then schedule ADO actions for immediate execution on this set of objects.

The following procedures support the two usage modes. Before describing the procedures, we introduce the notion of an ADO task as an entity that helps to track a particular evaluation or (an evaluation and execution) of ADO policies. A particular ADO task could be in one of the following states.

- Inactive
- Active
- Completed

97.2 DBMS_ILM Security Model

The DBMS_ILM package runs under invoker's rights.

97.3 DBMS_ILM Constants

The table in this topic lists the constants used by the DBMS_ILM package.

Table 97-1 DBMS_ILM Constants

Constant	Value	Type	Description
ILM_ALL_POLICIES	'ALL POLICIES'	VARCHAR2(20)	Selects all ADO policies on an object
ILM_EXECUTION_OFFLINE	1	NUMBER	Specifies that the object may be offline while ADO action is performed
ILM_EXECUTION_ONLINE	2	NUMBER	Specifies that the object should be online while ADO action is performed
SCOPE_DATABASE	1	NUMBER	Selects all ADO policies in the database
SCOPE_SCHEMA	2	NUMBER	Selects all ADO policies in the current schema
SCHEDULE_IMMEDIATE	1	NUMBER	Schedules ADO task for immediate execution
ARCHIVE_STATE_ACTIVE	'0'	VARCHAR2(1)	Represents the value of the ORA_ARCHIVE_STATE column of a row-archival enabled table that would make the row active
ARCHIVE_STATE_ARCHIVED	'1'	VARCHAR2(1)	Represents the value of the ORA_ARCHIVE_STATE column of a row-archival enabled table that would make the row inactive

97.4 DBMS_ILM Exceptions

The table in this topic lists the exceptions raised by the DBMS_ILM package.

Table 97-2 DBMS_ILM Exceptions

Exception	Error Code	Description
INVALID_ARGUMENT_VALUE	38327	Invalid argument value
INVALID_ILM_DICTIONARY	38328	Inconsistent dictionary state
INTERNAL_ILM_ERROR	38329	Internal error
INSUFFICIENT_PRIVILEGES	38330	Insufficient privileges

97.5 Summary of DBMS_ILM Subprograms

Thi table lists and describes the `DBMS_ILM` package subprograms.

Table 97-3 DBMS_ILM Package Subprograms

Subprogram	Description
ADD_TO_ILM Procedure	Adds the object specified through the argument to a particular ADO task and evaluates the ADO policies on this object
ARCHIVESTATENAME Function	Returns the value of the `ORA_ARCHIVE_STATE` column of a row-archival enabled table
EXECUTE_ILM Procedure	Executes an ADO task.
EXECUTE_ILM_TASK Procedure	Executes an ADO task that has been evaluated previously
PREVIEW_ILM Procedure	Evaluates all ADO policies in the scope specified by means of an argument
REMOVE_FROM_ILM Procedure	Removes the object specified through the argument from a particular ADO task
STOP_ILM Procedure	Stops ADO-related jobs created for a particular ADO task

97.5.1 ADD_TO_ILM Procedure

This procedure adds the object specified through the argument to a particular ADO task and evaluates the ADO policies on this object.

The procedure can only be executed on an ADO task in an inactive state. The results of the ADO policy evaluation on this object can be viewed using the appropriate views depending on role and access (`USER_ILMTASKS` or `DBA_ILMTASKS`, `USER_ILMEVALUATIONDETAILS` or `DBA_ILMEVALUATIONDETAILS`, `USER_ILMRESULTS` or `DBA_ILMRESULTS`).

Syntax

```
DBMS_ILM.ADD_TO_ILM (
   task_id           IN    NUMBER,
   owner             IN    VARCHAR2,
   object_name       IN    VARCHAR2,
   subobject_name    IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 97-4 *ADD_TO_ILM Procedure Parameters*

Parameter	Description
`task_id`	Identifies a particular ADO task
`owner`	Owner of the object
`object_name`	Name of the object
`subobject_name`	Name of the subobject (partition name in the case of partitioned tables)

97.5.2 ARCHIVESTATENAME Function

This function returns the value of the `ORA_ARCHIVE_STATE` column of a row-archival enabled table.

Syntax

```
DBMS_ILM.ARCHIVESTATENAME (
   value      IN  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 97-5 *ARCHIVESTATENAME Function Parameters*

Parameter	Description
`value`	Value for which the archive state name is to be returned

Usage Notes

Returns `ARCHIVE_STATE_ACTIVE` for `0`, `ARCHIVE_STATE_ARCHIVED` for others

> **See Also:**
>
> "Using In-Database Archiving" in *Oracle Database VLDB and Partitioning Guide*

97.5.3 EXECUTE_ILM Procedure

This procedure executes an ADO task.

There are two overloads to this procedure. The first overload executes an ADO task for a set of objects without having evaluated them previously. The second overload executes ADO policies for a specific object.

Syntax

```
DBMS_ILM.EXECUTE_ILM (
   task_id             OUT    NUMBER,
   ilm_scope           IN     NUMBER DEFAULT SCOPE_SCHEMA,
   execution_mode      IN     NUMBER DEFAULT ILM_EXECUTION_ONLINE);

DBMS_ILM.EXECUTE_ILM (
   owner               IN     VARCHAR2,
   object_name         IN     VARCHAR2,
   task_id             OUT    NUMBER,
   subobject_name      IN     VARCHAR2 DEFAULT NULL,
   policy_name         IN     VARCHAR2 DEFAULT ILM_ALL_POLICIES,
   execution_mode      IN     NUMBER DEFAULT ILM_EXECUTION_ONLINE);
```

Parameters

Table 97-6 *EXECUTE_ILM Procedure Parameters*

Parameter	Description
task_id	Identifies a particular ADO task
ilm_scope	Determines the set of objects considered for ADO execution. The default is to consider only the objects in the schema.
execution_mode	Whether the ADO task be executed online (ILM_EXECUTION_ONLINE) or offline (ILM_EXECUTION_OFFLINE)
owner	Owner of the object
object_name	Name of the object
subobject_name	Name of the subobject (partition name in the case of partitioned tables)
policy_name	Name of the ADO policy to be evaluated on the object. The package constant ILM_ALL_POLICIES should be used if all ADO policies on an object should be evaluated.

Usage Notes

- The EXECUTE_ILM procedure can be used by users who want more control of when ADO is performed, and who do not want to wait until the next maintenance window.
- The procedure executes like a DDL in that it auto commits before and after the ADO task and related jobs are created.

97.5.4 EXECUTE_ILM_TASK Procedure

This procedure executes an ADO task that has been evaluated previously and moves it to an active state.

Syntax

```
DBMS_ILM.EXECUTE_ILM_TASK (
   task_id             IN     NUMBER,
   execution_mode      IN     NUMBER DEFAULT ILM_EXECUTION_ONLINE);
   execution_schedule  IN     NUMBER DEFAULT SCHEDULE_IMMEDIATE);
```

Parameters

Table 97-7 *EXECUTE_ILM_TASK Procedure Parameters*

Parameter	Description
task_id	Identifies a particular ADO task
execution_mode	Whether the ADO task be executed online (ILM_EXECUTION_ONLINE) or offline (ILM_EXECUTION_OFFLINE)
execution_schedule	Identifies when the ADO task should be executed.Currently, the only choice available is immediate scheduling of ADO jobs

97.5.5 PREVIEW_ILM Procedure

This procedure evaluates the ADO policies on the objects specified using the `ILM_SCOPE` argument.

It returns a number as `task_id` which identifies a particular ADO task. This can be used to view the results of the policy evaluation in the appropriate views depending on role and access (`USER_ILMTASKS` or `DBA_ILMTASKS`, `USER_ILMEVALUATIONDETAILS` or `DBA_ILMEVALUATIONDETAILS`, `USER_ILMRESULTS` or `DBA_ILMRESULTS`).

The `PREVIEW_ILM` procedure leaves the ADO task in an inactive state. Once you have previewed the results, you can add or delete objects to this task.

Syntax

```
DBMS_ILM.PREVIEW_ILM (
   task_id           OUT    NUMBER,
  ilm_scope          IN     NUMBER DEFAULT SCOPE_SCHEMA);
```

Parameters

Table 97-8 *PREVIEW_ILM Procedure Parameters*

Parameter	Description
`task_id`	Identifies a particular ADO task
`ilm_scope`	Identifies the scope of execution. Should be either `SCOPE_DATABASE` or `SCOPE_SCHEMA` as described in Constants

97.5.6 REMOVE_FROM_ILM Procedure

This procedure removes the object specified through the argument from a particular ADO task.

The procedure can only be executed on an ADO task in an inactive state.

Syntax

```
DBMS_ILM.REMOVE_FROM_ILM (
   task_id           IN    NUMBER,
   owner             IN    VARCHAR2,
   object_name       IN    VARCHAR2,
   subobject_name    IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 97-9 *REMOVE_FROM_ILM Procedure Parameters*

Parameter	Description
`task_id`	Identifies a particular ADO task
`owner`	Owner of the object
`object_name`	Name of the object
`subobject_name`	Name of the subobject (partition name in the case of partitioned tables)

97.5.7 STOP_ILM Procedure

This procedure terminates ILM ADO jobs associated to a particular task Id or job name.

Syntax

```
DBMS_ILM.STOP_ILM (
   task_id               IN         NUMBER,
   p_drop_running_jobs   IN         BOOLEAN  DEFAULT FALSE),
   p_jobname             IN         VARCHAR2 DEFAULT NULL);
```

Parameters

Table 97-10 *STOP_ILM Procedure Parameters*

Parameter	Description
`task_id`	Number that uniquely identifies a particular ADO task
`p_drop_running_jobs`	Determines whether running jobs are dropped
`p_jobname`	Name of job to be terminated

98
DBMS_ILM_ADMIN

The `DBMS_ILM_ADMIN` package provides an interface to customize Automatic Data Optimization (ADO) policy execution. In combination with partitioning and compression, ADO policies can be used to help implement an Information Lifecycle Management (ILM) strategy.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_ILM_ADMIN Subprograms

> **See Also:**
>
> - *Oracle Database VLDB and Partitioning Guide* for information about managing Automatic Data Optimization (ADO) with this package
> - DBMS_ILM
> - DBMS_HEAT_MAP

98.1 DBMS_ILM_ADMIN Overview

To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification. You can also use Automatic Data Optimization (ADO) to automate the compression and movement of data between different tiers of storage within the database.

98.2 DBMS_ILM_ADMIN Security Model

This package runs under definer's rights. The user requires DBA privileges.

98.3 DBMS_ILM_ADMIN Constants

The table in this topic describes constants used by the `DBMS_ILM_ADMIN` package.

The value column refers to the numeric or character value that the constants resolve to.

Table 98-1 DBMS_ILM_ADMIN Constants

Constant	Value	Type	Description
ABS_JOBLIMIT	12	NUMBER	Specifies the absolute number of concurrent ILM ADO jobs.

Table 98-1 (Cont.) DBMS_ILM_ADMIN Constants

Constant	Value	Type	Description
DEG_PARALLEL	10	NUMBER	Decides the degree of parallelism to be used for ADO jobs
ENABLED	7	NUMBER	Provides a way to turn background ADO off or on
EXECUTION_INTERVAL	1	NUMBER	Determines the frequency with which ADO background evaluation occurs. Specified in minutes.
EXECUTION_MODE	4	NUMBER	Controls whether ADO execution is online, offline. The value for this parameter should either be DBMS_ILM_ADMIN.ILM_EXECUTION_OFFLINE or DBMS_ILM_ADMIN.ILM_EXECUTION_ONLINE.
HEAT_MAP_SEG_LOOKUP	8	NUMBER	Index scan done
HEAT_MAP_SEG_READ	1	NUMBER	Segment read done
HEAT_MAP_SEG_SCAN	4	NUMBER	Full table scan done
HEAT_MAP_SEG_WRITE	2	NUMBER	Segment write done
JOB_SIZELIMIT	13	NUMBER	Specifies the size (in megabytes) of the data that is processed by a single ILM ADO row level compression job.
JOBLIMIT	5	NUMBER	Controls the upper limit on number of ILM ADO jobs at any time. The maximum number of concurrent ADO jobs is JOBLIMIT*(number of instances)*(number of CPUs per instance).
POLICY_TIME	11	NUMBER	Decides if ADO policies are treated as though they are specified in seconds rather than days. Can take value ILM_POLICY_IN_SECONDS (treat policy time in seconds) or ILM_POLICY_IN_DAYS (treat policy time in days - default).
RETENTION_TIME	2	NUMBER	Controls the amount of time ADO history should be maintained. Specified in days.
TBS_PERCENT_FREE	9	NUMBER	Decides the targeted tablespace storage through ADO actions as a percentage of tablespace quota.
TBS_PERCENT_USED	8	NUMBER	Decides when a tablespace is considered full. Specified as a percentage of tablespace quota.

The DBMS_ILM_ADMIN package uses the constants as parameter values shown in Table 98-2.

Table 98-2 DBMS_ILM_ADMIN Constants Used as Parameter Values

Constant	Value	Type	Description
AUTO_OPTIMIZE_ENABLED	15	NUMBER	Indicates whether automatic storage compression is enabled.

Table 98-2 (Cont.) DBMS_ILM_ADMIN Constants Used as Parameter Values

Constant	Value	Type	Description
AUTO_OPTIMIZE_INACTIVITY_THRESHOLD	14	NUMBER	The period of inactivity that will determine that there are no modifications. This can be specified using DBMS_ILM_ADMIN.CUSTOMIZE_ILM. The default value is 1440 minutes, which is 1 day.
ILM_DISABLED	2	NUMBER	Indicates automatic ADO policy evaluation and execution is disabled
ILM_ENABLED	1	NUMBER	Indicates automatic ADO policy evaluation and execution is enabled
ILM_EXECUTION_OFFLINE	1	NUMBER	Specifies that the object may be offline while ADO action is performed.
ILM_EXECUTION_ONLINE	2	NUMBER	Specifies that the object should be online while ADO action is performed
ILM_POLICY_IN_DAYS	0	NUMBER	Indicates policy is specified in days. This is the default.

Table 98-2 (Cont.) DBMS_ILM_ADMIN Constants Used as Parameter Values

Constant	Value	Type	Description
ILM_POLICY_IN_SECONDS	1	NUMBER	Indicates policy unit is changed from days to seconds. This could be used to test ADO policy evaluation quickly instead of waiting for the policy duration.

Note:

- Setting ILM_POLICY_IN_SECONDS does not compress the blocks within the specified seconds.
- Setting ILM_POLICY_IN_SECONDS is for test ADO and should not be set in the production environment.

98.4 Summary of DBMS_ILM_ADMIN Subprograms

This table lists and briefly describes the DBMS_ILM_ADMIN package subprograms.

Table 98-3 DBMS_ILM_ADMIN Package Subprograms

Subprogram	Description
CLEAR_HEAT_MAP_ALL Procedure	Deletes all rows except the dummy row
CLEAR_HEAT_MAP_TABLE Procedure	Clears all or some statistics for the heat map table, deleting rows for a given table or segment which match a given pattern, or all such rows
CUSTOMIZE_ILM Procedure	Customizes environment for ILM execution by specifying the values for ILM execution related parameters
DISABLE_ILM Procedure	Turns off all background ILM scheduling
ENABLE_AUTO_OPTIMIZE Procedure	Enables Auto Compression for all Hybrid Columnar Compression objects in the PDB.
ENABLE_ILM Procedure	Turns on all background ILM scheduling
IGNORE_AUTO_OPTIMIZE_CRITERIA Procedure	Ignores the inactivity threshold so that uncompressed loads will be eligible for background auto optimization immediately.
SET_HEAT_MAP_ALL Procedure	Updates or inserts heat map rows for all tables
SET_HEAT_MAP_START Procedure	Sets the start date for collecting heat map data
SET_HEAT_MAP_TABLE Procedure	Updates or inserts a row for the specified table or segment

98.4.1 CLEAR_HEAT_MAP_ALL Procedure

This procedure deletes all rows in `HEAT_MAP_STAT$` except the dummy row.

Syntax

```
DBMS_ILM_ADMIN.CLEAR_HEAT_MAP_ALL;
```

98.4.2 CLEAR_HEAT_MAP_TABLE Procedure

This procedure clears all or some statistics for the heat map table, deleting rows for a given table or segment which match a given pattern, or all such rows.

Syntax

```
DBMS_ILM_ADMIN.CLEAR_HEAT_MAP_TABLE  (
   owner                  IN  VARCHAR2,
   tablename              IN  VARCHAR2,
   partition              IN  VARCHAR2 default '',
   access_date            IN  DATE DEFAULT NULL,
   segment_access_summary IN  NUMBER DEFAULT NULL);
```

Parameters

Table 98-4 *CLEAR_HEAT_MAP_TABLE Procedure Parameters*

Parameter	Description
`owner`	Table owner
`tablename`	Table name

Table 98-4 (Cont.) *CLEAR_HEAT_MAP_TABLE Procedure Parameters*

Parameter	Description
partition	Name of the subobject, defaults to NULL
access_date	Date for the entry in HEAT_MAP_STAT$ to be removed
segment_access_summary	Summary of segment access constants indicating access operations performed on the segment

98.4.3 CUSTOMIZE_ILM Procedure

This procedure customizes environment for ILM execution by specifying the values for ILM execution related parameters. These values take effect for the next background scheduling.

Syntax

```
DBMS_ILM_ADMIN.CUSTOMIZE_ILM  (
   parameter            IN       NUMBER,
   value                IN       NUMBER);
```

Parameters

Table 98-5 *CUSTOMIZE_ILM Procedure Parameters*

Parameter	Description
parameter	One of the parameter constants defined in DBMS_ILM_ADMIN package
value	Value of parameter

98.4.4 DISABLE_AUTO_OPTIMIZE Procedure

Disables Auto Compression for all Hybrid Columnar Compression objects in the PDB.

Syntax

After you issue enable_auto_optimize, auto compression will be enabled. New direct loads will be in an uncompressed format and the data will be compressed gradually in the background. For example, after five direct loads the user issues disable_auto_optimize. If the background auto compression task was only able to compress three of them (because of autotask time constraints, data not yet cold, etc.), the remaining two will be in an uncompressed format. The data from these two uncompressed loads will be compressed in the background even after the user issues disable_auto_optimize, but new direct loads will now be compressed.

```
DBMS_ILM_ADMIN.DISABLE_AUTO_OPTIMIZE;
```

98.4.5 DISABLE_ILM Procedure

This procedure turns off all background ILM scheduling.

Syntax

```
DBMS_ILM_ADMIN.DISABLE_ILM;
```

98.4.6 ENABLE_AUTO_OPTIMIZE Procedure

Enables Auto Compression for all Hybrid Columnar Compression objects in the PDB.

When Auto Compression is enabled, direct loads into a Hybrid Columnar Compression (HCC) object would use the uncompressed format to achieve faster loads. The system will then wait until there are no modifications to the newly loaded data, for the duration of the specified inactivity threshold. At that point, the data from the uncompressed direct load will be gradually HCC compressed using a background Auto Compression task.

Syntax

```
DBMS_ILM_ADMIN.ENABLE_AUTO_OPTIMIZE;
```

98.4.7 ENABLE_ILM Procedure

This procedure turns on all background ILM scheduling.

Syntax

```
DBMS_ILM_ADMIN.ENABLE_ILM;
```

98.4.8 IGNORE_AUTO_OPTIMIZE_CRITERIA Procedure

Ignores the inactivity threshold so that uncompressed loads will be eligible for background auto optimization immediately.

Syntax

```
DBMS_ILM_ADMIN.IGNORE_AUTO_OPTIMIZE_CRITERIA;
```

98.4.9 SET_HEAT_MAP_ALL Procedure

This procedure sets an HTTP request header. The request header is sent to the Web server as soon as it is set.

Syntax

```
DBMS_ILM_ADMIN.SET_HEAT_MAP_ALL  (
   access_date            IN DATE,
   segment_access_summary IN NUMBER);
```

Parameters

Table 98-6 *SET_HEAT_MAP_ALL Procedure Parameters*

Parameter	Description
access_date	Date for the entry in HEAT_MAP_STAT$ to be added
segment_access_summary	Summary of segment access constants indicating access operations performed on the segment

98.4.10 SET_HEAT_MAP_START Procedure

This procedure sets the start date for collecting heat map data.

Syntax

```
DBMS_ILM_ADMIN.SET_HEAT_MAP_START  (
   start_date  IN   DATE);
```

Parameters

Table 98-7 *SET_HEAT_MAP_START Procedure Parameters*

Parameter	Description
start_date	Indicates the new date from which all statistics are valid

98.4.11 SET_HEAT_MAP_TABLE Procedure

This procedure updates or inserts a row for the specified table or segment.

Syntax

```
DBMS_ILM_ADMIN.SET_HEAT_MAP_TABLE  (
   owner                  IN VARCHAR2,
   tablename              IN VARCHAR2,
   partition              IN VARCHAR2 DEFAULT '',
   access_date            IN DATE DEFAULT NULL,
   segment_access_summary IN NUMBER DEFAULT NULL);
```

Parameters

Table 98-8 *SET_HEAT_MAP_TABLE Procedure Parameters*

Parameter	Description
owner	Table owner
tablename	Table name
partition	Name of the subobject, defaults to NULL
access_date	Date for the entry in HEAT_MAP_STAT$ to be added
segment_access_summary	Summary of segment access constants indicating access operations performed on the segment

99
DBMS_IMMUTABLE_TABLE

Immutable tables are read-only tables that protect data against unauthorized modification. They also prevent against accidental data modifications that may be caused by human errors. The `DBMS_IMMUTABLE_TABLE` package allows you to delete the expired rows in an immutable table and add interval partitioning.

This chapter contains the following topics:

- DBMS_IMMUTABLE_TABLE Overview
- DBMS_IMMUTABLE_TABLE Security Model
- Summary of DBMS_IMMUTABLE_TABLE Subprograms

> **See Also:**
>
> - *Oracle Database Administrator's Guide*
> - *Oracle Database SQL Language Reference*
> - *Oracle Database Reference*

99.1 DBMS_IMMUTABLE_TABLE Overview

Immutable tables are `read-only` tables that protect data against unauthorized modification. Immutable tables also prevent against accidental data modifications that may be caused by human errors. You must specify a retention period for the immutable table and for rows within the immutable table. An immutable table can be dropped if it contains no rows or if the specified retention period for the table has elapsed. Rows can be deleted from an immutable table only after the specified retention period for rows in the table has elapsed. Except for increasing the retention periods and adding interval partitioning, you cannot modify the definition of an immutable table with this package.

99.2 DBMS_IMMUTABLE_TABLE Security Model

The `DBMS_IMMUTABLE_TABLE` package is owned by `SYS` and is installed as part of database installation. The routines in the package are run with invoker's rights and hence run with the privileges of the current user.

99.3 Summary of DBMS_IMMUTABLE_TABLE Subprograms

This table lists the `DBMS_IMMUTABLE_TABLE` subprograms in alphabetical order and briefly describes them.

Table 99-1 DBMS_IMMUTABLE_TABLE Package Subprograms

Subprogram	Description
ADD_INTERVAL_PARTITIONING Procedure	This procedure adds interval partitioning to an existing, non-partitioned, V1 or V2 immutable table.
DELETE_EXPIRED_ROWS Procedure	This procedure deletes the expired rows.

99.3.1 ADD_INTERVAL_PARTITIONING Procedure

This procedure adds interval partitioning to an existing, non-partitioned, V1 or V2 immutable table.

Syntax

```
DBMS_IMMUTABLE_TABLE.ADD_INTERVAL_PARTITIONING(
    schema_name               IN    VARCHAR2,
    table_name                IN    VARCHAR2,
    interval_number           IN    NUMBER,
    interval_frequency        IN    VARCHAR2,
    first_high_timestamp      IN    TIMESTAMP);
```

Parameters

Table 99-2 ADD_INTERVAL_PARTITIONING Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the immutable table.
`interval_number`	Sets how often the database creates partitions for the immutable table.
`interval_frequency`	Sets the frequency for the value that was set in the `interval_number` setting. Supported values are `YEAR`, `MONTH`, `DAY`, `HOUR`, and `MINUTE`.
`first_high_timestamp`	A timestamp that determines the upper boundary of the first partition in the immutable table.

Usage Notes

- Composite partitioning (that is, sub-partitioning) is not supported with the above interval partitioning.

99.3.2 DELETE_EXPIRED_ROWS Procedure

This procedure deletes some or all of the expired rows from the immutable table. This procedure commits before deleting any expired rows and commits after deleting any expired rows.

Syntax

```
DBMS_IMMUTABLE_TABLE.DELETE_EXPIRED_ROWS(
   schema_name                 IN VARCHAR2,
   table_name                  IN VARCHAR2,
```

```
  before_timestamp             IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
  number_of_rows_deleted       OUT NUMBER);
```

Parameters

Table 99-3 DELETE_EXPIRED_ROWS Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the immutable table.
before_timestamp	If the parameter is NULL, all expired rows in the table are deleted. If the parameter is not NULL and older than the timestamp calculated based on current time and row retention time, rows with timestamps less than the parameter value are deleted. If the parameter is younger than the timestamp calculated based on the current time and row retention time, the calculated timestamp is used, and all expired rows are deleted. The default value is NULL.
number_of_rows_deleted	The number of rows deleted.

100
DBMS_INMEMORY

The DBMS_INMEMORY package provides an interface for In-Memory Column Store (IM column store) functionality.

This chapter contains the following topics:

- DBMS_INMEMORY Overview
- DBMS_INMEMORY Security Model
- Summary of DBMS_INMEMORY Subprograms

100.1 DBMS_INMEMORY Overview

This package contains procedures for populating and repopulating the IM column store, and for dropping IM expressions from a specified table.

IM Population and Repopulation

In-Memory population (population) occurs when the database reads existing row-format data from disk, transforms it into columnar format, and then stores it in the IM column store. Only objects with the INMEMORY attribute are eligible for population.

Population, which transforms existing data on disk into columnar format, is different from repopulation, which loads new data into the IM column store. Repopulation occurs automatically after their columnar data undergo significant DML activity.

When an object has the INMEMORY attribute and a priority other than NONE, the database gradually populates the object in the IM column store according to an internal priority queue. Objects with priority of NONE are populated only when they undergo a full scan.

The DBMS_INMEMORY.POPULATE procedure forces immediate population of an object. The DBMS_INMEMORY.REPOPULATE procedure forces immediate repopulation of an object.

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about IM population

IM Expressions

IM expressions populate frequently evaluated query expressions in the IM column store for subsequent reuse. An IM expression is materialized as a hidden virtual column, prefixed with the string SYS_IME, and is accessed in the same way as a non-virtual column.

When you use DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS, the database adds the 20 hottest expressions to their respective tables as SYS_IME columns and applies the default INMEMORY column compression clause. If any SYS_IME columns that were added during a previous invocation are no longer in the latest expression list, then the database changes their attribute to NO INMEMORY.

The maximum number of SYS_IME columns for a table, regardless of whether the attribute is INMEMORY or NO INMEMORY, is 50. After the 50 expression limit is reached for a table, the database will not add new SYS_IME columns. To make space for new expressions, you must manually drop SYS_IME columns with the DBMS_INMEMORY.IME_DROP_EXPRESSIONS or DBMS_INMEMORY_ADMIN.IME_DROP_ALL_EXPRESSIONS procedures.

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about IM expressions

100.2 DBMS_INMEMORY Security Model

The DBMS_INMEMORY package subprograms execute with invoker's rights.

The POPULATE and REPOPULATE procedures require the invoking user to have SELECT privileges on the specified object. For IME_DROP_EXPRESSIONS, the invoking user must have ALTER TABLE privileges on the specified table.

100.3 Summary of DBMS_INMEMORY Subprograms

This table lists and briefly describes the DBMS_INMEMORY package subprograms.

Table 100-1 DBMS_INMEMORY Package Subprograms

Subprogram	Description
IME_DROP_EXPRESSIONS Procedure	Drops a specified set of SYS_IME virtual columns from a table
POPULATE Procedure	Forces population of the specified table
REPOPULATE Procedure	Forces repopulation of the specified table
SEGMENT_DEALLOCATE_VERSIONS Procedure	Deallocates non-current IMCUs in the IM column store

100.3.1 IME_DROP_EXPRESSIONS Procedure

This procedure drops a specified set of SYS_IME virtual columns from a table.

Syntax

```
DBMS_INMEMORY.IME_DROP_EXPRESSIONS(
   schema_name     IN    VARCHAR2,
   table_name      IN    VARCHAR2,
   column_name     IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 100-2 IME_DROP_EXPRESSIONS Procedure Parameters

Parameter	Description
schema_name	The name of the schema that contains the In-Memory table
table_name	The name of the In-Memory table that contains the SYS_IME columns
column_name	The name of the SYS_IME column. By default this value is null, which specifies all SYS_IME columns in this table.

Usage Notes

Typical reasons for dropping SYS_IME columns are space and performance. The maximum number of SYS_IME columns for a table, regardless of whether the attribute is INMEMORY or NO INMEMORY, is 50. After the 50-expression limit is reached for a table, the database will not add new SYS_IME columns. To make space for new expressions, you must manually drop SYS_IME columns with the DBMS_INMEMORY.IME_DROP_EXPRESSIONS or DBMS_INMEMORY_ADMIN.IME_DROP_ALL_EXPRESSIONS procedures.

To drop a specified SYS_IME column or all SYS_IME columns in the requested table, use DBMS_INMEMORY.IME_DROP_EXPRESSIONS . To populate these segments again, either invoke the DBMS_INMEMORY.POPULATE procedure, or perform a full table scan.

100.3.2 POPULATE Procedure

This procedure forces population of the specified table, partition, or subpartition into the IM column store.

Syntax

```
DBMS_INMEMORY.POPULATE(
   schema_name      IN    VARCHAR2,
   table_name       IN    VARCHAR2,
   subobject_name   IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 100-3 *POPULATE Procedure Parameters*

Parameter	Description
schema_name	Name of schema
table_name	Name of table
subobject_name	Partition or subpartition

100.3.3 REPOPULATE Procedure

This procedure forces repopulation of a table, partition, or subpartition that is currently populated in the IM column store.

Syntax

```
DBMS_INMEMORY.REPOPULATE(
   schema_name      IN    VARCHAR2,
   table_name       IN    VARCHAR2,
   subobject_name   IN    VARCHAR2 DEFAULT NULL,
   force            IN    BOOLEAN DEFAULT FALSE);
```

Parameters

Table 100-4 REPOPULATE Procedure Parameters

Parameter	Description
schema_name	Name of the schema that owns the object.
table_name	Name of the table requiring repopulation.
subobject_name	Name of the partition or subpartition. If null, then repopulate the entire table.
force	Whether to repopulate all IMCUs in the segment, just as in initial population. The following values are possible for the force parameter: • FALSE — The database repopulates only IMCUs containing modified rows. This is the default. • TRUE — The database drops the segment, and then rebuilds it. The database increments the statistics and performs all other tasks related to initial population. For example, IMCU 1 contains rows 1 to 500,000, and IMCU 2 contains rows 500,001 to 1,000,000. A statement modifies row 600,000. When force is FALSE, the database only repopulates IMCU 2. When force is TRUE, the database repopulates both IMCUs. Consider further that the INMEMORY_VIRTUAL_COLUMNS initialization parameter is set to ENABLE, and an application creates a new virtual column. When force is FALSE, the database only repopulates IMCU 2 with the new column. When force is TRUE, the database repopulates both IMCUs with the new column.

100.3.4 SEGMENT_DEALLOCATE_VERSIONS Procedure

This procedure deallocates non-current IMCUs in the IM column store.

Syntax

```
DBMS_INMEMORY.SEGMENT_DEALLOCATE_VERSIONS(
   SCHEMA_NAME      IN  VARCHAR2,
   TABLE_NAME       IN  VARCHAR2,
   PARTITION_NAME   IN  VARCHAR2 DEFAULT NULL,
   SPCPRESSURE      IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 100-5 SEGMENT_DEALLOCATE_VERSIONS Procedure Parameters

Parameter	Description
schema_name	Name of the schema that owns the object.
table_name	Name of the table requiring repopulation.
partition_name	Name of the partition or subpartition. If null, then repopulate the entire table.
spcpressure	Whether to force deallocation of non-current IMCUs (TRUE), or wait for the database to deallocate them automatically. By default, the database deallocates non-current IMCUs every two minutes.

Usage Notes

During repopulation, the IM column store maintains both the current IMCU and non-current IMCU. This mechanism, which is called *double buffering*, ensures that queries do not decrease performance because an IMCU is unavailable during repopulation. After repopulation completes, the IM column store retains the non-current IMCU for a short time (2 minutes by default) to optimize queries with older SCNs. Typically, the default behavior is sufficient. However, you can force deallocation of non-current IMCUs by using the SEGMENT_DEALLOCATE_VERSIONS procedure.

Example

The following program forces deallocation of non-current IMCUs for the products table:

```
BEGIN
  DBMS_INMEMORY.SEGMENT_DEALLOCATE_VERSIONS(
    schema_name => 'SH'
  , table_name  => 'PRODUCTS'
  , spcpressure => TRUE );
END;
```

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about double buffering in the IM column store

101
DBMS_INMEMORY_ADMIN

DBMS_INMEMORY_ADMIN provides interfaces for managing an In-Memory FastStart (IM FastStart) area and In-Memory Expressions (IM expressions).

This chapter contains the following topics:

- DBMS_INMEMORY_ADMIN Overview
- DBMS_INMEMORY_ADMIN Security Model
- DBMS_INMEMORY_ADMIN Operational Notes
- Summary of DBMS_INMEMORY_ADMIN Subprograms

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about Oracle Database In-Memory features

101.1 DBMS_INMEMORY_ADMIN Overview

This package provides interfaces for managing In-Memory Expressions (IM expressions) and the In-Memory FastStart (IM FastStart) area.

IM Expressions

Analytic queries often contain complex expressions or calculations that consume significant CPU and memory during execution. Use IME_CAPTURE_EXPRESSIONS to identify these frequently used ("hot") expressions and IME_POPULATE_EXPRESSIONS to populate them in the IM column store. By using IM expressions, the database avoids repeated computations and improves performance.

The database represents IM expressions as system-generated virtual columns. The name of an IM virtual column begins with SYS_IME. You can also use DBMS_INMEMORY_ADMIN.IME_DROP_ALL_EXPRESSIONS and DBMS_INMEMORY.IME_DROP_EXPRESSIONS to remove existing SYS_IME columns.

The DBA_IM_EXPRESSIONS view shows the SYS_IME columns that have the INMEMORY attribute. After using the IME_CAPTURE_EXPRESSIONS procedure, you can query this view to see the hot expressions added to different tables in the database.

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about IM expressions

IM FastStart Area

The IM FastStart area stores data that optimizes the population of the IM column store when the database restarts. Because the database reads columnar data directly from persistent storage without needing to compress or format it, population is faster when a database instance restarts.

When you enable IM FastStart for the IM column store, you must specify an ASSM tablespace for the IM FastStart area. The tablespace stores the data in a SecureFiles LOB named `SYSDBIMFS_LOGSEG$`. The `SYSAUX` tablespace stores the metadata. When data is populated or repopulated in the IM column store, the database automatically writes the data to the IM FastStart area. You cannot manually force a write. If you specify an object as `NO INMEMORY`, then the database removes it from the IM FastStart area.

When the IM FastStart area is under space pressure, the database automatically drops the oldest 15% of segments and continues saving columnar data. If space is unavailable, then the database stops writing to the IM FastStart area.

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about IM expressions

Automatic In-Memory

Automatic In-Memory uses access tracking and column statistics to manage objects in the IM column store. If the IM column store is full, and if other more frequently accessed segments would benefit from population in the IM column store, then the IM column store evicts inactive segments. If the IM column store is configured to hold all `INMEMORY` segments, however, then Automatic In-Memory takes no action.

By default, Automatic In-Memory checks usage statistics for the past 31 days. You can change the current setting by supplying the `AIM_STATWINDOW_DAYS` parameter to `DBMS_INMEMORY_ADMIN.AIM_SET_PARAMETER`.

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about Automatic In-Memory

Database In-Memory Wait on Populate

The `POPULATE_WAIT` function initiates population of all `INMEMORY` objects that have a priority greater than or equal to the specified priority, and then returns a status value for the population. A user-specified interval specifies the maximum time that the function waits before returning the value to the caller.

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn more about the wait on populate feature

101.2 DBMS_INMEMORY_ADMIN Security Model

This package requires administrator privileges. Package subprograms execute with invoker's rights.

101.3 DBMS_INMEMORY_ADMIN Operational Notes

It is possible for a `DBMS_INMEMORY_ADMIN` FastStart operation to fail or be interrupted.

In a failure or interruption scenario, the following rules determine which subprograms you can use:

- If `FASTSTART_ENABLE` does not succeed, then the only permitted operation is re-executing `FASTSTART_ENABLE`.
- If `FASTSTART_MIGRATE_STORAGE` does not succeed, then the only permitted operation is re-executing `FASTSTART_MIGRATE_STORAGE`.
- If `FASTSTART_DISABLE` does not succeed, then all `DBMS_INMEMORY_ADMIN` operations are permitted.

101.4 Summary of DBMS_INMEMORY_ADMIN Subprograms

This table lists the `DBMS_INMEMORY_ADMIN` subprograms and briefly describes them.

Table 101-1 DBMS_INMEMORY_ADMIN Package Subprograms

Subprogram	Description
AIM_GET_PARAMETER Procedure	This procedure obtains the current values for parameters that control Automatic In-Memory.
AIM_SET_PARAMETER Procedure	The procedure customizes the execution environment of Automatic In-Memory
FASTSTART_DISABLE Procedure	This procedure disables the In-Memory FastStart (IM FastStart) feature.
FASTSTART_ENABLE Procedure	This procedure enables IM FastStart and assigns a tablespace.
FASTSTART_MIGRATE_STORAGE Procedure	This procedure moves all IM FastStart data and metadata from the existing tablespace to the specified new tablespace.
GET_FASTSTART_TABLESPACE Function	This function returns the name of the tablespace that is currently designated for IM FastStart.
IME_CAPTURE_EXPRESSIONS Procedure	This procedure captures the 20 most frequently accessed ("hottest") expressions in the database in the specified time interval.
IME_CLOSE_CAPTURE_WINDOW Procedure	This procedure signals the end of the current expression capture window.
IME_DROP_ALL_EXPRESSIONS Procedure	This procedure drops all `SYS_IME` expression virtual columns in the database.
IME_GET_CAPTURE_STATE Procedure	This procedure returns the current capture state of the expression capture window and the timestamp of the most recent modification.

Table 101-1 (Cont.) DBMS_INMEMORY_ADMIN Package Subprograms

Subprogram	Description
IME_OPEN_CAPTURE_WINDOW Procedure	This procedure signals the beginning of an expression capture window.
IME_POPULATE_EXPRESSIONS Procedure	This procedure forces the population of expressions captured in the latest invocation of `DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS`.
POPULATE_WAIT Function	Initiates population of all `INMEMORY` objects that have a priority greater than or equal to the specified priority, and sets a timeout interval within which population must occur

101.4.1 AIM_GET_PARAMETER Procedure

This procedure obtains the current values for parameters that control Automatic In-Memory.

Syntax

```
DBMS_INMEMORY_ADMIN.AIM_GET_PARAMETER(
   parameter    IN    NUMBER,
   value        OUT   NUMBER);
```

Parameters

Parameter	Description
`parameter`	Specifies a predefined constant that controls Automatic In-Memory. The only valid constant is `AIM_STATWINDOW_DAYS`, which specifies the number of days in the sliding statistics window. Automatic In-Memory uses this duration to filter statistics for `INMEMORY` objects as part of its algorithms. For example, if the duration is set to 7 days, then Automatic In-Memory considers only statistics of the past 7 days for its algorithms. The default is `1`.
`value`	Specifies the value assigned to `AIM_STATWINDOW_DAYS`.

Example 101-1 Getting the Number of Days in the Statistics Window

The following code prints the number of days in the statistics window to the screen:

```
VARIABLE b_statwin NUMBER

BEGIN

DBMS_INMEMORY_ADMIN.AIM_GET_PARAMETER(DBMS_INMEMORY_ADMIN.AIM_STATWINDOW_DAYS,
 :b_statwin);
END;
/

PRINT b_statwin
```

Sample output appears below:

```
B_STATWIN
---------
       14
```

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn how to use `AIM_GET_PARAMETER`

101.4.2 AIM_SET_PARAMETER Procedure

The procedure customizes the execution environment of Automatic In-Memory.

Syntax

```
DBMS_INMEMORY_ADMIN.AIM_SET_PARAMETER(
   parameter    IN    NUMBER,
   value        IN    NUMBER);
```

Parameters

Parameter	Description
`parameter`	Specifies a predefined constant that controls Automatic In-Memory. The only valid constant is `AIM_STATWINDOW_DAYS`, which specifies the number of days in the sliding statistics window. The default is `1`.
`value`	Assigns the value assigned to `AIM_STATWINDOW_DAYS`.

Example 101-2 Setting the Number of Days in the Statistics Window

The following example gets the current number of days in the window, sets it to 14, and then prints the value to the screen:

```
VARIABLE b_statwin NUMBER

BEGIN

DBMS_INMEMORY_ADMIN.AIM_GET_PARAMETER(DBMS_INMEMORY_ADMIN.AIM_STATWINDOW_DAYS,
 :b_statwin);
END;
/

PRINT b_statwin

BEGIN

DBMS_INMEMORY_ADMIN.AIM_SET_PARAMETER(DBMS_INMEMORY_ADMIN.AIM_STATWINDOW_DAYS,
 14);
END;
/
```

```
BEGIN

DBMS_INMEMORY_ADMIN.AIM_GET_PARAMETER(DBMS_INMEMORY_ADMIN.AIM_STATWINDOW_DAYS,
 :b_statwin);
END;
/

PRINT b_statwin
```

Sample output appears below:

```
 B_STATWIN
----------
         1

 B_STATWIN
----------
        14
```

> **See Also:**
>
> *Oracle Database In-Memory Guide* to learn how to use `AIM_GET_PARAMETER`

101.4.3 FASTSTART_DISABLE Procedure

This procedure disables the In-Memory FastStart (IM FastStart) feature.

Syntax

```
DBMS_INMEMORY_ADMIN.FASTSTART_DISABLE();
```

Security Model

Administrator privileges are required to execute this procedure.

Usage Notes

When you execute the procedure, the database executes the following actions:

1. Waits until all IM FastStart operations complete
2. Disables the IM FastStart feature, and performs the following operations:
   - Cleans the IM FastStart area
   - Deletes IM FastStart metadata stored in the `SYSAUX` tablespace
   - Releases the IM FastStart tablespace (but does not delete it)

This procedure does not interrupt or affect any concurrent IM column store operations.

Examples

The following PL/SQL program disables the IM FastStart feature:

```
EXEC DBMS_INMEMORY_ADMIN.FASTSTART_DISABLE;
```

The following query shows that the LOB for the IM FastStart tablespace has been deleted (sample output included):

```
COL OWNER FORMAT a5
COL SEGMENT_NAME FORMAT a20
SELECT   l.OWNER, l.SEGMENT_NAME, SUM(s.BYTES)/1024/1024 MB
FROM     DBA_LOBS l, DBA_SEGMENTS s
WHERE    l.SEGMENT_NAME = s.SEGMENT_NAME
AND      l.TABLESPACE_NAME = 'FS_TBS'
GROUP BY l.OWNER, l.SEGMENT_NAME;

no rows selected
```

101.4.4 FASTSTART_ENABLE Procedure

This procedure enables In-Memory FastStart (IM FastStart), and designates a tablespace for the IM FastStart (FastStart) area.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

Syntax

```
DBMS_INMEMORY_ADMIN.FASTSTART_ENABLE(
   tbs_name    IN    VARCHAR2,
   nologging   IN    BOOLEAN DEFAULT TRUE);
```

Parameters

Table 101-2 FASTSTART_ENABLE Procedure Parameters

Parameter	Description
`tbs_name`	The name of the ASSM tablespace for the FastStart area.
`nologging`	The logging mode of the LOB created for the FastStart area. If the `nologging` parameter is set to `FALSE`, then the database creates the FastStart LOB with the `LOGGING` option. If set to `TRUE` (default), then the database creates the LOB with the `NOLOGGING` option.

Security Model

Administrator privileges are required to execute this procedure.

Usage Notes

To enable IM FastStart, the ASSM tablespace specified in `FASTSTART_ENABLE` must exist, and the `SYSAUX` tablespace must be online. Only one FastStart tablespace can exist for every PDB or non-CDB. The specified tablespace must have enough space to store data for the IM column store, and it must not contain any other data before it is designated for the FastStart area. Oracle recommends sizing the tablespace at least twice of the size of the `INMEMORY_SIZE` initialization parameter.

The database does not create the FastStart area on disk until the IM column store is populated. After population, the data periodically saves the columnar data (but not metadata such as the transaction journal) to the FastStart area, which is represented on disk as the `SYSDBIMFS_LOBSEG$` segment. The database stores the FastStart metadata in the `SYSAUX` tablespace. In an Oracle Real Application Clusters (Oracle RAC) environment, IM FastStart data is shared across all nodes.

> **Note:**
>
> IM FastStart is not supported in a standby database instance.

Whereas the initial loading of IMCUs into memory is expensive and CPU-bound, an IM FastStart tablespace requires intermittent I/O. The database periodically writes columnar data to the IM FastStart area. If a database instance must restart, then Oracle Database reads the columnar data directly from the IM FastStart area rather than reconstructing the IMCUs from scratch. No compression or formatting of the columnar data is required.

Examples

This example creates `fs_tbs` as an ASSM tablespace, and then uses `FASTSTART_ENABLE` to specify this tablespace as the IM FastStart area:

```
CREATE TABLESPACE fs_tbs
  DATAFILE 'fs_tbs.dbf' SIZE 500M
  EXTENT MANAGEMENT LOCAL
  SEGMENT SPACE MANAGEMENT AUTO;

EXEC DBMS_INMEMORY_ADMIN.FASTSTART_ENABLE('fs_tbs');
```

The following query shows that the IM FastStart LOB was created (sample output included):

```
COL OWNER FORMAT a5
COL SEGMENT_NAME FORMAT a20
SELECT   l.OWNER, l.SEGMENT_NAME, SUM(s.BYTES)/1024/1024 MB
FROM     DBA_LOBS l, DBA_SEGMENTS s
WHERE    l.SEGMENT_NAME = s.SEGMENT_NAME
AND      l.TABLESPACE_NAME = 'FS_TBS'
GROUP BY l.OWNER, l.SEGMENT_NAME;
```

```
OWNER SEGMENT_NAME         MB
----- -------------------- ----------
SYS   SYSDBIMFS_LOBSEG$    .125
```

101.4.5 FASTSTART_MIGRATE_STORAGE Procedure

This procedure moves the In-Memory FastStart (IM FastStart) data and catalogs from the current tablespace to a new tablespace.

Syntax

```
DBMS_INMEMORY_ADMIN.FASTSTART_MIGRATE_STORAGE(
   tbs_name    IN    VARCHAR2 );
```

Parameters

Table 101-3 FASTSTART_MIGRATE_STORAGE Procedure Parameters

Parameter	Description
tbs_name	The name of the new ASSM tablespace for the IM FastStart area.

Security Model

DBA privileges are required to execute this procedure.

Usage Notes

When you execute the procedure, the database executes the following actions:

1. Waits until all IM FastStart operations complete
2. Disables the IM FastStart feature
3. Copies IM FastStart data and metadata to the new tablespace, leaving the old tablespace intact
4. Re-enables IM FastStart the feature

Examples

The following program obtains the name of the IM FastStart tablespace, if one exists, and prints the result (sample output included):

```
VARIABLE b_fstbs VARCHAR2(20)
BEGIN
   :b_fstbs := DBMS_INMEMORY_ADMIN.GET_FASTSTART_TABLESPACE;
END;
/
PRINT b_fstbs

B_FSTBS
----------------------------
FS_TBS
```

The following statements create a new tablespace named `fs_tbs2`, and then migrate the IM FastStart area to this tablespace:

```
CREATE TABLESPACE fs_tbs2
  DATAFILE 'fs_tbs2.dbf' SIZE 500M
  EXTENT MANAGEMENT LOCAL
  SEGMENT SPACE MANAGEMENT AUTO;

EXEC DBMS_INMEMORY_ADMIN.FASTSTART_MIGRATE_STORAGE('fs_tbs2');
```

The following program prints the name of the current IM FastStart tablespace (sample output included):

```
BEGIN
   :b_fstbs := DBMS_INMEMORY_ADMIN.GET_FASTSTART_TABLESPACE;
END;
/
PRINT b_fstbs

B_FSTBS
----------------------------
FS_TBS2
```

101.4.6 GET_FASTSTART_TABLESPACE Function

This function returns the tablespace assigned to In-Memory FastStart (IM FastStart). If the feature is disabled, then the function returns `NOT ENABLED`.

Syntax

```
DBMS_INMEMORY_ADMIN.GET_FASTSTART_TABLESPACE();
```

Security Model

DBA privileges are required to execute this function.

Examples

This program obtains the name of the IM FastStart tablespace, if one exists, and prints the result:

```
VARIABLE b_fstbs VARCHAR2(20)
BEGIN
   :b_fstbs := DBMS_INMEMORY_ADMIN.GET_FASTSTART_TABLESPACE;
END;
/
PRINT b_fstbs

B_FSTBS
----------------------------
NOT ENABLED
```

101.4.7 IME_CAPTURE_EXPRESSIONS Procedure

This procedure captures the 20 most frequently accessed (“hottest”) expressions in the database in the specified time interval.

Syntax

```
DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS(
   snapshot     IN     VARCHAR2);
```

Parameters

Table 101-4 IME_CAPTURE_EXPRESSIONS Procedure Parameters

Parameter	Description
snapshot	Specifies a snapshot that defines the time interval in which expression statistics are considered. You can specify any of the following values: • CUMULATIVE The database considers all expression statistics since the creation of the database. • CURRENT The database considers only expression statistics from the past 24 hours. • WINDOW The database considers statistics for expressions tracked in the most recent expression capture window. The database adds hidden virtual columns for expressions tracked in the most recent window. If the capture window is currently open, then the database considers all expressions tracked in the current window up until this point, and then materializes the hottest expressions. To list the expressions that have been tracked in the current window, query DBA_EXPRESSION_STATISTICS with SNAPSHOT='WINDOW'.

Usage Notes

When you invoke this procedure, the database queries the Expression Statistics Store (ESS), and considers only expressions on tables that are at least partially populated in the IM column store. The database adds the 20 hottest expressions to their respective tables as hidden virtual columns, prefixed with the string SYS_IME, and applies the default INMEMORY column compression clause. If any SYS_IME columns added during a previous invocation are no longer in the latest top 20 list, then the database marks them as NO INMEMORY.

> **Note:**
>
> Executing the IME_CAPTURE_EXPRESSIONS procedure on a standby database has no effect.

The maximum number of SYS_IME columns for a table, regardless of whether the attribute is INMEMORY or NO INMEMORY, is 50. After the limit is reached for a table, the database will not add new SYS_IME columns. To make space for new expressions, you must manually drop SYS_IME

columns with the IME_DROP_ALL_EXPRESSIONS Procedure or IME_DROP_EXPRESSIONS Procedure.

The 50-expression limit for each table, which includes both `INMEMORY` and `NO INMEMORY` expressions, is different from the 20-expression limit for the database, which includes only `INMEMORY` expressions. For example, if 20 tables are populated in the IM column store, then each table might each have 1 `SYS_IME` column with the `INMEMORY` attribute, and 49 `SYS_IME` columns with the `NO INMEMORY` attribute.

IM expressions and virtual columns are stored in In-Memory structured called In-Memory Expression Units (IMEUs). Every IMEU is linked to a parent In-Memory Compression Unit (IMCU) from which it inherits compression characteristics.

ESS information is stored in the data dictionary and exposed in the `DBA_EXPRESSION_STATISTICS` view. This view shows the metadata that the optimizer has collected in the ESS. IM expressions are exposed as system-generated virtual columns, prefixed by the string `SYS_IME`, in the `DBA_IM_EXPRESSIONS` view.

Example 101-3 Capturing Expressions in a User-Defined Window

This example demonstrates use of the `WINDOW` capture mode. Your goal is to open and close an expression capture window, and then capture all expressions that the database tracked during this window. You perform the following steps:

1. Open an expression capture window, generate expressions, and then close the window:

```
EXEC DBMS_INMEMORY_ADMIN.IME_OPEN_CAPTURE_WINDOW();
-- Generate expressions for the database to track
EXEC DBMS_INMEMORY_ADMIN.IME_CLOSE_CAPTURE_WINDOW();
```

2. Query `DBA_EXPRESSION_STATICS` (sample output included):

```
COL OWNER FORMAT A6
COL TABLE_NAME FORMAT A9
COL COUNT FORMAT 99999
COL CREATED FORMAT A10
COL EXPRESSION_TEXT FORMAT A29

SELECT OWNER, TABLE_NAME, EVALUATION_COUNT AS COUNT,
       CREATED, EXPRESSION_TEXT
FROM   DBA_EXPRESSION_STATISTICS
WHERE  SNAPSHOT = 'WINDOW'
AND    OWNER = 'SH';

OWNER  TABLE_NAM  COUNT CREATED    EXPRESSION_TEXT
------ --------- ------ ---------- ------------------------
SH     SALES       4702 09-OCT-17  "QUANTITY_SOLD"
SH     SALES       4702 09-OCT-17  "QUANTITY_SOLD"*"AMOUNT_SOLD"
SH     SALES       4702 09-OCT-17  "PROD_ID"
SH     SALES       4702 09-OCT-17  "CUST_ID"
SH     SALES       4702 09-OCT-17  "CHANNEL_ID"
SH     SALES       4702 09-OCT-17  "AMOUNT_SOLD"
```

The preceding query shows both the columns tracked in the ESS and the expressions captured during the window for queries in the `sh` schema. During the most recent window, the database captured one expression: `QUANTITY_SOLD*AMOUNT_SOLD`.

3. Use `IME_CAPTURE_EXPRESSIONS` to make the database consider all expressions in the current window for materialization:

```
EXEC DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS('WINDOW');
```

4. Query `DBA_IM_EXPRESSIONS` (sample output included):

```
COL OWNER FORMAT a6
COL TABLE_NAME FORMAT a9
COL COLUMN_NAME FORMAT a25
SET LONG 50
SET LINESIZE 150

SELECT OWNER, TABLE_NAME, COLUMN_NAME, SQL_EXPRESSION
FROM   DBA_IM_EXPRESSIONS;

OWNER  TABLE_NAM COLUMN_NAME               SQL_EXPRESSION
------ --------- ------------------------- ------------------------------
SH     SALES     SYS_IME000100000025201B   "QUANTITY_SOLD"*"AMOUNT_SOLD"
```

   The preceding output shows all virtual columns that were added to the table and marked `INMEMORY` as part of the latest `IME_CAPTURE_EXPRESSIONS` invocation. The database gradually populates the captured expressions into the IM column store when it repopulates different IMCUs of the table.

5. Execute the following procedure to explicitly force a population of all captured IM expressions:

```
EXEC DBMS_INMEMORY_ADMIN.IME_POPULATE_EXPRESSIONS();
```

   Note that you can populate IM expressions from a specific table by executing the `DBMS_INMEMORY.REPOPULATE` procedure with the `force` parameter set to `TRUE`.

Example 101-4 Capturing Expressions for the Past Day

The following program captures expressions tracked during the last 24 hours:

```
EXEC DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS('CURRENT');
```

101.4.8 IME_CLOSE_CAPTURE_WINDOW Procedure

This procedure signals the end of the current expression capture window.

Syntax

```
DBMS_INMEMORY_ADMIN.IME_CLOSE_CAPTURE_WINDOW();
```

Usage Notes

On invocation of this procedure, the optimizer saves all gathered statistics to disk, and essentially freezes the expressions tracked in the window. The database preserves the statistics captured in this window until a new expression capture window is opened, at which point the database purges the statistics captured in the previous window.

Example 101-5 Example

This example opens an expression capture window, and then issues `IME_CAPTURE_EXPRESSIONS('WINDOW')` so that the database considers all expressions in the current window for materialization. Finally, the example closes the window.

```
EXEC DBMS_INMEMORY_ADMIN.IME_OPEN_CAPTURE_WINDOW();
-- Generate expressions so that the database can track them
EXEC DBMS_INMEMORY_ADMIN.IME_CLOSE_CAPTURE_WINDOW();
EXEC DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS('WINDOW');
```

101.4.9 IME_DROP_ALL_EXPRESSIONS Procedure

This procedure drops all `SYS_IME` expression virtual columns in the database.

Syntax

```
DBMS_INMEMORY_ADMIN.IME_DROP_ALL_EXPRESSIONS();
```

Usage Notes

The `IME_DROP_ALL_EXPRESSIONS` procedure drops all `SYS_IME` columns from all tables, regardless of whether they have the `INMEMORY` attribute. In effect, the procedure acts as a database-wide reset button.

Using `IME_DROP_ALL_EXPRESSIONS` triggers a drop of all IMEUs *and* IMCUs for segments that have `SYS_IME` columns. For example, if 50 populated tables have one `SYS_IME` column each, then `IME_DROP_ALL_EXPRESSIONS` removes all 50 tables from the IM column store. To populate these segments again, you must use the `DBMS_INMEMORY.POPULATE` procedure or perform a full table scan.

101.4.10 IME_GET_CAPTURE_STATE Procedure

This procedure returns the current capture state of the expression capture window and the timestamp of the most recent modification.

Syntax

```
DBMS_INMEMORY_ADMIN.IME_GET_CAPTURE_STATE(
   p_capture_state  OUT  VARCHAR2,
   p_last_modified  OUT  TIMESTAMP);
```

Parameters

Parameter	Description
`p_capture_state`	Describes the current state of the expression capture window. The following states are possible: • `OPEN` — Indicates that the window is open. • `CLOSED` — Indicates that the window is closed. • `DEFAULT` — Indicates that the window has not been used. It is equivalent to the `CLOSED` state.
`p_last_modified`	Indicates the timestamp of the most recent action.

Usage Notes

This procedure is useful for avoiding conflicting calls for IME_OPEN_CAPTURE_WINDOW Procedure and IME_CLOSE_CAPTURE_WINDOW Procedure. For example, if the current expression capture window state is `OPEN`, then you cannot open another window, and if the window state is `CLOSED`, then you cannot close a window.

Example 101-6 Determining the State of an Expression Capture Window

This example opens an expression capture window, and then determines its capture state.

```
EXEC DBMS_INMEMORY_ADMIN.IME_OPEN_CAPTURE_WINDOW();

VARIABLE b_state VARCHAR2(25)
VARIABLE b_time  VARCHAR2(10)
EXECUTE DBMS_INMEMORY_ADMIN.IME_GET_CAPTURE_STATE(:b_state, :b_time)
PRINT b_state b_time
```

The following sample output indicates that an expression capture window is currently open:

```
B_STATE
-------------------------------------------------
OPEN

B_TIME
-------------------------------------------------
09-OCT-17
```

101.4.11 IME_OPEN_CAPTURE_WINDOW Procedure

This procedure signals the beginning of an expression capture window.

Syntax

```
DBMS_INMEMORY_ADMIN.IME_OPEN_CAPTURE_WINDOW();
```

Security Model

Administrator privileges are required to execute this procedure.

Usage Notes

On invocation of this procedure, the optimizer begins a new window snapshot and starts tracking expressions that occur within this window. An expression capture window is global across all instances in an Oracle RAC database.

Conflicting actions are not permitted. For example, in an Oracle RAC database, opening expression capture window on instance 1 at time t0 and attempting to open another expression capture window on instance 2 at time t1 before closing the first window is a conflicting action. To obtain the current capture state and reduce the potential for conflicting procedure invocations, use the IME_GET_CAPTURE_STATE Procedure.

Example

This following program opens an expression capture window:

```
EXEC DBMS_INMEMORY_ADMIN.IME_OPEN_CAPTURE_WINDOW();
```

101.4.12 IME_POPULATE_EXPRESSIONS Procedure

This procedure forces the population of expressions captured in the latest invocation of `DBMS_INMEMORY_ADMIN.IME_CAPTURE_EXPRESSIONS`.

Syntax

```
DBMS_INMEMORY_ADMIN.IME_POPULATE_EXPRESSIONS();
```

Usage Notes

If you do not invoke this procedure, then the database gradually repopulates `SYS_IME` columns when their parent IMCUs are repopulated. If a table is not repopulated, then any new `SYS_IME` columns captured by the `IME_CAPTURE_EXPRESSIONS` procedure are not populated. `IME_POPULATE_EXPRESSIONS` solves this problem by forcing population.

Internally, the procedure invokes `DBMS_INMEMORY.REPOPULATE` for all tables that have `SYS_IME` columns with the `INMEMORY` attribute. To populate `SYS_IME` columns in a specified subset of tables, use `DBMS_INMEMORY.REPOPULATE` instead of `IME_POPULATE_EXPRESSIONS`.

101.4.13 POPULATE_WAIT Function

This function initiates population of all `INMEMORY` objects that have a priority greater than or equal to the specified priority, and returns a status value for the population. A user-specified interval specifies the maximum time that the function waits before returning the value to the caller.

Syntax

```
DBMS_INMEMORY_ADMIN.POPULATE_WAIT(
   priority    IN    VARCHAR2 DEFAULT 'LOW',
   percentage  IN    NUMBER   DEFAULT 100,
   timeout     IN    NUMBER   DEFAULT 99999999,
   force       IN    BOOLEAN DEFAULT FALSE)
RETURN VARCHAR2;
```

Parameters

Table 101-5 POPULATE_WAIT Function Parameters

Parameter	Description
`priority`	Specifies that the database populate all `INMEMORY` objects with the specified priority setting or higher. The default priority is `LOW`. `NONE` is considered lowest priority. If you set to `priority` to `NONE`, then this function waits for all `INMEMORY` objects to populate.

Table 101-5 (Cont.) POPULATE_WAIT Function Parameters

Parameter	Description
percentage	Specifies the percentage of population required for the function to consider population to be complete. The default is 100. For example, if percentage is 50 and priority is NONE, and if 50% of the INMEMORY objects are populated in the IM column store, then the function returns the value 0 (population successful).
timeout	Specifies the number of seconds that must pass before the function returns -1, which indicates that the populate operation timed out. The default is 99999999 seconds, which is 115.74 days. Assume that timeout is 600, priority is LOW, and percentage is 100. If 10 minutes pass, but all PRIORITY LOW objects are not yet fully populated, then the function returns -1.
force	Specifies that the database should drop all INMEMORY segments that have a priority greater than or equal than the specified priority, and then repopulate these segments. The default is FALSE. Assume that the INMEMORY attribute applies to the sales table, which is partitioned. Only half the sales partitions are currently populated in the IM column store. If you execute POPULATE_WAIT with force set to TRUE, then the database drops all sales segments, and then repopulates them.

Return Values

The following table describes the possible return values for POPULATE_WAIT. The function returns the values 0, 1, 2, and 3 only if the condition is met before the end of the interval specified by timeout. For example, if timeout is 600, then the function returns 1 only if an out-of-memory error occurs before 600 seconds pass. The function returns -1 only if the end of the timeout interval occurs *before* the database completes the requested operation.

Table 101-6 Return Values for POPULATE_WAIT

Constant	Value	Description
POPULATE_TIMEOUT	-1	The function timed out while waiting for population to complete. Existing population jobs continue running in the background after -1 is returned. Reissuing POPULATE_TIMEOUT after -1 is returned reinitiates population; segments that are already populated are not dropped.
POPULATE_SUCCESS	0	All objects that met the priority criteria were populated to the specified percentage of completion.
POPULATE_OUT_OF_MEMORY	1	The In-Memory pool had insufficient memory to populate the objects that met the priority criteria to the specified percentage of completion.
POPULATE_NO_INMEMORY_OBJECTS	2	No INMEMORY objects met the specified priority criteria.

Table 101-6 (Cont.) Return Values for POPULATE_WAIT

Constant	Value	Description
POPULATE_INMEMORY_SIZE_ZERO	3	The In-Memory column store is not enabled.

Usage Notes

Sample use cases for ensuring that objects are populated include:

- When the database is closed, open the database with `STARTUP RESTRICT` so that only administrators can access the database, and then execute `POPULATE_WAIT` with the desired timeout setting. If `POPULATE_WAIT` returns `-1`, indicating a timeout, then reexecute `POPULATE_WAIT`. When the function returns `0`, disable the restricted session so that non-administrative users can query the database.
- Block database connections by using services or an application tier technique. When no analytic indexes exists, and when the application depends on the IM column store to provide reasonable performance, these techniques prevent runaway queries.

Example 101-7 Specifying a Timeout Interval for In-Memory Population

In this example, the database contains a number of In-Memory tables with a variety of priority settings. Your goal is to populate *every* In-Memory table to 100% completion in a restricted database session, and then disable the restricted session so that the application can be guaranteed of querying only the In-Memory representations.

Assume that the database is shut down. In SQL*Plus, you connect to an idle instance as `SYSDBA`, and then execute the following command (sample output included):

```
SQL> STARTUP RESTRICT
ORACLE instance started.

Total System Global Area 1157624280 bytes
Fixed Size                  8839640 bytes
Variable Size             754974720 bytes
Database Buffers           16777216 bytes
Redo Buffers                7933952 bytes
In-Memory Area            369098752 bytes
Database mounted.
Database opened.
```

The database is open, but is accessible only to administrative users. You execute the following statements in SQL*Plus (sample output shown in bold):

```
VARIABLE b_pop_status NUMBER

SELECT DBMS_INMEMORY_ADMIN.POPULATE_WAIT(
         priority   => 'NONE' ,
         percentage => 100    ,
         timeout    => 300    )
  INTO b_pop_status
FROM   DUAL;
```

```
PRINT b_pop_status
-1
```

After 5 minutes, the function returns the number `-1`. This code indicates that the function timed out while waiting for population to complete. 5 minutes is not long enough to populate all `INMEMORY` tables. You re-execute the `SELECT` statement, specifying a 30-minute timeout:

```
SELECT DBMS_INMEMORY_ADMIN.POPULATE_WAIT(
         priority   => 'NONE' ,
         percentage => 100    ,
         timeout    => 1800   )
  INTO b_pop_status
FROM   DUAL;

PRINT b_pop_status
0
```

After 8 minutes, the function returns the number `0`. This code indicates that all tables are completely populated. You now disable the restricted session so that the application can start query In-Memory objects with full confidence that only In-Memory representations will be accessed:

```
ALTER SYSTEM DISABLE RESTRICTED SESSION;
```

102
DBMS_INMEMORY_ADVISE

This package enables you to determine if a workload can benefit from Oracle's Database In-Memory feature.

102.1 DBMS_INMEMORY_ADVISE Overview

DBMS_INMEMORY_ADVISE provides tools to determine the benefit of the Database In-Memory feature to the workflow on a database.

DBMS_INMEMORY_ADVISE perform its analysis on the workload of the current database. You can do the same analysis on another database workload by importing and examining AWR snapshots.

102.2 Summary of DBMS_INMEMORY_ADVISE Subprograms

In Oracle Database 19c, DBMS_INMEMORY_ADVISE includes a single subprogram: IS_INMEMORY_ELIGIBLE.

Table 102-1 DBMS_INMEMORY_ADVISE Subprograms

Subprogram	Description
IS_INMEMORY_ELIGIBLE Procedure	Examines AWR snapshots to determine if a given workload will benefit or not benefit from the Database In-Memory feature.

102.3 IS_INMEMORY_ELIGIBLE Procedure

This procedure examines Automatic Workload Repository (AWR) snapshots from a database to determine the workload's eligibility for Database In-Memory technology.

Many workloads benefit from Database In-Memory, however some may not. `IS_INMEMEMORY_ELIGIBLE` determines whether or not a workload would benefit from Database In-Memory and should be considered eligible or ineligible. Eligibility is gauged by the percentage of analytical activity in the workload. If you are planning to implement Database In-Memory, you can use this tool to quickly identify and filter out databases that are ineligible; those where analytic activity is low and where you would see no substantive gain from the use of Database In-Memory. You can then focus your Database In-Memory deployment on databases whose workload includes more analytic activity and could therefore benefit substantially.

Criterion Used by IS_INMEMEMORY_ELIGIBLE

The determination that a workflow is eligible or ineligible for Database In-Memory is based on a system-defined threshold. If the percentage of the workload involved analytic activity is below that threshold, then Database In-Memory will not noticeably improve performance. The greater the level of activity above that target analytic workload percentage threshold, the greater is the potential benefit that may be seen from use of Database In-Memory.

IS_INMEMORY_ELIGIBLE Procedure Syntax

IS_INMEMORY_ELIGIBLE has three overloads that let you choose how to define the scope of the eligibility test. The output parameters are the same in all three. The optional input parameter SNAP_DBID is also common.

The input parameters are mutually exclusive to each overload as shown in the syntax below. You cannot combine them in a single procedure call.

Syntax 1: Input as Most Recent *n* Days, Including Today

```
DBMS_INMEMORY_ADVISE.IS_INMEMORY_ELIGIBLE (
    TIME_WINDOW_DAYS   IN NUMBER,
    INMEM_ELIGIBLE     OUT BOOLEAN,
    ANALYSIS_SUMMARY   OUT VARCHAR2,
    SNAP_DBID          IN NUMBER DEFAULT NULL
);
```

Syntax 2: Input as a Snapshot Range

```
DBMS_INMEMORY_ADVISE.IS_INMEMORY_ELIGIBLE (
    START_SNAP_ID     IN NUMBER,
    END_SNAP_ID       IN NUMBER,
    INMEM_ELIGIBLE    OUT BOOLEAN,
    ANALYSIS_SUMMARY  OUT VARCHAR2,
    SNAP_DBID         IN NUMBER DEFAULT NULL
);
```

Syntax 3: Input as Time Interval (Start Time and End Time)

```
DBMS_INMEMORY_ADVISE.IS_INMEMORY_ELIGIBLE (
   START_TIME        IN TIMESTAMP,
   END_TIME          IN TIMESTAMP,
   INMEM_ELIGIBLE    OUT BOOLEAN,
   ANALYSIS_SUMMARY  OUT VARCHAR2,
   SNAP_DBID         IN NUMBER DEFAULT NULL
);
```

Parameters

Table 102-2 IS_INMEMORY_ELIGIBLE Procedure Parameters

Parameter	Description
TIME_WINDOW_DAYS	All snapshots occurring within the last *n* number of days, including today.
START_TIME	Specifies the start of a timespan for snapshot analysis. All AWR snapshots created after START_TIME, up to and included END_TIME are examined.
END_TIME	Specifies the end of a timespan for snapshot analysis. All AWR snapshots created prior to END_TIME, back to and including START_TIME are examined.

Table 102-2 (Cont.) IS_INMEMORY_ELIGIBLE Procedure Parameters

Parameter	Description
START_SNAP_ID	The first snapshot in a sequence of snapshots. This is the starting point of the workload to be analyzed.
END_SNAP_ID	The end of a sequence of snapshots.
INMEM_ELIGIBLE	The determination of whether or not the workflow is eligible for (will benefit from) Database In-Memory. Possible values: `True`, `False`.
ANALYSIS_SUMMARY	The analysis that explains the INMEM_ELIGIBLE evaluation. Example: `Observed Analytic Workload Percentage is 3% is less than target Analytic Workload Percentage 20%` Workloads with a percentage of analytical activity less than target are considered ineligible.
SNAP_DBID	Database ID of a database from where AWR data has been exported. This enables you to analyze the workload on another database. If you are testing for eligibility on the current database, exclude this parameter.

Usage Notes

Within a given workload, the level of analytical activity may vary over a series of snapshots. Depending on which subset of all available snapshots are selected for testing, the analytic workload percentage calculated by `IS_INMEMORY_ELIGIBLE` may correspondingly vary. So you may want to do several evaluations with different input parameters. Using the `TIME_WINDOW_DAYS` parameter gives you a quick summary of the workload based on that time window. Then you can drill down into specific time periods, using the parameters based on snapshot IDs to get a perspective on patterns in the workload.

If you are considering whether or not to enable Database In-Memory on your databases, do the following:

1. Run the In-Memory Eligibility Test on candidate databases to find out which databases can or cannot effectively use Database In-Memory.
2. Run the In-Memory Advisor on any Oracle 19c database, except those that the In-Memory Eligibility Test has determined are ineligible.

Examples

Example 1: Testing for eligibility within the time window of the last *n* days as input.

```
set serverout on;
set lines 200;

DECLARE
    inmem_eligible BOOLEAN;
    analysis_summary VARCHAR2(200);
BEGIN

    dbms_inmemory_advise.is_inmemory_eligible(1, inmem_eligible,
analysis_summary);
```

```
    DBMS_OUTPUT.PUT_LINE(
            CASE WHEN inmem_eligible = TRUE
                THEN 'Eligibile for In Memory'
                ELSE 'Not Eligible for In Memory'
            END
    );

    DBMS_OUTPUT.PUT_LINE(analysis_summary);

END;
/
```

Example 2: Testing for eligibility with a range of snapshot IDs as input.

```
set serverout on;
set lines 200;
DECLARE
    inmem_eligible BOOLEAN;
    analysis_summary VARCHAR2(200);
BEGIN

    dbms_inmemory_advise.is_inmemory_eligible(355, 356, inmem_eligible,
analysis_summary);

    DBMS_OUTPUT.PUT_LINE(
            CASE WHEN inmem_eligible = TRUE
                THEN 'Eligibile for In Memory'
                ELSE 'Not Eligible for In Memory'
            END
    );

    DBMS_OUTPUT.PUT_LINE(analysis_summary);

END;
/
```

Example 3: Testing for eligibility with a time interval as input. All snapshots beginning and ending within the time interval are included.

```
set serverout on;
set lines 200;
DECLARE
    inmem_eligible BOOLEAN;
    analysis_summary VARCHAR2(120);
BEGIN

    dbms_inmemory_advise.is_inmemory_eligible(TO_TIMESTAMP('21-JUN-23
05.22.27.262 PM', 'DD-Mon-RR HH:MI:SS.FF AM'),
                                              TO_TIMESTAMP('21-JUN-23
05.26.04.446 PM', 'DD-Mon-RR HH:MI:SS.FF AM'),
                                              inmem_eligible,
analysis_summary
    );
```

```
    DBMS_OUTPUT.PUT_LINE(
            CASE WHEN inmem_eligible = TRUE
                THEN 'Eligibile for In Memory'
                ELSE 'Not Eligible for In Memory'
            END
    );

    DBMS_OUTPUT.PUT_LINE(analysis_summary);

END;
/
```

103
DBMS_IOT

The DBMS_IOT package creates a table into which references to the chained rows for an index-organized table can be placed using the ANALYZE command. DBMS_IOT can also create an exception table into which references to the rows of an index-organized table that violate a constraint can be placed during the enable_constraint operation.

DBMS_IOT is not loaded during database installation. To install DBMS_IOT, run dbmsiotc.sql, available in the ADMIN directory.

This chapter contains the following topics:

- Summary of DBMS_IOT Subprograms

> **Note:**
>
> With the introduction of logical-rowids for IOTs with Oracle Database Release 8.1, you no longer need to use the procedures contained in this package which is retained for backward compatibility only. It is however required for servers running with Oracle Database Release 8.0.

103.1 Summary of DBMS_IOT Subprograms

This table lists and briefly describes the DBMS_IOT subprograms.

Table 103-1 ***DBMS_IOT Package Subprograms***

Subprogram	Description
BUILD_CHAIN_ROWS_TABLE Procedure	Creates a table into which references to the chained rows for an index-organized table can be placed using the ANALYZE command
BUILD_EXCEPTIONS_TABLE Procedure	Creates an exception table into which rows of an index-organized table that violate a constraint can be placed

103.1.1 BUILD_CHAIN_ROWS_TABLE Procedure

This procedure creates a table into which references to the chained rows for an index-organized table can be placed using the ANALYZE command.

Syntax

```
DBMS_IOT.BUILD_CHAIN_ROWS_TABLE (
   owner               IN VARCHAR2,
   iot_name            IN VARCHAR2,
   chainrow_table_name IN VARCHAR2 default 'IOT_CHAINED_ROWS');
```

Parameters

Table 103-2 BUILD_CHAIN_ROWS_TABLE Procedure Parameters

Parameter	Description
owner	Owner of the index-organized table.
iot_name	Index-organized table name.
chainrow_table_name	Intended name for the chained-rows table.

Usage Notes

You should create a separate chained-rows table for each index-organized table to accommodate its primary key.

Examples

```
CREATE TABLE l(a char(16),b char(16), c char(16), d char(240),
PRIMARY KEY(a,b,c)) ORGANIZATION INDEX pctthreshold 10 overflow;
EXECUTE DBMS_IOT.BUILD_CHAIN_ROWS_TABLE('SYS','L','LC');
```

A chained-row table is created with the following columns:

```
Column Name                    Null?    Type
------------------------------ -------- ----
OWNER_NAME                              VARCHAR2(30)
TABLE_NAME                              VARCHAR2(30)
CLUSTER_NAME                            VARCHAR2(30)
PARTITION_NAME                          VARCHAR2(30)
SUBPARTITION_NAME                       VARCHAR2(30)
HEAD_ROWID                              ROWID
TIMESTAMP                               DATE
A                                       CHAR(16)
B                                       CHAR(16)
C                                       CHAR(16)
```

103.1.2 BUILD_EXCEPTIONS_TABLE Procedure

This procedure creates an exception table.

Rows of an index-organized table that violate a constraint can be placed into this table during the execution of the following SQL statements:

- ALTER TABLE ... ENABLE CONSTRAINT ... EXCEPTIONS INTO
- ALTER TABLE ... ADD CONSTRAINT ... EXCEPTIONS INTO

Syntax

```
DBMS_IOT.BUILD_EXCEPTIONS_TABLE (
   owner                IN VARCHAR2,
   iot_name             IN VARCHAR2,
   exceptions_table_name IN VARCHAR2 default 'IOT_EXCEPTIONS');
```

Parameters

Table 103-3 BUILD_EXCEPTIONS_TABLE Procedure Parameters

Parameter	Description
owner	Owner of the index-organized table.
iot_name	Index-organized table name.
exceptions_table_name	Intended name for exception-table.

Usage Notes

You should create a separate exception table for each index-organized table to accommodate its primary key.

Examples

```
EXECUTE DBMS_IOT.BUILD_EXCEPTIONS_TABLE('SYS','L','LE');
```

An exception table for the preceding index-organized table with the following columns:

```
Column Name                    Null?    Type
------------------------------ -------- ----
ROW_ID                                  VARCHAR2(30)
OWNER                                   VARCHAR2(30)
TABLE_NAME                              VARCHAR2(30)
CONSTRAINT                              VARCHAR2(30)
A                                       CHAR(16)
B                                       CHAR(16)
C                                       CHAR(16)
```

104
DBMS_JAVA

The DBMS_JAVA package provides a PL/SQL interface for accessing database functionality from Java.

For a complete description of this package within the context of DBMS_JAVA, see DBMS_JAVA in the *Oracle Database Java Developer's Guide*

105
DBMS_JOB

The DBMS_JOB package schedules and manages jobs in the job queue.

> **Note:**
>
> The DBMS_JOB package has been superseded by the DBMS_SCHEDULER package, and support for DBMS_JOB might be removed in future releases of Oracle Database. In particular, if you are administering jobs to manage system load, you are encouraged to disable DBMS_JOB by revoking the package execution privilege for users.
>
> For more information, see DBMS_SCHEDULER and "Moving from DBMS_JOB to DBMS_SCHEDULER" in the *Oracle Database Administrator's Guide*.

This chapter contains the following topics:

- Security Model
- Operational Notes
- Summary of DBMS_JOB Subprograms

105.1 DBMS_JOB Security Model

DBMS_JOB uses the same security policies as DBMS_SCHEDULER. You must have the CREATE JOB privilege to use DBMS_JOB.

Jobs cannot be altered or deleted other than jobs owned by the user. This is true for all users including those users granted DBA privileges.

You can execute procedures that are owned by the user for which the user is explicitly granted EXECUTE. However, procedures for which the user is granted the execute privilege through roles cannot be executed.

Note that, once a job is started and running, there is no easy way to stop the job.

105.2 DBMS_JOB Operational Notes

These notes describe stopping a job, and working with Oracle Real Application Clusters.

Stopping a Job

Note that, once a job is started and running, there is no easy way to stop the job.

Working with Oracle Real Application Clusters

DBMS_JOB supports multi-instance execution of jobs. By default jobs can be executed on any instance, but only one single instance will execute the job. In addition, you can force instance binding by binding the job to a particular instance. You implement instance binding by

specifying an instance number to the instance affinity parameter. Note, however, that in Oracle Database 10g Release 1 (10.1) instance binding is not recommended. Service affinity is preferred. This concept is implemented in the DBMS_SCHEDULER package.

The following procedures can be used to create, alter or run jobs with instance affinity. Note that not specifying affinity means any instance can run the job.

- DBMS_JOB.SUBMIT
- DBMS_JOB.INSTANCE
- DBMS_JOB.CHANGE
- DBMS_JOB.RUN

DBMS_JOB.SUBMIT

To submit a job to the job queue, use the following syntax:

```
DBMS_JOB.SUBMIT(
   job       OUT    BINARY_INTEGER,
   what      IN     VARCHAR2,
   next_date IN     DATE DEFAULT SYSDATE,
   interval  IN     VARCHAR2 DEFAULT 'NULL',
   no_parse  IN     BOOLEAN DEFAULT FALSE,
   instance  IN     BINARY_INTEGER DEFAULT ANY_INSTANCE,
   force     IN     BOOLEAN DEFAULT FALSE);
```

Use the parameters `instance` and `force` to control job and instance affinity. The default value of `instance` is 0 (zero) to indicate that any instance can execute the job. To run the job on a certain instance, specify the `instance` value. Oracle displays error `ORA-23319` if the `instance` value is a negative number or `NULL`.

The `force` parameter defaults to `false`. If `force` is `TRUE`, any positive integer is acceptable as the job instance. If `force` is `FALSE`, the specified instance must be running, or Oracle displays error number `ORA-23428`.

DBMS_JOB.INSTANCE

To assign a particular instance to execute a job, use the following syntax:

```
DBMS_JOB.INSTANCE(  JOB IN BINARY_INTEGER,
  instance                IN BINARY_INTEGER,
  force                   IN BOOLEAN DEFAULT FALSE);
```

The `FORCE` parameter in this example defaults to `FALSE`. If the instance value is 0 (zero), job affinity is altered and any available instance can execute the job despite the value of force. If the `INSTANCE` value is positive and the `FORCE` parameter is `FALSE`, job affinity is altered only if the specified instance is running, or Oracle displays error `ORA-23428`.

If the `force` parameter is `TRUE`, any positive integer is acceptable as the job instance and the job affinity is altered. Oracle displays error ORA-23319 if the `instance` value is negative or `NULL`.

DBMS_JOB.CHANGE

To alter user-definable parameters associated with a job, use the following syntax:

```
DBMS_JOB.CHANGE(  JOB IN BINARY_INTEGER,
what                  IN VARCHAR2 DEFAULT NULL,
next_date             IN DATE DEFAULT NULL,
interval              IN VARCHAR2 DEFAULT NULL,
```

```
   instance             IN BINARY_INTEGER DEFAULT NULL,
   force                IN BOOLEAN DEFAULT FALSE );
```

Two parameters, `instance` and `force`, appear in this example. The default value of `instance` is `null` indicating that job affinity will not change.

The default value of `force` is `FALSE`. Oracle displays error `ORA-23428` if the specified instance is not running and error `ORA-23319` if the `instance` number is negative.

DBMS_JOB.RUN

The `force` parameter for `DBMS_JOB.RUN` defaults to `FALSE`. If `force` is `TRUE`, instance affinity is irrelevant for running jobs in the foreground process. If force is `FALSE`, the job can run in the foreground only in the specified instance. Oracle displays error `ORA-23428` if force is `FALSE` and the connected instance is the incorrect instance.

```
DBMS_JOB.RUN(
   job    IN BINARY_INTEGER,
   force  IN BOOLEAN DEFAULT FALSE);
```

105.3 Summary of DBMS_JOB Subprograms

This table lists the `DBMS_JOB` subprograms and briefly describes them.

Table 105-1 DBMS_JOB Package Subprograms

Subprogram	Description
BROKEN Procedure	Disables job execution
CHANGE Procedure	Alters any of the user-definable parameters associated with a job
INSTANCE Procedure	Assigns a job to be run by a instance
INTERVAL Procedure	Alters the interval between executions for a specified job
NEXT_DATE Procedure	Alters the next execution time for a specified job
REMOVE Procedure	Removes specified job from the job queue
RUN Procedure	Forces a specified job to run
SUBMIT Procedure	Submits a new job to the job queue
USER_EXPORT Procedures	Re-creates a given job for export, or re-creates a given job for export with instance affinity
WHAT Procedure	Alters the job description for a specified job

105.3.1 BROKEN Procedure

This procedure sets the broken flag. Broken jobs are never run.

Syntax

```
DBMS_JOB.BROKEN (
   job       IN  BINARY_INTEGER,
   broken    IN  BOOLEAN,
   next_date IN  DATE DEFAULT SYSDATE);
```

Parameters

Table 105-2 BROKEN Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
broken	Sets the job as broken or not broken. TRUE sets it as broken; FALSE sets it as not broken.
next_date	Next date when the job will be run.

> **Note:**
>
> If you set job as broken while it is running, Oracle resets the job's status to normal after the job completes. Therefore, only execute this procedure for jobs that are not running.

Usage Notes

- Your job will not be available for processing by the job queue in the background until it is committed.
- If a job fails 16 times in a row, Oracle automatically sets it as broken and then stops trying to run it.

105.3.2 CHANGE Procedure

This procedure changes any of the fields a user can set in a job.

Syntax

```
DBMS_JOB.CHANGE (
   job       IN  BINARY_INTEGER,
   what      IN  VARCHAR2,
   next_date IN  DATE,
   interval  IN  VARCHAR2,
   instance  IN  BINARY_INTEGER DEFAULT NULL,
   force     IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 105-3 CHANGE Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
what	PL/SQL procedure to run.
next_date	Next date when the job will be run.
interval	Date function; evaluated immediately before the job starts running.

Table 105-3 (Cont.) CHANGE Procedure Parameters

Parameter	Description
instance	When a job is submitted, specifies which instance can run the job. This defaults to NULL, which indicates that instance affinity is not changed.
force	If this is FALSE, then the specified instance (to which the instance number change) must be running. Otherwise, the routine raises an exception. If this is TRUE, then any positive integer is acceptable as the job instance.

Usage Notes

- Your job will not be available for processing by the job queue in the background until it is committed.
- The parameters instance and force are added for job queue affinity. Job queue affinity gives users the ability to indicate whether a particular instance or any instance can run a submitted job.
- If the parameters what, next_date, or interval are NULL, then leave that value as it is.

Example

```
BEGIN
   DBMS_JOB.CHANGE(14144, null, null, 'sysdate+3');
   COMMIT;
END;
```

105.3.3 INSTANCE Procedure

This procedure changes job instance affinity.

Syntax

```
DBMS_JOB.INSTANCE (
   job         IN BINARY_INTEGER,
   instance    IN BINARY_INTEGER,
   force       IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 105-4 INSTANCE Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
instance	When a job is submitted, a user can specify which instance can run the job.
force	If this is TRUE, then any positive integer is acceptable as the job instance. If this is FALSE (the default), then the specified instance must be running; otherwise the routine raises an exception.

Usage Notes

Your job will not be available for processing by the job queue in the background until it is committed.

105.3.4 INTERVAL Procedure

This procedure changes how often a job runs.

Syntax

```
DBMS_JOB.INTERVAL (
   job       IN  BINARY_INTEGER,
   interval  IN  VARCHAR2);
```

Parameters

Table 105-5 INTERVAL Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
interval	Date function, evaluated immediately before the job starts running.

Usage Notes

- If the job completes successfully, then this new date is placed in next_date. interval is evaluated by plugging it into the statement select interval into next_date from dual;
- The interval parameter must evaluate to a time in the future. Legal intervals include:

Interval	Description
'sysdate + 7'	Run once a week.
'next_day(sysdate,''TUESDAY'')'	Run once every Tuesday.
'null'	Run only once.

- If interval evaluates to NULL and if a job completes successfully, then the job is automatically deleted from the queue.
- Your job will not be available for processing by the job queue in the background until it is committed.

105.3.5 NEXT_DATE Procedure

This procedure changes when an existing job next runs.

Syntax

```
DBMS_JOB.NEXT_DATE (
   job       IN  BINARY_INTEGER,
   next_date IN  DATE);
```

Parameters

Table 105-6 NEXT_DATE Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
next_date	Date of the next refresh: it is when the job will be automatically run, assuming there are background processes attempting to run it.

Usage Notes

Your job will not be available for processing by the job queue in the background until it is committed.

105.3.6 REMOVE Procedure

This procedure removes an existing job from the job queue. This currently does not stop a running job.

Syntax

```
DBMS_JOB.REMOVE (
   job       IN  BINARY_INTEGER );
```

Parameters

Table 105-7 REMOVE Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.

Usage Notes

Your job will not be available for processing by the job queue in the background until it is committed.

Example

```
BEGIN
   DBMS_JOB.REMOVE(14144);
   COMMIT;
END;
```

105.3.7 RUN Procedure

This procedure runs job JOB now. It runs it even if it is broken.

Running the job recomputes next_date. See data dictionary view USER_JOBS or DBA_JOBS.

Syntax

```
DBMS_JOB.RUN (
   job       IN  BINARY_INTEGER,
   force     IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 105-8 RUN Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
force	If this is TRUE, then instance affinity is irrelevant for running jobs in the foreground process. If this is FALSE, then the job can be run in the foreground only in the specified instance.

Example

```
EXECUTE DBMS_JOB.RUN(14144);
```

> **WARNING:**
>
> This re-initializes the current session's packages.

Exceptions

An exception is raised if force is FALSE, and if the connected instance is the wrong one.

105.3.8 SUBMIT Procedure

This procedure submits a new job. It chooses the job from the sequence sys.jobseq.

Syntax

```
DBMS_JOB.SUBMIT (
   job       OUT BINARY_INTEGER,
   what      IN  VARCHAR2,
   next_date IN  DATE DEFAULT SYSDATE,
   interval  IN  VARCHAR2 DEFAULT 'null',
   no_parse  IN  BOOLEAN DEFAULT FALSE,
   instance  IN  BINARY_INTEGER DEFAULT any_instance,
   force     IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 105-9 SUBMIT Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view

Table 105-9 (Cont.) SUBMIT Procedure Parameters

Parameter	Description

what	PL/SQL text o the job to be run. This must be a valid PL/SQL statement or block of code. For example, to run a stored procedure P, you could pass the string P; (with the semi-colon) to this routine. The SQL that you submit in the what parameter is wrapped in the following PL/SQL block:
	
	DECLARE
	 job BINARY_INTEGER := :job;
	 next_date DATE := :mydate;
	 broken BOOLEAN := FALSE;
	BEGIN
	 WHAT
	 :mydate := next_date;
	 IF broken THEN :b := 1; ELSE :b := 0; END IF;
	END;
	
	Ensure that you include the ; semi-colon with the statement.

next_date	Next date when the job will be run.
interval	Date function that calculates the next time to run the job. The default is NULL. This must evaluate to a either a future point in time or NULL.
no_parse	A flag. The default is FALSE. If this is set to FALSE, then Oracle parses the procedure associated with the job. If this is set to TRUE, then Oracle parses the procedure associated with the job the first time that the job is run. For example, if you want to submit a job before you have created the tables associated with the job, then set this to TRUE.
instance	When a job is submitted, specifies which instance can run the job.
force	If this is TRUE, then any positive integer is acceptable as the job instance. If this is FALSE (the default), then the specified instance must be running; otherwise the routine raises an exception.

Usage Notes

- Your job will not be available for processing by the job queue in the background until it is committed.
- The parameters instance and force are added for job queue affinity. Job queue affinity gives users the ability to indicate whether a particular instance or any instance can run a submitted job.

Example

This submits a new job to the job queue. The job calls the procedure DBMS_DDL.ANALYZE_OBJECT to generate optimizer statistics for the table DQUON.ACCOUNTS. The statistics are based on a sample of half the rows of the ACCOUNTS table. The job is run every 24 hours:

```
VARIABLE jobno number;
BEGIN
   DBMS_JOB.SUBMIT(:jobno,
      'dbms_ddl.analyze_object(''TABLE'',
      ''DQUON'', ''ACCOUNTS'',
      ''ESTIMATE'', NULL, 50);'
      SYSDATE, 'SYSDATE + 1');
```

```
   COMMIT;
END;
/
Statement processed.
print jobno
JOBNO
----------
14144
```

105.3.9 USER_EXPORT Procedures

There are two overloaded procedures. The first produces the text of a call to re-create the given job. The second alters instance affinity (8*i* and after) and preserves the compatibility.

Syntax

```
DBMS_JOB.USER_EXPORT (
   job    IN     BINARY_INTEGER,
   mycall IN OUT VARCHAR2);

DBMS_JOB.USER_EXPORT (
   job      IN     BINARY_INTEGER,
   mycall   IN OUT VARCHAR2,
   myinst   IN OUT VARCHAR2);
```

Parameters

Table 105-10 USER_EXPORT Procedure Parameter

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
mycall	Text of a call to re-create the given job.
myinst	Text of a call to alter instance affinity.

105.3.10 WHAT Procedure

This procedure changes what an existing job does, and replaces its environment.

Syntax

```
DBMS_JOB.WHAT (
   job       IN  BINARY_INTEGER,
   what      IN  VARCHAR2);
```

Parameters

Table 105-11 WHAT Procedure Parameters

Parameter	Description
job	System-assigned ID of the job being run. To find this ID, query the JOB column of the USER_JOBS or DBA_JOBS view.
what	PL/SQL procedure to run.

Usage Notes

- Your job will not be available for processing by the job queue in the background until it is committed.
- Some legal values of `what` (assuming the routines exist) are:
  - `'myproc(''10-JAN-82'', next_date, broken);'`
  - `'scott.emppackage.give_raise(''JENKINS'', 30000.00);'`
  - `'dbms_job.remove(job);'`

106
DBMS_JSON_DUALITY

The package `DBMS_JSON_DUALITY` simplifies the creation and access of JSON-relational duality views. Given existing set of JSON documents, a corresponding relational schema and JSON-relational duality views are automatically created. In addition, the data can also be automatically imported from the JSON collection. The package also provides functions for non-interactive transactions.

This chapter contains the following topics:

- DBMS_JSON_DUALITY Overview
- DBMS_JSON_DUALITY Security
- Summary of DBMS_JSON_Duality Subprograms

Related Topics

-

106.1 DBMS_JSON_DUALITY Overview

Package `DBMS_JSON_DUALITY` provides subprograms for converting JSON collections to JSON-relational duality views. The package also provides subprograms for importing data into the duality views.

JSON documents are imported from document databases into a set of tables with a single JSON column. This can be done using existing tools provided by Oracle and document Databases. Following list points out to the appropriate sections that helps in loading JSON data and migrating data from MongoDB to Oracle Database:

- Loading External JSON Data
- Loading JSON Documents
- Migrate Application Data from MongoDB to Oracle Database

The `DBMS_JSON_DUALITY` package contains a set of procedures and functions that allows you to find, create relational schema, import JSON documents to duality views and perform non-interactive transactions on the imported duality views.

1. **JSON-To-Duality Converter :** Execute one of the appropriate JSON-To-Duality Migrator PL/SQL functions listed below to find and create the relational schema for the input collections (JSON tables).
   - `INFER_SCHEMA` : Infers the JSON schema that represents all of the input document sets.
   - `GENERATE_SCHEMA`: Produces the code to create the required database objects for each duality view.
   - `INFER_AND_GENERATE_SCHEMA`: performs both operations of `INFER_SCHEMA` and `GENERATE_SCHEMA`.
2. **JSON-To-Duality Importer :**
   This feature is a PL/SQL procedure `IMPORT` to import application data from sets of JSON documents into duality views defined using the JSON-To-Duality Converter feature. Based on the relational schema created by the JSON-To-Duality Converter, input data gets

decomposed and normalized into relational tables automatically. Any data that cannot be successfully imported is logged in an error log with the reason for the error and suggestions to fix the issue, if possible.

> **Note:**
>
> Use the importer feature only when you are satisfied with the relational schema recommended by the JSON Relational Duality Converter feature.

3. **Non-Interactive Transactions**
   This feature is a set of PL/SQL procedures that allows you to perform a multi-operational transaction on duality views. The following are the non-interactive procedures from package `DBMS_JSON_DUALITY`:
   - `BEGIN_TRANSACTION` : Begins a multi-operational transaction.
   - `REGISTER`: Check that the ETAG value of a document as last read matches that of the document in the database at the start of the transaction.
   - `COMMIT_TRANSACTION`: Commits a multi-operational transaction.

106.2 DBMS_JSON_DUALITY Security

This section lists the security restrictions related to `DBMS_JSON_DUALITY` package.

In terms of security, JSON-To-Duality Migrator has the following restrictions:

- It can not be run in `SYS` schema.
- The migrator requires "select any table" privilege if `targetSchema` is not current schema. Otherwise, duality view creation fails while attempting to access the tables in `targetSchema`.

  Example:

```
grant select any table to targetSchema;
```

106.3 Summary of DBMS_JSON_Duality Subprograms

This table lists the `DBMS_JSON_DUALITY` subprograms and briefly describes them.

DBMS_JSON_DUALITY Package Subprograms

Subprogram	Description
INFER_SCHEMA Function	Takes as input a set of JSON collections (tables with a single JSON column) and configuration parameters formatted as a JSON document. Returns the validated and normalized relational schema representing the collections as a JSON document
GENERATE_SCHEMA Function	Takes as input a normalized relational schema represented as a JSON document and returns the DDL script to generate the schema.

Subprogram	Description
INFER_AND_GENERATE_SCHEMA Function	Takes as input a set of JSON collections (tables with a single JSON column) and configuration parameters formatted as a JSON document. Returns the DDL script to generate the validated and normalized relational schema representing the collections as a CLOB.
IMPORT Procedure	Imports data into the tables underlying the duality view(s). Logs any documents that fail import into an error log.
BEGIN_TRANSACTION Procedure	Starts a special multi-op transaction.
COMMIT_TRANSACTION Procedure	Commits the multi-op transaction.
REGISTER Procedure	Registers each document to be modified in the transaction, ensuring that the document has not been modified since the last read.

106.3.1 INFER_SCHEMA Function

This section describes the syntax, input and output formats of the `INFER_SCHEMA` function.

Overview

This function takes tables with a single JSON column and their configuration parameters formatted as a JSON document as input. The input fields of the JSON configuration document are described in the section Input Configuration Parameters. It returns the validated and normalized relational schema representing the collections as a JSON document.

Syntax

```
FUNCTION INFER_SCHEMA(config in JSON)
RETURN JSON;
```

Example 106-1 Example for INFER_SCHEMA Function

```
schema_sql :=
    dbms_json_duality.infer_schema(
    json('
{
  "tableNames"           : ["EMP", "DEPT"],
  "viewNames"            : ["EMPLOYEES", "DEPARTMENTS"],
  "useFlexFields"        : true,
  "updatability"         : true,
  "sourceSchema"         : "HR",
  "targetSchema"         : "HR",
  "tablespace"           : "HR_TBS",
  "ingestLimit"          : 100000,
  "outputFormat"         : "executable",
  "minFrequency"         : 10,
  "softnessThreshold"    : 98
}'));

   dbms_output.put_line(schema_sql);
```

Input Configuration Fields

Parameter	Value Type
tableNames	JSON array of strings of the table names of JSON collection to be processed
viewNames	JSON array of strings of the output duality view names to be generated
useFlexFields	Boolean parameter which indicates whether flex fields should be used in the duality views. • When `true`, for each duality view, a flex column is added to each table that directly underlies the top-level fields of an object in the supported documents. • Default value for this parameter is `true`.
updatability	Boolean parameter which indicates whether the duality views created by the migrator should be updatable by default. • When `true`, the migrator will mark all duality views as updatable (i.e., annotations will be set for maximum updatability of the duality view). • When `false`, all duality views created will be read-only. • Default value for this parameter is `true`.
sourceSchema	This string parameter defines the schema name for the input tables. • Default value for this parameter is the current user schema.
targetSchema	This string parameter defines the schema name for the output tables and duality views. • There is no default value for this parameter. • If this parameter is not specified, a string whose value is the name of the database schema that will own the output database views."
tablespace	This string parameter defines the `tablespace` name to use for all tables created by the migrator. • This parameter has no default. • If this parameter is not specified, then the `tablespace` for tables created by the migrator will remain unspecified.
ingestLimit	This numeric parameter defines the maximum number of documents that will be analyzed in each JSON collection to infer the relational schema. • The default value of this parameter is 100000.

Parameter	Value Type
outputFormat	This string parameter defines the format of the output DDL script. • The possible values for this parameter are "standalone" and "executable". • The "executable" format creates a DDL script output which can be directly executed using the "execute immediate" PL/SQL construct. • The "standalone" format creates a DDL script which needs to be copied into a SQL file and executed independently. It cannot be executed using the "execute immediate" PL/SQL construct. • The default value of this parameter is "executable". **Note:** For cases where the output DDL statements can be larger than 32k characters, the user MUST use the "standalone" option.
minFrequency	This numeric parameter defines the threshold in percentage for identifying high entropy fields. • The default value for this parameter is 5, meaning that fields which appear in less than 5% of documents in a collection will be pruned from the schema.
softnessThreshold	This numeric parameter defines the threshold in percentage for allowable 'softness' in identifying functional dependencies in the input data. • The default value of this parameter is 99, meaning 1% of data could have missing or incorrect information. – The value of this parameter can be set based on the 'dirtiness' of input data. The dirtier (containing more errors) the data, the smaller the value of this parameter should be.

106.3.2 GENERATE_SCHEMA Function

This section describes the syntax, input and output formats of the `GENERATE_SCHEMA` function.

Overview

This function takes a normalized relational schema represented as a JSON document as input and returns the DDL script to generate the schema.

Syntax

```
FUNCTION GENERATE_SCHEMA(er_schema in JSON)
  RETURN CLOB;
```

Example 106-2 Example for GENERATE_SCHEMA Function

```
schema_example :=dbms_json_duality.generate_schema(schema_sql);
   dbms_output.put_line(schema_example);
```

> **Note:**
>
> The `schema_sql` is the output from `INFER_SCHEMA` function. The details of `INFER_SCHEMA` function is available here.

106.3.3 INFER_AND_GENERATE_SCHEMA Function

This section describes the syntax, input and output formats of the `INFER_AND_GENERATE_SCHEMA` function.

Overview

This function takes a set of JSON collections (tables with a single JSON column) and configuration parameters formatted as a JSON document as input. The input fields of the JSON configuration document are identical to the `INFER_SCHEMA` function and listed here. It returns the DDL script to generate the validated and normalized relational schema representing the collections as a CLOB.

Syntax

```
FUNCTION INFER_AND_GENERATE_SCHEMA(config in JSON)
  RETURN CLOB;
```

Example 106-3 Example for INFER_AND_GENERATE_SCHEMA Function

```
schema_sql :=dbms_json_duality.infer_and_generate_schema(json('
{
  "tableNames"           : ["EMP", "DEPT"],
  "viewNames"            : ["EMPLOYEES", "DEPARTMENTS"],
  "useFlexFields"        : true,
  "updatability"         : true,
  "sourceSchema"         : "HR",
  "targetSchema"         : "HR",
  "tablespace"           : "HR_TBS",
  "ingestLimit"          : 100000,
  "outputFormat"         : "executable",
  "minFrequency"         : 10,
  "softnessThreshold"    : 98
}'));

dbms_output.put_line(schema_sql);
```

106.3.4 IMPORT Procedure

This section describes the syntax, input and output formats of the `IMPORT` procedure.

Overview

This procedure imports data from a table with a single JSON column into a Duality View. It also logs any documents that fail import into an error log.

> **Note:**
>
> In the process of normalization, data may be transformed, cast to different data types, and truncated to honor maximum size limits. In addition, any data not conforming to the resulting schema may be rejected during the import phase. Oracle recommends to verify that all data has been successfully imported by running verification tests and looking at the error log for resolving import errors. An easy way to ensure imported data is valid is by running a `json_diff` between the input collection and output duality view and checking that the duality view data only contains additional fields and no fields have been removed or modified.
>
> ```
> select json_diff(input_collection.data, output_duality_view.data)
> from
> input_collection c, output_duality_view d where c.data.pk = d.data.pk;
> ```

Syntax

```
PROCEDURE IMPORT(table_owner_name   in VARCHAR2 DEFAULT NULL,
                 table_name         in VARCHAR2,
                 view_owner_name    in VARCHAR2 DEFAULT NULL,
                 view_name          in VARCHAR2,
                 err_log_owner_name in VARCHAR2 DEFAULT NULL,
                 err_log_name       in VARCHAR2 DEFAULT NULL,
                 reject_limit       in NUMBER   DEFAULT NULL);
```

Input Parameter Fields

Parameter	Value Type
`table_owner_name`	The name of the owner of the input table with a single JSON column. You can specify the owner in `table_name`. Otherwise, the schema of the current connected user is used.
`table_name`	The name of the input table with a single JSON column.
`view_owner_name`	The name of the owner of the output duality view. You can specify the owner in `veiw_name`.

Parameter	Value Type
	Otherwise, the schema of the current connected user is used
view_name	The name of the output duality view.
err_log_owner_name	The name of the owner of the error log to use for document that cannot be imported successfully. You can specify the owner in err_log_name. Otherwise, the schema of the current connected user is used.
err_log_name	The name of the error log to use.The default value is NULL, so no errors will be logged by default.
reject_limit	Specifies the maximum number of errors that can be logged before the operation is aborted. The default (NULL) is unlimited, so the operation will never be aborted by default when error_log_name is specified.

Example 106-4 Example for IMPORT Procedure

```
exec dbms_errorlog.create_error_log(dml_table_name => 'COLLECTION_DUALITY',
err_log_table_name => 'ERR_LOG', skip_unsupported => TRUE);
exec dbms_json_duality.import(table_name => 'COLLECTION', view_name =>
'COLLECTION_DUALITY',
err_log_name => 'ERR_LOG');
```

106.3.5 BEGIN_TRANSACTION Procedure

This section describes the syntax of the non-interactive procedure named BEGIN_TRANSACTION

Overview

The typical work-flow for a non-interactive transaction involves the following procedures in the below-mentioned sequence:

- Begin the transaction
- Register the documents
- Perform queries or DMLs on the duality views
- Commit the transaction

The BEGIN_TRANSACTION procedure starts a special multi-operation transaction with the following properties:

- The multi-operation transaction provides repeatable reads, i.e., all reads run as of the snapshot when the transaction began.
- Locks are only taken for rows that are modified, unmodified rows remain unlocked throughout the transaction

Syntax

```
procedure begin_transaction;
```

Example 106-5 Example for BEGIN_TRANSACTION Procedure

```
dbms_json_duality.begin_transaction()
```

> **Note:**
>
> The `BEGIN_TRANSACTION` procedure is normally followed by `REGISTER` procedure. The detailed information about `REGISTER` procedure can be found here.

106.3.6 COMMIT_TRANSACTION Procedure

This section describes the syntax of the non-interactive procedure named `COMMIT_TRANSACTION`

Overview

This procedure commits the special multi-operation transaction. It validates all modified documents in the transaction, i.e., ensures that no other session has modified these documents.

Syntax

```
procedure commit_transaction;
```

Example 106-6 Example for COMMIT_TRANSACTION Procedure

`COMMIT_TRANSACTION` procedure is normally the last step in the non-interactive transaction after queries are performed on the duality views.

```
dbms_json_duality.commit_transaction();
```

106.3.7 REGISTER Procedure

This section describes the syntax of the non-interactive procedure named `REGISTER`

Overview

This procedure verifies that the current `etag` of the object is unchanged from a prior read. If the `etag` doesn't match the expected `etag`, and error is thrown from this function. The register procedure relies on object ids (oid) to identify objects/documents to register and validate. The object id can be obtained by querying the `resid` hidden column from the duality view. Its value is simply a concatenated binary encoding of all the primary key columns of the root table in the duality view.

Syntax

```
 procedure register(view_name in VARCHAR2,  /* duality view name
*/
                    oid       in RAW,       /* duality view obj identifier
*/
                    etag      in RAW);      /* document etag */
```

Example 106-7 Example for REGISTER Procedure

This example assumes that a duality view `team_dv` is already created. The detailed information about creating the duality views are explained in Creating Duality Views topic in the Overview of JSON-Relational Duality Views. Obtaining the `RESID` and `ETAG` is mandatory step prior to calling the `REGISTER` procedure. The following shows an example code to obtain the `RESID` and `ETAG` of the duality view `team_dv`:

```
SELECT RESID, DATA FROM team_dv dv
  WHERE dv.DATA.name LIKE 'Mercedes%';
```

This produces an output that includes the `oid` (`RESID`) and the `ETAG` information as shown below:

```
RESID
-----
DATA
----
FB03C2040400
{"_id" : 303,
 "_metadata":
  {"etag" : "039A7874ACEE6B6709E06E42E4DC6355",
   "asof" : "00000000001BE239"},
 "name" : "Mercedes",
 ...}
```

> **Note:**
>
> `REGISTER` procedure is normally performed after executing the `BEGIN_TRANSACTION` procedure. The details of `BEGIN_TRANSACTION` procedure is provided here.

```
BEGIN
  DBMS_JSON_DUALITY.begin_transaction();
  DBMS_JSON_DUALITY.register('team_dv',
                             hextoraw('FB03C2040400'),
                             hextoraw('039A7874ACEE6B6709E06E42E4DC6355'));
  ..........................................
  ****Perform the updating (DML) operations****
  ..........................................
```

```
  DBMS_JSON_DUALITY.commit_transaction();
END
```

107
DBMS_JSON

The `DBMS_JSON` package provides an interface for data-guide operations.

This chapter contains the following topics:

- DBMS_JSON Overview
- DBMS_JSON Security Model
- DBMS_JSON Constants
- Summary of DBMS_JSON Subprograms

107.1 DBMS_JSON Overview

Package `DBMS_JSON` provides subprograms for manipulating JavaScript Object Notation (JSON) data that is stored in Oracle Database.

107.2 DBMS_JSON Security Model

`PUBLIC` is granted the `EXECUTE` privilege on package `DBMS_JSON`. Its subprograms execute with invoker's rights privileges.

107.3 DBMS_JSON Constants

The DBMS_JSON package uses these constants to define the JSON schema types and data-guide formatting options.

Table 107-1 DBMS_JSON Constants Defined for JSON Data-Guide Formatting

Name	Value	Description
FORMAT_FLAT	2	Display flat format

Table 107-1 (Cont.) DBMS_JSON Constants Defined for JSON Data-Guide Formatting

Name	Value	Description
FORMAT_HIERARCHICAL	1	Display hierarchical format. **Note:** If a field X in the JSON data has different data types in different documents (number/ boolean/string), FORMAT_HIER ARCHICAL will propagate the type to the common denominator (usually string)
FORMAT_SCHEMA	3	Display JSON schema format **Note:** If a field X in the JSON data has different data types in different documents (number/ boolean/string), FORMAT_SCH EMA will list all encountered types in a JSON Array. Oracle recommends the use of FORMAT_SCH EMA if you want to use the generated schema for validation.
PRETTY	1	Use appropriate indention to improve readability

Table 107-2 DBMS_JSON Constants for JSON Schema Types

Name	Type	Value	Description
TYPE_ARRAY	NUMBER(2)	6	A JSON array
TYPE_BOOLEAN	NUMBER(2)	2	A JSON boolean
TYPE_GEOJSON	NUMBER(2)	7	Geographic JSON data
TYPE_NULL	NUMBER(2)	1	The JSON NULL value
TYPE_NUMBER	NUMBER(2)	3	A JSON number
TYPE_OBJECT	NUMBER(2)	5	A JSON object
TYPE_STRING	NUMBER(2)	4	A JSON string
TYPE_BINARY	NUMBER(2)	17	Oracle extended JSON type binary
TYPE_DATE	NUMBER(2)	13	Oracle extended JSON type date
TYPE_DOUBLE	NUMBER(2)	12	Oracle extended JSON type double
TYPE_DSINTERVAL	NUMBER(2)	16	Oracle extended JSON type day-second interval
TYPE_FLOAT	NUMBER(2)	11	Oracle extended JSON type float
TYPE_TIMESTAMP	NUMBER(2)	14	Oracle extended JSON type timestamp
TYPE_YMINTERVAL	NUMBER(2)	15	Oracle extended JSON type year-month interval

Table 107-3 DBMS_JSON Constants for mvrefreshmode Parameter

Name	Type	Value	Description
MV_REFRESH_ON_STATEMENT	NUMBER(2)	1	Creates the materialized view with refresh on statement.
MV_REFRESH_ON_COMMIT	NUMBER(2)	2	Creates the materialized view with refresh on commit.
MV_REFRESH_ON_DEMAND	NUMBER(2)	3	Creates the materialized view with refresh on demand.

See Also:

JSON Developer's Guide

107.4 Summary of DBMS_JSON Subprograms

This table lists the `DBMS_JSON` subprograms and briefly describes them.

DBMS_JSON Package Subprograms

Subprogram	Description
ADD_VIRTUAL_COLUMNS Procedure	Add virtual columns based on data-guide information. This has no effect when running on the shard catalog server — no virtual column is added.
CREATE_VIEW Procedure	Create a view with relational columns and scalar JSON fields as specified in a data guide.
CREATE_VIEW_ON_PATH Procedure	Create a view based on data-guide information, with relational columns, top-level scalar types, and fully expanded sub-tree under a given path. When running on the shard catalog server this raises an error stating that the data guide is empty.
DROP_VIRTUAL_COLUMNS Procedure	Drop virtual columns created by procedure `add_virtual_columns`. This has no effect when running on the shard catalog server.
GET_INDEX_DATAGUIDE Function	Get JSON data guide from a data guide-enabled JSON search index. When running on the shard catalog server this returns a single empty row as result.
GET_VIEW_SQL Function	Get the data definition language (DDL) statement for creating a view without actually creating the view.
JSON_TYPE_CONVERTIBLE_CHECK Procedure	Check whether existing data stored as JSON text can be migrated to the JSON data type.
RENAME_COLUMN Procedure	Set the preferred name for a view column or a virtual column creating using a data guide. This has no effect when running on the shard catalog server.

> **Note:**
>
> In the context of sharding, each individual shard maintains its own data-guide information, which is obtained from the JSON documents stored in that shard. When running on individual shard, procedures in this package that use data-guide information use only the information that is maintained for that shard.

107.4.1 ADD_VIRTUAL_COLUMNS Procedure

This procedure adds virtual columns based on the data guide.

The virtual column name is the value of `o:preferred_vc_name` in the data guide. The procedure ignores JSON objects, arrays, and fields under arrays in the data guide. Before it

adds virtual columns, procedure `ADD_VIRTUAL_COLUMNS` first drops any existing virtual columns that were projected from fields in the same JSON column by a previous invocation of `ADD_VIRTUAL_COLUMNS` or by data-guide change-trigger procedure `add_vc` (in effect, it does what procedure `DBMS_JSON.DROP_VIRTUAL_COLUMNS` does).

> **See Also:**
>
> - DROP_VIRTUAL_COLUMNS Procedure
> - *Oracle Database JSON Developer's Guide*

Syntax

```
DBMS_JSON.ADD_VIRTUAL_COLUMNS (
    tablename               IN  VARCHAR2,
    jcolname                IN  VARCHAR2,
    dataguide               IN  CLOB,
    resolvenameconflicts    IN  BOOLEAN  DEFAULT TRUE,
    colnameprefix           IN VARCHAR2  DEFAULT NULL,
    mixedcasecolumns        IN BOOLEAN   DEFAULT FALSE);
```

For the following signature you must have a data guide-enabled search index on the JSON column. This is not needed for the previous signature.

```
DBMS_JSON.ADD_VIRTUAL_COLUMNS (
    tablename  IN  VARCHAR2,
    jcolname   IN  VARCHAR2,
    frequency      NUMBER    DEFAULT 0,
    hidden         BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 107-4 ADD_VIRTUAL_COLUMNS Procedure Parameters

Parameter	Description
`tablename`	Name of the table containing JSON column `jcolname`.
`jcolname`	Name of the JSON column in table `tablename` that contains the data from which to create the virtual column.
`dataguide`	The data guide. When `o:hidden` in the data guide for a particular JSON field is set to `TRUE`, the corresponding virtual column is added as a hidden column. The default value of `o:hidden` is `FALSE`.
`resolvenameconflicts`	By default, this parameter is set to `TRUE`. The procedure automatically resolves the virtual column name conflicts by appending a sequence number. If it is set to `FALSE`, then in the event of any conflicts among `o:preferred_column_name`, an error is raised.
`colnameprefix`	By default, the virtual column name is the same as the JSON field name. This parameter allows you to add a prefix to the virtual column names.
`mixedcasecolumns`	By default, the virtual column names are case sensitive. If this parameter value is set to `FALSE`, the virtual column names become non-case-sensitive.

Table 107-4 (Cont.) ADD_VIRTUAL_COLUMNS Procedure Parameters

Parameter	Description
`frequency`	Sets the minimum frequency threshold to display JSON columns. A frequency of 0 means display all JSON columns. Also, all JSON columns are displayed if statistics have not been collected, effectively overriding any value set by this parameter.
`hidden`	`TRUE` means the added virtual column is hidden; `FALSE` means it is not. The default is `FALSE`.

Usage Notes

Procedure `DBMS_STATS.GATHER_STATS` collects statistics in the data guide. If the frequency statistic has not been collected, frequency is `NULL`. Setting the frequency to a value greater than zero means do not include columns for which there are no frequency statistics collected (statistic is `NULL`), unless `DBMS_STATS.GATHER_STATS` has never been executed. In that case, the frequency parameter is ignored and all columns are displayed in the view.

107.4.2 CREATE_VIEW Procedure

This procedure creates a view with relational columns, using scalar JSON fields as specified in the data guide. A data guide-enabled JSON search index is not required for this procedure; the data guide is passed to the procedure.

> **See Also:**
>
> *Oracle Database JSON Developer's Guide*

Syntax

```
PROCEDURE CREATE_VIEW (
   viewname                VARCHAR2,
   tablename               VARCHAR2,
   jcolname                VARCHAR2,
   dataguide               CLOB,
   resourcepath            VARCHAR2 DEFAULT NULL,
   materialize             BOOLEAN  DEFAULT FALSE,
   mvrefreshmode           NUMBER   DEFAULT MV_REFRESH_ON_STATEMENT,
   path                    VARCHAR2 DEFAULT '$',
   resolvenameconflicts    BOOLEAN  DEFAULT TRUE,
   colnameprefix           VARCHAR2 DEFAULT NULL,
   mixedcasecolumns        BOOLEAN DEFAULT FALSE);
```

Parameters

Table 107-5 DBMS_JSON.CREATE_VIEW Procedure Parameters

Parameter	Description
`viewname`	Name of the view.
`tablename`	Name of the table containing JSON column `jcolname`.

Table 107-5 (Cont.) DBMS_JSON.CREATE_VIEW Procedure Parameters

Parameter	Description
jcolname	Name of the JSON column in table tablename that is used to create the view.
dataguide	The data guide.
resourcepath	This parameter is for internal use. Value of this parameter is always NULL.
materialize	The value of this parameter is boolean and indicates if the view is materialized or not.
mvrefreshmode	When materialize is true, this parameter specifies the materialized view refresh mode. For more information on materialized view refresh mode options, see DBMS_JSON Constants.
path	The path of the JSON field to be expanded. It uses JSON path-expression syntax. It expands the descendants under the specified path, and creates view columns for each scalar value in the resulting sub-tree. The path $ creates a view starting from the JSON document root.
resolvenameconflicts	By default, this parameter is set to TRUE. The procedure automatically resolves the virtual column name conflicts by appending a sequence number. If it is set to FALSE, then in the event of any conflicts among o:preferred_column_name, an error is raised.
colnameprefix	By default, the view column name is the same as the JSON field name. This parameter allows users to provide a prefix to prepend to the view column names.
mixedcasecolumns	By default, the view column names are case sensitive. You can use this parameter to change the case sensitivity behavior of the view column names.

107.4.3 CREATE_VIEW_ON_PATH Procedure

This procedure creates a view with relational columns, using top-level scalar values and the scalar values in the expanded sub-tree under a given path. The JSON column must have a data guide-enabled search index.

> **See Also:**
>
> *Oracle Database JSON Developer's Guide*

Syntax

```
PROCEDURE CREATE_VIEW_ON_PATH(
   viewname VARCHAR2,
   tablename VARCHAR2,
   jcolname VARCHAR2,
   path VARCHAR2,
   frequency NUMBER DEFAULT 0);
```

Parameters

Table 107-6 CREATE_VIEW_ON_PATH Procedure Parameters

Parameter	Description
`viewname`	Name of the view.
`tablename`	Name of the table containing JSON column `jcolname`.
`jcolname`	Name of the JSON column in table `tablename` that is used to create the view. The column must have a data guide-enabled JSON search index, or else an error is raised.
`path`	The path of the JSON field to be expanded. It uses JSON path-expression syntax. It expands the descendants under the specified path, and creates view columns for each scalar value in the resulting sub-tree. The path `$` creates a view starting from the JSON document root.
`frequency`	The minimum frequency threshold for displaying the JSON columns. A frequency of 0 means display all JSON columns. All JSON columns are also displayed if statistics have not been collected, effectively overriding any value set by this parameter. The view only displays JSON fields with frequency greater than the given `frequency`. It does not display JSON fields added after collecting statistics if the given frequency is greater than 0, if their statistic columns are `NULL`.

107.4.4 DROP_VIRTUAL_COLUMNS Procedure

Drop all virtual columns that were added using PL/SQL procedure `DBMS_JSON.add_virtual_columns` or using data-guide change-trigger procedure `add_vc`.

> **See Also:**
>
> - ADD_VIRTUAL_COLUMNS Procedure
> - *Oracle Database JSON Developer's Guide*

Syntax

```
PROCEDURE DROP_VIRTUAL_COLUMNS(
   tablename VARCHAR2,
   jcolname VARCHAR2);
```

Parameters

Table 107-7 DBMS_JSON.DROP_VIRTUAL_COLUMNS Procedure Parameters

Parameter	Description
`tablename`	Name of the table containing JSON column `jcolname`.
`jcolname`	Name of the JSON column in table `tablename`.

107.4.5 GET_INDEX_DATAGUIDE Function

`GET_INDEX_DATAGUIDE` gets JSON data guide from data guide-enabled JSON search index.

> **See Also:**
>
> *Oracle Database JSON Developer's Guide*

Syntax

```
FUNCTION GET_INDEX_DATAGUIDE(
   tablename VARCHAR2,
   jcolname VARCHAR2,
   format NUMBER,
   pretty NUMBER DEFAULT 0)
   RETURN CLOB;
```

Parameters

Table 107-8 DBMS_JSON.GET_INDEX_DATAGUIDE Procedure Parameters

Parameter	Description
`tablename`	Name of the table containing JSON column `jcolname`.
`jcolname`	Name of the JSON column in table `tablename` that has a data guide-enabled JSON search index.
`format`	The data-guide format: • `FORMAT_HIERARCHICAL` — hierarchical format • `FORMAT_FLAT` — flat format
`pretty`	A value of `DBMS_JSON.PRETTY` means pretty-print the data guide, using indention to improve readability.

Example 107-1 Example Get Data Guide in Hierarchical Pretty Format

This example returns the data guide in hierarchical format.

```
SELECT DBMS_JSON.GET_INDEX_DATAGUIDE('T1', 'PO',
DBMS_JSON.FORMAT_HIERARCHICAL, DBMS_JSON.PRETTY)
FROM DUAL;
```

107.4.6 GET_VIEW_SQL Function

This function returns the creating view DDL without actually creating the view. A data guide-enabled JSON search index is not required for this function; the data guide is passed to the function.

> **See Also:**
>
> *Oracle Database JSON Developer's Guide*

Syntax

```
FUNCTION GET_VIEW_SQL (
   viewname                VARCHAR2,
   tablename               VARCHAR2,
   jcolname                VARCHAR2,
   dataguide               CLOB,
   materialize             BOOLEAN  DEFAULT FALSE,
   mvrefreshmode           NUMBER   DEFAULT MV_REFRESH_ON_STATEMENT,
   path                    VARCHAR2 DEFAULT '$',
   resolvenameconflicts    BOOLEAN  DEFAULT TRUE,
   colnameprefix           VARCHAR2 DEFAULT NULL,
   mixedcasecolumns        BOOLEAN DEFAULT TRUE)
   RETURN CLOB;
```

Parameters

Table 107-9 DBMS_JSON.GET_VIEW_SQL Function Parameters

Parameter	Description
`viewname`	Name of the view.
`tablename`	Name of the table containing JSON column `jcolname`.
`jcolname`	Name of the JSON column in table `tablename` that is used to create the view.
`dataguide`	The data guide.
`materialize`	The value of this parameter is Boolean and indicates if the view is materialized or not.
`mvrefreshmode`	When materialize is true, this parameter specifies the materialized view refresh mode. For more information on materialized view refresh mode options, see DBMS_JSON Constants.
`path`	The path of the `JSON` field to be expanded. It uses `JSON` path-expression syntax. It expands the descendants under the specified path, and creates view columns for each scalar value in the resulting sub-tree. The path $ creates a view starting from the `JSON` document root.
`resolvenameconflicts`	By default, this parameter is set to `TRUE`. The function automatically resolves the virtual column name conflicts by appending a sequence number. If it is set to `FALSE`, then in the event of any conflicts among `o:preferred_column_name`, an error is raised.

Table 107-9 (Cont.) DBMS_JSON.GET_VIEW_SQL Function Parameters

Parameter	Description
`colnameprefix`	By default, the view column name is the same as the `JSON` field name. This parameter allows users to provide a prefix to prepend to the view column names.
`mixedcasecolumns`	By default, the view column names are case sensitive. You can use this parameter to change the case sensitivity behavior of the view column names.

Usage Notes

- If Wide Tables are enabled for the database (`MAX_COLUMNS=EXTENDED`):
  - When `viewname` is `NULL`, the function returns only the select statement of the view DDL and it can select more than 4096 columns.
  - When `viewname` is not `NULL`, the function returns create view DDL and it selects at most 4096 columns.
  - As one `json_table` can only produce at most 4096 columns, the function will split paths into joins among multiple `json_tables` if the paths are more than 4096, when `viewname` is `NULL`.
- If Wide Tables are not enabled for the database (`MAX_COLUMNS=STANDARD`):
  - When `viewname` is `NULL`, the function returns only the select statement of the view DDL and it can select more than 1000 columns.
  - When `viewname` is not `NULL`, the function returns create view DDL and it selects at most 1000 columns.
  - As one `json_table` can only produce at most 1000 columns, the function will split paths into joins among multiple `json_tables` if the paths are more than 1000, when `viewname` is `NULL`.

> **See Also:**
>
> *Oracle Database Reference* for more information on the `MAX_COLUMNS` initialization parameter

107.4.7 JSON_TYPE_CONVERTIBLE_CHECK Procedure

This procedure checks whether existing data stored as JSON text can be migrated to the JSON data type.

> **See Also:**
>
> *Oracle Database JSON Developer's Guide*

Syntax

```
PROCEDURE JSON_TYPE_CONVERTIBLE_CHECK (
   owner                VARCHAR2,
   tableName            VARCHAR2,
   columnName           VARCHAR2,
   statusTableName      VARCHAR2,
   fastCheck            BOOLEAN DEFAULT FALSE,
   appendStatus         BOOLEAN DEFAULT FALSE);
```

Parameters

Table 107-10 DBMS_JSON.JSON_TYPE_CONVERTIBLE_CHECK Procedure Parameters

Parameter	Description
`owner`	Name of the owner of the table.
`tableName`	Name of the table.
`columnName`	Name of the column in the table `tableName` that contains data to convert to the JSON type.
`statusTableName`	Name of the table to use to add the tracking status of the operation. This table might already exist or it might need to be created. If it already exists, the procedure verifies that it complies with the expected shape.
`fastCheck`	The value of this optional parameter is Boolean. If this parameter is set to `TRUE`, the `is_json` check constraint is run to verify that the input data is convertible. If the parameter is `FALSE`, then the `oson` constructor is run to verify that the input data is convertible. The default value is `FALSE`.
`appendStatus`	The value of this optional parameter is Boolean. If this parameter is set to `TRUE`, the status table will not be truncated. If the parameter is `FALSE`, then if the status table already exists, it will be truncated. In either case, if the status table does not already exist, it is created and will then contain only new data from running the procedure. The default value is `FALSE`.

107.4.8 RENAME_COLUMN Procedure

This procedure sets the preferred name for a JSON column, to be used by the create view, or add virtual columns procedure.

> **See Also:**
>
> *Oracle Database JSON Developer's Guide*

Syntax

```
PROCEDURE RENAME_COLUMN(
   tablename VARCHAR2,
   jcolname VARCHAR2,
```

```
   path VARCHAR2,
   type NUMBER,
   preferred_name VARCHAR2);
```

Parameters

Table 107-11 RENAME_COLUMN Procedure Parameters

Parameter	Description
tablename	Name of the table containing JSON column jcolname.
jcolname	Name of the JSON column in table tablename. It must have a data guide-enabled JSON search index, or else an error is raised.
path	Path to the JSON field on which to set the preferred column name.
type	The type of the JSON field targeted by path. Two JSON fields can have the same path if they are of different types. Possible values: • TYPE_NULL • TYPE_STRING • TYPE_NUMBER • TYPE_BOOLEAN • TYPE_OBJECT • TYPE_ARRAY
preferred_name	Preferred name for the JSON field specified by path. If there is a name conflict, a system generated name is used instead.

Example 107-2 Example Renaming a Column

This example renames a field to item_name.

```
EXEC DBMS_JSON.RENAME_COLUMN('T1', 'PO', '$.purchaseOrder.items.name',
DBMS_JSON.TYPE_STRING, 'item_name');
```

108
DBMS_JSON_SCHEMA

The `DBMS_JSON_SCHEMA` package provides subprograms for validating and generating JavaScript Object Notation (JSON) schemas that are stored in Oracle Database.

This chapter contains the following topics:

- DBMS_JSON_SCHEMA Security Model
- DBMS_JSON_SCHEMA Constants
- Summary of DBMS_JSON_SCHEMA Subprograms

108.1 DBMS_JSON_SCHEMA Security Model

The `DBMS_JSON_SCHEMA` package is owned by the `SYS` account.

The `PUBLIC` account is granted the `EXECUTE` privilege on the `DBMS_JSON_SCHEMA` package.

The `DBMS_JSON_SCHEMA` package is defined with the `AUTHID CURRENT_USER` clause.

The user who invokes the `DBMS_JSON_SCHEMA.DESCRIBE()` function must have sufficient privileges on the object that is being described.

108.2 DBMS_JSON_SCHEMA Constants

The DBMS_JSON_SCHEMA package uses the following constants for schema validation.

Table 108-1 DBMS_JSON_SCHEMA Constants for Schema Validation

Name	Value	Description
RAISE_ERROR	1	Raise errors for invalid schemas.
RAISE_NONE	0	Do not raise errors for invalid schemas.
SCHEMA_INVALID	0	The schema is not valid.
SCHEMA_VALID	1	The schema is valid.

> **See Also:**
>
> JSON Developer's Guide

108.3 Summary of DBMS_JSON_SCHEMA Subprograms

This table lists the DBMS_JSON_SCHEMA subprograms and briefly describes them.

DBMS_JSON_SCHEMA Package Subprograms

Subprogram	Description
DESCRIBE Function	Creates a schema from a specified object (table, view, object type; or a synonym resolving to a table, view, duality view, object or collection type, or domain).
IS_SCHEMA_VALID Function	Check the validity of a JSON schema.
IS_VALID Function	Check the validity of JSON data for a specified schema (function).
IS_VALID Procedure	Check the validity of JSON data for a specified schema (procedure).
VALIDATE_REPORT Function	Reads the error report from the result of a schema validation.

108.3.1 DESCRIBE Function

This procedure creates a schema from a specified object.

> **See Also:**
>
> - *Oracle Database JSON Developer's Guide*

Syntax

```
FUNCTION DESCRIBE(
   object_name IN VARCHAR2,
   owner_name  IN VARCHAR2  DEFAULT NULL,
   column_name IN VARCHAR2  DEFAULT NULL)
RETURN JSON;
```

Parameters

Table 108-2 DESCRIBE Function Parameters

Parameter	Description
`object_name`	The object to use to create the schema. The object can be a table, view, object type, or a synonym resolving to a table, view, duality view, object type, collection type, or domain.
`owner_name`	The name of the user who the schema is created for. If the `owner_name` value is NULL, the table is created in the current user's schema.
`column_name`	If a column name is specified, the schema returned is only for the column in the table or view. The `column_name` argument can be used only with table or view objects.

108.3.2 IS_SCHEMA_VALID Function

This function checks a schema for validity.

> **See Also:**
>
> - *Oracle Database JSON Developer's Guide*

Syntax

```
FUNCTION IS_SCHEMA_VALID(
   json_data    IN    JSON)
RETURN PLS_INTEGER;
```

Parameters

Table 108-3 DESCRIBE Function Parameters

Parameter	Description
json_data	The schema to check for validity.

The function returns 0 if the schema is not valid.

108.3.3 IS_VALID Function

This function checks the validity of JSON data for a specified schema.

> **See Also:**
>
> - *Oracle Database JSON Developer's Guide*

Syntax

```
FUNCTION IS_VALID(
   json_data    IN    JSON,
   json_schema IN    JSON,
   raise_err    IN    PLS_INTEGER
   DEFAULT DBMS_JSON_SCHEMA.RAISE_NONE)
RETURN PLS_INTEGER;
```

Parameters

Table 108-4 IS_VALID Function Parameters

Parameter	Description
json_data	The JSON data to validate.
json_schema	The schema to use to validate the data.
raise_err	• RAISE_NONE — Do not raise an error for invalid schemas. • RAISE_ERROR — Raise an error for invalid schemas.

If the schema is valid, the SCHEMA_VALID constant, 1, is returned. If the schema is not valid, the SCHEMA_INVALID constant, 0, is returned, unless the raise_err value is RAISE_ERROR, in which case an error is raised.

108.3.4 IS_VALID Procedure

This procedure checks the validity of JSON data for a specified schema.

> **See Also:**
>
> - *Oracle Database JSON Developer's Guide*

Syntax

```
PROCEDURE IS_VALID(
   json_data   IN  JSON,
   json_schema IN  JSON,
   result      OUT BOOLEAN,
   errors      OUT JSON);
```

Parameters

Table 108-5 IS_VALID Procedure Parameters

Parameter	Description
json_data	The JSON data to validate.
json_schema	The schema to use to validate the data.
result	The result of the validation, which is TRUE if the JSON data is valid and FALSE if not.
errors	A list of error messages that indicate why the JSON data was found invalid when validated against the specified schema.

108.3.5 VALIDATE_REPORT Function

This function reads the error report from the result of a schema validation.

> **See Also:**
>
> - *Oracle Database JSON Developer's Guide*

Syntax

```
FUNCTION VALIDATE_REPORT(
   json_data   IN   VARCHAR2,
   json_schema IN   VARCHAR2)
RETURN JSON;
```

Parameters

Table 108-6 VALIDATE_REPORT Function Parameters

Parameter	Description
`json_data`	The JSON data to validate.
`json_schema`	The schema to use to validate the data.

109
DBMS_KAFKA

The `DBMS_KAFKA` package provides a PL/SQL interface for enabling Oracle SQL access to topics in Kafka clusters.

Users granted `READ` access to an Oracle SQL access to Kafka (OSAK) cluster can use the `DBMS_KAFKA` package to create applications that query Kafka data from Oracle Database views and tables.

- DBMS_KAFKA Overview
- DBMS_KAFKA LOADING Mode
- DBMS_KAFKA Global Temporary Tables
- DBMS_KAFKA Streaming Mode
- DBMS_KAFKA Seekable Mode
- ADD_PARTITIONS
- DROP_ALL_APPS
- ENABLE_VIEW_QUERY
- SET_TRACING

109.1 DBMS_KAFKA Overview

The `DBMS_KAFKA` packages enable you to access and process Kafka data with Oracle SQL access to Kafka (OSAK).

To enable applications to consume Kafka data, you use the `DBMS_KAFKA` packages. Before you can use Kafka resources, an administrator for Kafka must have registered and enabled access to a Kafka cluster.

To access the Kafka topics, you create an Oracle SQL Access to Kafka (OSAK) application. You first have to decide what mode to use.

- Loading mode: Use to load data from a Kafka topic into an Oracle Database table.
- Streaming mode: Use to read sequentially through a Kafka topic.
- Seekable mode: Use to randomly access a Kafka topic between starting and ending timestamps.

You then create the application using the appropriate OSAK package:

- `CREATE_LOAD_APP`: Creates an application that can be used in Loading mode.
- `CREATE_STREAMING_APP`: Creates an application that can be used in Streaming mode.
- `CREATE_SEEKABLE_APP`: Creates an application that can be used in Seekable mode.

Other DBMS_KAFKA packages enable you to manage the Kafka data.

The following is an overview of the procedure for each package

Loading Data Into a Load Application

- Use `DBMS_KAFKA.CREATE_LOAD_APP` to create an Oracle SQL Access to Kafka Load application
- Optionally, use `DBMS_KAFKA_INIT_OFFSET_TS` or `DBMS_KAFKA_INIT_OFFSET` to set the first Kafka record that you want to be read.
- Run `LOOP` until done.
  - Use `DBMS_KAFKA.EXECUTE_LOAD_APP` to load Kafka data starting from where a previous read left off to the current high water mark.
- Use `DBMS_KAFKA.DROP_LOAD_APP` to drop the load application after you are finished with the Kafka data.

Loading Data Into a Streaming Mode Application

To query Kafka data in Streaming mode to read sequentially through a Kafka topic, the procedure is as follows:

1. Use `DBMS_KAFKA.CREATE_STREAMING_APP` to create the Oracle SQL Access to Kafka streaming application.
2. Optionally, use `DBMS_KAFFA_INIT_OFFSET_TS` or `DBMS_KAFKA_INIT_OFFSET` to set the first Kafka record that you want to be read.
3. Run `LOOP` on the data in SQL:
   a. Call `DBMS_KAFKA.LOAD_TEMP_TABLE` to load the global temporary table with the next set of rows from Kafka
   b. Use `SELECT` from the OSAK global temporary table.
   c. Process the data retrieved
   d. If the processing was successful, then use `DBMS_KAFKA.UPDATE_OFFSET` to advance to the next set of Kafka records.
   e. Use `COMMIT` to commit the offset tracking information.
4. When finished with the application, use `DBMS_KAFKA.DROP_STREAMING_APP` to drop the application.

Loading Data into a Seekable Mode Application

To query Kafka data in Seekable mode, so that you can access Kafka records between two timestamps, an overview of the procedure is as follows:

1. `DBMS_KAFKA.CREATE_SEEKABLE_APP` Use to create the Oracle SQL Access to Kafka seekable application
2. Run `LOOP` on the Kafka data in SQL:
   a. Use `DBMS_KAFKA.SEEK_OFFSET_TS` to seek to a defined window of time in a Kafka topic.
   b. Call `DBMS_KAFKA.LOAD_TEMP_TABLE` to load a global temporary table with the set of rows from Kafka that you want to analyze.
   c. Use `SELECT` from the OSAK global temporary table.
   d. Process the data.
3. When done with the application, use `DBMS_KAFKA.DROP_SEEKABLE_APP` to drop the application.

109.2 DBMS_KAFKA LOADING Mode

Use the `DBMS_KAFKA LOADING` mode packages to load Kafka data incrementally into Oracle Database.

Loading procedures enable you to load available Kafka records into Oracle Database, which can then serve as a data warehouse for that data. You can then combine that data with Oracle Database tables for analytics.

An application declares that it is a loading application by calling the PL/SQL procedure `DBMS_KAFKA.CREATE_LOAD_APP` to initialize state for subsequent calls to `DBMS_KAFKA.EXECUTE_LOAD_APP`. `DBMS_KAFKA.CREATE_LOAD_APP` creates a single view over all partitions of the topic. An application can optionally call the `DBMS_KAFKA.INIT_OFFSET[_TS]` procedure to set the starting point in Kafka topic partitions.

The `DBMS_KAFKA.EXECUTE_LOAD_APP` procedure is called in an application loop to load data from where the previous call left off to the current high water mark of the Kafka topic. This procedure runs in an autonomous transaction.

When you are finished working with the Kafka data, you can remove the application by using `DBMS_KAFKA.DROP_LOAD_APP`.

109.2.1 CREATE_LOAD_APP

This procedure creates an Oracle SQL Access to Kafka Load application that retrieves data from all partitions in a Kafka topic to load that Kafka data into an Oracle Database table. It also creates, if not already present, a metadata view that is used to inspect the Kafka cluster for live topic and partition information regarding the Kafka topic. This view is created once, and serves all applications that are sharing the same cluster. This model is restrictive in that only one application instance is allowed to call `DBMS_KAFKA.EXECUTE_LOAD_APP` for the created `LOAD` application.

Parameters

Parameter	Description
`cluster_name`	The name of a registered Oracle SQL access to Kafka cluster that has the topic that you want to associate with this application. Case-insensitive. The registered cluster names can be obtained from the OSAK Administrator, by using the following statement: `SELECT cluster_name FROM sys.user_kafka_clusters;`
`application_name`	The application name. This parameter is also used as the Kafka group that can read the topic. Case-insensitive.
`topic_name`	The topic name in the Kafka cluster whose contents you want to retrieve. Case-sensitive.

Parameter	Description
options	Includes a list of properties formatted as a JSON document. Options are described in more detail in the topic "DBMS_KAFKA OPTIONS Passed to CREATE_XXX_APP".

Examples

Example 109-1 CREATE_LOAD_APP Procedure for Oracle SQL Access to Kafka

In the following example, a load application called ExampleApp is created for data from the Kafka cluster ExampleCluster, using the Kafka topic my-company-app-event1 and the option ProducerRecord.

```
PROCEDURE CREATE_LOAD_APP (
                           ExampleCluster         IN  VARCHAR2,
                           ExampleApp             IN  VARCHAR2,
                           my-company-app-event1  IN  VARCHAR2,
                           ProducerRecord         IN  CLOB
```

Related Topics

- DBMS_KAFKA OPTIONS Passed to CREATE_xxx_APP

109.2.2 DROP_LOAD_APP

The DBMS_KAFKA_ADM procedure DROP_LOAD_APP drops the Oracle SQL for Kafka (OSAK) LOAD application, and removes related metadata.

Syntax

```
EXEC DROP_LOAD_APP (
                          cluster_name     IN VARCHAR2,
                          application_name IN VARCHAR2
  );
```

Parameters

Table 109-1 DROP_LOAD_APP Procedure Parameters for DBMS_KAFKA_ADM.

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.
application_name	Name of an existing application associated with the Kafka cluster. Case-insensitive.

Usage Notes

Use this procedure to drop an Oracle SQL access to Kafka (OSAK) application when you no longer want to load data from a Kafka topic into an Oracle Database table.

Examples

Suppose you have completed your work with the application called `ExampleApp` with data from the Kafka cluster `ExampleCluster`. Use this procedure to drop the application:

```
EXEC DROP_LOAD_APP (
                          ExampleCluster   IN VARCHAR2,
                          ExampleApp       IN VARCHAR2
  );
```

109.2.3 EXECUTE_LOAD_APP

The `DBMS_KAFKA` procedure `EXECUTE_LOAD_APP` loads a user table from a dedicated Oracle SQL access to Kafka (OSAK) view. To use this procedure, you must previously have created a load application with `CREATE_LOAD_APP`.

Syntax

```
PROCEDURE EXECUTE_LOAD_APP (
                             cluster_name        IN  VARCHAR2,
                             application_name    IN  VARCHAR2,
                             target_table        IN  VARCHAR2,
                             records_loaded      OUT INTEGER,
                             parallel_hint      IN INTEGER DEFAULT 0
);
```

Parameters

Table 109-2 EXECUTE_LOAD_APPS Procedure Parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of a registered Oracle SQL access to Kafka cluster that has the topic that you want associated with this application. Case-insensitive. The registered cluster names can be obtained from the OSAK Administrator, or by using or by using the following SQL statement: `SELECT cluster_name from SYS.USER_KAFKA_CLUSTERS;`
application_name	The name of an existing application associated with the Kafka cluster Case-insensitive.
target_table	A target table in the Oracle Database that will be loaded with data from a Kafka topic. It must be consistent with some or all of the columns retrieved from the Kafka cluster topic by the OSAK view created by the `LOAD` operation.
records_loaded	(OUT) The number of Kafka records loaded

Table 109-2 (Cont.) EXECUTE_LOAD_APPS Procedure Parameters for DBMS_KAFKA

Parameter	Description
parallel_hint	(IN) (Optional) The degree of parallelism to use when loading the application that maps exclusively to a particular OSAK view. If a parallel hint is not specified, or it is less than or equal to 1, then parallelism is not used to load the table. **Note:** Only use a parallel hint when `PARALLEL_DEGREE_POLICY` for the user session or the system is set to `MANUAL`. For all other policies (for example, `AUTO`), no parallel hint should be passed. If parallel hint exceeds the granule count of an OSAK view, then an exception will be raised.

Usage Notes

Oracle SQL access to Kafka (OSAK) views are used transparently by PL/SQL calls to load records from a topic in a Kafka cluster previously created by the `CREATE_LOAD_APP` procedure. Applications call `EXECUTE_LOAD_APP` to load an Oracle Database user table from a dedicated OSAK view.

> **Note:**
>
> Only one user can query an OSAK temporary table using a single database application instance serving the OSAK application. However by using `EXECUTE_LOAD_APP`, you load Kafka data into a permanent database table that is accessible by multiple applications.

Each call to the `EXECUTE_LOAD_APP` procedure reads new records from the Kafka topic, and inserts these records into the Oracle Database table. `EXECUTE_LOAD_APP` also advances offsets of all Kafka partitions, so that the next time `EXECUTE_LOAD_APP` is run, it will insert new rows. Using EXECUTE_LOAD_APP enables you to perform incremental loads, so that you can update the Oracle Database table with updates to the Kafka topic. Because that Kafka data is moved into standard Oracle Database tables, they become available for processing and analysis by multiple applications.

Examples

Suppose you have completed your initial cluster definition for the Kafka cluster `ExampleCluster`, and registered the cluster. Next, you use this procedure to load data from the Kafka cluster `ExampleCluster` into the application `ExampleLoadApp`:

```
DECLARE
    v_records_inserted INTEGER;
BEGIN
    SYS.DBMS_KAFKA.EXECUTE_LOAD_APP (

'ExampleCluster',

'ExampleLoadApp',
```

```
'ExampleLoadTable',

v_records_inserted);
END;
```

109.3 DBMS_KAFKA Global Temporary Tables

Use the `DBMS_KAFKA LOAD_TEMP_TABLE` mode packages to load Kafka data into a temporary table, from which the Kafka data view in Oracle Database is created.

The `DBMS_KAFKA.LOAD_TEMP_TABLE` procedure is called in an application loop to load data into an Oracle SQL access to Kafka (OSAK) application. For both `STREAMING` and `SEEKABLE` applications, you use `CREATE_APP_xxx` (where xxx is either `STREAMING` or `SEEKABLE`) to create the application. Next, you use an application loop while calling `LOAD_TEMP_TABLE` for the application, and process the data loaded into the temporary tables.

109.3.1 LOAD_TEMP_TABLE

The `DBMS_KAFKA` procedure `LOAD_TEMP_TABLE` selects all data from an Oracle SQL access to Kafka view into a temporary table.Use this procedure to create an Oracle SQL access to Kafka (OSAK) dedicated temporary table that an application can use with SQL queries to analyze the data, or to or join with Oracle Database tables.

Syntax

```
FUNCTION LOAD_TEMP_TABLE(
                   temporary-table-name   IN VARCHAR2
  ) RETURN INTEGER;
```

Parameters

Table 109-3 LOAD_TEMP_TABLE Procedure Parameters for DBMS_KAFKA

Parameter	Description
temporary-table-name	Name of the temporary table that you want to create Type: VARCHAR Case-insensitive.
parallel-hint	(Optional) the degree of parallelism when loading the global temporary table that maps exclusively to a particular OSAK view. If a parallel hint is not specified, or it is less than or equal to 1, then parallelism is not used to load the table. **Note:** Only use a parallel hint when `PARALLEL_DEGREE_POLICY` for the user session or the system is set to `MANUAL`. For all other policies (for example, `AUTO`), no parallel hint should be passed. If parallel hint exceeds the granule count of an OSAK view, then an exception will be raised.

Usage Notes

Oracle SQL access to Kafka (OSAK) views are used transparently by PL/SQL calls to load records from Kafka topics into either a dedicated OSAK global temporary table or into a user-defined table. `STREAMING` and `SEEKABLE` applications call `LOAD_TEMP_TABLE` which loads an OSAK global temporary table from a dedicated OSAK view, while `LOAD` applications call `EXECUTE_LOAD_APP` to load a user table from a dedicated OSAK view.

> **Caution:**
>
> OSAK views are dedicated views, which can serve only one application instance. They are not shared by other instances of the application, and they should not be queried by a generalized tool (for example, SQL*Plus or SQL Developer). Concurrent access of an OSAK view can cause a race condition where a tool could inadvertently read more rows than the dedicated application, and inadvertently advance Kafka offsets beyond what the dedicated application has read. As a result, the application can then fail to obtain from the Kafka topic records that should have been processed.

Examples

Suppose you want to load and process records associated with a streaming application called `ExampleApp` from Kafka cluster `ExampleCluster`. By default, the OSAK views are configured to read from the earliest record present to last record currently published when they are initially created. After you create the streaming application with `DBMS_KAFKA.CREATE_STREAMING_APP`, the following is an example Kafka data processing loop for that streaming application:

```
BEGIN
   LOOP
       SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_EXAMPLEAPP_0);
       FOR kafka_record IN (
             SELECT kafka_offset offset
                    FROM ORA$DKVGTT_EXAMPLECLUSTER_EXAMPLEAPP_0)
       LOOP
              SYS.DBMS_OUTPUT.PUT_LINE ('Processing record: ' ||
kafka_record.offset);
              --application logic to process the Kafka records
       END LOOP;
       IF (application logic was successful) THEN
            --Update internal metadata to confirm Kafka records were
successfully processed
            SYS.DBMS_KAFKA.UPDATE_OFFSET

('ORA$DKV_EXAMPLECLUSTER_EXAMPLEAPP_0');
            COMMIT;
       ELSE
            --add your application logic to correct for any failures
       END IF;
   END LOOP;
END;
```

109.4 DBMS_KAFKA Streaming Mode

Use the `DBMS_KAFKA STREAMING` mode packages to stream data from a Kafka topic into an Oracle SQL access to Kafka (OSAK) global temporary table from a dedicated OSAK view. To use this package, you must first use `LOAD_TEMP_TABLE` to create the dedicated OSAK view. Use `STREAMING` mode for applications that require access to Kafka topics in a sequential manner from the beginning, or from a specific starting point in a Kafka topic. This mode allows a SQL query using an OSAK temporary table to access Kafka records sequentially in an application processing loop. An application declares that it is a streaming application by calling the PL/SQL procedure `DBMS_KAFKA.CREATE_STREAMING_APP` to initialize state for subsequent queries of OSAK views.

109.4.1 CREATE_STREAMING_APP

The `DBMS_KAFKA` procedure `CREATE_STREAMING_APP` creates an Oracle SQL access to Kafka (OSAK) streaming application. The application includes a set of dedicated OSAK global temporary tables and OSAK views that are used for retrieving new and unread records from partitions in a Kafka topic. It also creates, if not already present, a metadata view that is used to inspect the Kafka cluster for live topic and partition information regarding the Kafka topic. This view is created once, and serves all applications that are sharing the same cluster.

Syntax

```
PROCEDURE CREATE_STREAMING_APP (
                                  cluster_name          IN  VARCHAR2,
                                  application_name      IN  VARCHAR2,
                                  topic_name            IN  VARCHAR2,
                                  options               IN  CLOB,
                                  view_count            IN  INTEGER DEFAULT 1
  );
```

Parameters

Table 109-4 CREATE_STREAMING_APP Procedure Parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of a registered Oracle SQL access to Kafka cluster that has the topic that you want associated with this application. Case-insensitive. The registered cluster names can be obtained from the OSAK Administrator, or by using or by using the following SQL statement: `SELECT cluster_name from SYS.USER_KAFKA_CLUSTERS;`
application_name	The name of the application. Also used as the Kafka group that will read the topic. Case-insensitive.
topic_name	The topic name in the Kafka cluster whose contents will be retrieved.

Table 109-4 (Cont.) CREATE_STREAMING_APP Procedure Parameters for DBMS_KAFKA

Parameter	Description
options	Includes a list of properties formatted as a JSON document. Options are described in more detail in the topic "DBMS_KAFKA OPTIONS Passed to CREATE_XXX_APP".
view_count	(OPTIONAL) Identifies the number of Oracle SQL access to Kafka (OSAK) view pairs to create. Valid values are 1 to *N*, where *N* is the number of Kafka partitions in a topic, or 0, which defaults to *N*. (Default is 1).

Usage Notes

Each OSAK view is exclusively used by one instance of an Oracle SQL access to Kafka application. Each application instance call populates the view with Kafka rows. The application then can then run one or more SQL queries against the content in the OSAK view A STREAMING application is different from a LOAD or SEEKING application in that it can choose how many OSAK views need to be created. The number of OSAK views must be between 1 and *N* where *N* is the number of partitions in the Kafka topic.

As with other types of OSAK applications, each application instance exclusively queries one unique OSAK temporary table. Each OSAK view includes the cluster name, the application name, and an application instance identifier (ID). Creating multiple application instances enables applications to scale out and divide the workload of analyzing Kafka data across application instances running concurrently on one or more threads, processes, or systems.

The number of Kafka partitions bound to a specific OSAK view and its associated OSAK global temporary table will vary depending upon how many views are created and how many partitions exist. If *N* OSAK view/temporary table pairs are created, then the application user must have been allocated at least *N* sessions per user, so that *N* application instances can run concurrently.

Examples

Suppose you want to create a set of four views for a streaming application called `ExampleApp`, streamed from a Kafka topic called `ExampleTopic` in the Kafka cluster `ExampleCluster` that has four partitions, where each view is associated with one partition. You can enter the following statement:

```
DECLARE
   v_options  VARCHAR2;
BEGIN
  v_options := '{"fmt" : "DSV", "reftable" : "user_reftable_name"}';
  SYS.DBMS_KAFKA.CREATE_STREAMING_APP (
                                     'ExampleCluster',
                                     'ExampleApp',
                                     'ExampleTopic',
                                      v_options,
                                      4);
END;
/
```

Alternatively, to create one view for an application that is associated with all four partitions of the topic, you enter the following statement:

```
DECLARE
v_options VARCHAR2;
BEGIN
              v_options := '{"fmt" : "DSV", "reftable" :
"user_reftable_name"}';
 SYS.DBMS_KAFKA.CREATE_STREAMING_APP (

'ExampleCluster',

'ExampleApp',

'ExampleTopic',
                                                        v_options,
                                                        1);

END;
/
```

Related Topics

- DBMS_KAFKA OPTIONS Passed to CREATE_xxx_APP

109.4.2 DROP_STREAMING_APP

The `DBMS_KAFKA` procedure `DROP_STREAMING_APP` drops the streaming application.This function removes the Oracle SQL access to Kafka (OSAK) view, and drops all associated database objects.

Syntax

```
PROCEDURE DROP_STREAMING_APP (
                                cluster_name     IN VARCHAR2,
                                application_name IN VARCHAR2
  );
```

Parameters

Table 109-5 DROP_STREAMING_APP Procedure Parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.
application_name	The name of an existing application using the Kafka cluster. Case-insensitive.

Usage Notes

When you are done with an application, you use this function to drop an OSAK view associated with an application.

Examples

Suppose you have completed your work with the Kafka cluster `ExampleCluster`, which was being used by Streaming application `ExampleApp`. You can then use this procedure to drop the application::

```
EXEC SYS.DBMS_KAFKA.DROP_STREAMING_APP (
                                        'ExampleCluster', 'ExampleApp');
```

109.4.3 INIT_OFFSET

The `DBMS_KAFKA_ADM` procedure `INIT_OFFSET` enables you to select a particular offset as the starting point for reading Kafka data. Use this option when you want to select a particular starting point in the data to load, instead of loading data from the first record available.

Syntax

```
PROCEDURE INIT_OFFSET (
         view_name IN VARCHAR2,
         record_count IN INTEGER,
         water_mark IN VARCHAR2 ;
```

Parameters

Table 109-6 INIT_OFFSET Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing OSAK view (VARCHAR2)
record_count	record number (INTEGER)
water_mark	Watermark (VARCHAR) The high or low watermark that indicates the desired relative positioning. Values are restricted to `WATER_MARK_HIGH` (`'WMH'`) or `WATER_MARK_LOW` (`'WM'`) constants. Default: `WATER_MARK_HIGH`

Usage Notes

The `INIT_OFSET` procedure enables `STREAMING` or `LOAD` applications to start reading current records after the application has been created without being forced to first read a backlog of old records that are no longer interesting.

Use this function to specify a starting point based on the difference (delta) number of records from either the high or low water mark of every partition where you want your application to read data, instead of starting the read point from the beginning of the data records in the Kafka topic.

Examples

Suppose that you want your application to restart the processing with the last 100 records available and continue on from that point. You can achieve this result by running the following

procedure before the application is restarted, or as part of the application logic before data retrieval loop:

```
SYS.DBMS_KAFKA.INIT_OFFSET (

'ORA$DKV_EXAMPLECLUSTER_EXAMPLEAPP_0',
                                        100,
SYS.DBMS_KAFKA.WATER_MARK_HIGH);
```

109.4.4 INIT_OFFSET_TS (Milliseconds Since Epoch)

The `DBMS_KAFKA` procedure `INIT_OFFSET_TS` using milliseconds since epoch specifies a starting offset using a timestamp.

Using milliseconds since epoch initializes the starting offset related to a timestamp for each Kafka partition belonging to the OSAK view. INIT_OFFSET_TS would typically be called at the outset of a new application instance dedicated to processing the view or recovering after an application instance shutdown or failure.

.

Syntax

```
PROCEDURE INIT_OFFSET_TS (
                            view_name           IN VARCHAR2,
                            start_timestamp_ms  IN INTEGER);
```

Parameters

Table 109-7 INIT_OFFSET_TS (Millseconds Since Epoch) Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Kafka cluster (`VARCHAR`) Case-insensitive.
start_timestamp	The timestamp of the offset from which you want to start your application (INTEGER). The first record returned will have a timestamp equal to the timestamp provided, or the nearest timestamp greater than the timestamp provided.

Usage Notes

`INIT_OFFSET_TS` using milliseconds since a Kafka epoch initializes the starting offset related to a timestamp for each Kafka partition belonging to the OSAK view. A typical use case is to call `INIT_OFFSET_TS` with a specified epoch starting point at the outset of a new application instance dedicated to processing the view, or to recovering after an application instance shutdown or failure.

This procedure serves to position the processing of Kafka topic records to a point that is relatively current, potentially skipping unprocessed older records in the Kafka partitions.

> **Note:**
>
> Be aware that the time between initializing the offset and the first fetch can be delayed. During this gap in time, it is possible that the record for the chosen offset can be deleted due to either the record exceeding the Kafka retention time, or to the record being explicitly removed.

Examples

Suppose that you want to calculate and return the number of milliseconds since the epoch time from the input `TIMESTAMP`. The timestamp is considered to be in the session's timezone unless the timezone is provided. You provide the parameter `datetime` (an integer) Timestamp to convert to milliseconds since epoch time. The parameter timezone (integer) is optional, providing the timezone of the timestamp. If you do not specify a timezone, then the timezone defaults to the session's timezone:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS ('ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                               1593829800,
                               1593833400);

SYS.DBMS_KAFKA.LOAD_TEMP_TABLE ('ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0');
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
```

109.4.5 INIT_OFFSET_TS (Timestamp with Separate Timezone Parameter)

The `DBMS_KAFKA` procedure `INIT_OFFSET_TS` using a TIMESTAMP with a string parameter specifying the TIMESTAMP TIME ZONE. specifies a starting position to return dataUsing a timestamp specified with a timestamp with time zone positions the processing of Kafka data to a timezone related to a timestamp for each Kafka partition belonging to the Oracle SQL access to Kafka (OSAK) view. `INIT_OFFSET_TS` would typically be called at the outset of a new application instance dedicated to processing the view or recovering after an application instance shutdown or failure.

Syntax

```
PROCEDURE INIT_OFFSET_TS (
                           view_name           IN VARCHAR2,
                           start_timestamp     IN TIMESTAMP,
                           timezone            IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 109-8 INIT_OFFSET_TS Timestamp with Time Zone Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Kafka cluster (`VARCHAR`) Case-insensitive.

Table 109-8 (Cont.) INIT_OFFSET_TS Timestamp with Time Zone Procedure Parameters for DBMS_KAFKA

Parameter	Description
start_timestamp	The timestamp of the offset from which you want to start your application (INTEGER). The first record returned will have a timestamp equal to the timestamp provided, or the nearest timestamp greater than the timestamp provided.
timezone	The timezone of the timestamp (INTEGER). If no value is provided, then the default is to use the session timezone.

Usage Notes

`INIT_OFFSET_TS` using a timestamp with a separate time zone parameter initializes the starting offset related to a timestamp and a time zone for each Kafka partition belonging to the OSAK view. A typical use case is to call `INIT_OFFSET_TS` with a timestamp with a separate time zone parameter at the outset of a new application instance dedicated to processing the view, or to recovering after an application instance shutdown or failure.

This procedure serves to position the processing of Kafka topic records to a point that is relatively current, potentially skipping unprocessed older records in the Kafka partitions.

> **Note:**
>
> Be aware that the time between initializing the offset and the first fetch can be delayed. During this gap in time, it is possible that the record for the chosen offset can be deleted due to either the record exceeding the Kafka retention time, or to the record being explicitly removed.

Examples

Suppose that you want to select a later time to process data from a Kafka topic than at the first record available. For example, if a data center was down for maintenance for a weekend, and you only want to process new Kafka data generated after the data center was reopened at 6 P.M. (18:00:00), then you want to start the data record timestamped after 6 PM. To achieve this with an OSAK view, run the following procedure before the application is restarted, or as part of the application logic before the data retrieval loop:

```
SYS.DBMS_KAFKA.INIT_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_EXAMPLEAPP_0',
                                          TO_DATE ('2023/07/05 18:00:00',

'YYYY/MM/DD HH:MI:SS'))
                                          timestamp=1603507387101;
```

109.4.6 INIT_OFFSET_TS (Timestamp with Time Zone)

The `DBMS_KAFKA` procedure `INIT_OFFSET_TS` using timestamp with timezone specifies a starting offset using a timestamp.Using a timestamp specified with a time zone initializes the starting

offset related to a timestamp for each Kafka partition belonging to the Oracle SQL access to Kafka (OSAK) view. `INIT_OFFSET_TS` would typically be called at the outset of a new application instance dedicated to processing the view or recovering after an application instance shutdown or failure.

Syntax

```
PROCEDURE INIT_OFFSET_TS (
                            view_name           IN VARCHAR2,
                            start_timestamp     IN TIMESTAMP WITH TIME ZONE);
```

Parameters

Table 109-9 INIT_OFFSET_TS Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Kafka cluster (`VARCHAR`) Case-insensitive.
start_timestamp	The timestamp of the offset from which you want to start your application (INTEGER). The first record returned will have a timestamp equal to the timestamp provided, or the nearest timestamp greater than the timestamp provided.
timezone	(Optional). The timezone of the timestamp (INTEGER). If not provided, then it defaults to the session timezone

Usage Notes

`INIT_OFFSET_TS` initializes the starting offset related to a timestamp for each Kafka partition belonging to the OSAK view. A typical use case is to call INIT_OFFSET_TS at the outset of a new application instance dedicated to processing the view, or to recovering after an application instance shutdown or failure.

This procedure serves to position the processing of Kafka topic records to a point that is relatively current, potentially skipping unprocessed older records in the Kafka partitions.

> **Note:**
>
> Be aware that the time between initializing the offset and the first fetch can be delayed. During this gap in time, it is possible that the record for the chosen offset can be deleted due to either the record exceeding the Kafka retention time, or to the record being explicitly removed.

Examples

Timestamp with specified time zone

Suppose that you want to select a later time to process data from a Kafka topic than at the first record available. For example, if a data center was down for maintenance for a weekend, and you only want to process new Kafka data generated after the data center was reopened at 6 P.M. (18:00:00), then you want to start the data record timestamped after 6pm. To achieve this

with an OSAK view, run the following procedure before the application is restarted, or as part of the application logic before the data retrieval loop:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS ('ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                               TO_TIMESTAMP ('2020/07/04 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                               TO_TIMESTAMP ('2020/07/04 03:30:00',
'YYYY/MM/DD HH:MI:SS'),
                               'UTC');

SYS.DBMS_KAFKA.LOAD_TEMP_TABLE ('ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0');
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
```

Timestamp without a specified time zone

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS ('ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                               TO_TIMESTAMP_TZ ('2020/07/04 02:30:00 -8:00',
'YYYY/MM/DD HH:MI:SS TZH:TZM',
                               TO_TIMESTAMP_TZ ('2020/07/04 03:30:00 -8:00',
'YYYY/MM/DD HH:MI:SS TZH:TZM'));

SYS.DBMS_KAFKA.LOAD_TEMP_TABLE ('ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0');
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
```

109.4.7 UPDATE_OFFSET

The `DBMS_KAFKA` procedure `UPDATE_OFFSET` update the last Kafka offsets read so the next pass in the loop will retrieve and process new unread Kafka records. `UPDATE_OFFSET` transparently advances Kafka partition offsets of the Kafka group ID for all of the partitions belonging to the Oracle SQL access to Kafka (OSAK) view, so that for every call to `DBMS_KAFKA.LOAD_TEMP_TABLE`, a new set of unread Kafka records is retrieved and processed.

Syntax

```
PROCEDURE UPDATE_OFFSET (view_name IN VARCHAR2);
```

Parameters

Table 109-10 UPDATE_OFFSET Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	The name of an OSAK view Case-insensitive.

Usage Notes

`UPDATE_OFFSET` advances the Kafka offset read point to the next set of Kafka records if the processing of the earlier records was successful. `UPDATE_OFFSET` also initiates an Oracle Database transaction, if a transaction is not already started, and records the last offsets in metadata tables. For this reason, applications should commit a transaction after each call to `UPDATE_OFFSET`. Because OSAK manages offsets within an Oracle Database transaction, in which ACID (Atomicity, Consistency, Isolation, Durability) is preserved in the transaction, no

records are lost or reread. If the transaction fails to complete successfully, then offsets are not advanced, and the application will pick up where it previously off off when it resumes.

Examples

A typical use case is to advance an offset read to the next set of Kafka records after the previous set has been processed successfully, such as in a Streaming mode application. For example:

```
BEGIN
   LOOP
       SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_EXAMPLEAPP_0);
       FOR kafka_record IN (
             SELECT kafka_offset offset
                    FROM ORA$DKVGTT_EXAMPLECLUSTER_EXAMPLEAPP_0)
       LOOP
              SYS.DBMS_OUTPUT.PUT_LINE ('Processing record: ' ||
kafka_record.offset);
              --application logic to process the Kafka records
       END LOOP;
       IF (application logic was successful) THEN
            --Update internal metadata to confirm Kafka records were
successfully processed
            SYS.DBMS_KAFKA.UPDATE_OFFSET

('ORA$DKV_EXAMPLECLUSTER_EXAMPLEAPP_0');
            COMMIT;
       ELSE
            --application logic to correct what failed
       END IF;
   END LOOP;
END;
```

You can also use it as part of the process of access an OSAK view directly to retrieve updates. For example, suppose you have a Streaming mode application that is monitoring average temperatures over time in a laboratory called `Lab1`. You can use `UPDATE_OFFSET` to retrieve the average temperature since it was last checked:

```
EXEC DBMS_KAFKA.ENABLE_VIEW_QUERY(

'ORA$DKV_EXAMPLECLUSTER_STREAMINGAPP_0');
SELECT AVG(temperature) FROM ORA$DKV_EXAMPLECLUSTER_STREAMINGAPP_0
                                                    WHERE sensor_name =
'LAB1_SENSOR';
EXEC DBMS_KAFKA.UPDATE_OFFSET(

'ORA$DKV_EXAMPLECLUSTER_STREAMINGAPP_0');
COMMIT;
```

109.5 DBMS_KAFKA Seekable Mode

Use the `DBMS_KAFKA SEEKABLE` mode packages to read Kafka records that exist between two points in time.
An application declares that it is a loading application by calling the PL/SQL procedure `DBMS_KAFKA.CREATE_SEEKABLE_APP` to set up and create a seekable Oracle SQL access to Kafka (OSAK) view and temporary table.

Use the Seekable mode view with `SEEK_OFFSET_TS` to search through the Kafka data to find a set of data whose timestamps exist between two points in time. You can then use `LOAD_TEMP_TABLE` to load that data range into an OSAK temporary table, and run one or more application queries against the OSAK global temporary table containing the Kafka records in that timestamp range. When you have completed your queries, you can then use `DROP_SEEKABLE_APP` to drop the application.

.

109.5.1 CREATE_SEEKABLE_APP

The `DBMS_KAFKA` procedure `CREATE_SEEKABLE_APP` Creates one Oracle SQL access to Kafka (OSAK) view and an associated global temporary table You can use `CREATE_SEEKABLE_APP` to seek and load Kafka records between a particular timeframe. Use of this procedure is restricted to one application instance to perform seek operations.

Syntax

```
PROCEDURE CREATE_SEEKABLE_APP (
                                cluster_name          IN  VARCHAR2,
                                application_name      IN  VARCHAR2,
                                topic_name            IN  VARCHAR2,
                                options               IN  CLOB
```

Parameters

Parameter	Description
cluster_name	The name of a registered Oracle SQL access to Kafka cluster that has the topic that you want to associate with this application. Case-insensitive. The registered cluster names can be obtained from the OSAK Administrator, by using the following statement: `SELECT cluster_name FROM sys.user_kafka_clusters;`
application_name	The application name. This parameter is also used as the Kafka group that can read the topic. Case-insensitive.
topic_name	The topic name in the Kafka cluster whose contents you want to retrieve. Case-sensitive.

Parameter	Description
options	Includes a list of properties formatted as a JSON document. Options are described in more detail in the topic "DBMS_KAFKA OPTIONS Passed to CREATE_XXX_APP".

Examples

Example 109-2 CREATE_SEEKABLE_APP Procedure for Oracle SQL Access to Kafka

```
DECLARE
   v_options  VARCHAR2;
BEGIN
  v_options := '{"fmt" : "DSV", "reftable" : "user_reftable_name"}';
  SYS.DBMS_KAFKA.CREATE_SEEKABLE_APP (
                                     'ExampleCluster',
                                     'ExampleApp',
                                     'ExampleTopic',
END;
/
```

Related Topics

- DBMS_KAFKA OPTIONS Passed to CREATE_xxx_APP

109.5.2 DROP_SEEKABLE_APP

The DBMS_KAFKA procedure DROP_SEEKABLE_APP drops the seeking application. This function removes the Oracle SQL access to Kafka (OSAK) view, and its related metadata.

Syntax

```
PROCEDURE DROP_SEEKABLE_APP (
                               cluster_name     IN VARCHAR2,
                               application_name IN VARCHAR2
  );
```

Parameters

Table 109-11 DROP_SEEKABLE_APP Procedure Parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.
application_name	The name of an existing application using the Kafka cluster. Case-insensitive.

Usage Notes

When you are done with an application, you use this function to drop an OSAK view associated with an application.

Examples

Suppose you have completed your work with the Kafka cluster `ExampleCluster`, which was being used by the Seekable application `ExampleApp`. You can then use this procedure to drop the application:

```
EXEC SYS.DBMS_KAFKA.DROP_SEEKABLE_APP (
                                        'ExampleCluster', 'ExampleApp');
```

109.5.3 SEEK_OFFSET_TS (Timestamp with Separate Timezone Parameter)

The `DBMS_KAFKA` procedure `SEEK_OFFSET_TS` using a timestamp with a timezone parameter positions an Oracle SQL access to Kafka (OSAK) view to start reading Kafka records between two Kafka epoch timestamps that you specify.Use this procedure to specify a timeframe for records within a Kafka topic that you want to seek, as specified by epoch time in milliseconds.

Syntax

```
PROCEDURE SEEK_OFFSET_TS(
                view_name           IN VARCHAR2,
                start_timestamp_ms  IN INTEGER,
                end_timestamp_ms    IN INTEGER,
                timezone  IN INTEGER);
```

Parameters

Table 109-12 SEEK_OFFSET_TS (Timestamp with Separate Timezone Parameter)Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Kafka cluster (`VARCHAR`) Case-insensitive.
start_timestamp	The timestamp of the offset from which you want to start your application (INTEGER). The last record returned will have a timestamp equal to the timestamp provided, or the nearest timestamp greater than the timestamp provided.

Table 109-12 (Cont.) SEEK_OFFSET_TS (Timestamp with Separate Timezone Parameter)Procedure Parameters for DBMS_KAFKA

Parameter	Description
end_timestamp	The timestamp of the offset from which you want to end your application (INTEGER). The last record returned will have a timestamp equal to the timestamp provided, or the nearest timestamp less than the timestamp provided. The range of records returned is inclusive of the start timestamp and exclusive on the end timestamp. For example, a range of 2:00PM to 3:00PM returns records with timestamps in the range of 2:00:00.000000 PM to 2:59:59.999999PM This range enables the application to search from 2:00 to 3:00, and then 3:00 to 4:00, without having overlapping records between the sets.
timezone	The timezone of both of the timestamp arguments (INTEGER). If no value is provided, then the default is to use the session timezone.

Usage Notes

`SEEK_OFFSET_TS` seeks Kafka topic records that exist within a specific timeframe that you select, specified by Kafka epoch time in milliseconds. The last record returned is the last record published.

The purpose of `SEEK_OFFSET_TS` is to position an OSAK view to start reading Kafka records within a given timeframe, as defined by the epoch timestamps that you specify. If the window of records exceeds the range of actual records in a Kafka topic, this procedure will return whatever records do exist.

For example:

- If the timestamps are either both below the low water mark, or both above the high water mark, then no records will be returned.
- If the start timestamp is below the low water mark, then the first record returned will be the low water mark.
- If the end timestamp is above the high water mark, then the last record returned will be the difference of the high water mark (HWM), minus 1 (HWM - 1). For example, in a new topic has 100 records, the offset range is from 0 to 99, and the HWM is set at 100.

> **Note:**
>
> The data retrieved by a SEEKABLE application using Kafka epoch timeframes as the specification for retrieval cannot include any outlier records that are delivered after the load of the OSAK global temporary table, even if they have a timestamp that is within the timestamp window provided.

Examples

Suppose that you want to investigate issues that have occurred in the past. If the data is still present in the Kafka steam, then you can create a Seekable application by calling

DBMS_KAFKA.CREATE_SEEKABLE_APP. You can then call the SEEK_OFFSET_TS procedure to request the OSAK view to retrieve a range of data records. For example, if an IT consultant was informed that a production issue occurred around 3:00 in the morning, then the consultant can use the following procedure to load the temporary table, and then select to retrieve an hour's worth of data around that time:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                                          TO_DATE ('2023/04/02 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                                          TO_DATE ('2023/04/02 03:30:00',
'YYYY/MM/DD HH:MI:SS'));
SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0);
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
```

Another use case might be if an application with sequential access to a Kafka stream detected a potential anomaly, and was configured to insert a row into an anomaly table. The anomaly table would include the Kafka timestamp, as well as any other data it was configured to record. In that case, another application can then use this information to retrieve records around the suspected record to see if there were any other issues. To achieve this goal, run the following procedure, load the temporary table, and then select and apply application logic to the results:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                                          TO_DATE ('2020/07/04 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                                          TO_DATE ('2020/07/04 03:30:00',
'YYYY/MM/DD HH:MI:SS'));

SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0);
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
--application logic
```

109.5.4 SEEK_OFFSET_TS (Timestamp in Milliseconds)

The DBMS_KAFKA procedure SEEK_OFFSET_TS (timestamp in milliseconds) seeks Kafka records in a topic that exist between a window of time specified by epoch time in milliseconds.

Syntax

```
PROCEDURE SEEK_OFFSET_TS(
                view_name           IN VARCHAR2,
                start_timestamp_ms  IN INTEGER,
                end_timestamp_ms    IN INTEGER);
```

Parameters

Table 109-13 SEEK_OFFSET_TS (Timestamp in Milliseconds) Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Oracle SQL access to Kafka (OSAK) view. Case-insensitive.
start_timestamp_ms	The timestamp in milliseconds where you want to seek to the first record
end_timestamp_ms	The timestamp in milliseconds where you want to seek to the last record

Usage Notes

`SEEK_OFFSET_TS` seeking a timestamp range in milliseconds)seeks Kafka topic records that exist within a starting timestamp and ending timestamp window that you select, specified by Kafka epoch time in milliseconds. The last record returned is the last record published.

The purpose of `SEEK_OFFSET_TS` using milliseconds to specify the window is to position an OSAK view to start reading Kafka records within a given timeframe, as defined by the epoch timestamps that you specify. If the window of records exceeds the range of actual records in a Kafka topic, this procedure will return whatever records do exist.

For example:

- If the timestamps are either both below the low water mark, or both above the high water mark, then no records will be returned.
- If the start timestamp is below the low water mark, then the first record returned will be the low water mark.
- If the end timestamp is above the high water mark, then the last record returned will be the difference of the high water mark (HWM), minus 1 (HWM - 1). For example, in a new topic has 100 records, the offset range is from 0 to 99, and the HWM is set at 100.

> **Note:**
>
> Kafka record timestamps are either assigned by Kafka (that is, by transaction time), or are assigned by the application (that is, valid or decision time). The data retrieved by a SEEKABLE application will not include outlier records that were delivered after the load of the OSAK global temporary table that have a timestamp that is within the timestamp window provided.

Examples

Suppose that you want to investigate issues that have occurred in the past. If the data is still present in the Kafka steam, then you can create a Seekable application by calling `DBMS_KAFKA.CREATE_SEEKABLE_APP`. You can then call the `SEEK_OFFSET_TS` procedure to request the OSAK view to retrieve a range of data records. For example, if an IT consultant was informed that a production issue occurred around 3:00 in the morning, then the consultant

can use the following procedure to load the temporary table, and then select to retrieve an hour's worth of data around that time:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                                          TO_DATE ('2023/04/02 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                                          TO_DATE ('2023/04/02 03:30:00',
'YYYY/MM/DD HH:MI:SS'));
SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0);
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
```

Another use case might be if an application with sequential access to a Kafka stream detected a potential anomaly, and was configured to insert a row into an anomaly table. The anomaly table would include the Kafka timestamp, as well as any other data it was configured to record. In that case, another application can then use this information to retrieve records around the suspected record to see if there were any other issues. To achieve this goal, run the following procedure, load the temporary table, and then select and apply application logic to the results:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                                          TO_DATE ('2020/07/04 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                                          TO_DATE ('2020/07/04 03:30:00',
'YYYY/MM/DD HH:MI:SS'));

SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0);
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
--application logic
```

109.5.5 SEEK_OFFSET_TS (Timestamp with Timezone)

The `DBMS_KAFKA` procedure `SEEK_OFFSET_TS` seeks Kafka records in a topic that exist between a window of time specified by `TIMESTAMP WITH TIME ZONE`.Use this procedure to position an Oracle SQL access to Kafka (OSAK) view to start reading Kafka records between two specified `TIMESTAMP WITH TIME ZONE` timestamps that you specify.

Syntax

```
PROCEDURE SEEK_OFFSET_TS(
                view_name           IN VARCHAR2,
                start_timestamp_ms  IN INTEGER,
                end_timestamp_ms    IN INTEGER);
```

Parameters

Table 109-14 SEEK_OFFSET_TS Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Kafka view name (VARCHAR) Case-insensitive.
start_timestamp	The timestamp where you want to seek the first record (INTEGER)
end_timestamp	The timestamp where you want to seek to the last record (INTEGER).

Usage Notes

The purpose of SEEK_OFFSET_TS defined by start and end timestamps is to position an OSAK view to start reading Kafka records within a given timeframe, as defined by the epoch timestamps that you specify. If the window of records exceeds the range of actual records in a Kafka topic, this procedure will return whatever records do exist.

For example:

- If the timestamps are either both below the low water mark, or both above the high water mark, then no records will be returned.
- If the start timestamp is below the low water mark, then the first record returned will be the low water mark.
- If the end timestamp is above the high water mark, then the last record returned will be the difference of the high water mark (HWM), minus 1 (HWM - 1). For example, in a new topic has 100 records, the offset range is from 0 to 99, and the HWM is set at 100.

> **Note:**
>
> Kafka record timestamps are either assigned by Kafka (that is, transaction time), or are assigned by applications (that is, valid or decision time).The data retrieved by a SEEKABLE application using Kafka epoch timeframes as the specification for retrieval cannot include any outlier records that are delivered after the load of the OSAK global temporary table, even if they have a timestamp that is within the timestamp window provided.

Examples

Suppose that you want to investigate issues that have occurred in the past. If the data is still present in the Kafka steam, then you can create a Seekable application by calling DBMS_KAFKA.CREATE_SEEKABLE_APP. You can then call the SEEK_OFFSET_TS procedure to request the OSAK view to retrieve a range of data records. For example, if an IT consultant was informed that a production issue occurred around 3:00 in the morning, then the consultant can use the following procedure to load the temporary table, and then select to retrieve an hour's worth of data around that time:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
```

```
                                          TO_DATE ('2023/04/02 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                                          TO_DATE ('2023/04/02 03:30:00',
'YYYY/MM/DD HH:MI:SS'));
SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0);
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
```

Another use case might be if an application with sequential access to a Kafka stream detected a potential anomaly, and was configured to insert a row into an anomaly table. The anomaly table would include the Kafka timestamp, as well as any other data it was configured to record. In that case, another application can then use this information to retrieve records around the suspected record to see if there were any other issues. To achieve this goal, run the following procedure, load the temporary table, and then select and apply application logic to the results:

```
SYS.DBMS_KAFKA.SEEK_OFFSET_TS (

'ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0',
                                          TO_DATE ('2020/07/04 02:30:00',
'YYYY/MM/DD HH:MI:SS',
                                          TO_DATE ('2020/07/04 03:30:00',
'YYYY/MM/DD HH:MI:SS'));

SYS.DBMS_KAFKA.LOAD_TEMP_TABLE

(ORA$DKVGTT_EXAMPLECLUSTER_SEEKABLEAPP_0);
SELECT <columns> FROM ORA$DKV_EXAMPLECLUSTER_SEEKABLEAPP_0;
--application logic
```

109.6 ADD_PARTITIONS

The `DBMS_KAFKA_ADM` procedure `ADD_PARTITIONS` adds additional Kafka partitions to an existing set of Oracle SQL access to Kafka (OSAK) views.

Syntax

```
PROCEDURE ADD_PARTITIONS (
                      cluster_name         IN  VARCHAR2,
                      application_name     IN  VARCHAR2
```

Parameters

Table 109-15 ADD_PARTITIONS procedure parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.

Table 109-15 (Cont.) ADD_PARTITIONS procedure parameters for DBMS_KAFKA

Parameter	Description
application_name	Name of an existing application associated with the Kafka cluster. Case-insensitive.

Usage Notes

The semantics for `ADD_PARTITIONS` are similar to the `CREATE_xxx_APP` calls, except it preserves state information about existing Kafka topic partitions (for example, `committed_offset`), and binds new partitions either to either existing views, or to new OSAK views.

> **Caution:**
>
> Before calling `ADD_PARTITIONS`, all application instances for this application must be shut down. After `ADD_PARTITIONS` is called successfully, the application instances can be restarted. If the application is a streaming application configured to create one view instance per partition, then new views will be created. Using these views will require starting additional application instances, each dedicated to processing exclusively one of the new views created.

If no partitions have been added to the topic, then the procedure will terminate successfully without altering the existing views.

When an OSAK view is created, and the partitions are assigned, the view manages a list of partitions. For example, suppose an OSAK view was previously assigned to read partitions 0,1,2,3 in a Kafka topic. If `ADD_PARTITIONS` later adds a new partition to this view (partition 16, for example), then the OSAK view is now configured to fetch Kafka records from partitions 0,1,2,3, and 16.

This procedure runs both DDL and DML, which are transactional. This call should only be run outside of any existing transactional context.

Examples

```
BEGIN
  SYS.DBMS_KAFKA.ADD_PARTITIONS ('ExampleCluster',
                                                    'ExampleApp');
END;
/
```

109.7 DROP_ALL_APPS

The DBMS_KAFKA procedure DROP_ALL_APPS drops all applications for the Kafka cluster.Use this procedure to confirm that a connection can be established with the configured security information. The function returns the state of the cluster.

Syntax

```
PROCEDURE DROP_ALL_APPS (
                            \cluster_name    IN VARCHAR2
);
```

Parameters

Table 109-16 DROP_ALL_APPS Procedure Parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.

Usage Notes

The DROP_ALL_APPS procedure runs both Data Definition Language (DDL) and Data Manipulation Language (DML) changes, which are transactional. Only run this call outside of any existing transactional context.

> **Note:**
>
> This is an autonomous transaction.

Examples

Suppose that one or more of the OSAK applications no longer exist in the Kafka cluster ExampleCluster. You can then drop all of the applications in that cluster by using the following statement:

```
EXEC SYS.DBMS_KAFKA.DROP_ALL_APPS ('ExampleCluster');
```

109.8 ENABLE_VIEW_QUERY

The DBMS_KAFKA procedure ENABLE_VIEW_QUERY sets a context within the current Oracle Database session, enabling an application to query a view. Use this procedure to enable an application to query an Oracle SQL access to Kafka (OSAK) view directly.

Syntax

```
PROCEDURE ENABLE_VIEW_QUERY(view_name IN VARCHAR2);
```

Parameters

Table 109-17 ENABLE_VIEW_QUERY Procedure Parameters for DBMS_KAFKA

Parameter	Description
view_name	Name of an existing Kafka view Case-insensitive.

Usage Notes

Only one query against the OSAK view should be done within a transaction. Multiple queries of an OSAK view cannot guarantee repeatable read behavior. Because Kafka is a streaming service, it is possible that new records can be published to a Kafka topic in between the `SELECT` queries. If new records are published, then this can result in additional rows being seen by the application.

> **Caution:**
>
> This is an advanced procedure. Oracle strongly recommends that this call is only used by developers who have a deep understanding of the Oracle SQL access to Kafka (OSAK) processing model and are knowledgeable about debugging Oracle access plans. This includes understanding how Kafka offsets are advanced by OSAK after an OSAK view is queried.

In general, Oracle recommends that you use `LOAD_TEMP_TABLE` or `EXECUTE_LOAD_TABLE`. These procedures query the Kafka views transparently, and load them a single time into either a global temporary table, or into an Oracle Database table for further analysis using SQL.

Examples

Enabling a user to query an OSAK view directly should only be used in two cases:

1. When a query is only referencing the OSAK view in a `FROM` clause to scan a Kafka stream. For example, suppose you want to access an OSAK view directly and retrieve the average temperature of a sensor in laboratory 1 since it was last checked. You can use the following query:

```
EXEC DBMS_KAFKA.ENABLE_VIEW_QUERY(

'ORA$DKV_EXAMPLECLUSTER_STREAMINGAPP_0');
SELECT AVG(temperature)
FROM ORA$DKV_EXAMPLECLUSTER_STREAMINGAPP_0WHERE sensor_name =
'LAB1_SENSOR';
EXEC DBMS_KAFKA.UPDATE_OFFSET(
                                'ORA$DKV_EXAMPLECLUSTER_STREAMINGAPP_0');
COMMIT;
```

2. When a query creates a simple table join between the OSAK view and an Oracle Database table, where the OSAK view is forced to be the outer table of a join by using the `ORDERED` hint in a query. This case ensures that the Kafka data is retrieved only once. For example:

```
SELECT /*+ ORDERED *\/ COUNT(*)
FROM thermostat_spec s, ora$dkv_thermostat_0 t
WHERE s.device_type = t.device_type
AND  t.current_setting >
 t.temperature_setting +
s.device_max_variation;
```

109.9 SET_TRACING

The `DBMS_KAFKA` procedure `SET_TRACING` enables or disables debug level tracing for the external table driver code associated with an Oracle SQL access to Kafka (OSAK) application.This function generates logging output to the trace file for the session.

Syntax

```
PROCEDURE SET_TRACING(
                        cluster_name        IN VARCHAR2,
                        application_name IN VARCHAR2,
                        enable                   IN BOOLEAN);
```

Parameters

Table 109-18 DROP_SEEKABLE_APP Procedure Parameters for DBMS_KAFKA

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.
application_name	The name of an existing application using the Kafka cluster. Case-insensitive.
enable	`[true |false]` Default: `false`. Set `true` to enable the debug output, or `false` to disable the debug output

Usage Notes

Use to generate debug logging output to the trace file for the session. Set to `true` to enable the debug output, or `false` to disable the debug output.

> **Note:**
>
> The following event must already be enabled for the database:
>
> ```
> event="39431 trace name context forever, level 1" # Enable external
> table debug tracing
> ```

Examples

Suppose you want to start tracing on the Kafka cluster `ExampleCluster`, which was being used by the application `ExampleApp`. You can then use this procedure:

```
EXEC SYS.DBMS_KAFKA.SET_TRACING('KAFKACLUS1', 'KAFKACLUS1', True);
```

109.10 DBMS_KAFKA OPTIONS Passed to CREATE_xxx_APP

The options parameters for the `DBMS_KAFKA.CREATE_LOAD_APP`, `CREATE_STREAMING_APP`, and `CREATE_SEEKABLE_APP` packages are provided to the packages in a JSON document
The `DBMS_KAFKA` options can be used with each of the `DBMS_KAFKA` packages to create and define Oracle SQL access to Kafka (OSAK) applications.

> **Note:**
>
> 54 bytes are reserved for the maximum Kafka metadata chunk size in the buffer. For example, a bufsize of 1M can hold a record of 1M – 54 bytes. Bufsize defaults to 1000 KB with a max of 10,000 KB, so the largest Kafka record size supported is 10000 kb – 54.

Table 109-19 Options for DBMS_KAFKA

Option Name	Description
`avrodecimaltype`	Specifies the representation of a decimal stored in the byte array. Valid values: `int`, `integer`, `str`, `string` Default: If this parameter is not used, then an Avro decimal column is read assuming byte arrays store the numerical representation of the values (that is, default to `int`), which is as the Avro specification defines. Only allowed if the `fmt` option is specified as `AVRO`. Related access parameter: `com.oracle.bigdata.avro.decimaltpe`
`avroschema`	JSON document representing the schema of the Kafka value payload. The schema must define an Avro record type. There is no default value. A value must be specified if the `fmt` option is specified as `AVRO`

Table 109-19 (Cont.) Options for DBMS_KAFKA

Option Name	Description
blankasnull	When set to true, loads fields consisting of spaces as null. Values: [true | false] Default: false Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.blankasnull
bufsize	Sets the buffer size in kilobytes for large record reads. Set this value if you need to read records that are greater than the default buffer size. Default (in kilobytes): 1000 Related access parameter: com.oracle.bigdata.buffersize
conversionerrs	If a row has data type conversion errors, then the related columns are stored as null, or if the row is rejected. Values: [reject_record | store_null] Default: store_null Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.conversionerrors
datefmt	Specifies the date format in the source file. The value auto checks for the following formats. J, MM-DD-YYYYBC, MM-DD-YYYY, YYYYMMDD HHMISS, YYMMDD HHMISS, YYYY.DDD, YYYY-MM-DD Default: yyyy-mm-dd hh24:mi:ss Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.dateformat
escapedby	Specifies the character used to escape any embedded field terminators or line terminators in the value for fields. The character value must be wrapped in single-quotes. For example: '\' If the parameter is not specified, then there is no value. Only allowed if fmt option is specified as DSV Related access parameter: com.oracle.bigdata.csv.rowformat.fields.escapedby
fmt	Format of the value payload of a Kafka record. Values: [DSV | JSON | AVRO] There is no default value. The format must be specified in the options passed. Related access parameter: com.oracle.bigdata.kafka.format
jsondt	Store JSON in varchar2 (max size) or clob to support large JSON records. Values: [varchar2 | clob] Default: varchar2 Only allowed if the fmt option is specified as JSON.

Table 109-19 (Cont.) Options for DBMS_KAFKA

Option Name	Description
nulldefinedas	Specifies the character used to indicate the value of a field is NULL. If the parameter is not specified, then there is no value. Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.csv.rowformat.nulldefinedas
quote	Specifies the quote character for the fields. When specified, the characters defined as the quote characters are removed during loading. Valid values: character Default: Null, meaning no quote character Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.quote
rejectlmt	The operation errors out after a specified number of rows (*number*, an integer) are rejected. This operation only applies when rejecting records due to conversion errors. To allow all records to be processed even if conversion errors exist, pass the value 'unlimited'. Valid values: [*number* | 'unlimited'] Default: 0, which means that no conversion errors are allowed.
removequotes	Removes any quotes that are around any field in the source file. Values: [true | false] Default: false Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.removequotes
separator	Specifies the character used to separate the field values. The character value must be wrapped in single quotes. For example: '|' Default: ',' Related access parameter: com.oracle.bigdata.csv.rowformat.fields.terminator
terminator	Specifies the character used to separate the record values. The character value must be wrapped in single quotes. For example: '|' Default: '\n' Related access parameter: com.oracle.bigdata.csv.rowformat.lines.terminator
trimspaces	Specifies how the leading and trailing spaces of the fields are trimmed. Valid values: rtrim, ltrim, notrim, ltrim, ldrtrim Default: notrim Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.trimspaces
truncatecol	If the data in the file is too long for a field, then this option truncates the value of the field rather than rejecting the row or setting the field to NULL. Values: [true | false] Default: false Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.truncatecol

Table 109-19 (Cont.) Options for DBMS_KAFKA

Option Name	Description
tsfmt	Specifies the timestamp format in the source file. The value auto checks for the following formats: YYYY-MM-DD HH:MI:SS.FF, YYYY-MM-DD HH:MI:SS.FF3, MM/DD/YYYY HH:MI:SS.FF3 Default: yyyy-mm-dd hh24:mi:ss.ff Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.timestampformat
tslzfmt	Specifies the timestamp with local timezone format in the source file. The value auto checks for the following formats: DD Mon YYYY HH:MI:SS.FF TZR, MM/DD/YYYY HH:MI:SS.FF TZR, YYYY-MM-DD HH:MI:SS+/-TZR, YYYY-MM-DD HH:MI:SS.FF3, DD.MM.YYYY HH:MI:SS TZR Default: yyyy-mm-dd hh24:mi:ss.ff Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.timestampltzformat
tstzfmt	Specifies the timestamp with timezone format in the source file. The value auto checks for the following formats: DD Mon YYYY HH:MI:SS.FF TZR, MM/DD/YYYY HH:MI:SS.FF TZR, YYYY-MM-DD HH:MI:SS+/-TZR, YYYY-MM-DD HH:MI:SS.FF3, DD.MM.YYYY HH:MI:SS TZR Default: yyyy-mm-dd hh24:mi:ss.ff Only allowed if the fmt option is specified as DSV Related access parameter: com.oracle.bigdata.timestamptzformat

Example 109-3 AVRO Schema Record Type

The following is an example of an AVRO data schema. In this case, the type of AVRO data defined is a record, and the data is a record of sensor values for a temperature monitor as monitored over a set period of time.

```
{
  "type" : "record",
  "name" : "sensor_value",
  "namespace" : "example.sensor",
  "fields" : [ {
    "name" : "EventTime",
    "type" : "long",
    "logicalType" : "timestamp-millis"
  }, {
    "name" : "IotDeviceType",
    "type" : "int"
  }, {
    "name" : "IotDeviceUnitId",
    "type" : "int"
  }, {
    "name" : "TempSetting",
    "type" : "double"
```

```
  }, {
    "name" : "TempReading",
    "type" : "double"
  } ]
}
```

110
DBMS_KAFKA_ADM

The `DBMS_KAFKA_ADM` package provides a PL/SQL interface to create cluster definitions, which you can then use to grant access to Kafka cluster data for applications.

Administrators granted the `OSAK_ADMIN_ROLE` can use `DBMS_KAFKA_ADM` package to create applications that query Kafka data from Oracle Database views and tables.

- DBMS_KAFKA_ADM Overview
- DBMS_KAFKA_ADM Security Model
- DBMS_KAFKA_ADM Constants
- Summary of DBMS_KAFKA_ADM Procedures
- CHECK_CLUSTER
- DEREGISTER_CLUSTERDISABLE_CLUSTER
- DISABLE_CLUSTER
- ENABLE_CLUSTER
- REGISTER_CLUSTER
- UPDATE_CLUSTER_INFO

110.1 DBMS_KAFKA_ADM Overview

The `DBMS_KAFKA_ADM` package enables you to manage the Oracle SQL access to a Kafka cluster.

To set up access to Kafka topics, you require the privileges granted with `OSAK_ADMIN_ROLE`. You use these privileges to provide a cluster definition that helps you to manage resources when connecting to Kafka clusters. The ability of users to share a cluster definition enables the underlying framework to share connections to the clusters. Sharing connections conserves memory and network resources.

110.2 DBMS_KAFKA_ADM Security Model

Oracle recommends that you grant the `OSAK_ADMIN_ROLE` in the target Oracle Database to an administrator user for Oracle SQL Access to Kafka.

Using the `OSAK_ADMIN_ROLE` simplifies the process of granting system privileges required to create and administer Oracle SQL access to Kafka applications. Also, administrators with the `OSAK_ADMIN_ROLE` grant access for other users to the Oracle Database object that has the relevant Kafka view and table for which they need access.

In addition to the `OSAK_ADMIN_ROLE`, an administrator for Oracle SQL access to Kafka applications must have the following system privileges:

- `CREATE SESSION`
- `ALTER SESSION`

- `CREATE CREDENTIAL`, to create a Kafka SASL-SSL (Simple Authentication and Security Layer) password or OSS (Oracle Streaming Service) authToken
- `CREATE ANY DIRECTORY`, to create cluster access and cluster configuration directory
- `DROP ANY DIRECTORY`, to drop cluster access and cluster configuration directory
- `READ ON` privileges to `sys.dbms_kafka_clusters`
- `READ ON` privileges to `sys.dbms_kafka_applications`
- `READ ON` privileges to `sys.dbms_kafka_messages`
- `EXECUTE ON` privileges to `sys.dbms_kafka_adm`

110.3 DBMS_KAFKA_ADM Constants

The `DBMS_KAFKA_ADM` package is a stateful package that uses the constants described here as part of the declared package state in the package specification.
The `DBMS_KAFKA_ADM` constants include constants with literal initial values and constants that you specify in your package build.

Kafka Providers

The provider of the Kafka server can be either Apache, or Oracle Cloud Infrastructure (OCI) Oracle Streaming Service (OSS).

```
KAFKA_PROVIDER_APACHE CONSTANT VARCHAR2(6) := 'APACHE';
KAFKA_PROVIDER_OSS CONSTANT VARCHAR2(3) := 'OSS';
```

States of a cluster definition

```
STATE_CONNECTED CONSTANT INTEGER := 0;
STATE_MAINTENANCE CONSTANT INTEGER := 1;
STATE_BROKEN CONSTANT INTEGER := 2;
STATE_DEREGISTERED CONSTANT INTEGER := 3;
```

Kafka connection modes

Kafka connection modes specify one of the following connection mode constants:

- High Throughput

  If the connection mode (`connmode`) is `high_throughput`, then the applications associated with this connection require all of the data to be delivered as fast as possible. This connection mode is optimal for use from STREAMING and LOADING applications.

  For example:

  ```
  CONNECTION_MODE_HI_THRU CONSTANT VARCHAR2(15) := 'high_throughput';
  ```

- Low Latency

  If the connection mode (`connmode`) is `low_latency`, then the underlying layer attempts to return the first rows as fast as possible. This connection mode is most useful for applications that use only a small amount of Kafka data for each load.

For example:

```
CONNECTION_MODE_LO_LAT CONSTANT VARCHAR2(11) := 'low_latency';
```

- Options

```
OPT_CONNECTION_MODE CONSTANT VARCHAR2(30) := 'connmode';
```

110.4 Summary of DBMS_KAFKA_ADM Procedures

`DBMS_KAFKA_ADM` procedures enable you to create, delete, and configure Oracle SQL access to Kafka.

Kafka providers

Kafka provider constants are used to define the Kafka cluster. Because the security models are different between an Apache Kafka-based cluster and an Oracle Streaming Service-based cluster, you must define to which type of Kafka environment you are connecting, so that the correct secure connection parameters can be passed.
For example:

```
KAFKA_PROVIDER_APACHE CONSTANT VARCHAR2(6) := 'APACHE';
KAFKA_PROVIDER_OSS CONSTANT VARCHAR2(3) := 'OSS';
```

States of a cluster definition

```
STATE_CONNECTED CONSTANT INTEGER := 0;
STATE_MAINTENANCE CONSTANT INTEGER := 1;
STATE_BROKEN CONSTANT INTEGER := 2;
STATE_DEREGISTERED CONSTANT INTEGER := 3;
```

Kafka connection modes

Kafka connection modes specify the Kafka rules that Oracle SQL access to Kafka follows to read a batch of records. Depending on the requirements for your application, you can specify either a low latency rule or a high throughput rule, depending on which option is the more efficient use of resources. You should consider tradeoffs between throughput and latency when you create your connections. The latency between the time that a record is generated and the time that you want that data made available to your application should determine the whether you make low latency your priority, or high throughput your priority.

If your application requires data access in near real-time, so that your application can react to data within seconds or a minute of its being generated by the Kafka cluster, then you should choose low latency connections. With a low latency connection, the Oracle SQL Kafka connection buffers are set to a smaller size, and the reads to the Kafka topic are made frequently, so that the delay between the time a record is generated to the Kafka topic and consumed by the application using Oracle SQL access to Kafka is within the performance requirements of your application. Each read from the Kafka topic is expected to consist of a relatively few number of records, and the number of reads from the Kafka cluster are high. This form of Kafka connection is particularly suited to updating an application with near real-time concurrent transactions, or where the application must monitor and react to changes quickly. However, this configuration can be resource-intensive, as the number of transactions are relatively high for the amount of data received in each transaction.

If your application does not have near real-time requirements, so that your application can delay updates in favor of processing a large batch of records, in a cadence of 15 minutes, or an hour, or every 8 hours, then a high throughput connection can be a more efficient choice. With a connection configured for high throughput, each batch of Kafka data is streamed more infrequently, but the amount of data streamed in each transaction is larger. Accordingly, the Oracle SQL access to Kafka buffer size is larger, and the network I/O and memory allocations are configured for these larger batches. Because these transactions are streamed less frequently, the overall amount of server resources consumed by the Oracle SQL access to Kafka transaction is less

- Low Latency

```
CONNECTION_MODE_LO_LAT CONSTANT VARCHAR2(11) := 'low_latency';
```

- High Throughput

```
CONNECTION_MODE_HI_THRU CONSTANT VARCHAR2(15) := 'high_throughput';
```

- Options

```
OPT_CONNECTION_MODE CONSTANT VARCHAR2(30) := 'connmode';
```

This table lists the DBMS_KAFKA_ADM procedures, and briefly describes them.

Table 110-1 DBMS_KAFKA_ADM Package Procedures

Subprogram	Description
CHECK_CLUSTER	Tests the Kafka cluster connectivity.
DEREGISTER_CLUSTER	Deregisters the Kafka cluster.
DISABLE_CLUSTER	Disables the Kafka cluster.
ENABLE_CLUSTER	Enables the Kafka cluster.
REGISTER_CLUSTER	Registers the Kafka cluster.
UPDATE_CLUSTER_INFO	Updates the host server, or Kafka cluster options, or both.

110.5 CHECK_CLUSTER

The `DBMS_KAFKA_ADM` procedure `CHECK_CLUSTER` tests the Kafka cluster connectivity.
Use this procedure to confirm that a connection can be established with the configured security information. The function returns the state of the cluster.

Syntax

```
FUNCTION CHECK_CLUSTER(
                   cluster_name    IN VARCHAR2
  ) RETURN INTEGER;
```

Parameters

Table 110-2 CHECK_CLUSTER procedure parameters for DBMS_KAFKA_ADM

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.

Usage Notes

As the Oracle SQL access to Kafka cluster administrator, you should check the connectivity of the cluster. Checking the connectivity ensures that the bootstrap server list and the security related configuration that you have provided for the Kafka cluster is proper, and that a connection can be successfully established. You can use this function to test the connection when you run `REGISTER_CLUSTER`, and the security configuration has been established. After the cluster connection has been verified, you can make the cluster definition available to be used by user applications.

You also use this function if users have reported issues when calling `DBMS_KAFKA.LOAD_TEMP_TABLE` and `DBMS_KAFKA.EXECUTE_LOAD_APP`, both of which select from Oracle SQL access to Kafka views.

Examples

Suppose you have completed your initial cluster definition for the cluster ExampleCluster, and registered the cluster. Next, you use this procedure to check the configuration:

```
EXEC SYS.DBMS_OUTPUT.PUT_LINE (
           SYS.DBMS_KAFKA_ADM.CHECK_CLUSTER ('ExampleCluster'));
```

110.6 DEREGISTER_CLUSTER

The `DBMS_KAFKA_ADM` procedure `DEREGISTER_CLUSTER` deregisters the Kafka cluster.

> **Note:**
>
> This procedure will only deregister and remove the cluster if there are no user applications associated with the cluster

Syntax

```
PROCEDURE DEREGISTER_CLUSTER (
                    cluster_name        IN VARCHAR2,
                    forced              IN BOOLEAN DEFAULT FALSE
  );
```

Parameters

Table 110-3 DEREGISTER_CLUSTER Procedure Parameters for DBMS_KAFKA_ADM

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.
forced	(Optional) Marks the cluster as deregistered even if there are user applications still defined. Default: `FALSE`. Values: `[TRUE|FALSE]`

Usage Notes

When a cluster definition is no longer needed, the OSAK Administrator can remove the cluster definition.

If the `forced` parameter is passed as `TRUE`, then the cluster's state is marked as deregistered. Applications that are still associated with this cluster will no longer be functional. When the last application is dropped, this deregistered cluster definition will also be removed.

Examples

In the following example, an unused cluster definition for the Kafka cluster ExampleCluster is deregistered:

```
EXEC SYS.DBMS_KAFKA_ADM.DEREGISTER_CLUSTER ('ExampleCluster');
```

Suppose that there are still user applications (views and tables) associated with the cluster ExampleCluster, but you want to deregister the cluster. In the following example, the cluster ExampleCluster is deregistered with the forced parameter set to TRUE:

```
exec DBMS_KAFKA_ADM.DEREGISTER_CLUSTER(
    cluster_name => 'ExampleCluster',
    forced => TRUE);
```

110.7 DISABLE_CLUSTER

The `DBMS_KAFKA_ADM` procedure `DISABLE_CLUSTER` disables the Kafka cluster.

> **Note:**
>
> When this procedure completes successfully, the state of the application changes to `MAINTENANCE`.

Syntax

```
PROCEDURE DISABLE_CLUSTER(
        cluster_name        IN VARCHAR2
);
```

Parameters

Table 110-4 DISABLE_CLUSTER Procedure Parameters for DBMS_KAFKA_ADM

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.

Usage Notes

A disabled cluster will prevent the cluster security information from being retrieved, thus preventing Kafka connections from being created or accessed.

Before you begin maintenance on a Kafka cluster, an Oracle SQL access to Kafka administrator should disable access to the cluster by using the `DBMS_KAFKA_ADM.DISABLE_CLUSTER` procedure. After the maintenance is completed, you can reeanble the cluster access by using the `DBMS_KAFKA_ADM.ENABLE_CLUSTER` procedure.

This procedure raises an exception if the state cannot be changed to `MAINTENANCE`

Examples

In the following example, the cluster definition for ExampleCluster is temporarily disabled. In this state, views cannot connect to or retrieve data from the Kafka cluster:

```
EXEC SYS.DBMS_KAFKA_ADM.DISABLE_CLUSTER ('ExampleCluster');
```

110.8 ENABLE_CLUSTER

The `DBMS_KAFKA_ADM` procedure `ENABLE_CLUSTER` enables the Kafka cluster.

> **Note:**
>
> When this procedure completes successfully, the state of the application changes to `CONNECTED`.

Syntax

```
PROCEDURE ENABLE_CLUSTER(
                           cluster_name    IN VARCHAR2
  );
```

Parameters

Table 110-5 ENABLE_CLUSTER Procedure Parameters for DBMS_KAFKA_ADM

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive.

Usage Notes

After you complete maintenance on a Kafka cluster, an Oracle SQL access to Kafka administrator should enable access to the cluster by using the `DBMS_KAFKA_ADM.ENABLE_CLUSTER` procedure.

This procedure raises an exception if the state cannot be changed to `CONNECTED`.

Examples

Suppose that you have disabled the cluster ExampleCluster to perform maintenance, and you now want to reenable the ability of views to connect to and retrieve data from that Kafka cluster. Enter the following command:

```
EXEC SYS.DBMS_KAFKA_ADM.ENABLE_CLUSTER ('ExampleCluster');
```

110.9 REGISTER_CLUSTER

The `DBMS_KAFKA_ADM` procedure `REGISTER_CLUSTER` registers the Kafka cluster and makes it available for access on Oracle Database.

Syntax

```
FUNCTION REGISTER_CLUSTER (
                      cluster_name             IN VARCHAR2,
                      startup_servers          IN VARCHAR2,
                      kafka_provider_provider  IN VARCHAR2,
                      cluster_access_dir       IN VARCHAR2,
                      credential_name          IN VARCHAR2 DEFAULT NULL,
                      cluster_config_dir       IN VARCHAR2 DEFAULT NULL,
                      cluster_description      IN VARCHAR2 DEFAULT NULL,
```

```
                        options                 IN CLOB DEFAULT NULL
)  RETURN INTEGER;
```

Parameters

Table 110-6 REGISTER_CLUSTER Procedure Parameters for DBMS_KAFKA_ADM

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive
bootstrap_servers	The startup servers of the Kafka cluster Case-sensitive
kafka_provider_*provider*	The provider of the Kafka server, which can be either Apache or Oracle Cloud Infrastructure (OCI) Oracle Streaming Service (OSS).. Syntax: DBMS_KAFKA_ADM.KAFKA_PROVIDER_*provider* Values for *provider*: [APACHE|OSS] Case-insensitive See also the DBMS_KAFKA_ADMconstant KAFKA_PROVIDER
cluster_access_dir	The Oracle directory object for determining access to this cluster Case-insensitive
credential_name	Credential name associated with the password to connect to Kafka. See SYS.DBMS_CREDENTIAL for possible name values and how to create a credential. Must be provided if Kafka connection requires a password, otherwise NULL
cluster_config_dir	The Oracle directory object containing the cluster configuration files required for secure clusters. The value is NULL for OSS or non- secure clusters. Case-insensitive
cluster_description	(Optional) A text description of the cluster. Maximum length: 400 characters.
options	(Optional) Cluster options. The format is a JSON-formatted document. For possible options, see SYS.DBMS_KAFKA_ADM.OPT_%

Usage Notes

REGISTER_CLUSTER automatically runs a CHECK_CLUSTER call after the cluster definition has been stored. The status returned from the CHECK_CLUSTER call is the returned value from the REGISTER_CLUSTER. This procedure raises an exception if the state cannot be determined.

At the end of the lifecycle for an Oracle SQL access to Kafka (OSAK) cluster, the administrator can remove the OSAK cluster by using the DBMS_KAFKA_ADM.DEREGISTER_CLUSTER function.

Examples

In the following example, the function request with REGISTER_CLUSTER registers the Kafka cluster KAFKACLUS1 that is on the server mykafkastartup, using Apache as the provider. The cluster access directory on the Oracle Database instance is OSAK_KAFKACLUS1_ACCESS,

accessed with the credential `KAFKACLUS1CRED1`, using the cluster configuration directory `OSAK_KAFKACLUS1_CONFIG`. The description for this cluster (which is optional), is "My test cluster kafkaclus1." A user (`examplekafka-user`) is granted `READ` access to the Kafka cluster data on `OSAK_KAFKACLUS1_ACCESS`,

```
SQL> select DBMS_KAFKA_ADM.REGISTER_CLUSTER ('KAFKACLUS1',
                                    'mykafkastartup-host:9092',
                                   DBMS_KAFKA_ADM.KAFKA_PROVIDER_APACHE,
                                   'OSAK_KAFKACLUS1_ACCESS'
                                   'KAFKACLUS1CRED1',
                                   'OSAK_KAFKACLUS1_CONFIG',
                                   'My test cluster kafkaclus1') from dual;

Output: Successful registration return 0 (zero)
SQL> DBMS_KAFKA_ADM_RE...
             0


        SQL> grant read on directory OSAK_KAFKACLUS1_ACCESS to examplekafka-
user;
```

You can also create non-secured cluster access. In the following example, a non-secure Kafka cluster access is registered using Oracle SQL access to Kafka (OSAK):

1. The cluster access database directory is created with an empty path. This directory is used to control which Oracle users can access the Kafka cluster:

   ```
   SQL> CREATE DIRECTORY
   OSAK_KAFKACLUS2_ACCESS AS '';
   ```

2. A cluster configuration operating system directory using the path structure *`Oracle base/`*osak/*cluster_name*/config is created, with the corresponding Oracle directory object, where the Oracle base is `/u01/app/oracle`, and the Kafka cluster name is `kafkaclus2`:

   ```
   $mkdir /u01/app/oracle/osak/kafkaclus2/config;
   .
   .
   .
   SQL> CREATE DIRECTORY OSAK_KAFKACLUS2_CONFIG AS 'u01/app/oracle/osak/
   kafkaclus2/config';
   ```

3. Create an empty `osakafka.properties` file, or create an `osakafka.properties` file with OSAK tuning or debugging properties, as you decide is required for your enterprise.
4. Register the Kafka cluster using `DBMS_KAKFA_ADM.REGISTER_CLUSTER()`. For example:

   ```
   SQL> select DBMS_KAFKA_ADM.REGISTER_CLUSTER (
             cluster_name  => 'KAFKACLUS2',
             bootstrap_servers =>'mykafkastartup-host:9092',
             kafka_provider => DBMS_KAFKA_ADM.KAFKA_PROVIDER_APACHE,
             cluster_access_dir => 'OSAK_KAFKACLUS2_ACCESS',
             credential_name => NULL,
             cluster_config_dir => 'OSAK_KAFKACLUS2_CONFIG',
              cluster_description =>  'My test cluster kafkaclus2',
   ```

```
            options  => NULL)
from dual;
Output: Successful registration return 0 (zero)
SQL> DBMS_KAFKA_ADM_RE…..
             0
```

5. Grant read access to users. In this example, user examplekafka-user2 is granted access:

```
SQL> grant read on directory
osak_kafkaclus2_access to examplekafka-user2;
```

110.10 UPDATE_CLUSTER_INFO

The `DBMS_KAFKA_ADM` procedure `UPDATE_CLUSTER_INFO` updates the Kafka host server, or the options for the Kafka cluster, or both.

Syntax

```
PROCEDURE UPDATE_CLUSTER_INFO(
                      cluster_name         IN VARCHAR2,
                      startup_servers    IN VARCHAR2 DEFAULT NULL,
                      options              IN CLOB DEFAULT NULL
 );
```

Parameters

Table 110-7 UPDATE_CLUSTER_INFO Procedure Parameters for DBMS_KAFKA_ADM

Parameter	Description
cluster_name	Name of an existing Kafka cluster Case-insensitive
startup_servers	The startup servers of the Kafka cluster Case-sensitive Because this is part of the Kafka source configuration, this string is treated as an opaque value (used as is), which is passed to the Kafka cluster connection logic.
options	(Optional) Cluster options. The format is a JSON-formatted document. For possible options, see `SYS.DBMS_KAFKA_ADM.OPT_%`

Usage Notes

The `UPDATE_CLUSTER_INFO` procedure updates the Kafka cluster definition, including the startup server list, or cluster options, or both. It also disconnects from the Kafka server. As the Oracle SQL access to Kafka (OSAK) Administrator, use this procedure if the Kafka cluster environment changes, and you need to change the cluster definition and configuration.

If an update fails, then the procedure raises an exception.

Examples

In the following example, the Kafka administrator has reconfigured the cluster to use the servers `newhost` and `host2`, and is now updating the list of startup servers for the Kafka cluster `ExampleCluster`:

```
EXEC SYS.DBMS_KAFKA_ADM.UPDATE_CLUSTER_INFO
('ExampleCluster',

'newhost:9092,host2:9092');
```

111
DBMS_LDAP

The DBMS_LDAP package lets you access data from LDAP servers.

For a complete description of this package within the context of Oracle Internet Directory, see `DBMS_LDAP` in the *Application Developer's Guide for Oracle Identity Management*.

112
DBMS_LDAP_UTL

The DBMS_LDAP_UTL package contains the Oracle Extension utility functions.

For a complete description of this package within the context of Oracle Internet Directory, see `DBMS_LDAP_UTL` in the *Application Developer's Guide for Oracle Identity Management*.

113
DBMS_LIBCACHE

The DBMS_LIBCACHE package consists of one subprogram that prepares the library cache on an Oracle instance by extracting SQL and PL/SQL from a remote instance and compiling this SQL locally without execution. The value of compiling the cache of an instance is to prepare the information the application requires to execute in advance of failover or switchover.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_LIBCACHE Subprograms

113.1 DBMS_LIBCACHE Overview

Compiling a shared cursor consists of open, parse, and bind operations, plus the type-checking and execution plan functions performed at the first execution. All of these steps are executed in advance by the package DBMS_LIBCACHE for SELECT statements.

The open and parse functions are executed in advance for PL/SQL and DML. For PL/SQL, executing the parse phase has the effect of loading all library cache heaps other than the MCODE.

113.2 DBMS_LIBCACHE Security Model

To execute DBMS_LIBCACHE you must directly access the same objects as do SQL statements. You can best accomplish this by utilizing the same user id as the original system on the remote system.

When there are multiple schema users, DBMS_LIBCACHE should be called for each.

Alternatively, DBMS_LIBCACHE may be called with the generic user PARSER. However, this user cannot parse the SQL that uses objects with access granted though roles. This is a standard PL/SQL security limitation.

113.3 Summary of DBMS_LIBCACHE Subprograms

The DBMS_LIBCACHE package includes the COMPILE_FROM_REMOTE procedure subprogram.

Table 113-1 ***DBMS_LIBCACHE Package Subprograms***

Subprogram	Description
COMPILE_FROM_REMOTE Procedure	Extracts SQL in batch from the source instance and compiles the SQL at the target instance

113.3.1 COMPILE_FROM_REMOTE Procedure

This procedure extracts SQL in batch from the source instance and compiles the SQL at the target instance.

Syntax

```
DBMS_LIBCACHE.COMPILE_FROM_REMOTE (
   p_db_link                IN     dbms_libcache$def.db_link%type,
   p_username               IN     VARCHAR2 default null,
   p_threshold_executions   IN     NATURAL  default 3,
   p_threshold_sharable_mem IN     NATURAL  default 1000,
   p_parallel_degree        IN     NATURAL  default 1);
```

Parameters

Table 113-2 COMPILE_FROM_REMOTE Procedure Parameters

Parameter	Description
p_db_link	Database link to the source name (mandatory). The database link pointing to the instance that will be used for extracting the SQL statements. The user must have the role SELECT_ON_CATALOG at the source instance. For improved security, the connection may use a password file or LDAP authentication. The database link is mandatory only for releases with dbms_libcache$def.ACCESS_METHOD = DB_LINK_METHOD
p_instance_name	(Reserved for future use). The name of the instance that will be used for extracting the SQL statements. The instance name must be unique for all instances excluding the local instance. The name is not case sensitive.
p_username	Source username (default is all users). The name of the username that will be used for extracting the SQL statements. The username is an optional parameter that is used to ensure the parsing user id is the same as that on the source instance. For an application where users connect as a single user_id, for example APPS, APPS is the parsing user_id that is recorded in the shared pool. To select only SQL statements parsed by APPS, enter the string 'APPS' in this field. To also select statements executed by batch, repeat the executing the procedure with the schema owner, for example GL. If the username is supplied, it must be valid. The name is not case sensitive.
p_threshold_executions	The lower bound for the number of executions, below which a SQL statement will not be selected for parsing. This parameter is optional. It allows the application to extract and compile statements with executions, for example, greater than 3. The default value is 1. This means SQL statements that have never executed, including invalid SQL statements, will not be extracted.
p_threshold_sharable_mem	The lower bound for the size of the shared memory consumed by the cursors on the source instance. Below this value a SQL statement will not be selected for parsing. This parameter is optional. It allows the application to extract and compile statements with shared memory for example, greater than 10000 bytes.

Table 113-2 (Cont.) COMPILE_FROM_REMOTE Procedure Parameters

Parameter	Description
p_parallel_degree	The number of parallel jobs that execute to complete the parse operation. These tasks are spawned as parallel jobs against a sub-range of the SQL statements selected for parsing. This parameter is reserved for parallel compile jobs which are currently not implemented.

114
DBMS_LOB

The DBMS_LOB package provides subprograms to operate on BLOBs, CLOBs, NCLOBs, BFILEs, and temporary LOBs. You can use DBMS_LOB to access and manipulate specific parts of a LOB or complete LOBs.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Datatypes
- Operational Notes
- Rules and Limits
- Exceptions
- Summary of DBMS_LOB Subprograms

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide*

114.1 DBMS_LOB Overview

DBMS_LOB can read and modify BLOBs, CLOBs, and NCLOBs; it provides read-only operations for BFILEs. The bulk of the LOB operations are provided by this package.

114.2 DBMS_LOB Security Model

This package must be created under SYS. Operations provided by this package are performed under the current calling user, not under the package owner SYS.

Any DBMS_LOB subprogram called from an anonymous PL/SQL block is executed using the privileges of the current user.

When creating the procedure, users can set the AUTHID to indicate whether they want definer's rights or invoker's rights. For example:

```
CREATE PROCEDURE proc1 AUTHID DEFINER ...
```

or

```
CREATE PROCEDURE proc1 AUTHID CURRENT_USER ...
```

> **See Also:**
>
> For more information on AUTHID and privileges, see *Oracle Database PL/SQL Language Reference*

You can provide secure access to BFILEs using the DIRECTORY feature discussed in BFILENAME function in the*Oracle Database SecureFiles and Large Objects Developer's Guide* and the *Oracle Database SQL Language Reference*.

For information about the security model pertaining to temporary LOBs, see Operational Notes.

114.3 DBMS_LOB Constants

This topic describes the constants used by the DBMS_LOB package

These are shown in following tables:

- Table 114-1
- Table 114-2
- Table 114-3
- Table 114-4
- Table 114-5
- Table 114-6

Table 114-1 DBMS_LOB Constants - Basic

Constant	Type	Value	Description
CALL	PLS_INTEGER	12	Create the TEMP LOB with call duration
FILE_READONLY	BINARY_INTEGER	0	Open the specified BFILE read-only
LOB_READONLY	BINARY_INTEGER	0	Open the specified LOB read-only
LOB_READWRITE	BINARY_INTEGER	1	Open the specified LOB read-write
LOBMAXSIZE	INTEGER	18446744073709551615	Maximum size of a LOB in bytes
SESSION	PLS_INTEGER	10	Create the TEMP LOB with session duration

Table 114-2 DBMS_LOB Constants - Option Types

Constant	Definition	Value	Description
OPT_COMPRESS	BINARY_INTEGER	1	Set/Get the SECUREFILE compress option value
OPT_DEDUPLICATE	BINARY_INTEGER	4	Set/Get the SECUREFILE Deduplicate option value

Table 114-2 (Cont.) DBMS_LOB Constants - Option Types

Constant	Definition	Value	Description
OPT_ENCRYPT	BINARY_INTEGER	2	Get the SECUREFILE encrypt option value

Table 114-3 DBMS_LOB Constants - Option Values

Constant	Definition	Value	Description
COMPRESS_OFF	BINARY_INTEGER	0	For SETOPTIONS Procedures, set compress off; for GETOPTIONS Functions, compress is off
COMPRESS_ON	BINARY_INTEGER	1	For SETOPTIONS Procedures, set compress on; for GETOPTIONS Functions, compress is on
DEDUPLICATE_OFF	BINARY_INTEGER	0	For SETOPTIONS Procedures, set deduplicate is off; for GETOPTIONS Functions, deduplicate is off
DEDUPLICATE_ON	BINARY_INTEGER	4	For SETOPTIONS Procedures, set deduplicate is on; for GETOPTIONS Functions, deduplicate is on
ENCRYPT_OFF	BINARY_INTEGER	0	For GETOPTIONS Functions, encrypt is off
ENCRYPT_ON	BINARY_INTEGER	2	For GETOPTIONS Functions, encrypt is on

Table 114-4 DBMS_LOB Constants - DBFS State Value Types

Constant	Definition	Value	Description
DBFS_LINK_NEVER	PLS_INTEGER	0	LOB has never been archived
DBFS_LINK_NO	PLS_INTEGER	2	LOB was archived, but as been read back in to the RDBMS
DBFS_LINK_YES	PLS_INTEGER	1	LOB is currently archived

Table 114-5 DBMS_LOB Constants - DBFS Cache Flags

Constant	Definition	Value	Description
DBFS_LINK_CACHE	PLS_INTEGER	1	Put the LOB data to the archive, but keep the data in the RDBMS as a cached version
DBFS_LINK_NOCACHE	PLS_INTEGER	0	Put the LOB data to the archive, and remove the data from the RDBMS.

Table 114-6 DBMS_LOB Constants - Miscellaneous

Constant	Definition	Value	Description
CONTENTTYPE_MAX_SIZE	PLS_INTEGER	128	Maximum number of bytes allowed in the content type string

Table 114-6 (Cont.) DBMS_LOB Constants - Miscellaneous

Constant	Definition	Value	Description
DBFS_LINK_PATH_MAX_SIZE	PLS_INTEGER	1024	The maximum length of DBFS pathnames

114.4 DBMS_LOB Datatypes

The table in this topic describes the datatypes used by DBMS_LOB.

.

Table 114-7 Datatypes Used by DBMS_LOB

Type	Description
BLOB	Source or destination binary LOB.
RAW	Source or destination RAW buffer (used with BLOB).
CLOB	Source or destination character LOB (including NCLOB).
VARCHAR2	Source or destination character buffer (used with CLOB and NCLOB).
INTEGER	Specifies the size of a buffer or LOB, the offset into a LOB, or the amount to access.
BFILE	Large, binary object stored outside the database.

The DBMS_LOB package defines no special types.

An NCLOB is a CLOB for holding fixed-width and varying-width, multibyte national character sets.

The clause ANY_CS in the specification of DBMS_LOB subprograms for CLOBs enables the CLOB type to accept a CLOB or NCLOB locator variable as input.

114.5 DBMS_LOB Operational Notes

All DBMS_LOB subprograms work based on LOB locators. For the successful completion of DBMS_LOB subprograms, you must provide an input locator that represents a LOB that already exists in the database tablespaces or external file system.

See also Chapter 1 of *Oracle Database SecureFiles and Large Objects Developer's Guide*

Starting from 12.2 release, you can select a persistent LOB locator from a remote table into a local variable. The remote column can be of type BLOB, CLOB, or NCLOB. You cannot select BFILE from a remote table. The LOB variable that refers to the LOB value in a remote table is called a remote locator.

All the DBMS_LOB APIs other than the ones that are meant for BFILEs will now accept and support operations on remote LOB locators. All the APIs that take in two locators must have both LOBs collocated at one database.

> **See Also:**
>
> *Distributed LOBs* chapter in *Oracle Database SecureFiles and Large Objects Developer's Guide*.

To use LOBs in your database, you must first use SQL data definition language (DDL) to define the tables that contain LOB columns.

- Internal LOBs
- External LOBs
- Temporary LOBs

Internal LOBs

To populate your table with internal LOBs after LOB columns are defined in a table, you use the SQL data manipulation language (DML) to initialize or populate the locators in the LOB columns.

External LOBs

For an external LOB (BFILE) to be represented by a LOB locator, you must:

- Ensure that a `DIRECTORY` object representing a valid, existing physical directory has been defined, and that physical files (the LOBs you plan to add) exist with read permission for the database. If your operating system uses case-sensitive path names, then be sure you specify the directory in the correct format.
- Pass the `DIRECTORY` object and the filename of the external LOB you are adding to the `BFILENAME` function to create a LOB locator for your external LOB.

Once you have completed these tasks, you can insert or update a row containing a LOB column using the specified LOB locator.

After the LOBs are defined and created, you can then `SELECT` from a LOB locator into a local PL/SQL LOB variable and use this variable as an input parameter to `DBMS_LOB` for access to the LOB value.

For details on the different ways to do this, See *Oracle Database SecureFiles and Large Objects Developer's Guide*

Temporary LOBs

The database supports the definition, creation, deletion, access, and update of temporary LOBs. Your temporary tablespace stores the temporary LOB data. Temporary LOBs are not permanently stored in the database. Their purpose is mainly to perform transformations on LOB data.

For temporary LOBs, you must use the OCI, PL/SQL, or another programmatic interface to create or manipulate them. Temporary LOBs can be either `BLOBs`, `CLOBs`, or `NCLOBs`.

A temporary LOB is empty when it is created. By default, all temporary LOBs are deleted at the end of the session in which they were created. If a process dies unexpectedly or if the database crashes, then temporary LOBs are deleted, and the space for temporary LOBs is freed.

There is also an interface to let you group temporary LOBs together into a logical bucket. The duration represents this logical store for temporary LOBs. Each temporary LOB can have

separate storage characteristics, such as CACHE/ NOCACHE. There is a default store for every session into which temporary LOBs are placed if you don't specify a specific duration. Additionally, you are able to perform a free operation on durations, which causes all contents in a duration to be freed.

There is no support for consistent read (CR), undo, backup, parallel processing, or transaction management for temporary LOBs. Because CR and roll backs are not supported for temporary LOBs, you must free the temporary LOB and start over again if you encounter an error.

Because CR, undo, and versions are not generated for temporary LOBs, there is potentially a performance impact if you assign multiple locators to the same temporary LOB. Semantically, each locator should have its own copy of the temporary LOB.

A copy of a temporary LOB is created if the user modifies the temporary LOB while another locator is also pointing to it. The locator on which a modification was performed now points to a new copy of the temporary LOB. Other locators no longer see the same data as the locator through which the modification was made. A deep copy was not incurred by permanent LOBs in these types of situations, because CR snapshots and version pages enable users to see their own versions of the LOB cheaply.

You can gain pseudo-REF semantics by using pointers to locators in OCI and by having multiple pointers to locators point to the same temporary LOB locator, if necessary. In PL/SQL, you must avoid using more than one locator for each temporary LOB. The temporary LOB locator can be passed by reference to other procedures.

Because temporary LOBs are not associated with any table schema, there are no meanings to the terms in-row and out-of-row temporary LOBs. Creation of a temporary LOB instance by a user causes the engine to create and return a locator to the LOB data. The PL/SQL DBMS_LOB package, PRO*C/C++, OCI, and other programmatic interfaces operate on temporary LOBs through these locators just as they do for permanent LOBs.

There is no support for client side temporary LOBs. All temporary LOBs reside in the server.

Temporary LOBs do not support the EMPTY_BLOB or EMPTY_CLOB functions that are supported for permanent LOBs. The EMPTY_BLOB function specifies the fact that the LOB is initialized, but not populated with any data.

A temporary LOB instance can only be destroyed by using OCI or the DBMS_LOB package by using the appropriate FREETEMPORARY or OCIDurationEnd statement.

A temporary LOB instance can be accessed and modified using appropriate OCI and DBMS_LOB statements, just as for regular permanent internal LOBs. To make a temporary LOB permanent, you must explicitly use the OCI or DBMS_LOB COPY command, and copy the temporary LOB into a permanent one.

Security is provided through the LOB locator. Only the user who created the temporary LOB is able to see it. Locators are not expected to be able to pass from one user's session to another. Even if someone did pass a locator from one session to another, they would not access the temporary LOBs from the original session. Temporary LOB lookup is localized to each user's own session. Someone using a locator from somewhere else is only able to access LOBs within his own session that have the same LOB ID. Users should not try to do this, but if they do, they are not able to affect anyone else's data.

The database keeps track of temporary LOBs for each session in a v$ view called V$TEMPORARY_LOBS, which contains information about how many temporary LOBs exist for each session. V$ views are for DBA use. From the session, the database can determine which user owns the temporary LOBs. By using V$TEMPORARY_LOBS in conjunction with DBA_SEGMENTS, a DBA can see how much space is being used by a session for temporary LOBs. These tables

can be used by DBAs to monitor and guide any emergency cleanup of temporary space used by temporary LOBs.

The following notes are specific to temporary LOBs:

1. All functions in `DBMS_LOB` return `NULL` if any of the input parameters are `NULL`. All procedures in `DBMS_LOB` raise an exception if the LOB locator is input as `NULL`.
2. Operations based on `CLOBs` do not verify if the character set IDs of the parameters (`CLOB` parameters, `VARCHAR2` buffers and patterns, and so on) match. It is the user's responsibility to ensure this.
3. Data storage resources are controlled by the DBA by creating different temporary tablespaces. DBAs can define separate temporary tablespaces for different users, if necessary.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information on `NOCOPY` syntax

114.6 DBMS_LOB Rules and Limits

This topic describes general DBMS_LOB rules and limits, rules and limits specific to external files (BFILEs), and maximum LOB and buffer sizes.

General Rules and Limits

- Oracle Database does not support constraints on columns or attributes whose type is a LOB, with the following exceptions:
  - `NOT NULL` constraints are supported for an LOB column or attribute `IS JSON FORMAT JSON` constraints trigger a JSON validation when you close a LOB.
  - Oracle does not recommend that you use `DBMS_LOB` to modify LOBs under `IS JSON FORMAT JSON` constraints. Because LOBs under `IS JSON FORMAT JSON` constraints are not validated, you can only check that the LOB content is a valid JSON. However, if required, Oracle recommends that you wrap a LOB operation inside the Open/Close API to perform write operations on LOBs under the `IS JSON FORMAT JSON` constraints. This validates that the LOB content has valid JSON when you close the LOB. If the LOB is not open, every write operation validates that the LOB content is a valid JSON, and hence, you will not able to write the JSON in pieces.
- The following rules apply in the specification of subprograms in this package:
  - `newlen`, `offset`, and `amount` parameters for subprograms operating on `BLOBs` and `BFILEs` must be specified in terms of *bytes*.
  - `newlen`, `offset`, and `amount` parameters for subprograms operating on `CLOBs` must be specified in terms of *characters*.

  In multi-byte character sets, it is not possible to interpret these offsets correctly.
- A subprogram raises an `INVALID_ARGVAL` exception if the following restrictions are not followed in specifying values for parameters (unless otherwise specified):
  1. Only positive, absolute offsets from the beginning of `LOB` data are permitted: Negative offsets from the tail of the LOB are not permitted.

2. Only positive, nonzero values are permitted for the parameters that represent size and positional quantities, such as `amount`, `offset`, `newlen`, `nth`, and so on. Negative offsets and ranges observed in SQL string functions and operators are not permitted.
3. The value of `offset`, `amount`, `newlen`, `nth` must not exceed the value `lobmaxsize` 18446744073709551615 (2^{64}-1) in any `DBMS_LOB` subprogram.
4. For `CLOB`s in a database with a multibyte database character set and for NCLOBs, the maximum value for these parameters must not exceed `trunc(lobmaxsize/2)=9223372036854775807` characters.

- PL/SQL language specifications stipulate an upper limit of 32767 bytes (not characters) for `RAW` and `VARCHAR2` parameters used in `DBMS_LOB` subprograms. For example, if you declare a variable to be:

  ```
  charbuf VARCHAR2(3000)
  ```

  Then, `charbuf` can hold 3000 single byte characters or 1500 2-byte fixed width characters. This has an important consequence for `DBMS_LOB` subprograms for `CLOB`s and `NCLOB`s.

- The `%CHARSET` clause indicates that the form of the parameter with `%CHARSET` must match the form of the `ANY_CS` parameter to which it refers.

  For example, in `DBMS_LOB` subprograms that take a `VARCHAR2` buffer parameter, the form of the `VARCHAR2` buffer must match the form of the `CLOB` parameter. If the input LOB parameter is of type `NCLOB`, then the buffer must contain `NCHAR` data. Conversely, if the input LOB parameter is of type `CLOB`, then the buffer must contain `CHAR` data.

  For `DBMS_LOB` subprograms that take two `CLOB` parameters, both `CLOB` parameters must have the same form; that is, they must both be `NCLOB`s, or they must both be `CLOB`s.

- If the value of `amount` plus the `offset` exceeds the maximum LOB size allowed by the database, then access exceptions are raised.

  Under these input conditions, read subprograms, such as `READ`, `COMPARE`, `INSTR`, and `SUBSTR`, read until `End of Lob/File` is reached. For example, for a `READ` operation on a `BLOB` or `BFILE`, if the user specifies offset value of 3 GB and an amount value of 2 GB on a LOB that is 4GB in size, then `READ` returns only 1GB (4GB-3GB) bytes.

- Functions with `NULL` or invalid input values for parameters return a `NULL`. Procedures with `NULL` values for destination LOB parameters raise exceptions.

- Operations involving patterns as parameters, such as `INSTR` do not support regular expressions or special matching characters (such as `%` in the `LIKE` operator in SQL) in the `pattern` parameter or substrings.

- The `End Of` LOB condition is indicated by the `READ` procedure using a `NO_DATA_FOUND` exception. This exception is raised only upon an attempt by the user to read beyond the end of the LOB. The `READ` buffer for the last read contains 0 bytes.

- For consistent LOB updates, you must lock the row containing the destination LOB before making a call to any of the procedures (mutators) that modify LOB data.

- Unless otherwise stated, the default value for an `offset` parameter is 1, which indicates the first byte in the `BLOB` or `BFILE` data, and the first character in the `CLOB` or `NCLOB` value. No default values are specified for the `amount` parameter — you must input the values explicitly.

- You must lock the row containing the destination internal LOB before calling any subprograms that modify the LOB, such as `APPEND`, `COPY`, `ERASE`, `TRIM`, or `WRITE`. These subprograms do not implicitly lock the row containing the LOB.

Rules and Limits Specific to External Files (BFILEs)

- The subprograms `COMPARE`, `INSTR`, `READ`, `SUBSTR`, `FILECLOSE`, `FILECLOSEALL` and `LOADFROMFILE` operate only on an *opened* `BFILE` locator; that is, a successful `FILEOPEN` call must precede a call to any of these subprograms.
- For the functions `FILEEXISTS`, `FILEGETNAME` and `GETLENGTH`, a file's open/close status is unimportant; however, the file must exist physically, and you must have adequate privileges on the `DIRECTORY` object and the file.
- `DBMS_LOB` does not support any concurrency control mechanism for `BFILE` operations.
- In the event of several open files in the session whose closure has not been handled properly, you can use the `FILECLOSEALL` subprogram to close all files opened in the session and resume file operations from the beginning.
- If you are the creator of a `DIRECTORY`, or if you have system privileges, then use the `CREATE OR REPLACE`, `DROP`, and `REVOKE` statements in SQL with extreme caution.

  If you, or other grantees of a particular directory object, have several open files in a session, then any of the preceding commands can adversely affect file operations. In the event of such abnormal termination, your only choice is to invoke a program or anonymous block that calls `FILECLOSEALL`, reopen your files, and restart your file operations.
- All files opened during a user session are implicitly closed at the end of the session. However, Oracle strongly recommends that you close the files after *both* normal and abnormal termination of operations on the `BFILE`.

  In the event of normal program termination, proper file closure ensures that the number of files that are open simultaneously in the session remains less than `SESSION_MAX_OPEN_FILES`.

  In the event of abnormal program termination from a PL/SQL program, it is imperative that you provide an exception handler that ensures closure of all files opened in that PL/SQL program. This is necessary because after an exception occurs, only the exception handler has access to the `BFILE` variable in its most current state.

  After the exception transfers program control outside the PL/SQL program block, all references to the open `BFILEs` are lost. The result is a larger open file count which may or may not exceed the `SESSION_MAX_OPEN_FILES` value.

  For example, consider a `READ` operation past the end of the `BFILE` value, which generates a `NO_DATA_FOUND` exception:

```
-- This assumes a directory 'DDD' whose path is already known
DECLARE
       fil BFILE:= bfilename('DDD', 'filename.foo');
       pos INTEGER;
       amt BINARY_INTEGER;
       buf RAW(40);
BEGIN
       SELECT ad_graphic INTO fil FROM print_media WHERE product_id = 3106;
       dbms_lob.open(fil, dbms_lob.lob_readonly);
       amt := 40; pos := 1 + dbms_lob.getlength(fil); buf := '';
       dbms_lob.read(fil, amt, pos, buf);
       dbms_output.put_line('Read F1 past EOF: '||
           utl_raw.cast_to_varchar2(buf));
       dbms_lob.close(fil);
END;

ORA-01403: no data found
```

```
ORA-06512: at "SYS.DBMS_LOB", line 373
ORA-06512: at line 10
```

After the exception has occurred, the `BFILE` locator variable file goes out of scope, and no further operations on the file can be done using that variable. Therefore, the solution is to use an exception handler:

```
DECLARE
     fil BFILE;
     pos INTEGER;
     amt BINARY_INTEGER;
     buf RAW(40);
BEGIN
     SELECT ad_graphic INTO fil FROM print_media WHERE product_id = 3106;
     dbms_lob.open(fil, dbms_lob.lob_readonly);
     amt := 40; pos := 1 + dbms_lob.getlength(fil); buf := '';
     dbms_lob.read(fil, amt, pos, buf);
     dbms_output.put_line('Read F1 past EOF: '||
          utl_raw.cast_to_varchar2(buf));
     dbms_lob.close(fil);
     exception
     WHEN no_data_found
     THEN
       BEGIN
         dbms_output.put_line('End of File reached. Closing file');
         dbms_lob.fileclose(fil);
         -- or dbms_lob.filecloseall if appropriate
       END;
END;
     /

Statement processed.
End of File reached. Closing file
```

In general, you should ensure that files opened in a PL/SQL block using `DBMS_LOB` are closed before normal or abnormal termination of the block.

Maximum LOB Size

The maximum size for LOBs supported by the database is equal to the value of the blocksize of the tablespace the LOB column resides in times the value 2^{32}-1 (4294967295). This allows for a maximum LOB size ranging from 8 terabytes to 128 terabytes.

Maximum Buffer Size

The maximum buffer size, 32767 bytes.

For `BLOBs`, where buffer size is expressed in bytes, the number of bytes cannot exceed 32767.

For `CLOBs` or `NCLOBs`, where buffer size is expressed in characters, the number of characters cannot result in a buffer larger than 32767 bytes. For example, if you are using fixed-width, two-byte characters, then specifying 20000 characters is an error (20000*2 = 40000, which is greater than 32767).

114.7 DBMS_LOB Exceptions

The table in the topic describes the exceptions for DBMS_LOB.

Table 114-8 DBMS_LOB Exceptions

Exception	Code	Description
ACCESS_ERROR	22925	You are trying to write too much data to the LOB: LOB size is limited to 4 gigabytes.
CONTENTTYPE_TOOLONG	43859	The length of the contenttype string exceeds the defined maximum. Modify the length of the contenttype string and retry the operation.
CONTENTTYPEBUF_WRONG	43862	The length of the contenttype buffer is less than defined constant. Modify the length of the contenttype buffer and retry the operation.
INVALID_ARGVAL	21560	The argument is expecting a non-NULL, valid value but the argument value passed in is NULL, invalid, or out of range.
INVALID_DIRECTORY	22287	The directory used for the current operation is not valid if being accessed for the first time, or if it has been modified by the DBA since the last access.
NO_DATA_FOUND	1403	ENDOFLOB indicator for looping read operations. This is not a hard error.
NOEXIST_DIRECTORY	22285	The directory leading to the file does not exist.
NOPRIV_DIRECTORY	22286	The user does not have the necessary access privileges on the directory or the file for the operation.
OPEN_TOOMANY	22290	The number of open files has reached the maximum limit.
OPERATION_FAILED	22288	The operation attempted on the file failed.
QUERY_WRITE	14553	Cannot perform a LOB write inside a query or PDML parallel execution server
SECUREFILE_BADLOB	43856	A non-SECUREFILE LOB type was used in a SECUREFILE only call
SECUREFILE_BADPARAM	43857	An invalid argument was passed to a SECUREFILE subprogram
SECUREFILE_MARKERASED	43861	The mark provided to a FRAGMENT_* operation has been deleted
SECUREFILE_OUTOFBOUNDS	43883	Attempted to perform a FRAGMENT_* operation past the LOB end
UNOPENED_FILE	22289	The file is not open for the required operation to be performed.
VALUE_ERROR	6502	PL/SQL error for invalid values to subprogram's parameters.

114.8 Summary of DBMS_LOB Subprograms

This table lists the DBMS_LOB subprograms and briefly describes them.

Table 114-9 DBMS_LOB Package Subprograms

Subprogram	Description
APPEND Procedures	Appends the contents of the source LOB to the destination LOB
CLOB2FILE Procedure	Writes the content of a CLOB into a file.
CLOSE Procedure	Closes a previously opened internal or external LOB
COMPARE Functions	Compares two entire LOBs or parts of two LOBs

Table 114-9 (Cont.) DBMS_LOB Package Subprograms

Subprogram	Description
CONVERTTOBLOB Procedure	Reads character data from a source `CLOB` or `NCLOB` instance, converts the character data to the specified character, writes the converted data to a destination `BLOB` instance in binary format, and returns the new offsets
CONVERTTOCLOB Procedure	Takes a source `BLOB` instance, converts the binary data in the source instance to character data using the specified character, writes the character data to a destination `CLOB` or `NCLOB` instance, and returns the new offsets
COPY Procedures	Copies all, or part, of the source LOB to the destination LOB
COPY_DBFS_LINK Procedures	Copies the DBFS Link in the source LOB to the destination LOB
COPY_FROM_DBFS_LINK	Retrieves the data for the LOB from the DBFS store
CREATETEMPORARY Procedures	Creates a temporary `BLOB` or `CLOB` and its corresponding index in the user's default temporary tablespace
DBFS_LINK_GENERATE_PATH Functions	Returns a unique file path name for use in creating a DBFS Link
ERASE Procedures	Erases all or part of a LOB
FILECLOSE Procedure	Closes the file
FILECLOSEALL Procedure	Closes all previously opened files
FILEEXISTS Function	Checks if the file exists on the server
FILEGETNAME Procedure	Gets the directory object name and file name
FILEISOPEN Function	Checks if the file was opened using the input `BFILE` locators
FILEOPEN Procedure	Opens a file
FRAGMENT_DELETE Procedure	Deletes the data at the specified offset for the specified length from the LOB
FRAGMENT_INSERT Procedures	Inserts the specified data (limited to 32K) into the LOB at the specified offset
FRAGMENT_MOVE Procedure	Moves the amount of bytes (`BLOB`) or characters (`CLOB`/`NCLOB`) from the specified offset to the new offset specified
FRAGMENT_REPLACE Procedures	Replaces the data at the specified offset with the specified data (not to exceed 32k)
FREETEMPORARY Procedures	Frees the temporary `BLOB` or `CLOB` in the default temporary tablespace
GET_DBFS_LINK Functions	Returns the DBFS Link path associated with the specified SecureFile
GET_DBFS_LINK_STATE Procedures	Retrieves the current DBFS Link state of the specified SecureFile
GET_LOB_DEDUPLICATION_RATIO Function	Returns the deduplication ratio, which indicates that amount of space you can save by enabling deduplication.
GETCHUNKSIZE Functions	Returns the amount of space used in the LOB chunk to store the LOB value
GETCONTENTTYPE Functions	Returns the content ID string previously set by means of the SETCONTENTTYPE Procedure
GETLENGTH Functions	Gets the length of the LOB value

Table 114-9 (Cont.) DBMS_LOB Package Subprograms

Subprogram	Description
GETOPTIONS Functions	Obtains settings corresponding to the `option_type` field for a particular LOB
GET_STORAGE_LIMIT Function	Returns the storage limit for LOBs in your database configuration
INSTR Functions	Returns the matching position of the *nth* occurrence of the pattern in the LOB
ISOPEN Functions	Checks to see if the LOB was already opened using the input locator
ISREMOTE Function	Checks to see if the LOB is local to the database or if it belongs to a remote database.
ISSECUREFILE Function	Returns `TRUE` if the LOB locator passed to is for a SecureFiles LOB, otherwise, returns `FALSE`
ISTEMPORARY Functions	Checks if the locator is pointing to a temporary LOB
LOADBLOBFROMFILE Procedure	Loads `BFILE` data into an internal `BLOB`
LOADCLOBFROMFILE Procedure	Loads `BFILE` data into an internal `CLOB`
MOVE_TO_DBFS_LINK Procedures	Writes the specified SecureFile data to the DBFS store
OPEN Procedures	Opens a LOB (internal, external, or temporary) in the indicated mode
READ Procedures	Reads data from the LOB starting at the specified offset
SET_DBFS_LINK Procedures	Links the specified SecureFile to the specified path name. It does not copy the data to the path
SETCONTENTTYPE Procedure	Sets the content type string for the data in the LOB
SETOPTIONS Procedures	Enables CSCE features on a per-LOB basis, overriding the default LOB column settings
SUBSTR Functions	Returns part of the LOB value starting at the specified offset
TRIM Procedures	Trims the LOB value to the specified shorter length
WRITE Procedures	Writes data to the LOB from a specified offset
WRITEAPPEND Procedures	Writes a buffer to the end of a LOB

114.8.1 APPEND Procedures

This procedure appends the contents of a source internal LOB to a destination LOB. It appends the complete source LOB.

Syntax

```
DBMS_LOB.APPEND (
   dest_lob IN OUT  NOCOPY BLOB,
   src_lob  IN             BLOB);

DBMS_LOB.APPEND (
   dest_lob IN OUT  NOCOPY CLOB  CHARACTER SET ANY_CS,
   src_lob  IN             CLOB  CHARACTER SET dest_lob%CHARSET);
```

Parameters

Table 114-10 APPEND Procedure Parameters

Parameter	Description
dest_lob	Locator for the internal LOB to which the data is to be appended.
src_lob	Locator for the internal LOB from which the data is to be read.

Exceptions

Table 114-11 APPEND Procedure Exceptions

Exception	Description
VALUE_ERROR	Either the source or the destination LOB is NULL.
QUERY_WRITE	Cannot perform a LOB write inside a query or PDML parallel execution server

Usage Notes

- It is not mandatory that you wrap the LOB operation inside the Open/Close interfaces. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.

  If you do not wrap the LOB operation inside the Open/Close API, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. If a LOB is under IS JSON FORMAT JSON constraints, and you do not wrap the LOB operations inside the Open/Close API, it checks that the LOB content is a valid JSON for every write operation on the LOB.
- If APPEND is called on a LOB that has been archived, it implicitly gets the LOB before the first byte is written
- If APPEND is called on a SecureFiles LOB that is a DBFS Link, an exception is thrown.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.2 CLOB2FILE Procedure

This procedure writes the content of a CLOB into a bfile. This procedure gets called from the deprecated dbms_xslprocessor.clob2file internally.

Syntax

```
DBMS_LOB.CLOB2FILE(
   src_cl      IN  CLOB,
```

```
   file_loc    IN  VARCHAR2,
   file_name   IN  VARCHAR2,
   csid        IN  NUMBER   := 0,
   open_mode   IN  VARCHAR2 :='wb');
```

Parameters

Table 114-12 CLOB2FILE Procedure Parameters

Parameter	Description
`src_cl`	Source CLOB locator to write into a file
`file_loc`	Directory object name where the file is located
`file_name`	File name
`csid`	Character set id of the CLOB locator • Must be a valid Oracle id; otherwise returns an error • If the value is `0`, then the content of the output file will be in the database character set
`open_mode`	The mode to open the output file in. • `wb` — write byte mode, overwrites the file The default value is `wb`.

114.8.3 CLOSE Procedure

This procedure closes a previously opened internal or external LOB.

Syntax

```
DBMS_LOB.CLOSE (
   lob_loc    IN OUT NOCOPY BLOB);

DBMS_LOB.CLOSE (
   lob_loc    IN OUT NOCOPY CLOB CHARACTER SET ANY_CS);

DBMS_LOB.CLOSE (
   file_loc   IN OUT NOCOPY BFILE);
```

Parameters

Table 114-13 CLOSE Procedure Parameters

Parameter	Description
`lob_loc`	LOB locator. For more information, see Operational Notes.

Exceptions

No error is returned if the `BFILE` exists but is not opened. An error is returned if the LOB is not open.

Usage Notes

`CLOSE` requires a round-trip to the server for both internal and external LOBs. For internal LOBs, `CLOSE` triggers other code that relies on the close call, and for external LOBs (`BFILEs`), `CLOSE` actually closes the server-side operating system file.

It is not mandatory that you wrap all LOB operations inside the Open/Close interfaces. However, if you open a LOB, you must close it before you commit the transaction; an error is produced if you do not. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.

It is an error to commit the transaction before closing all opened LOBs that were opened by the transaction. When the error is returned, the openness of the open LOBs is discarded, but the transaction is successfully committed. Hence, all the changes made to the LOB and non-LOB data in the transaction are committed, but the domain and function-based indexes are not updated. If this happens, you should rebuild the functional and domain indexes on the LOB column.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.4 COMPARE Functions

This function compares two entire LOBs or parts of two LOBs.

Syntax

```
DBMS_LOB.COMPARE (
   lob_1            IN BLOB,
   lob_2            IN BLOB,
   amount           IN INTEGER := DBMS_LOB.LOBMAXSIZE,
   offset_1         IN INTEGER := 1,
   offset_2         IN INTEGER := 1)
  RETURN INTEGER;

DBMS_LOB.COMPARE (
   lob_1            IN CLOB  CHARACTER SET ANY_CS,
   lob_2            IN CLOB  CHARACTER SET lob_1%CHARSET,
   amount           IN INTEGER := DBMS_LOB.LOBMAXSIZE,
   offset_1         IN INTEGER := 1,
   offset_2         IN INTEGER := 1)
  RETURN INTEGER;

DBMS_LOB.COMPARE (
   lob_1            IN BFILE,
   lob_2            IN BFILE,
   amount           IN INTEGER,
   offset_1         IN INTEGER := 1,
   offset_2         IN INTEGER := 1)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(COMPARE, WNDS, WNPS, RNDS, RNPS);
```

Parameters

Table 114-14 COMPARE Function Parameters

Parameter	Description
lob_1	LOB locator of first target for comparison.
lob_2	LOB locator of second target for comparison.
amount	Number of bytes (for BLOBs) or characters (for CLOBs/NCLOBSs) to compare.
offset_1	Offset in bytes or characters on the first LOB (origin: 1) for the comparison.
offset_2	Offset in bytes or characters on the second LOB (origin: 1) for the comparison.

Return Values

- INTEGER: 0 if the comparison succeeds, nonzero if not.
- NULL, if any of amount, offset_1 or offset_2 is not a valid LOB offset value. A valid offset is within the range of 1 to LOBMAXSIZE inclusive.

Usage Notes

- You can only compare LOBs of the same datatype (LOBs of BLOB type with other BLOBs, and CLOBs with CLOBs, and BFILEs with BFILEs). For BFILEs, the file must be already opened using a successful FILEOPEN operation for this operation to succeed.
- COMPARE returns 0 if the data exactly matches over the range specified by the offset and amount parameters. COMPARE returns -1 if the first CLOB is less than the second, and 1 if it is greater.
- For fixed-width *n*-byte CLOBs, if the input amount for COMPARE is specified to be greater than (DBMS_LOB.LOBMAXSIZE/*n*), then COMPARE matches characters in a range of size (DBMS_LOB.LOBMAXSIZE/*n*), or Max(length(clob1), length(clob2)), whichever is lesser.
- If COMPARE is called on any LOB that has been archived, it implicitly gets the LOB before the compare begins.
- If COMPARE() is called on a SecureFiles LOB that is a DBFS Link, the linked LOB is streamed from DBFS, if possible, otherwise an exception is thrown.

Exceptions

Table 114-15 COMPARE Function Exceptions for BFILE operations

Exception	Description
UNOPENED_FILE	File was not opened using the input locator.
NOEXIST_DIRECTORY	Directory does not exist.
NOPRIV_DIRECTORY	You do not have privileges for the directory.
INVALID_DIRECTORY	Directory has been invalidated after the file was opened.
INVALID_OPERATION	File does not exist, or you do not have access privileges on the file.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.5 CONVERTTOBLOB Procedure

This procedure reads character data from a source `CLOB` or `NCLOB` instance, converts the character data to the character set you specify, writes the converted data to a destination `BLOB` instance in binary format, and returns the new offsets.

You can use this interface with any combination of persistent or temporary LOB instances as the source or destination.

Syntax

```
DBMS_LOB.CONVERTTOBLOB(
  dest_lob       IN OUT     NOCOPY  BLOB,
  src_clob       IN         CLOB CHARACTER SET ANY_CS,
  amount         IN         INTEGER,
  dest_offset    IN OUT     INTEGER,
  src_offset     IN OUT     INTEGER,
  blob_csid      IN         NUMBER,
  lang_context   IN OUT     INTEGER,
  warning        OUT        INTEGER);
```

Parameters

Table 114-16 CONVERTTOBLOB Procedure Parameters

Parameter	Description
`dest_lob`	LOB locator of the destination LOB instance.
`src_clob`	LOB locator of the source LOB instance.
`amount`	Number of characters to convert from the source LOB. If you want to copy the entire LOB, pass the constant `DBMS_LOB.LOBMAXSIZE`. If you pass any other value, it must be less than or equal to the size of the LOB.
`dest_offset`	`(IN)` Offset in bytes in the destination LOB for the start of the write. Specify a value of 1 to start at the beginning of the LOB. `(OUT)` The new offset in bytes after the end of the write.
`src_offset`	`(IN)` Offset in characters in the source LOB for the start of the read. `(OUT)` Offset in characters in the source LOB right after the end of the read.
`blob_csid`	Desired character set ID of the converted data.
`lang_context`	`(IN)` Language context, such as shift status, for the current conversion. `(OUT)` The language context at the time when the current conversion is done. This information is returned so you can use it for subsequent conversions without losing or misinterpreting any source data. For the very first conversion, or if do not care, use the default value of zero.

Table 114-16 (Cont.) CONVERTTOBLOB Procedure Parameters

Parameter	Description
warning	(OUT) Warning message. This parameter indicates when something abnormal happened during the conversion. You are responsible for checking the warning message. Currently, the only possible warning is — inconvertible character. This occurs when the character in the source cannot be properly converted to a character in destination. The default replacement character (for example, '?') is used in place of the inconvertible character. The return value of this error message is defined as the constant warn_inconvertible_char in the DBMS_LOB package.

Usage Notes

Preconditions

Before calling the CONVERTTOBLOB procedure, the following preconditions must be met:

- Both the source and destination LOB instances must exist.
- If the destination LOB is a persistent LOB, the row must be locked. To lock the row, select the LOB using the FOR UPDATE clause of the SELECT statement.

Constants and Defaults

All parameters are required. You must pass a variable for each OUT or IN OUT parameter. You must pass either a variable or a value for each IN parameter.

Table 114-17 gives a summary of typical values for each parameter. The first column lists the parameter, the second column lists the typical value, and the last column describes the result of passing the value. Note that constants are used for some values. These constants are defined in the dbmslob.sql package specification file.

Table 114-17 DBMS_LOB.CONVERTTOBLOB Typical Values

Parameter	Value	Description
amount	LOBMAXSIZE (IN)	convert the entire file
dest_offset	1 (IN)	start from the beginning
src_offset	1 (IN)	start from the beginning
blob_csid	DEFAULT_CSID (IN)	default CSID, use same CSID as source LOB
lang_context	DEFAULT_LANG_CTX (IN)	default language context
warning	NO_WARNING (OUT)	no warning message, success
	WARN_INCONVERTIBLE_CHAR (OUT)	character in source cannot be properly converted

General Notes

- You must specify the desired character set for the destination LOB in the blob_csid parameter. You can pass a zero value for blob_csid. When you do so, the database assumes that the desired character set is the same as the source LOB character set.

- You must specify the offsets for both the source and destination LOBs, and the number of characters to copy from the source LOB. The `amount` and `src_offset` values are in characters and the `dest_offset` is in bytes. To convert the entire LOB, you can specify `LOBMAXSIZE` for the `amount` parameter.
- `CONVERTTOBLOB` gets the source and/or destination LOBs as necessary prior to conversion and write of the data.

Exceptions

Table 114-18 gives possible exceptions this procedure can throw. The first column lists the exception string and the second column describes the error conditions that can cause the exception.

Table 114-18 CONVERTTOBLOB Procedure Exceptions

Exception	Description
`VALUE_ERROR`	Any of the input parameters are `NULL` or `INVALID`.
`INVALID_ARGVAL`	One or more of the following: - `src_offset` or `dest_offset` < 1. - `src_offset` or `dest_offset` > `LOBMAXSIZE`. - `amount` < 1. - `amount` > `LOBMAXSIZE`.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for more information on using LOBs in application development

114.8.6 CONVERTTOCLOB Procedure

This procedure takes a source `BLOB` instance, converts the binary data in the source instance to character data using the character set you specify, writes the character data to a destination `CLOB` or `NCLOB` instance, and returns the new offsets.

You can use this interface with any combination of persistent or temporary LOB instances as the source or destination.

Syntax

```
DBMS_LOB.CONVERTTOCLOB(
   dest_lob       IN OUT NOCOPY  CLOB CHARACTER SET ANY_CS,
   src_blob       IN             BLOB,
   amount         IN             INTEGER,
   dest_offset    IN OUT         INTEGER,
   src_offset     IN OUT         INTEGER,
   blob_csid      IN             NUMBER,
   lang_context   IN OUT         INTEGER,
   warning        OUT            INTEGER);
```

Parameters

Table 114-19 CONVERTTOCLOB Procedure Parameters

Parameter	Description
dest_lob	LOB locator of the destination LOB instance.
src_blob	LOB locator of the source LOB instance.
amount	Number of bytes to convert from the source LOB. If you want to copy the entire BLOB, pass the constant DBMS_LOB.LOBMAXSIZE. If you pass any other value, it must be less than or equal to the size of the BLOB.
dest_offset	(IN) Offset in characters in the destination LOB for the start of the write. Specify a value of 1 to start at the beginning of the LOB. (OUT) The new offset in characters after the end of the write. This offset always points to the beginning of the first complete character after the end of the write.
src_offset	(IN) Offset in bytes in the source LOB for the start of the read. (OUT) Offset in bytes in the source LOB right after the end of the read.
blob_csid	The character set ID of the source data
lang_context	(IN) Language context, such as shift status, for the current conversion. (OUT) The language context at the time when the current conversion is done. This information is returned so you can use it for subsequent conversions without losing or misinterpreting any source data. For the very first conversion, or if do not care, use the default value of zero.
warning	Warning message. This parameter indicates when something abnormal happened during the conversion. You are responsible for checking the warning message. Currently, the only possible warning is — inconvertible character. This occurs when the character in the source cannot be properly converted to a character in destination. The default replacement character (for example, '?') is used in place of the inconvertible character. The return value of this error message is defined as the constant warn_inconvertible_char in the DBMS_LOB package.

Usage Notes

Preconditions

Before calling the CONVERTTOCLOB procedure, the following preconditions must be met:

- Both the source and destination LOB instances must exist.
- If the destination LOB is a persistent LOB, the row must be locked before calling the CONVERTTOCLOB procedure. To lock the row, select the LOB using the FOR UPDATE clause of the SELECT statement.

Constants and Defaults

All parameters are required. You must pass a variable for each OUT or IN OUT parameter. You must pass either a variable or a value for each IN parameter.

Table 114-20 gives a summary of typical values for each parameter. The first column lists the parameter, the second column lists the typical value, and the last column describes the result

of passing the value. Note that constants are used for some values. These constants are defined in the `dbmslob.sql` package specification file.

Table 114-20 DBMS_LOB.CONVERTTOCLOB Typical Values

Parameter	Value	Description
`amount`	`LOBMAXSIZE (IN)`	convert the entire file
`dest_offset`	`1 (IN)`	start from the beginning
`src_offset`	`1 (IN)`	start from the beginning
`csid`	`DEFAULT_CSID (IN)`	default `CSID`, use destination `CSID`
`lang_context`	`DEFAULT_LANG_CTX (IN)`	default language context
`warning`	`NO_WARNING (OUT)`	no warning message, success
	`WARN_INCONVERTIBLE_CHAR (OUT)`	character in source cannot be properly converted

General Notes

- You must specify the desired character set for the source LOB in the `blob_csid` parameter. You can pass a zero value for `blob_csid`. When you do so, the database assumes that the desired character set is the same as the destianation LOB character set.
- You must specify the offsets for both the source and destination LOBs, and the number of characters to copy from the source LOB. The `amount` and `src_offset` values are in bytes and the `dest_offset` is in characters. To convert the entire LOB, you can specify `LOBMAXSIZE` for the `amount` parameter.
- `CONVERTTOCLOB` gets the source and/or destination LOBs as necessary prior to conversion and write of the data.

Exceptions

Table 114-21 CONVERTTOCLOB Procedure Exceptions

Exception	Description
`VALUE_ERROR`	Any of the input parameters are `NULL` or `INVALID`.
`INVALID_ARGVAL`	One or more of the following: - `src_offset` or `dest_offset` < 1. - `src_offset` or `dest_offset` > `LOBMAXSIZE`. - `amount` < 1. - `amount` > `LOBMAXSIZE`.

See Also:

Oracle Database SecureFiles and Large Objects Developer's Guide for more information on using LOBs in application development

114.8.7 COPY Procedures

This procedure copies all, or a part of, a source internal LOB to a destination internal LOB. You can specify the offsets for both the source and destination LOBs, and the number of bytes or characters to copy.

Syntax

```
DBMS_LOB.COPY (
  dest_lob    IN OUT NOCOPY BLOB,
  src_lob     IN            BLOB,
  amount      IN            INTEGER,
  dest_offset IN            INTEGER := 1,
  src_offset  IN            INTEGER := 1);

DBMS_LOB.COPY (
  dest_lob    IN OUT NOCOPY CLOB  CHARACTER SET ANY_CS,
  src_lob     IN            CLOB  CHARACTER SET dest_lob%CHARSET,
  amount      IN            INTEGER,
  dest_offset IN            INTEGER := 1,
  src_offset  IN            INTEGER := 1);
```

Parameters

Table 114-22 COPY Procedure Parameters

Parameter	Description
dest_lob	LOB locator of the copy target.
src_lob	LOB locator of source for the copy.
amount	Number of bytes (for BLOBs) or characters (for CLOBs) to copy.
dest_offset	Offset in bytes or characters in the destination LOB (origin: 1) for the start of the copy.
src_offset	Offset in bytes or characters in the source LOB (origin: 1) for the start of the copy.

Exceptions

Table 114-23 COPY Procedure Exceptions

Exception	Description
VALUE_ERROR	Any of the input parameters are NULL or invalid.
INVALID_ARGVAL	Either: - src_offset or dest_offset < 1 - src_offset or dest_offset > LOBMAXSIZE - amount < 1 - amount > LOBMAXSIZE
QUERY_WRITE	Cannot perform a LOB write inside a query or PDML parallel execution server

Usage Notes

- If the offset you specify in the destination LOB is beyond the end of the data currently in this LOB, then zero-byte fillers or spaces are inserted in the destination `BLOB` or `CLOB` respectively. If the offset is less than the current length of the destination LOB, then existing data is overwritten.
- It is not an error to specify an amount that exceeds the length of the data in the source LOB. Thus, you can specify a large amount to copy from the source LOB, which copies data from the `src_offset` to the end of the source LOB.
- It is not mandatory that you wrap the LOB operation inside the Open/Close interfaces. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.
- If you do not wrap the LOB operation inside the Open/Close API, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. If a LOB is under `IS JSON FORMAT JSON` constraints, and you do not wrap the LOB operations inside the Open/Close API, it checks that the LOB content is a valid JSON for every write operation on the LOB.
- Prior to copy, the source and destination LOBs are retrieved, if they are currently archived. For a complete over-write, the destination LOB is not retrieved.
- If the source LOB is a DBFS Link, the data is streamed from DBFS, if possible, otherwise an exception is thrown. If the destination LOB is a DBFS Link, an exception is thrown.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.8 COPY_DBFS_LINK Procedures

This procedure copies the DBFS Link in the source LOB to the destination LOB.

Syntax

```
DBMS_LOB.COPY_DBFS_LINK (
  lob_loc_dst    IN OUT BLOB,
  lob_loc_src    IN     BLOB,
  flags          IN     PLS_INTEGER DEFAULT DBFS_LINK_NOCACHE);

DBMS_LOB.COPY_DBFS_LINK (
  lob_loc_dst    IN OUT CLOB CHARACTER SET ANY_CS,
  lob_loc_src    IN     CLOB CHARACTER SET ANY_CS,
  flags          IN     PLS_INTEGER DEFAULT DBFS_LINK_NOCACHE);
```

Parameters

Table 114-24 COPY_DBFS_LINK Procedure Parameters

Parameter	Description
lob_loc_dst	LOB to be made to reference the same storage data as lob_loc_src
lob_loc_src	LOB from which to copy the reference
flags	Options to COPY_DBFS_LINK: • DBFS_LINK_NOCACHE specifies to only copy the DBFS Link • DBFS_LINK_CACHE specifies to copy the DBFS Link and read the data into the database LOB specified by lob_loc_dst so that the data is cached

Exceptions

Table 114-25 COPY_DBFS_LINK Procedure Exceptions

Exception	Description
SECUREFILE_BADLOB	Either lob_loc_src or lob_loc_dst is not a SECUREFILE
INVALID_ARGVAL	lob_loc_src LOB has not been archived
ORA-01555	If the source LOB has been retrieved, never archived, or if the LOB has been migrated in and out (modified or not) since the locator was gotten.

114.8.9 COPY_FROM_DBFS_LINK

This procedure retrieves the archived SecureFiles LOB data from the DBFS HSM store and to the database.

Syntax

```
DBMS_LOB.COPY_FROM_DBFS_LINK (
  lob_loc       IN OUT BLOB);

DBMS_LOB.COPY_FROM_DBFS_LINK (
  lob_loc       IN OUT CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 114-26 COPY_FROM_DBFS_LINK Procedure Parameters

Parameter	Description
lob_loc	LOB to be retrieved from the archive

Usage Note

COPY_FROM_DBFS_LINK does not remove the underlying DBFS file.

If the LOB is successfully retrieved, COPY_FROM_DBFS_LINK silently returns success.

Exceptions

Table 114-27 COPY_FROM_DBFS_LINK Procedure Exceptions

Exception	Description
SECUREFILE_BADLOB	lob_loc is not a SECUREFILE
ORA-01555	If the LOB has already been retrieved and has been modified since retrieval, if the LOB has been migrated in and out (modified or not) since the locator was retrieved

114.8.10 CREATETEMPORARY Procedures

This procedure creates a temporary BLOB or CLOB and its corresponding index in your default temporary tablespace.

Syntax

```
DBMS_LOB.CREATETEMPORARY (
   lob_loc IN OUT NOCOPY BLOB,
   cache   IN            BOOLEAN,
   dur     IN            PLS_INTEGER := DBMS_LOB.SESSION);

DBMS_LOB.CREATETEMPORARY (
   lob_loc IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   cache   IN            BOOLEAN,
   dur     IN            PLS_INTEGER := 10);
```

Parameters

Table 114-28 CREATETEMPORARY Procedure Parameters

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.
cache	Specifies if LOB should be read into buffer cache or not.
dur	1 of 2 predefined duration values (SESSION or CALL) which specifies a hint as to whether the temporary LOB is cleaned up at the end of the session or call. If dur is omitted, then the session duration is used.

See Also:

- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure
- *Oracle Database PL/SQL Language Reference* for more information about NOCOPY and passing temporary lobs as parameters

114.8.11 DBFS_LINK_GENERATE_PATH Functions

This subprogram returns a unique file path name for use in creating a DBFS Link.

Syntax

```
DBMS_LOB.DBFS_LINK_GENERATE_PATH (
  lob_loc       IN BLOB,
  storage_dir   IN VARCHAR2)
 RETURN VARCHAR2;

DBMS_LOB.DBFS_LINK_GENERATE_PATH (
  lob_loc       IN CLOB CHARACTER SET ANY_CS,
  storage_dir   IN VARCHAR2)
 RETURN VARCHAR2;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(dbfs_link_generate_path,
       WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-29 DBFS_LINK_GENERATE_PATH Function Parameters

Parameter	Description
lob_loc	LOB to be retrieved from DBFS
storage_dir	DBFS directory that will be the parent directory of the file

Exceptions

Table 114-30 DBFS_LINK_GENERATE_PATH Function Exceptions

Exception	Description
SECUREFILE_WRONGTYPE	lob_loc is not a SECUREFILE

Usage Notes

Returns a globally unique file pathname that can be used for archiving. This is guaranteed to be globally unique across all calls to this function for different LOBs and versions of that LOB. It is always the same for the same LOB and version.

114.8.12 ERASE Procedures

This procedure erases an entire internal LOB or part of an internal LOB.

Syntax

```
DBMS_LOB.ERASE (
   lob_loc           IN OUT   NOCOPY   BLOB,
   amount            IN OUT   NOCOPY   INTEGER,
   offset            IN                INTEGER := 1);

DBMS_LOB.ERASE (
```

```
lob_loc          IN OUT   NOCOPY   CLOB CHARACTER SET ANY_CS,
amount           IN OUT   NOCOPY   INTEGER,
offset           IN                INTEGER := 1);
```

Parameters

Table 114-31 ERASE Procedure Parameters

Parameter	Description
lob_loc	Locator for the LOB to be erased.For more information, see Operational Notes.
amount	Number of bytes (for BLOBs or BFILES) or characters (for CLOBs or NCLOBs) to be erased.
offset	Absolute offset (origin: 1) from the beginning of the LOB in bytes (for BLOBs) or characters (CLOBs).

Usage Notes

- When data is erased from the middle of a LOB, zero-byte fillers or spaces are written for BLOBs or CLOBs respectively.
- The actual number of bytes or characters erased can differ from the number you specified in the amount parameter if the end of the LOB value is reached before erasing the specified number. The actual number of characters or bytes erased is returned in the amount parameter.
- ERASE gets the LOB if it is archived, unless the erase covers the entire LOB.
- If the LOB to be erased is a DBFS Link, an exception is thrown.

> **Note:**
>
> The length of the LOB is not decreased when a section of the LOB is erased. To decrease the length of the LOB value, see the "TRIM Procedures".

Exceptions

Table 114-32 ERASE Procedure Exceptions

Exception	Description
VALUE_ERROR	Any input parameter is NULL.
INVALID_ARGVAL	Either: - amount < 1 or amount > LOBMAXSIZE - offset < 1 or offset > LOBMAXSIZE
QUERY_WRITE	Cannot perform a LOB write inside a query or PDML parallel execution server

Usage Notes

It is not mandatory that you wrap the LOB operation inside the Open/Close interfaces. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before

performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.

If you do not wrap the LOB operation inside the Open/Close API, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. If a LOB is under `IS JSON FORMAT JSON` constraints, and you do not wrap the LOB operations inside the Open/Close API, it checks that the LOB content is a valid JSON for every write operation on the LOB.

> **See Also:**
>
> - TRIM Procedures
> - *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.13 FILECLOSE Procedure

This procedure closes a `BFILE` that has already been opened through the input locator.

> **Note:**
>
> The database has only read-only access to `BFILEs`. This means that `BFILEs` cannot be written through the database.

Syntax

```
DBMS_LOB.FILECLOSE (
    file_loc IN OUT NOCOPY BFILE);
```

Parameters

Table 114-33 FILECLOSE Procedure Parameters

Parameter	Description
`file_loc`	Locator for the `BFILE` to be closed.

Exceptions

Table 114-34 FILECLOSE Procedure Exceptions

Exception	Description
`VALUE_ERROR`	`NULL` input value for `file_loc`.
`UNOPENED_FILE`	File was not opened with the input locator.
`NOEXIST_DIRECTORY`	Directory does not exist.
`NOPRIV_DIRECTORY`	You do not have privileges for the directory.
`INVALID_DIRECTORY`	Directory has been invalidated after the file was opened.

Table 114-34 (Cont.) FILECLOSE Procedure Exceptions

Exception	Description
INVALID_OPERATION	File does not exist, or you do not have access privileges on the file.

See Also:

- "FILEOPEN Procedure"
- "FILECLOSEALL Procedure"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.14 FILECLOSEALL Procedure

This procedure closes all BFILEs opened in the session.

Syntax

```
DBMS_LOB.FILECLOSEALL;
```

Exceptions

Table 114-35 FILECLOSEALL Procedure Exception

Exception	Description
UNOPENED_FILE	No file has been opened in the session.

See Also:

- "FILEOPEN Procedure"
- "FILECLOSE Procedure"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.15 FILEEXISTS Function

This function finds out if a specified BFILE locator points to a file that actually exists on the server's file system.

Syntax

```
DBMS_LOB.FILEEXISTS (
   file_loc     IN    BFILE)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(FILEEXISTS, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-36 FILEEXISTS Function Parameter

Parameter	Description
file_loc	Locator for the BFILE.

Return Values

Table 114-37 FILEEXISTS Function Return Values

Return	Description
0	Physical file does not exist.
1	Physical file exists.

Exceptions

Table 114-38 FILEEXISTS Function Exceptions

Exception	Description
NOEXIST_DIRECTORY	Directory does not exist.
NOPRIV_DIRECTORY	You do not have privileges for the directory.
INVALID_DIRECTORY	Directory has been invalidated after the file was opened.

See Also:

- "FILEISOPEN Function".
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.16 FILEGETNAME Procedure

This procedure determines the directory object and filename, given a BFILE locator.

This function only indicates the directory object name and filename assigned to the locator, not if the physical file or directory actually exists.

The maximum constraint values for the dir_alias buffer is 30, and for the entire path name, it is 2000.

Syntax

```
DBMS_LOB.FILEGETNAME (
   file_loc   IN    BFILE,
   dir_alias  OUT   VARCHAR2,
   filename   OUT   VARCHAR2);
```

Parameters

Table 114-39 FILEGETNAME Procedure Parameters

Parameter	Description
file_loc	Locator for the BFILE
dir_alias	Directory object name
filename	Name of the BFILE

Exceptions

Table 114-40 FILEGETNAME Procedure Exceptions

Exception	Description
VALUE_ERROR	Any of the input parameters are NULL or INVALID.
INVALID_ARGVAL	dir_alias or filename are NULL.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.17 FILEISOPEN Function

This function finds out whether a BFILE was opened with the specified FILE locator.

Syntax

```
DBMS_LOB.FILEISOPEN (
   file_loc   IN    BFILE)
  RETURN INTEGER;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(fileisopen, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-41 FILEISOPEN Function Parameter

Parameter	Description
file_loc	Locator for the BFILE.

Return Values

INTEGER: 0 = file is not open, 1 = file is open

Usage Notes

If the input FILE locator was never passed to the FILEOPEN procedure, then the file is considered not to be opened by this locator. However, a different locator may have this file open. In other words, openness is associated with a specific locator.

Exceptions

Table 114-42 FILEISOPEN Function Exceptions

Exception	Description
NOEXIST_DIRECTORY	Directory does not exist.
NOPRIV_DIRECTORY	You do not have privileges for the directory.
INVALID_DIRECTORY	Directory has been invalidated after the file was opened.

> **See Also:**
>
> - "FILEEXISTS Function"
> - *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.18 FILEOPEN Procedure

This procedure opens a BFILE for read-only access. BFILE data may not be written through the database.

Syntax

```
DBMS_LOB.FILEOPEN (
   file_loc   IN OUT NOCOPY  BFILE,
   open_mode  IN             BINARY_INTEGER := file_readonly);
```

Parameters

Table 114-43 FILEOPEN Procedure Parameters

Parameter	Description
file_loc	Locator for the BFILE.
open_mode	File access is read-only.

Exceptions

Table 114-44 FILEOPEN Procedure Exceptions

Exception	Description
VALUE_ERROR	file_loc or open_mode is NULL.
INVALID_ARGVAL	open_mode is not equal to FILE_READONLY.
OPEN_TOOMANY	Number of open files in the session exceeds session_max_open_files.
NOEXIST_DIRECTORY	Directory associated with file_loc does not exist.
INVALID_DIRECTORY	Directory has been invalidated after the file was opened.
INVALID_OPERATION	File does not exist, or you do not have access privileges on the file.

See Also:

- "FILECLOSE Procedure"
- "FILECLOSEALL Procedure"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.19 FRAGMENT_DELETE Procedure

This procedure deletes the data at the specified offset for the specified length from the LOB without having to rewrite all the data in the LOB following the specified offset.

Syntax

```
DBMS_LOB.FRAGMENT_DELETE (
   lob_loc     IN OUT NOCOPY BLOB,
   amount      IN            INTEGER,
   offset      IN            INTEGER);

DBMS_LOB.FRAGMENT_DELETE (
   lob_loc     IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   amount      IN            INTEGER,
   offset      IN            INTEGER);
```

Parameters

Table 114-45 FRAGMENT_DELETE Procedure Parameters

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.
amount	Number of bytes (BLOB) or characters (CLOB/NCLOB) to be removed from the LOB
offset	Offset into the LOB in bytes (BLOB) or characters (CLOB/NCLOB) to begin the deletion

Exceptions

Table 114-46 FRAGMENT_DELETE Procedure Exceptions

Exception	Description
INVALID_ARGVAL	A parameter value was invalid
QUERY_WRITE	Cannot perform operation during a query
SECUREFILE_BADLOB	A non-SECUREFILE LOB was used in a SECUREFILE LOB only call
SECUREFILE_OUTOFBOUNDS	Attempted to perform a FRAGMENT_* operation past LOB end

114.8.20 FRAGMENT_INSERT Procedures

This procedure inserts the specified data (limited to 32K) into the LOB at the specified offset.

Syntax

```
DBMS_LOB.FRAGMENT_INSERT (
   lob_loc     IN OUT NOCOPY BLOB,
   amount      IN            INTEGER,
   offset      IN            INTEGER,
   buffer      IN            RAW);

DBMS_LOB.FRAGMENT_INSERT (
   lob_loc     IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   amount      IN            INTEGER,
   offset      IN            INTEGER,
   buffer      IN            VARCHAR2 CHARACTER SET lob_loc%CHARSET);
```

Parameters

Table 114-47 FRAGMENT_INSERT Procedure Parameters

Parameter	Description
lob_loc	LOB locator.For more information, see Operational Notes.
amount	Number of bytes (BLOB) or characters (CLOB/NCLOB) to be inserted into the LOB
offset	Offset into the LOB in bytes (BLOB) or characters (CLOB/NCLOB) to begin the insertion

Table 114-47 (Cont.) FRAGMENT_INSERT Procedure Parameters

Parameter	Description
buffer	Data to insert into the LOB

Exceptions

Table 114-48 FRAGMENT_INSERT Procedure Exceptions

Exception	Description
INVALID_ARGVAL	A parameter value was invalid
QUERY_WRITE	Cannot perform operation during a query
SECUREFILE_BADLOB	A non-SECUREFILE LOB was used in a SECUREFILE LOB only call
SECUREFILE_OUTOFBOUNDS	Attempted to perform a FRAGMENT_* operation past LOB end

Usage Notes

FRAGMENT_INSERT gets the LOB, if necessary, before performing operations on the LOB.

114.8.21 FRAGMENT_MOVE Procedure

This procedure moves the amount of bytes (BLOB) or characters (CLOB/NCLOB) from the specified offset to the new offset specified.

Syntax

```
DBMS_LOB.FRAGMENT_MOVE (
   lob_loc       IN OUT NOCOPY BLOB,
   amount        IN            INTEGER,
   src_offset    IN            INTEGER,
   dest_offset   IN            INTEGER);

DBMS_LOB.FRAGMENT_MOVE (
   lob_loc       IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   amount        IN            INTEGER,
   src_offset    IN            INTEGER,
   dest_offset   IN            INTEGER);
```

Parameters

Table 114-49 FRAGMENT_MOVE Procedure Parameters

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.
amount	Number of bytes (BLOB) or characters (CLOB/NCLOB) to be moved in the LOB
src_offset	Beginning offset into the LOB in bytes (BLOB) or characters (CLOB/NCLOB) to put the data
dest_offset	Beginning offset into the LOB in bytes (BLOB) or characters (CLOB/NCLOB) to remove the data

Exceptions

Table 114-50 FRAGMENT_MOVE Procedure Exceptions

Exception	Description
INVALID_ARGVAL	A parameter value was invalid
QUERY_WRITE	Cannot perform operation during a query
SECUREFILE_BADLOB	A non-SECUREFILE LOB was used in a SECUREFILE LOB only call
SECUREFILE_OUTOFBOUNDS	Attempted to perform a FRAGMENT_* operation past LOB end

Usage Notes

- All offsets are pre-move offsets.
- Offsets of more than 1 past the end of the LOB are not permitted.
- FRAGMENT_MOVE gets the LOB, if necessary, before performing operations on the LOB.

114.8.22 FRAGMENT_REPLACE Procedures

This procedure replaces the data at the specified offset with the specified data (not to exceed 32k).

Syntax

```
DBMS_LOB.FRAGMENT_REPLACE (
   lob_loc     IN OUT NOCOPY BLOB,
   old_amount  IN            INTEGER,
   new_amount  IN            INTEGER,
   offset      IN            INTEGER,
   buffer      IN            RAW);

DBMS_LOB.FRAGMENT_REPLACE (
   lob_loc     IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,   old_amount  IN
INTEGER,
   new_amount  IN           INTEGER,
   offset      IN           INTEGER,
   buffer      IN           VARCHAR2 CHARACTER SET lob_loc%CHARSET);
```

Parameters

Table 114-51 FRAGMENT_REPLACE Function Parameters

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.
old_amount	Number of bytes (BLOB) or characters (CLOB/NCLOB) to be replaced in the LOB
new_amount	Number of bytes (BLOB) or characters (CLOB/NCLOB) to written to the LOB
offset	Beginning offset into the LOB in bytes (BLOB) or characters (CLOB/NCLOB) to put the data
buffer	Data to insert into the LOB

Exceptions

Table 114-52 FRAGMENT_REPLACE Procedure Exceptions

Exception	Description
INVALID_ARGVAL	A parameter value was invalid
QUERY_WRITE	Cannot perform operation during a query
SECUREFILE_BADLOB	A non-SECUREFILE LOB was used in a SECUREFILE LOB only call
SECUREFILE_OUTOFBOUNDS	Attempted to perform a FRAGMENT_* operation past LOB end

Usage Notes

- Invoking this procedure is equivalent to deleting the old amount of bytes/characters at offset and then inserting the new amount of bytes/characters at offset.
- FRAGMENT_REPLACE gets the LOB, if necessary, before performing operations on the LOB.

114.8.23 FREETEMPORARY Procedures

This procedure frees the temporary BLOB or CLOB in the default temporary tablespace.

Syntax

```
DBMS_LOB.FREETEMPORARY (
   lob_loc  IN OUT  NOCOPY BLOB);

DBMS_LOB.FREETEMPORARY (
   lob_loc  IN OUT  NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 114-53 FREETEMPORARY Procedure Parameters

Parameter	Description
lob_loc	LOB locator.For more information, see Operational Notes.

Usage Notes

- When a new temporary LOB is created, and there is currently no temporary LOB in use with the same duration (session, call), a new temporary LOB segment is created. When the temporary LOB is freed, the space it consumed is released to the temporary segment. If there are no other temporary LOBs for the same duration, the temporary segment is also freed.
- After the call to FREETEMPORARY, the LOB locator that was freed is marked as invalid.
- If an invalid LOB locator is assigned to another LOB locator using OCILobLocatorAssign in OCI or through an assignment operation in PL/SQL, then the target of the assignment is also freed and marked as invalid.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.24 GET_DBFS_LINK Functions

This function returns the DBFS path name for the specified SecureFile LOB.

Syntax

```
DBMS_LOB.GET_DBFS_LINK (
  lob_loc             IN     BLOB,
  storage_path        OUT VARCHAR2(DBFS_LINK_PATH_MAX_SIZE),
  lob_length          OUT NUMBER);

DBMS_LOB.GET_DBFS_LINK (
  lob_loc             IN     CLOB CHARACTER SET ANY_CS,
  storage_path        OUT VARCHAR2(DBFS_LINK_PATH_MAX_SIZE),
  lob_length          OUT NUMBER);
```

Parameters

Table 114-54 GET_DBFS_LINK Function Parameters

Parameter	Description
lob_loc	LOB to be retrieved from DBFS
storage_path	Path where the LOB is stored in DBFS
lob_length	LOB length at the time of write to DBFS

Return Values

The Archive ID

Exceptions

Table 114-55 GET_DBFS_LINK Function Exceptions

Exception	Description
SECUREFILE_BADLOB	lob_loc is not a SECUREFILE
ORA-01555	The LOB has already been retrieved and has been modified since retrieval or the LOB has been migrated in and out (modified or not) since the locator was retrieved

114.8.25 GET_DBFS_LINK_STATE Procedures

`GET_DBFS_LINK_STATE` retrieves the current link state of the specified SecureFile.

Syntax

```
DBMS_LOB.GET_DBFS_LINK_STATE (
  lob_loc       IN BLOB,
  storage_path  OUT VARCHAR2(DBFS_LINK_PATH_MAX_SIZE),
  state         OUT NUMBER,
  cached        OUT BOOLEAN);

DBMS_LOB.GET_DBFS_LINK_STATE (
  lob_loc       IN CLOB CHARACTER SET ANY_CS,
  storage_path  OUT VARCHAR2(DBFS_LINK_PATH_MAX_SIZE),
  state         OUT NUMBER,
  cached        OUT BOOLEAN);
```

Parameters

Table 114-56 GET_DBFS_LINK_STATE Procedure Parameters

Parameter	Description
`lob_loc`	LOB to be retrieved from the archive
`storage_path`	Path where the LOB is stored in the DBFS HSM store
`state`	One of `DBFS_LINK_NEVER`, `DBFS_LINK_NO` or `DBFS_LINK_YES`
`cached`	If the LOB is archived and the data was specified to be cashed on put

Exceptions

Table 114-57 GET_DBFS_LINK_STATE Procedure Exceptions

Exception	Description
`SECUREFILE_BADLOB`	`lob_loc` is not a `SECUREFILE`

Usage Notes

- If the LOB has never been archived, state is set to `DBMS_LOB.DBFS_LINK_NEVER`. If the LOB has been archived, state is set to `DBMS_LOB.DBFS_LINK_YES`. If the LOB has been previously retrieved from the archive, state is set to `DBFS_LINK_NO`.
- If the LOB was archived, but the data was left in the RDBMS, `cached` is set to `TRUE`. If the data was removed after the link was created, `cached` is set to `FALSE`, and `NULL` if state is `DBMS_LOB.DBFS_LINK_NEVER`.

114.8.26 GET_LOB_DEDUPLICATION_RATIO Function

The `GET_LOB_DEDUPLICATION_RATIO` function estimates the storage space that you can save by enabling the deduplication feature for an existing SecureFile LOB and returns the deduplication ratio.

The deduplication ratio is estimated for the number of rows in the LOB column that you specify. For example, let's consider that the deduplication ratio is 2.33. It indicates that after you enable the deduplication feature, you can save around half of the space for the sampled rows in the LOB column.

Disclaimer: The deduplication ratio is an approximate value, which is calculated based on the sampled rows in the LOB column. The actual space that you save when you enable deduplication for the complete table may be different.

Syntax

```
DBMS_LOB.GET_LOB_DEDUPLICATION_RATIO (
   tablespacename        IN     VARCHAR2,
   tabowner              IN     VARCHAR2,
   tabname               IN     VARCHAR2,
   lobcolumnname         IN     VARCHAR2,
   partname              IN     VARCHAR2,
   dedup_ratio           OUT    NUMBER,
   subset_numrows        IN     NUMBER DEFAULT DEDUP_RATIO_LOB_MAXROWS
)
```

Parameters

Table 114-58 GET_DBFS_LINK_STATE Procedure Parameters

Parameter	Description
`tablespacename`	Name of the tablespace
`tabowner`	Owner of the table
`tabname`	Name of the table
`lobcolumnname`	Name of the LOB column for which you want to calculate the deduplication ratio
`partname`	In case of partitioned tables, enter the related partition name
`subset_numrows`	Number of rows sampled to estimate the deduplication ratio. By default, the deduplication ratio is calculated for all the rows.

Return Values

Returns the deduplication ratio, `dedup_ratio`, which indicates the space that you can save by enabling the deduplication feature.

Example

The following sample code calculates and returns the deduplication ratio for `C`, a LOB column, in `ACME_TABLE`, which is owned by `JOHN`.

```
DECLARE
 dedup_ratio    number;
 l_table_name   varchar2(128) = "ACME_TABLE";
 l_column_name  varchar2(3000) = "C";
 l_tablespace_name varchar2(128) := "TBS";
 l_owner varchar2(128) := "JOHN";
BEGIN
dedup_ratio := dbms_lob.GET_LOB_DEDUPLICATION_RATIO(
    l_tablespace_name,
    l_owner,
    l_table_name,
    l_column_name,
    '',
    dedup_ratio,
    -1);
dbms_output.put_line('Deduplication ratio: ' || dedup_ratio);
END;
```

Usage Notes

- The maximum number of LOBs that this function can process is 100000 or 1% of the total number of rows in the table, whichever is lesser.

114.8.27 GETCONTENTTYPE Functions

This procedure returns the content type string previously set by means of the SETCONTENTTYPE Procedure.

Syntax

```
DBMS_LOB.GETCONTENTTYPE (
   lob_loc  IN BLOB)
 RETURN VARCHAR2;

DBMS_LOB.GETCONTENTTYPE (
   lob_loc  IN CLOB CHARACTER SET ANY_CS)
 RETURN VARCHAR2;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(getcontenttype, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-59 GETCONTENTTYPE Function Parameters

Parameter	Description
`lob_loc`	LOB whose content type is to be retrieved

Return Values

The returned content type.

If the SecureFiles LOB does not have a `contenttype` associated with it, `GETCONTENTTYPE()` returns `NULL`.

Exceptions

Table 114-60 GETCONTENTTYPE Function Exceptions

Exception	Description
`SECUREFILE_BADLOB`	`lob_loc` is not a `SECUREFILE`

Related Topics

- SETCONTENTTYPE Procedure
  This procedure sets the content type string for the data in the LOB.

114.8.28 GET_STORAGE_LIMIT Function

This function returns the LOB storage limit for the specified LOB.

Syntax

```
DBMS_LOB.GET_STORAGE_LIMIT (
   lob_loc  IN CLOB CHARACTER SET ANY_CS)
 RETURN INTEGER;

DBMS_LOB.GET_STORAGE_LIMIT (
   lob_loc  IN BLOB)
 RETURN INTEGER;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(get_storage_limit, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-61 GET_STORAGE_LIMIT Function Parameters

Parameter	Description
`lob_loc`	LOB locator. For more information, see Operational Notes.

Return Value

The value returned from this function is the maximum allowable size for specified LOB locator. For `BLOB`s, the return value depends on the block size of the tablespace the LOB resides in and is calculated as (2^{32})-1 (4294967295) times the block size of the tablespace. For `CLOB`s/`NCLOB`s, the value returned is the(2^{32})-1 (4294967295) times the block size of the tablespace divided by the character width of the `CLOB`/`NCLOB`.

Usage

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for details on LOB storage limits

114.8.29 GETCHUNKSIZE Functions

When creating the table, you can specify the chunking factor, a multiple of tablespace blocks in bytes. This corresponds to the chunk size used by the LOB data layer when accessing or modifying the LOB value. Part of the chunk is used to store system-related information, and the rest stores the LOB value. This function returns the amount of space used in the LOB chunk to store the LOB value.

Syntax

```
DBMS_LOB.GETCHUNKSIZE (
   lob_loc IN BLOB)
  RETURN INTEGER;

DBMS_LOB.GETCHUNKSIZE (
   lob_loc IN CLOB CHARACTER SET ANY_CS)
  RETURN INTEGER;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(getchunksize, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-62 GETCHUNKSIZE Function Parameters

Parameter	Description
`lob_loc`	LOB locator. For more information, see Operational Notes.

Return Values

The return value is a usable chunk size in bytes.

Usage Notes

- With regard to basic LOB files, performance is improved if you enter read/write requests using a multiple of this chunk size. For writes, there is an added benefit, because LOB chunks are versioned, and if all writes are done on a chunk basis, then no extra or excess versioning is done or duplicated. You could batch up the `WRITE` until you have enough for a chunk, instead of issuing several `WRITE` calls for the same chunk.

  These tactics of performance improvement do not apply to SecureFiles.
- Note that chunk size is independent of LOB type (`BLOB`, `CLOB`, `NCLOB`, Unicode or other character set).

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.30 GETLENGTH Functions

This function gets the length of the specified LOB. The length in bytes or characters is returned.

The length returned for a `BFILE` includes the `EOF`, if it exists. Any 0-byte or space filler in the LOB caused by previous `ERASE` or `WRITE` operations is also included in the length count. The length of an empty internal LOB is 0.

Syntax

```
DBMS_LOB.GETLENGTH (
   lob_loc    IN  BLOB)
  RETURN INTEGER;

DBMS_LOB.GETLENGTH (
   lob_loc    IN  CLOB   CHARACTER SET ANY_CS)
  RETURN INTEGER;

DBMS_LOB.GETLENGTH (
   file_loc    IN  BFILE)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(GETLENGTH, WNDS, WNPS, RNDS, RNPS);
```

Parameters

Table 114-63 GETLENGTH Function Parameter

Parameter	Description
`file_loc`	The file locator for the LOB whose length is to be returned.

Return Values

The length of the LOB in bytes or characters as an `INTEGER`. `NULL` is returned if the input LOB is `NULL` or if the input `lob_loc` is `NULL`. An error is returned in the following cases for `BFILEs`:

- `lob_loc` does not have the necessary directory and operating system privileges
- `lob_loc` cannot be read because of an operating system read error

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.31 GETOPTIONS Functions

This function obtains compression, deduplication, and encryption settings corresponding to the `option_type` field for a particular LOB.

Syntax

```
DBMS_LOB.GETOPTIONS (
   lob_loc             IN     BLOB,
   option_types        IN     PLS_INTEGER)
 RETURN PLS_INTEGER;

DBMS_LOB.GETOPTIONS (
  lob_loc             IN     CLOB CHARACTER SET ANY_CS,
  option_types        IN     PLS_INTEGER)
RETURN PLS_INTEGER;
```

Parameters

Table 114-64 GETOPTIONS Function Parameter

Parameter	Description
`lob_loc`	Locator for the LOB to be examined. For more information, see Operational Notes.
`option_type`	See Table 114-2

Return Values

The return values are a combination of `COMPRESS_ON`, `ENCRYPT_ON` and `DEDUPLICATE_ON` (see Table 114-3) depending on which option types (see Table 114-2) are passed in.

Exceptions

Table 114-65 GETOPTIONS Procedure Exceptions

Exception	Description
`INVALID_ARGVAL`	A parameter value was invalid
`QUERY_WRITE`	Cannot perform operation during a query
`SECUREFILE_BADLOB`	A non-`SECUREFILE` LOB was used in a `SECUREFILE` LOB only call

Usage Notes

You cannot turn compression or deduplication on or off for a SecureFile column that does not have those features on. The GetOptions Functions and SETOPTIONS Procedures work on individual SecureFiles. You can turn off a feature on a particular SecureFile and turn on a feature that has already been turned off by SetOptions, but you cannot turn on an option that has not been given to the SecureFile when the table was created.

114.8.32 INSTR Functions

This function returns the matching position of the *nth* occurrence of the pattern in the LOB, starting from the offset you specify.

Syntax

```
DBMS_LOB.INSTR (
   lob_loc    IN   BLOB,
   pattern    IN   RAW,
   offset     IN   INTEGER := 1,
   nth        IN   INTEGER := 1)
  RETURN INTEGER;

DBMS_LOB.INSTR (
   lob_loc    IN   CLOB      CHARACTER SET ANY_CS,
   pattern    IN   VARCHAR2  CHARACTER SET lob_loc%CHARSET,
   offset     IN   INTEGER := 1,
   nth        IN   INTEGER := 1)
  RETURN INTEGER;

DBMS_LOB.INSTR (
   file_loc   IN   BFILE,
   pattern    IN   RAW,
   offset     IN   INTEGER := 1,
   nth        IN   INTEGER := 1)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(INSTR, WNDS, WNPS, RNDS, RNPS);
```

Parameters

Table 114-66 INSTR Function Parameters

Parameter	Description
`lob_loc`	Locator for the LOB to be examined. For more information, see Operational Notes.
`file_loc`	The file locator for the LOB to be examined.
`pattern`	Pattern to be tested for. The pattern is a group of `RAW` bytes for `BLOBs`, and a character string (`VARCHAR2`) for `CLOBs`.The maximum size of the pattern is 16383 bytes.
`offset`	Absolute offset in bytes (`BLOBs`) or characters (`CLOBs`) at which the pattern matching is to start. (origin: 1)
`nth`	Occurrence number, starting at 1.

Return Values

Table 114-67 INSTR Function Return Values

Return	Description
`INTEGER`	Offset of the start of the matched pattern, in bytes or characters. It returns 0 if the pattern is not found.

Table 114-67 (Cont.) INSTR Function Return Values

Return	Description
NULL	Either: -any one or more of the `IN` parameters was `NULL` or `INVALID`. -`offset` < 1 or `offset` > `LOBMAXSIZE`. -`nth` < 1. -`nth` > `LOBMAXSIZE`.

Usage Notes

The form of the `VARCHAR2` buffer (the `pattern` parameter) must match the form of the `CLOB` parameter. In other words, if the input LOB parameter is of type `NCLOB`, then the buffer must contain `NCHAR` data. Conversely, if the input LOB parameter is of type `CLOB`, then the buffer must contain `CHAR` data.

For `BFILEs`, the file must be already opened using a successful `FILEOPEN` operation for this operation to succeed.

Operations that accept `RAW` or `VARCHAR2` parameters for pattern matching, such as `INSTR`, do not support regular expressions or special matching characters (as in the case of SQL `LIKE`) in the pattern parameter or substrings.

Exceptions

Table 114-68 INSTR Function Exceptions for BFILES

Exception	Description
`UNOPENED_FILE`	File was not opened using the input locator.
`NOEXIST_DIRECTORY`	Directory does not exist.
`NOPRIV_DIRECTORY`	You do not have privileges for the directory.
`INVALID_DIRECTORY`	Directory has been invalidated after the file was opened.
`INVALID_OPERATION`	File does not exist, or you do not have access privileges on the file.

See Also:

- "SUBSTR Functions"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.33 ISOPEN Functions

This function checks to see if the LOB was already opened using the input locator. This subprogram is for internal and external LOBs.

Syntax

```
DBMS_LOB.ISOPEN (
   lob_loc IN BLOB)
  RETURN INTEGER;

DBMS_LOB.ISOPEN (
   lob_loc IN CLOB CHARACTER SET ANY_CS)
  RETURN INTEGER;

DBMS_LOB.ISOPEN (
   file_loc IN BFILE)
  RETURN INTEGER;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(isopen, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-69 ISOPEN Function Parameters

Parameter	Description
`lob_loc`	LOB locator. For more information, see Operational Notes.
`file_loc`	File locator.

Return Values

The return value is 1 if the LOB is open, 0 otherwise.

Usage Notes

For `BFILES`, openness is associated with the locator. If the input locator was never passed to `OPEN`, the `BFILE` is not considered to be opened by this locator. However, a different locator may have opened the `BFILE`. More than one `OPEN` can be performed on the same `BFILE` using different locators.

For internal LOBs, openness is associated with the LOB, not with the locator. If locator1 opened the LOB, then locator2 also sees the LOB as open. For internal LOBs, `ISOPEN` requires a round-trip, because it checks the state on the server to see if the LOB is indeed open.

For external LOBs (`BFILEs`), `ISOPEN` also requires a round-trip, because that's where the state is kept.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.34 ISREMOTE Function

This function checks to see if the LOB is local to the database or if it belongs to a remote database.

Syntax

```
DBMS_LOB.ISREMOTE (
   lob_loc IN BLOB)
  RETURN BOOLEAN;

DBMS_LOB.ISREMOTE (
   lob_loc IN CLOB CHARACTER SET ANY_CS)
  RETURN BOOLEAN;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(isremote, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-70 ISREMOTE Function Parameter

Parameter	Description
lob_loc	Locator for the LOB.

Return Values

BOOLEAN: TRUE for remote LOBs obtained over a database link; FALSE for LOBs obtained from local database

> **See Also:**
>
> - *Distributed LOBs* chapter in Database SecureFiles and Large Objects Developer's Guide for more details on the usage of this procedure.

114.8.35 ISSECUREFILE Function

This function returns TRUE if the LOB locator passed to it is for a SecureFile LOB. It returns FALSE otherwise.

Syntax

```
DBMS_LOB ISSECUREFILE(
    lob_loc    IN      BLOB)
  RETURN BOOLEAN;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(issecurefile, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-71 ISSECUREFILE Function Parameter

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.

Return Values

This function returns TRUE if the LOB locator passed to it is for a SecureFile LOB. It returns FALSE otherwise.

114.8.36 ISTEMPORARY Functions

This function determines whether a LOB instance is temporary.

Syntax

```
DBMS_LOB.ISTEMPORARY (
   lob_loc IN BLOB)
  RETURN INTEGER;

DBMS_LOB.ISTEMPORARY (
   lob_loc IN CLOB CHARACTER SET ANY_CS)
  RETURN INTEGER;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(istemporary, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 114-72 ISTEMPORARY Procedure Parameters

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.

Return Values

The return value is 1 if the LOB is temporary and exists; 0 if the LOB is not temporary or does not exist; NULL if the given locator is NULL.

Usage Notes

When you free a Temporary LOB with FREETEMPORARY, the LOB locator is not set to NULL. Consequently, ISTEMPORARY will return 0 for a locator that has been freed but not explicitly reset to NULL.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.37 LOADBLOBFROMFILE Procedure

This procedure loads data from BFILE to internal BLOB. This achieves the same outcome as LOADFROMFILE, and returns the new offsets.

Syntax

```
DBMS_LOB.LOADBLOBFROMFILE (
   dest_lob    IN OUT NOCOPY BLOB,
   src_bfile   IN            BFILE,
   amount      IN            INTEGER,
   dest_offset IN OUT        INTEGER,
   src_offset  IN OUT        INTEGER);
```

Parameters

Table 114-73 LOADBLOBFROMFILE Procedure Parameters

Parameter	Description
dest_lob	BLOB locator of the target for the load.
src_bfile	BFILE locator of the source for the load.
amount	Number of bytes to load from the BFILE. You can also use DBMS_LOB.LOBMAXSIZE to load until the end of the BFILE.
dest_offset	(IN) Offset in bytes in the destination BLOB (origin: 1) for the start of the write. (OUT) New offset in bytes in the destination BLOB right after the end of this write, which is also where the next write should begin.
src_offset	(IN) Offset in bytes in the source BFILE (origin: 1) for the start of the read. (OUT) Offset in bytes in the source BFILE right after the end of this read, which is also where the next read should begin.

Usage Notes

- You can specify the offsets for both the source and destination LOBs, and the number of bytes to copy from the source BFILE. The amount and src_offset, because they refer to the BFILE, are in terms of bytes, and the dest_offset is in bytes for BLOBs.
- If the offset you specify in the destination LOB is beyond the end of the data currently in this LOB, then zero-byte fillers or spaces are inserted in the destination BLOB. If the offset is less than the current length of the destination LOB, then existing data is overwritten.
- There is an error if the input amount plus offset exceeds the length of the data in the BFILE (unless the amount specified is LOBMAXSIZE which you can specify to continue loading until the end of the BFILE is reached).
- It is not mandatory that you wrap the LOB operation inside the OPEN/CLOSE operations. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.
- If you do not wrap the LOB operation inside the OPEN/CLOSE, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. Therefore, it is recommended that you enclose write operations to the LOB within the OPEN or CLOSE statement.

- LOADFROMFILE gets the destination LOB prior to the load unless the load covers the entire LOB.

Constants and Defaults

There is no easy way to omit parameters. You must either declare a variable for `IN/OUT` parameter or provide a default value for the `IN` parameter. Here is a summary of the constants and the defaults that can be used.

Table 114-74 Suggested Values of the Parameter

Parameter	Default Value	Description
`amount`	`DBMS_LOB.LOBMAXSIZE (IN)`	Load the entire file
`dest_offset`	`1 (IN)`	start from the beginning
`src_offset`	`1 (IN)`	start from the beginning

Constants defined in `DBMSLOB.SQL`

```
lobmaxsize                    CONSTANT INTEGER        := DBMS_LOB.LOBMAXSIZE;
```

Exceptions

Table 114-75 LOADBLOBFROMFILE Procedure Exceptions

Exception	Description
`VALUE_ERROR`	Any of the input parameters are `NULL` or `INVALID`.
`INVALID_ARGVAL`	Either: - `src_offset` or `dest_offset` < 1. - `src_offset` or `dest_offset` > `LOBMAXSIZE`. - `amount` < 1. - `amount` > `LOBMAXSIZE`.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.38 LOADCLOBFROMFILE Procedure

This procedure loads data from a `BFILE` to an internal `CLOB/NCLOB` with necessary character set conversion and returns the new offsets.

Syntax

```
DBMS_LOB.LOADCLOBFROMFILE (
   dest_lob       IN OUT NOCOPY   NOCOPY CLOB CHARACTER SET ANY_CS,
   src_bfile      IN              BFILE,
   amount         IN              INTEGER,
   dest_offset    IN OUT          INTEGER,
   src_offset     IN OUT          INTEGER,
```

```
bfile_csid      IN              NUMBER,
lang_context    IN OUT          INTEGER,
warning         OUT             INTEGER);
```

Parameters

Table 114-76 LOADCLOBFROMFILE Procedure Parameters

Parameter	Description
dest_lob	CLOB/NCLOB locator of the target for the load.
src_bfile	BFILE locator of the source for the load.
amount	Number of bytes to load from the BFILE. Use DBMS_LOB.LOBMAXSIZE of load until the end of the BFILE.
dest_offset	(IN) Offset in characters in the destination CLOB (origin: 1) for the start of the write. (OUT) The new offset in characters right after the end of this load, which is also where the next load should start. It always points to the beginning of the first complete character after the end of load. If the last character is not complete, offset goes back to the beginning of the partial character.
src_offset	(IN) Offset in bytes in the source BFILE (origin: 1) for the start of the read. (OUT) Offset in bytes in the source BFILE right after the end of this read, which is also where the next read should begin.
bfile_csid	Character set id of the source (BFILE) file.
lang_context	(IN) Language context, such as shift status, for the current load. (OUT) The language context at the time when the current load stopped, and what the next load should be using if continuing loading from the same source. This information is returned to the user so that they can use it for the continuous load without losing or misinterpreting any source data. For the very first load or if do not care, simply use the default 0. The details of this language context is hidden from the user. One does not need to know what it is or what's in it in order to make the call
warning	(OUT) Warning message. This indicates something abnormal happened during the loading. It may or may not be caused by the user's mistake. The loading is completed as required, and it's up to the user to check the warning message. Currently, the only possible warning is the inconvertible character. This happens when the character in the source cannot be properly converted to a character in destination, and the default replacement character (for example, '?') is used in place. The message is defined the constant value DBMS_LOB.WARN_INCONVERTIBLE_CHAR.

Usage Notes

You can specify the offsets for both the source and destination LOBs, and the number of bytes to copy from the source BFILE. The amount and src_offset, because they refer to the BFILE, are in terms of bytes, and the dest_offset is in characters for CLOBs.

If the offset you specify in the destination LOB is beyond the end of the data currently in this LOB, then zero-byte fillers or spaces are inserted in the destination CLOB. If the offset is less than the current length of the destination LOB, then existing data is overwritten.

There is an error if the input amount plus offset exceeds the length of the data in the BFILE (unless the amount specified is LOBMAXSIZE which you can specify to continue loading until the end of the BFILE is reached).

Note the following requirements:

- The destination character set is always the same as the database character set in the case of `CLOB` and national character set in the case of `NCLOB`.
- `csid=0` indicates the default behavior that uses database `csid` for `CLOB` and national `csid` for `NCLOB` in the place of source `csid`. Conversion is still necessary if it is of varying width
- It is not mandatory that you wrap the LOB operation inside the `OPEN/CLOSE` operations. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.

  If you do not wrap the LOB operation inside the `OPEN/CLOSE`, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. Therefore, it is recommended that you enclose write operations to the LOB within the `OPEN` or `CLOSE` statement.

The source `BFILE` can contain data in the Unicode character set. The Unicode standard defines many encoding schemes that provide mappings from Unicode characters to sequences of bytes. Table 114-77 lists Unicode encodings schemes supported by this subprogram.

Table 114-77 Supported Unicode Encoding Schemes

Encoding Scheme	Oracle Name	bfile_csid Value
UTF-8	AL32UTF8	873
UTF-16BE	AL16UTF16	2000
UTF-16LE	AL16UTF16LE	2002
CESU-8	UTF8	871
UTF-EBCDIC	UTFE	872
UTF-16	UTF16	1000

All three `UTF-16` encoding schemes encode Unicode characters as 2-byte unsigned integers. Integers can be stored in big-endian or in little-endian byte order. The `UTF-16BE` encoding scheme defines big-endian data. The `UTF-16LE` scheme defines little-endian data. The `UTF-16` scheme requires that the source `BFILE` contains the Byte Order Mark (BOM) character in the first two bytes to define the byte order. The BOM code is `0xFEFF`. If the code is stored as `{0xFE,0xFF}`, the data is interpreted as big-endian. If it is stored as `{0xFF,0xFE}`, the data is interpreted as little-endian.

In `UTF-8` and in `CESU-8` encodings the Byte Order Mark is stored as `{0xEF,0xBB, 0xBF}`. With any of the Unicode encodings, the corresponding BOM sequence at the beginning of the file is recognized and not loaded into the destination LOB.

Constants

Here is a summary of the constants and the suggested values that can be used.

Table 114-78 Suggested Values of the LOADCLOBFROMFILE Parameter

Parameter	Suggested Value	Description
amount	DBMS_LOB.LOBMAXSIZE (IN)	Load the entire file
dest_offset	1 (IN)	start from the beginning
src_offset	1 (IN)	start from the beginning
csid	0 (IN)	default csid, use destination csid
lang_context	0 (IN)	default language context
warning	0 (OUT)	no warning message, everything is ok

Constants defined in DBMSLOB.SQL

```
lobmaxsize                    CONSTANT INTEGER       := 18446744073709551615;
warn_inconvertible_char       CONSTANT INTEGER       := 1;
default_csid                  CONSTANT INTEGER       := 0;
default_lang_ctx              CONSTANT INTEGER       := 0;
no_warning                    CONSTANT INTEGER       := 0;
```

Exceptions

Table 114-79 LOADCLOBFROMFILE Procedure Exceptions

Exception	Description
VALUE_ERROR	Any of the input parameters are NULL or INVALID.
INVALID_ARGVAL	Either: - src_offset or dest_offset < 1. - src_offset or dest_offset > LOBMAXSIZE. - amount < 1. - amount > LOBMAXSIZE.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.39 MOVE_TO_DBFS_LINK Procedures

This procedure archives the specified LOB data (from the database) into the DBFS HSM Store.

Syntax

```
DBMS_LOB.MOVE_TO_DBFS_LINK (
  lob_loc       IN OUT BLOB,
  storage_path  IN     VARCHAR2(dbfs_link_path_max_size),
  flags         IN     BINARY INTEGER DEFAULT DBFS_LINK_NOCACHE);

DBMS_LOB.MOVE_TO_DBFS_LINK (
  lob_loc       IN OUT CLOB CHARACTER SET ANY_CS,
```

```
  storage_path  IN     VARCHAR2(dbfs_link_path_max_size),
  flags         IN     BINARY_INTEGER DEFAULT DBFS_LINK_NOCACHE);
```

Parameters

Table 114-80 MOVE_TO_DBFS_LINK Procedure Parameters

Parameter	Description
lob_loc	LOB to be archived
storage_path	Path where the LOB will be stored
flags	Either DBFS_LINK_CACHE or DBFS_LINK_NOCACHE. If DBFS_LINK_CACHE is specified, the LOB data continues to be stored in the RDBMS as well as being written to the DBFS store. DBFS_LINK_NOCACHE specifies that the LOB data should be deleted from the RDBMS once written to the DBFS.

Exceptions

Table 114-81 MOVE_TO_DBFS_LINK Procedure Exceptions

Exception	Description
SECUREFILE_BADLOB	lob_loc is not a SECUREFILE

Usage Notes

- If the LOB is already archived, the procedure silently returns as if the put was successful. In that case, if DBFS_LINK_NOCACHE is specified, or flags is defaulted, the LOB data is removed from the RDBMS.
- Calling this procedure multiple times on the same LOB with the same flags has no effect.
- Calling the procedure on a LOB that is already archived causes the LOB to be cached (DBFS_LINK_CACHE) or removed (DBFS_LINK_NOCACHE) according to the flag setting.

114.8.40 OPEN Procedures

This procedure opens a LOB, internal or external, in the indicated mode. Valid modes include read-only, and read/write.

Syntax

```
DBMS_LOB.OPEN (
   lob_loc   IN OUT NOCOPY BLOB,
   open_mode IN            BINARY_INTEGER);

DBMS_LOB.OPEN (
   lob_loc   IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   open_mode IN            BINARY_INTEGER);

DBMS_LOB.OPEN (
   file_loc  IN OUT NOCOPY BFILE,
   open_mode IN            BINARY_INTEGER := file_readonly);
```

Parameters

Table 114-82 OPEN Procedure Parameters

Parameter	Description
lob_loc	LOB locator. For more information, see Operational Notes.
open_mode	Mode in which to open. For BLOB and CLOB types, the mode can be either: LOB_READONLY or LOB_READWRITE. For BFILE types, the mode must be FILE_READONLY.

Usage Notes

> **Note:**
>
> If the LOB was opened in read-only mode, and if you try to write to the LOB, then an error is returned. BFILE can only be opened with read-only mode.

OPEN requires a round-trip to the server for both internal and external LOBs. For internal LOBs, OPEN triggers other code that relies on the OPEN call. For external LOBs (BFILEs), OPEN requires a round-trip because the actual operating system file on the server side is being opened.

It is not mandatory that you wrap all LOB operations inside the Open/Close interfaces. However, if you open a LOB, you must close it before you commit the transaction; an error is produced if you do not. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.

It is an error to commit the transaction before closing all opened LOBs that were opened by the transaction. When the error is returned, the openness of the open LOBs is discarded, but the transaction is successfully committed. Hence, all the changes made to the LOB and non-LOB data in the transaction are committed, but the domain and function-based indexes are not updated. If this happens, you should rebuild the functional and domain indexes on the LOB column.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.41 READ Procedures

This procedure reads a piece of a LOB, and returns the specified amount into the buffer parameter, starting from an absolute offset from the beginning of the LOB.

The number of bytes or characters actually read is returned in the amount parameter. If the input offset points past the End of LOB, then amount is set to 0, and a NO_DATA_FOUND exception is raised.

Syntax

```
DBMS_LOB.READ (
   lob_loc   IN             BLOB,
   amount    IN OUT  NOCOPY INTEGER,
   offset    IN             INTEGER,
   buffer    OUT            RAW);

DBMS_LOB.READ (
   lob_loc   IN             CLOB CHARACTER SET ANY_CS,
   amount    IN OUT  NOCOPY INTEGER,
   offset    IN             INTEGER,
   buffer    OUT            VARCHAR2 CHARACTER SET lob_loc%CHARSET);

DBMS_LOB.READ (
   file_loc   IN             BFILE,
   amount    IN OUT   NOCOPY INTEGER,
   offset    IN              INTEGER,
   buffer    OUT             RAW);
```

Parameters

Table 114-83 READ Procedure Parameters

Parameter	Description
lob_loc	Locator for the LOB to be read. For more information, see Operational Notes.
file_loc	The file locator for the LOB to be examined.
amount	Number of bytes (for BLOBs) or characters (for CLOBs) to read, or number that were read.
offset	Offset in bytes (for BLOBs) or characters (for CLOBs) from the start of the LOB (origin: 1).
buffer	Output buffer for the read operation.

Exceptions

Table 114-84 lists exceptions that apply to any LOB instance. Table 114-85 lists exceptions that apply only to BFILEs.

Table 114-84 READ Procedure Exceptions

Exception	Description
VALUE_ERROR	Any of lob_loc, amount, or offset parameters are NULL.
INVALID_ARGVAL	Either: - amount < 1 - amount > 32767 bytes (or the character equivalent) - offset < 1 - offset > LOBMAXSIZE - amount is greater, in bytes or characters, than the capacity of buffer.
NO_DATA_FOUND	End of the LOB is reached, and there are no more bytes or characters to read from the LOB: amount has a value of 0.

Table 114-85 READ Procedure Exceptions for BFILEs

Exception	Description
UNOPENED_FILE	File is not opened using the input locator.
NOEXIST_DIRECTORY	Directory does not exist.
NOPRIV_DIRECTORY	You do not have privileges for the directory.
INVALID_DIRECTORY	Directory has been invalidated after the file was opened.
INVALID_OPERATION	File does not exist, or you do not have access privileges on the file.

Usage Notes

- The form of the VARCHAR2 buffer must match the form of the CLOB parameter. In other words, if the input LOB parameter is of type NCLOB, then the buffer must contain NCHAR data. Conversely, if the input LOB parameter is of type CLOB, then the buffer must contain CHAR data.
- When calling DBMS_LOB.READ from the client (for example, in a BEGIN/END block from within SQL*Plus), the returned buffer contains data in the client's character set. The database converts the LOB value from the server's character set to the client's character set before it returns the buffer to the user.
- READ gets the LOB, if necessary, before the read.
- If the LOB is a DBFS LINK, data is streamed from DBFS, if possible, otherwise an exception is thrown.

> **See Also:**
>
> *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.42 SET_DBFS_LINK Procedures

This function links the specified SecureFile to the specified path name. It does not copy the data to the path.

Syntax

```
DBMS_LOB.SET_DBFS_LINK (
  lob_loc        IN OUT BLOB,
  archive_id     IN     RAW(1024));

DBMS_LOB.SET_DBFS_LINK(
  lob_loc_dst    IN OUT CLOB CHARACTER SET ANY_CS,
  archive_id     IN     RAW(1024));
```

Parameters

Table 114-86 SET_DBFS_LINK Procedure Parameters

Parameter	Description
lob_loc	LOB for which to store the reference value
archive_id	Archive ID as returned by calling either of the GET_DBFS_LINK Functions Functions

Exceptions

Table 114-87 SET_DBFS_LINK Procedure Exceptions

Exception	Description
SECUREFILE_BADLOB	lob_loc is not a SECUREFILE

114.8.43 SETCONTENTTYPE Procedure

This procedure sets the content type string for the data in the LOB.

Syntax

```
DBMS_LOB.SETCONTENTTYPE (
   lob_loc      IN OUT NOCOPY BLOB,
   contenttype  IN            VARCHAR2);

DBMS_LOB.SETCONTENTTYPE (
   lob_loc     IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   contenttype IN            VARCHAR2);
```

Parameters

Table 114-88 SETCONTENTTYPE Procedure Parameters

Parameter	Description
lob_loc	LOB to be assigned the content type
contenttype	String to be assigned

Exceptions

Table 114-89 SETCONTENTTYPE Procedure Exceptions

Exception	Description
SECUREFILE_BADLOB	lob_loc is not a SECUREFILE

Usage Notes

To clear an existing content type associated with a SECUREFILE, invoke SETCONTENTTYPE with contenttype set to empty string.

114.8.44 SETOPTIONS Procedures

This procedure enables/disables compression and deduplication on a per-LOB basis, overriding the default LOB column settings.

Syntax

```
DBMS_LOB.SETOPTIONS (
   lob_loc             IN     BLOB,
   option_types        IN     PLS_INTEGER,
   options             IN     PLS_INTEGER);

DBMS_LOB.SETOPTIONS (
  lob_loc             IN     CLOB CHARACTER SET ANY_CS,
  option_types        IN     PLS_INTEGER,
  options             IN     PLS_INTEGER);
```

Parameters

Table 114-90 SETOPTIONS Procedure Parameter

Parameter	Description
lob_loc	Locator for the LOB to be examined. For more information, see Operational Notes.
option_type	See Table 114-2
options	See Table 114-3

Exceptions

Table 114-91 SETOPTIONS Procedure Exceptions

Exception	Description
SECUREFILE_BADLOB	Unsupported object type for the operation
INVALID_ARGVAL	A parameter value was invalid
QUERY_WRITE	Cannot perform operation during a query

Usage Notes

- DBMS_LOB.SETOPTIONS cannot be used to enable or disable encryption on individual LOBs.
- You cannot turn the compression or deduplication features on or off for a SecureFile column if they were not turned when the table was created.

  The GETOPTIONS Functions and SETOPTIONS Procedures work on individual SecureFiles. You can turn off compression or deduplication on a particular SecureFiles LOB and turn on them on, *if* they have already been turned off by SETOPTIONS.
- This call incurs a round-trip to the server to make the changes persistent.

114.8.45 SUBSTR Functions

This function returns `amount` bytes or characters of a LOB, starting from an absolute `offset` from the beginning of the LOB.

For fixed-width `n`-byte `CLOBs`, if the input amount for `SUBSTR` is greater than (32767/`n`), then `SUBSTR` returns a character buffer of length (32767/`n`), or the length of the `CLOB`, whichever is lesser. For CLOBs in a varying-width character set, `n` is the maximum byte-width used for characters in the CLOB.

Syntax

```
DBMS_LOB.SUBSTR (
   lob_loc     IN    BLOB,
   amount      IN    INTEGER := 32767,
   offset      IN    INTEGER := 1)
  RETURN RAW;

DBMS_LOB.SUBSTR (
   lob_loc     IN    CLOB   CHARACTER SET ANY_CS,
   amount      IN    INTEGER := 32767,
   offset      IN    INTEGER := 1)
  RETURN VARCHAR2 CHARACTER SET lob_loc%CHARSET;

DBMS_LOB.SUBSTR (
   file_loc     IN    BFILE,
   amount      IN    INTEGER := 32767,
   offset      IN    INTEGER := 1)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(SUBSTR, WNDS, WNPS, RNDS, RNPS);
```

Parameters

Table 114-92 SUBSTR Function Parameters

Parameter	Description
`lob_loc`	Locator for the LOB to be read. For more information, see Operational Notes.
`file_loc`	The file locator for the LOB to be examined.
`amount`	Number of bytes (for `BLOBs`) or characters (for `CLOBs`) to be read.
`offset`	Offset in bytes (for `BLOBs`) or characters (for `CLOBs`) from the start of the LOB (origin: 1).

Return Values

Table 114-93 SUBSTR Function Return Values

Return	Description
`RAW`	Function overloading that has a `BLOB` or `BFILE` in parameter.
`VARCHAR2`	`CLOB` version.

Table 114-93 (Cont.) SUBSTR Function Return Values

Return	Description
NULL	Either: - any input parameter is NULL - amount < 1 - amount > 32767 - offset < 1 - offset > LOBMAXSIZE

Exceptions

Table 114-94 SUBSTR Function Exceptions for BFILE operations

Exception	Description
UNOPENED_FILE	File is not opened using the input locator.
NOEXIST_DIRECTORY	Directory does not exist.
NOPRIV_DIRECTORY	You do not have privileges for the directory.
INVALID_DIRECTORY	Directory has been invalidated after the file was opened.
INVALID_OPERATION	File does not exist, or you do not have access privileges on the file.

Usage Notes

- The form of the VARCHAR2 buffer must match the form of the CLOB parameter. In other words, if the input LOB parameter is of type NCLOB, then the buffer must contain NCHAR data. Conversely, if the input LOB parameter is of type CLOB, then the buffer must contain CHAR data.
- When calling DBMS_LOB.SUBSTR from the client (for example, in a BEGIN/END block from within SQL*Plus), the returned buffer contains data in the client's character set. The database converts the LOB value from the server's character set to the client's character set before it returns the buffer to the user.
- DBMS_LOB.SUBSTR will return 8191 or more characters based on the characters stored in the LOBs. If all characters are not returned as a consequence of the character byte size exceeding the available buffer, the user should either call DBMS_LOB.SUBSTR with a new offset to read the remaining characters, or call the subprogram on loop until all the data is extracted.
- SUBSTR gets the LOB, if necessary, before read.
- If the LOB is a DBFS Link, the data is streamed from DBFS, if possible, otherwise, an exception is thrown.

> **See Also:**
>
> - "INSTR Functions"
> - "READ Procedures"
> - *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.46 TRIM Procedures

This procedure trims the value of the internal LOB to the length you specify in the `newlen` parameter.

Specify the length in bytes for `BLOBs`, and specify the length in characters for `CLOBs`.

> **Note:**
>
> The `TRIM` procedure decreases the length of the LOB to the value specified in the `newlen` parameter.

If you attempt to `TRIM` an empty LOB, then nothing occurs, and `TRIM` returns no error. If the new length that you specify in `newlen` is greater than the size of the LOB, then an exception is raised.

Syntax

```
DBMS_LOB.TRIM (
   lob_loc         IN OUT  NOCOPY BLOB,
   newlen          IN             INTEGER);

DBMS_LOB.TRIM (
   lob_loc         IN OUT  NOCOPY CLOB CHARACTER SET ANY_CS,
   newlen          IN             INTEGER);
```

Parameters

Table 114-95 TRIM Procedure Parameters

Parameter	Description
`lob_loc`	Locator for the internal LOB whose length is to be trimmed. For more information, see Operational Notes.
`newlen`	New, trimmed length of the LOB value in bytes for `BLOBs` or characters for `CLOBs`.

Exceptions

Table 114-96 TRIM Procedure Exceptions

Exception	Description
VALUE_ERROR	lob_loc is NULL.
INVALID_ARGVAL	Either: - new_len < 0 - new_len > LOBMAXSIZE
QUERY_WRITE	Cannot perform a LOB write inside a query or PDML parallel execution server

Usage Notes

- It is not mandatory that you wrap the LOB operation inside the Open/Close interfaces. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.
- If you do not wrap the LOB operation inside the Open/Close API, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. If a LOB is under IS JSON FORMAT JSON constraints, and you do not wrap the LOB operations inside the Open/Close API, it checks that the LOB content is a valid JSON for every write operation on the LOB.
- TRIM gets the LOB, if necessary, before altering the length of the LOB, unless the new length specified is '0'

See Also:

- "ERASE Procedures"
- "WRITEAPPEND Procedures"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.47 WRITE Procedures

This procedure writes a specified amount of data into an internal LOB, starting from an absolute offset from the beginning of the LOB. The data is written from the buffer parameter.

WRITE replaces (overwrites) any data that already exists in the LOB at the offset, for the length you specify.

Syntax

```
DBMS_LOB.WRITE (
   lob_loc  IN OUT NOCOPY  BLOB,
   amount   IN             INTEGER,
```

```
   offset   IN             INTEGER,
   buffer   IN             RAW);

DBMS_LOB.WRITE (
   lob_loc  IN OUT  NOCOPY CLOB   CHARACTER SET ANY_CS,
   amount   IN             INTEGER,
   offset   IN             INTEGER,
   buffer   IN             VARCHAR2 CHARACTER SET lob_loc%CHARSET);
```

Parameters

Table 114-97 WRITE Procedure Parameters

Parameter	Description
lob_loc	Locator for the internal LOB to be written to. For more information, see Operational Notes
amount	Number of bytes (for BLOBs) or characters (for CLOBs) to write
offset	Offset in bytes (for BLOBs) or characters (for CLOBs) from the start of the LOB (origin: 1) for the write operation.
buffer	Input buffer for the write

Exceptions

Table 114-98 WRITE Procedure Exceptions

Exception	Description
VALUE_ERROR	Any of lob_loc, amount, or offset parameters are NULL, out of range, or INVALID.
INVALID_ARGVAL	Either: - amount < 1 - amount > 32767 bytes (or the character equivalent) - offset < 1 - offset > LOBMAXSIZE
QUERY_WRITE	Cannot perform a LOB write inside a query or PDML parallel execution server
SECUREFILE_OUTOFBOUNDS	Attempted to perform a write operation past the end of a LOB having FRAGMENT_* on it

Usage Notes

- There is an error if the input amount is more than the data in the buffer. If the input amount is less than the data in the buffer, then only amount bytes or characters from the buffer is written to the LOB. If the offset you specify is beyond the end of the data currently in the LOB, then zero-byte fillers or spaces are inserted in the BLOB or CLOB respectively.
- The form of the VARCHAR2 buffer must match the form of the CLOB parameter. In other words, if the input LOB parameter is of type NCLOB, then the buffer must contain NCHAR data. Conversely, if the input LOB parameter is of type CLOB, then the buffer must contain CHAR data.
- When calling DBMS_LOB.WRITE from the client (for example, in a BEGIN/END block from within SQL*Plus), the buffer must contain data in the client's character set. The database

converts the client-side buffer to the server's character set before it writes the buffer data to the LOB.

- It is not mandatory that you wrap the LOB operation inside the Open/Close interfaces. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.
- If you do not wrap the LOB operation inside the Open/Close API, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. If a LOB is under `IS JSON FORMAT JSON` constraints, and you do not wrap the LOB operations inside the Open/Close API, it checks that the LOB content is a valid JSON for every write operation on the LOB.
- `WRITE` gets the LOB, if necessary, before writing the LOB, unless the write is specified to overwrite the entire LOB.

See Also:

- "APPEND Procedures"
- "COPY Procedures"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

114.8.48 WRITEAPPEND Procedures

This procedure writes a specified amount of data to the end of an internal LOB. The data is written from the `buffer` parameter.

Syntax

```
DBMS_LOB.WRITEAPPEND (
   lob_loc IN OUT NOCOPY BLOB,
   amount  IN            INTEGER,
   buffer  IN            RAW);

DBMS_LOB.WRITEAPPEND (
   lob_loc IN OUT NOCOPY CLOB CHARACTER SET ANY_CS,
   amount  IN            INTEGER,
   buffer  IN            VARCHAR2 CHARACTER SET lob_loc%CHARSET);
```

Parameters

Table 114-99 WRITEAPPEND Procedure Parameters

Parameter	Description
`lob_loc`	Locator for the internal LOB to be written to. For more information, see Operational Notes
`amount`	Number of bytes (for `BLOBs`) or characters (for `CLOBs`) to write
`buffer`	Input buffer for the write

Usage Notes

There is an error if the input amount is more than the data in the buffer. If the input amount is less than the data in the buffer, then only amount bytes or characters from the buffer are written to the end of the LOB.

Exceptions

Table 114-100 WRITEAPPEND Procedure Exceptions

Exception	Description
VALUE_ERROR	Any of lob_loc, amount, or offset parameters are NULL, out of range, or INVALID.
INVALID_ARGVAL	Either: - amount < 1 - amount > 32767 bytes (or the character equivalent)
QUERY_WRITE	Cannot perform a LOB write inside a query or PDML parallel execution server

Usage Notes

- The form of the VARCHAR2 buffer must match the form of the CLOB parameter. In other words, if the input LOB parameter is of type NCLOB, then the buffer must contain NCHAR data. Conversely, if the input LOB parameter is of type CLOB, then the buffer must contain CHAR data.
- When calling DBMS_LOB.WRITEAPPEND from the client (for example, in a BEGIN/END block from within SQL*Plus), the buffer must contain data in the client's character set. The database converts the client-side buffer to the server's character set before it writes the buffer data to the LOB.
- It is not mandatory that you wrap the LOB operation inside the Open/Close interfaces. If you did not open the LOB before performing the operation, the functional and domain indexes on the LOB column are updated during the call. However, if you opened the LOB before performing the operation, you must close it before you commit the transaction. When an internal LOB is closed, it updates the functional and domain indexes on the LOB column.
- If you do not wrap the LOB operation inside the Open/Close API, the functional and domain indexes are updated each time you write to the LOB. This can adversely affect performance. If a LOB is under IS JSON FORMAT JSON constraints, and you do not wrap the LOB operations inside the Open/Close API, it checks that the LOB content is a valid JSON for every write operation on the LOB.
- WRITEAPPEND gets the LOB, if necessary, before appending to the LOB.

See Also:

- "APPEND Procedures"
- "COPY Procedures"
- "WRITE Procedures"
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for additional details on usage of this procedure

115
DBMS_LOCK

The `DBMS_LOCK` package provides an interface to Oracle Lock Management services.

You can request a lock of a specific mode, give it a unique name recognizable in another procedure in the same or another instance, change the lock mode, and release it.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Rules and Limits
- Operational Notes
- Summary of DBMS_LOCK Subprograms

> **See Also:**
>
> For more information, and an example of how to use the `DBMS_LOCK` package, *Oracle Database Development Guide*

115.1 DBMS_LOCK Overview

The `DBMS_LOCK` package has many beneficial uses.

These uses include the following:

- Providing exclusive access to a device, such as a terminal
- Providing application-level enforcement of read locks
- Detecting when a lock is released and cleanup after the application
- Synchronizing applications and enforcing sequential processing

115.2 DBMS_LOCK Security Model

There might be operating system-specific limits on the maximum number of total locks available. This *must* be considered when using locks or making this package available to other users. Consider granting the `EXECUTE` privilege only to specific users or roles.

A better alternative would be to create a cover package limiting the number of locks used and grant `EXECUTE` privilege to specific users. An example of a cover package is documented in the `DBMS_LOCK.SQL` package specification file. The abbreviations for these locks as they appear in Enterprise Manager monitors are in parentheses.

115.3 DBMS_LOCK Constants

The DBMS_LOCK package includes several constants to use when specifying parameter values.

These constants are shown in the following table.

Table 115-1 DBMS_LOCK Constants

Name	Alternate Name(s)	Type	Value	OEM Abbreviation	Description
NL_MODE	NuL1	INTEGER	1	-	-
SS_MODE	Sub Shared	INTEGER	2	ULRS	This can be used on an aggregate object to indicate that share locks are being acquired on subparts of the object.
SX_MODE	• Sub eXclusive • Row Exclusive Mode	INTEGER	3	ULRX	This can be used on an aggregate object to indicate that exclusive locks are being acquired on sub-parts of the object.
S_MODE	• Shared • Row Exclusive Mode • Intended Exclusive	INTEGER	4	ULRSX	-
SSX_MODE	• Shared Sub eXclusive • Share Row Exclusive Mode	INTEGER	5	-	This indicates that the entire aggregate object has a share lock, but some of the sub-parts may additionally have exclusive locks.
X_MODE	Exclusive	INTEGER	6	ULX	-

These are the various lock modes (nl -> "NuLl", ss -> "Sub Shared", sx -> "Sub eXclusive", s -> "Shared", ssx -> "Shared Sub eXclusive", x -> "eXclusive").

115.4 DBMS_LOCK. Rules and Limits

When another process holds "held", an attempt to get "get" succeeds or fails, based on the held mode and type of get.

The following table describes the results:

Table 115-2 Lock Compatibility

HELD MODE	GET NL	GET SS	GET SX	GET S	GET SSX	GET X
NL	Success	Success	Success	Success	Success	Success
SS	Success	Success	Success	Success	Success	Fail
SX	Success	Success	Success	Fail	Fail	Fail
S	Success	Success	Fail	Success	Fail	Fail
SSX	Success	Success	Fail	Fail	Fail	Fail
X	Success	Fail	Fail	Fail	Fail	Fail

```
maxwait  constant integer := 32767;
```

The constant `maxwait` waits forever.

115.5 DBMS_LOCK Operational Notes

User locks never conflict with Oracle locks because they are identified with the prefix "UL". You can view these locks using the Enterprise Manager lock monitor screen or the appropriate fixed views.

User locks are automatically released when a session terminates. The lock identifier is a number in the range of 0 to 1073741823.

Because a reserved user lock is the same as an Oracle lock, it has all the functionality of an Oracle lock, such as deadlock detection. Be certain that any user locks used in distributed transactions are released upon `COMMIT`, or an undetected deadlock may occur.

`DBMS_LOCK` is most efficient with a limit of a few hundred locks for each session. Oracle strongly recommends that you develop a standard convention for using these locks in order to avoid conflicts among procedures trying to use the same locks. For example, include your company name as part of your lock names.

115.6 Summary of DBMS_LOCK Subprograms

This table lists the `DBMS_LOCK` subprograms and briefly describes them.

Table 115-3 ***DBMS_LOCK Package Subprograms***

Subprogram	Description
ALLOCATE_UNIQUE Procedure	Allocates a unique lock ID to a named lock
ALLOCATE_UNIQUE_AUTONOMOUS Procedure	Allocates a unique lock ID to a named lock
CONVERT Function	Converts a lock from one mode to another
RELEASE Function	Releases a lock
REQUEST Function	Requests a lock of a specific mode.

115.6.1 ALLOCATE_UNIQUE Procedure

This procedure allocates a unique lock identifier (in the range of 1073741824 to 1999999999) a specified lock name. Lock identifiers are used to enable applications to coordinate their use of locks. This is provided because it may be easier for applications to coordinate their use of locks based on lock names rather than lock numbers.

Syntax

```
DBMS_LOCK.ALLOCATE_UNIQUE (
   lockname         IN  VARCHAR2,
   lockhandle       OUT VARCHAR2,
   expiration_secs  IN  INTEGER   DEFAULT 864000);
```

Parameters

Table 115-4 ALLOCATE_UNIQUE Procedure Parameters

Parameter	Description
lockname	Name of the lock for which you want to generate a unique ID. Do not use lock names beginning with ORA$; these are reserved for products supplied by Oracle.
lockhandle	Returns the handle to the lock ID generated by ALLOCATE_UNIQUE. You can use this handle in subsequent calls to REQUEST, CONVERT, and RELEASE. A handle is returned instead of the actual lock ID to reduce the chance that a programming error accidentally creates an incorrect, but valid, lock ID. This provides better isolation between different applications that are using this package. LOCKHANDLE can be up to VARCHAR2 (128). All sessions using a lock handle returned by ALLOCATE_UNIQUE with the same lock name are referring to the same lock. Therefore, do not pass lock handles from one session to another.
expiration_secs	Number of seconds to wait after the last ALLOCATE_UNIQUE has been performed on a specified lock, before permitting that lock to be deleted from the DBMS_LOCK_ALLOCATED table. The default waiting period is 10 days. You should not delete locks from this table. Subsequent calls to ALLOCATE_UNIQUE may delete expired locks to recover space.

Usage Notes

If you choose to identify locks by name, you can use ALLOCATE_UNIQUE to generate a unique lock identification number for these named locks.

The first session to call ALLOCATE_UNIQUE with a new lock name causes a unique lock ID to be generated and stored in the dbms_lock_allocated table. Subsequent calls (usually by other sessions) return the lock ID previously generated.

A lock name is associated with the returned lock ID for at least expiration_secs (defaults to 10 days) past the last call to ALLOCATE_UNIQUE with the specified lock name. After this time, the row in the dbms_lock_allocated table for this lock name may be deleted in order to recover space. ALLOCATE_UNIQUE performs a commit.

> **WARNING:**
>
> Named user locks may be less efficient, because Oracle uses SQL to determine the lock associated with a specified name.

Exceptions

ORA-20000, ORU-10003: Unable to find or insert lock <lockname> into catalog dbms_lock_allocated.

115.6.2 ALLOCATE_UNIQUE_AUTONOMOUS Procedure

This procedure allocates a unique lock identifier (in the range of 1073741824 to 1999999999) a specified lock name and is an autonomous version of the ALLOCATE_UNIQUE procedure. This procedure works exactly same as that of ALLOCATE_UNIQUE, except that the procedure will run as an autonomous transaction. Therefore the commits in ALLOCATE_UNIQUE_AUTONOMOUS procedure will not affect the calling procedure. The ALLOCATE_UNIQUE_AUTONOMOUS procedure is implemented in DB 12.1 and later releases.

Syntax

```
DBMS_LOCK.ALLOCATE_UNIQUE_AUTONOMOUS (
   lockname         IN  VARCHAR2,
   lockhandle       OUT VARCHAR2,
   expiration_secs  IN  INTEGER   DEFAULT 864000);
```

Parameters

Table 115-5 ALLOCATE_UNIQUE_AUTONOMOUS Procedure Parameters

Parameter	Description
lockname	Name of the lock for which you want to generate a unique ID. Do not use lock names beginning with ORA$; these are reserved for products supplied by Oracle.
lockhandle	Returns the handle to the lock ID generated by ALLOCATE_UNIQUE_AUTONOMOUS. You can use this handle in subsequent calls to REQUEST, CONVERT, and RELEASE. A handle is returned instead of the actual lock ID to reduce the chance that a programming error accidentally creates an incorrect, but valid, lock ID. This provides better isolation between different applications that are using this package. LOCKHANDLE can be up to VARCHAR2 (128). All sessions using a lock handle returned by ALLOCATE_UNIQUE_AUTONOMOUS with the same lock name are referring to the same lock. Therefore, do not pass lock handles from one session to another.
expiration_secs	Number of seconds to wait after the last ALLOCATE_UNIQUE_AUTONOMOUS has been performed on a specified lock, before permitting that lock to be deleted from the ALLOCATE_UNIQUE_AUTONOMOUS table. The default waiting period is 10 days. You should not delete locks from this table. Subsequent calls to ALLOCATE_UNIQUE_AUTONOMOUS may delete expired locks to recover space.

Usage Notes

If you choose to identify locks by name, you can use ALLOCATE_UNIQUE_AUTONOMOUS to generate a unique lock identification number for these named locks.

The first session to call ALLOCATE_UNIQUE_AUTONOMOUS with a new lock name causes a unique lock ID to be generated and stored in the dbms_lock_allocated table. Subsequent calls (usually by other sessions) return the lock ID previously generated.

A lock name is associated with the returned lock ID for at least `expiration_secs` (defaults to 10 days) past the last call to `ALLOCATE_UNIQUE_AUTONOMOUS` with the specified lock name. After this time, the row in the `dbms_lock_allocated` table for this lock name may be deleted in order to recover space. `ALLOCATE_UNIQUE_AUTONOMOUS` performs a commit.

> **WARNING:**
>
> Named user locks may be less efficient, because Oracle uses SQL to determine the lock associated with a specified name.

Exceptions

`ORA-20000`, `ORU-10003`: Unable to find or insert lock <`lockname`> into catalog `dbms_lock_allocated`.

115.6.3 CONVERT Function

This function converts a lock from one mode to another. `CONVERT` is an overloaded function that accepts either a user-defined lock identifier, or the lock handle returned by the `ALLOCATE_UNIQUE` procedure.

Syntax

```
DBMS_LOCK.CONVERT(
   id         IN INTEGER ||
   lockhandle IN VARCHAR2,
   lockmode   IN INTEGER,
   timeout    IN NUMBER DEFAULT MAXWAIT)
  RETURN INTEGER;
```

Parameters

Table 115-6 CONVERT Function Parameters

Parameter	Description
`id` or `lockhandle`	User assigned lock identifier, from 0 to 1073741823, or the lock handle, returned by `ALLOCATE_UNIQUE`, of the lock mode you want to change
`lockmode`	New mode that you want to assign to the specified lock. For the available modes and their associated integer identifiers, see Constants.
`timeout`	Number of seconds to continue trying to change the lock mode. If the lock cannot be converted within this time period, then the call returns a value of 1 (timeout).

Return Values

Table 115-7 CONVERT Function Return Values

Return Value	Description
`0`	Success

Table 115-7 (Cont.) CONVERT Function Return Values

Return Value	Description
1	Timeout
2	Deadlock
3	Parameter error
4	Don't own lock specified by `id` or `lockhandle`
5	Illegal lock handle

115.6.4 RELEASE Function

This function explicitly releases a lock previously acquired using the `REQUEST` function.

Locks are automatically released at the end of a session. `RELEASE` is an overloaded function that accepts either a user-defined lock identifier, or the lock handle returned by the `ALLOCATE_UNIQUE` procedure.

Syntax

```
DBMS_LOCK.RELEASE (
   id          IN INTEGER)
  RETURN INTEGER;

DBMS_LOCK.RELEASE (
   lockhandle IN VARCHAR2)
  RETURN INTEGER;
```

Parameters

Table 115-8 RELEASE Function Parameter

Parameter	Description
`id` or `lockhandle`	User assigned lock identifier, from 0 to 1073741823, or the lock handle, returned by `ALLOCATE_UNIQUE`, of the lock mode you want to change

Return Values

Table 115-9 RELEASE Function Return Values

Return Value	Description
0	Success
3	Parameter error
4	Do not own lock specified by `id` or `lockhandle`
5	Illegal lock handle

115.6.5 REQUEST Function

This function requests a lock with a specified mode.

`REQUEST` is an overloaded function that accepts either a user-defined lock identifier, or the lock handle returned by the `ALLOCATE_UNIQUE` procedure.

Syntax

```
DBMS_LOCK.REQUEST(
   id                 IN  INTEGER ||
   lockhandle         IN  VARCHAR2,
   lockmode           IN  INTEGER DEFAULT X_MODE,
   timeout            IN  INTEGER DEFAULT MAXWAIT,
   release_on_commit  IN  BOOLEAN DEFAULT FALSE)
  RETURN INTEGER;
```

The current default values, such as `X_MODE` and `MAXWAIT`, are defined in the `DBMS_LOCK` package specification.

Parameters

Table 115-10 REQUEST Function Parameters

Parameter	Description
`id` or `lockhandle`	User assigned lock identifier, from 0 to 1073741823, or the lock handle, returned by `ALLOCATE_UNIQUE`, of the lock mode you want to change
`lockmode`	Mode that you are requesting for the lock. For the available modes and their associated integer identifiers, see Constants.
`timeout`	Number of seconds to continue trying to grant the lock. If the lock cannot be granted within this time period, then the call returns a value of 1 (`timeout`).
`release_on_commit`	Set this parameter to `TRUE` to release the lock on commit or roll-back. Otherwise, the lock is held until it is explicitly released or until the end of the session.

Return Values

Table 115-11 REQUEST Function Return Values

Return Value	Description
0	Success
1	Timeout
2	Deadlock
3	Parameter error
4	Already own lock specified by `id` or `lockhandle`
5	Illegal lock handle

116
DBMS_LOGMNR

The DBMS_LOGMNR package, one of a set of LogMiner packages, contains the subprograms you use to initialize the LogMiner tool and to begin and end a LogMiner session.

> **Note:**
>
> The ability to create flat file dictionary dumps of pluggable databases (PDBs) is desupported in Oracle Database 21c.
>
> In previous releases, using a flat file dictionary was one means of mining the redo logs for the changes associated with a specific PDB whose data dictionary was contained within the flat file. This feature is now desupported. Starting with Oracle Database 21c, Oracle recommends that you call DBMS_LOGMNR.START_LOGMNR, and supply the system change number (SCN) or time range that you want to mine. The SCN or time range options of START_LOGMNR are enhanced to support mining of individual PDBs.

The CONTINUOUS_MINE functionality of the LogMiner package is obsolete. It was deprecated in Oracle Database 12c release 2 (12.2). There is no replacement functionality.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Views
- Operational Notes
- Summary of DBMS_LOGMNR Subprograms

> **See Also:**
>
> *Oracle Database Utilities* for information regarding LogMiner.

116.1 DBMS_LOGMNR Overview

Oracle LogMiner, which is part of Oracle Database, enables you to query online and archived redo log files through a SQL interface. The DBMS_LOGMNR package provides the majority of the tools needed to start and stop LogMiner and specify the redo log files of interest.
All changes made to user data or to the database dictionary are recorded in the Oracle redo log files, so that database recovery operations can be performed. You can take advantage of the data recorded in the redo log files to accomplish other tasks, such as the following::

- Pinpointing when a logical corruption to a database, such as errors made at the application level, may have begun
- Determining what actions you would have to take to perform fine-grained recovery at the transaction level.
- Performance tuning and capacity planning through trend analysis.
- Tracking any data manipulation language (DML) and data definition language (DDL) statements run on the database, the order in which they were executed, and who executed them.

> **See Also:**
>
> DBMS_LOGMNR_D for information on the package subprograms that extract a LogMiner dictionary and re-create LogMiner tables in alternate tablespaces

116.2 DBMS_LOGMNR Security Model

You must have the `EXECUTE_CATALOG_ROLE` role to use the `DBMS_LOGMNR` package.

116.3 DBMS_LOGMNR Constants

The `DBMS_LOGMNR` package defines several enumerated constants for specifying parameter values. Enumerated constants must be prefixed with the package name, for example, `DBMS_LOGMNR.NEW`.

> **Note:**
>
> The `continuous_mine` option for the `dbms_logmnr.start_logmnr` package is desupported in Oracle Database 19c (19.1), and is no longer available. The use of flat files is desupported for PDBs in Oracle Database 19c, and desupported entirely in later releases.

The following table describes the constants for the `ADD_LOGFILE options` flag in the `DBMS_LOGMNR` package.

Table 116-1 Constants for ADD_LOGFILE Options Flag

Constant	Description
`NEW`	Implicitly calls the `DBMS_LOGMNR.END_LOGMNR` procedure to end the current LogMiner session and then creates a new session. The new session starts a new list of redo log files to be analyzed, beginning with the redo log file you specify.
`ADDFILE`	Adds the specified redo log file to the list of redo log files to be analyzed. Any attempt to add a duplicate file raises an exception (`ORA-01289`). This is the default if no options flag is specified.

Table 116-2 describes the constants for the `START_LOGMNR` options flag in the `DBMS_LOGMNR` package.

Table 116-2 Constants for START_LOGMNR Options Flag

Constant	Description
COMMITTED_DATA_ONLY	If set, DML statements corresponding to committed transactions are returned. DML statements corresponding to a committed transaction are grouped together. Transactions are returned in their commit order. Transactions that are rolled back or in-progress are filtered out, as are internal redo records (those related to index operations, management, and so on). If this option is not set, all rows for all transactions (committed, rolled back, and in-progress) are returned in the order in which they are found in the redo logs (in order of SCN values).
SKIP_CORRUPTION	Directs a select operation on the V$LOGMNR_CONTENTS view to skip any corruptions in the redo log file being analyzed and continue processing. This option works only when a block in the redo log file (and not the header of the redo log file) is corrupt. You should check the INFO column in the V$LOGMNR_CONTENTS view to determine the corrupt blocks skipped by LogMiner. When a corruption in the redo log file is skipped, the OPERATION column contains the value CORRUPTED_BLOCKS, and the STATUS column contains the value 1343.
DDL_DICT_TRACKING	If the LogMiner dictionary in use is in the redo log files, LogMiner updates its internal dictionary if a DDL event occurs. This ensures that correct SQL_REDO and SQL_UNDO information is maintained for objects that are modified after the LogMiner internal dictionary is built. The database to which LogMiner is connected must be open. This option cannot be used in conjunction with the DICT_FROM_ONLINE_CATALOG option..
DICT_FROM_ONLINE_CATALOG	Directs LogMiner to use the current online database dictionary rather than a LogMiner dictionary contained in the redo log files being analyzed. This option cannot be used in conjunction with the DDL_DICT_TRACKING option. The database to which LogMiner is connected must be the same one that generated the redo log files. Expect to see a value of 2 in the STATUS column of the V$LOGMNR_CONTENTS view if the table definition in the database does not match the table definition in the redo log file.
DICT_FROM_REDO_LOGS	If set, LogMiner expects to find a LogMiner dictionary in the redo log files that were specified. The redo log files are specified with the DBMS_LOGMNR.ADD_LOGFILE procedure or with the DBMS_LOGMNR.START_LOGMNR procedure.
NO_SQL_DELIMITER	If set, the SQL delimiter (a semicolon) is not placed at the end of reconstructed SQL statements. This is helpful for applications that open a cursor and then execute the reconstructed statements.
NO_ROWID_IN_STMT	If set, the ROWID clause is not included in the reconstructed SQL statements. The redo log file may already contain logically unique identifiers for modified rows if supplemental logging is enabled. When using this option, you must be sure that supplemental logging was enabled in the source database at the appropriate level and that no duplicate rows exist in the tables of interest. LogMiner does not make any guarantee regarding the uniqueness of logical row identifiers.
PRINT_PRETTY_SQL	If set, LogMiner formats the reconstructed SQL statements for ease of reading. These reconstructed SQL statements are not executable.

Table 116-2 (Cont.) Constants for START_LOGMNR Options Flag

Constant	Description
STRING_LITERALS_IN_STMT	If set, SQL_REDO and SQL_UNDO use literals for numbers and datetime and interval column types.

To specify more than one option, use a plus sign (+) between them. For example:

```
EXECUTE DBMS_LOGMNR.START_LOGMNR(OPTIONS => -
   DBMS_LOGMNR.DDL_DICT_TRACKING + DBMS_LOGMNR.DICT_FROM_REDO_LOGS);
```

116.4 DBMS_LOGMNR Views

The DBMS_LOGMNR package uses the views listed under *Accessing Logminer Operational Information In Views* in *Oracle Database Utilities*.

> **See Also:**
>
> *Oracle Database Utilities*

116.5 DBMS_LOGMNR Operational Notes

A **LogMiner session** begins with a call to DBMS_LOGMNR.ADD_LOGFILE or DBMS_LOGMNR.START_LOGMNR (the former if you plan to specify log files explicitly; the latter if you plan to use continuous mining). The session ends with a call to DBMS_LOGMNR.END_LOGMNR.

Within a LogMiner session, you can specify the redo log files to be analyzed and the SCN or time range of interest; then you can issue SQL SELECT statements against the V$LOGMNR_CONTENTS view to retrieve the data of interest.

ADD_LOGFILE Procedure must be invoked before START_LOGMNR Procedure.

> **Note:**
>
> You must add log files before filtering. Continuous logging is no longer supported. If logfiles have not been added that match the time or the SCN that you provide, then DBMS_LOGMNR.START_LOGMNR fails with the error ORA-01291: missing logfile.

116.6 Summary of DBMS_LOGMNR Subprograms

This table lists the DBMS_LOGMNR subprograms and briefly describes them.

In a multitenant container database (CDB) some subprograms must be called from the root. There may be other differences as well. See the individual subprogram descriptions for details.

Table 116-3 DBMS_LOGMNR Package Subprograms

Subprogram	Description
ADD_LOGFILE Procedure	Adds a redo log file to the existing or newly created list of redo log files for LogMiner to process, so that if a new list is created, this marks the beginning of a LogMiner session
COLUMN_PRESENT Function	Call this function for any row returned from the `V$LOGMNR_CONTENTS` view to determine if undo or redo column values exist for the column specified by the `column_name` input parameter to this function
END_LOGMNR Procedure	Finishes a LogMiner session
MINE_VALUE Function	Call this function for any row returned from the `V$LOGMNR_CONTENTS` view to retrieve the undo or redo column value of the column specified by the `column_name` input parameter to this function
REMOVE_LOGFILE Procedure	Removes a redo log file from the list of redo log files for LogMiner to process
START_LOGMNR Procedure	Initializes the LogMiner utility and starts LogMiner (unless the session was already started with a call to `DBMS_LOGMNR.ADD_LOGFILE`)

116.6.1 ADD_LOGFILE Procedure

This procedure adds a file to an existing or newly created list of log files for LogMiner to process.

> **Note:**
>
> The `continuous_mine` option for the `dbms_logmnr.start_logmnr` package is desupported in Oracle Database 19c (19.1), and is no longer available. The export of the LogMiner dictionary to flat files is desupported for PDBs in Oracle Database 19c, and desupported entirely in later releases.

In a CDB, the `ADD_LOGFILE` procedure must be called from the root database. You must have the `LOGMINING` administrative privilege to use this procedure.

Syntax

```
DBMS_LOGMNR.ADD_LOGFILE (
   LogFileName     IN VARCHAR2,
   options         IN BINARY_INTEGER default ADDFILE );
```

Parameters

Table 116-4 ADD_LOGFILE Procedure Parameters

Parameter	Description
`LogFileName`	Specifies the name of the redo log file to add to the list of redo log files to be analyzed during this session.

Table 116-4 (Cont.) ADD_LOGFILE Procedure Parameters

Parameter	Description
options	Does one of the following: • Starts a new LogMiner session and a new list of redo log files for analysis (DBMS_LOGMNR.NEW) • Adds a file to an existing list of redo log files for analysis (DBMS_LOGMNR.ADDFILE) See Table 116-1.

Exceptions

Table 116-5 ADD_LOGFILE Procedure Exceptions

Exception	Description
ORA-01284	Specified file cannot be opened.
ORA-01287	Specified file is from a different database incarnation.
ORA-01289	Specified file has already been added to the list. Duplicate redo log files cannot be added.
ORA-01290	Specified file is not in the current list and therefore cannot be removed from the list.
ORA-01324	Specified file cannot be added to the list because there is a DB_ID mismatch.

Usage Notes

- The DBMS_LOGMNR.ADD_LOGFILE call from a PDB connection is not supported for adhoc users and returns error 65040.
- Dumping Flat File dictionary and mining using Flat File dictionary are not supported from a PDB connection in Oracle Database 19c, and desupported entirely in later releases.
- Before querying the V$LOGMNR_CONTENTS view, you must make a successful call to the DBMS_LOGMNR.START_LOGMNR procedure (within the current LogMiner session).
- The LogMiner session must be set up with a list of redo log files to be analyzed. Use the ADD_LOGFILE procedure to specify the list of redo log files to analyze.
- If you want to analyze more than one redo log file, you must call the ADD_LOGFILE procedure separately for each redo log file. The redo log files do not need to be registered in any particular order.
- Both archived and online redo log files can be mined.
- After you have added the first redo log file to the list, each additional redo log file that you add to the list must be associated with the same database and database RESETLOGS SCN as the first redo log file. (The database RESETLOGS SCN uniquely identifies each execution of an ALTER DATABASE OPEN RESETLOGS statement. When the online redo logs are reset, Oracle creates a new and unique incarnation of the database.)
- To analyze the redo log files from a different database (or a database incarnation with a different database RESETLOGS SCN) than that with which the current list of redo log files is associated, use the END_LOGMNR procedure to end the current LogMiner session, and then build a new list using the ADD_LOGFILE procedure.

- LogMiner matches redo log files by the log sequence number. Thus, two redo log files with different names but with the same log sequence number will return the ORA-01289 exception. For instance, the online counterpart of an archived redo log file has a different name from the archived redo log file, but attempting to register it with LogMiner after registering the archived counterpart will result in the ORA-01289 exception being returned.

116.6.2 COLUMN_PRESENT Function

This function is designed to be used in conjunction with the `MINE_VALUE` function.

If the `MINE_VALUE` function returns a `NULL` value, it can mean either:

- The specified column is not present in the redo or undo portion of the data.
- The specified column is present and has a `NULL` value.

To distinguish between these two cases, use the `COLUMN_PRESENT` function, which returns a `1` if the column is present in the redo or undo portion of the data. Otherwise, it returns a `0`.

Syntax

```
DBMS_LOGMNR.COLUMN_PRESENT (
     sql_redo_undo      IN  RAW,
     column_name        IN  VARCHAR2 default '') RETURN NUMBER;
```

Parameters

Table 116-6 COLUMN_PRESENT Function Parameters

Parameter	Description
`sql_redo_undo`	Specifies either the `REDO_VALUE` or the `UNDO_VALUE` column in the `V$LOGMNR_CONTENTS` view from which to extract data values. See the Usage Notes for more information.
`column_name`	Specifies the fully qualified name (`schema.table.column`) of the column for which this function will return information. In a CDB, the column name is specified as follows: `container_name:schema.table.column`

Return Values

Table 116-7 describes the return values for the `COLUMN_PRESENT` function. The `COLUMN_PRESENT` function returns 1 if the self-describing record (the first parameter) contains the column specified in the second parameter. This can be used to determine the meaning of `NULL` values returned by the `DBMS_LOGMNR.MINE_VALUE` function.

Table 116-7 Return Values for COLUMN_PRESENT Function

Return	Description
`0`	Specified column is not present in this row of `V$LOGMNR_CONTENTS`.
`1`	Column is present in this row of `V$LOGMNR_CONTENTS`.

Exceptions

Table 116-8 COLUMN_PRESENT Function Exceptions

Exception	Description
ORA-01323	Currently, a LogMiner dictionary is not associated with the LogMiner session. You must specify a LogMiner dictionary for the LogMiner session.
ORA-00904	Value specified for the column_name parameter is not a fully qualified column name.

Usage Notes

- To use the COLUMN_PRESENT function, you must have successfully started LogMiner.
- The COLUMN_PRESENT function must be invoked in the context of a select operation on the V$LOGMNR_CONTENTS view.
- The COLUMN_PRESENT function does not support LONG, LOB, ADT, or COLLECTION datatypes.
- The value for the sql_redo_undo parameter depends on the operation performed and the data of interest:
  - If an update operation was performed and you want to know what the value was prior to the update operation, specify UNDO_VALUE.
  - If an update operation was performed and you want to know what the value is after the update operation, specify REDO_VALUE.
  - If an insert operation was performed, typically you would specify REDO_VALUE (because the value of a column prior to an insert operation will always be NULL).
  - If a delete operation was performed, typically you would specify UNDO_VALUE (because the value of a column after a delete operation will always be NULL).

116.6.3 END_LOGMNR Procedure

This procedure finishes a LogMiner session. Because this procedure performs cleanup operations that may not otherwise be done, you must use it to properly end a LogMiner session. This procedure is called automatically when you log out of a database session or when you call DBMS_LOGMNR.ADD_LOGFILE and specify the NEW option.

Syntax

```
DBMS_LOGMNR.END_LOGMNR;
```

Exceptions

Table 116-9 END_LOGMNR Procedure Exception

Exception	Description
ORA-01307	No LogMiner session is currently active. The END_LOGMNR procedure was called without adding any log files or before the START_LOGMNR procedure was called

116.6.4 MINE_VALUE Function

This function facilitates queries based on a column's data value.

This function takes two arguments. The first one specifies whether to mine the redo (`REDO_VALUE`) or undo (`UNDO_VALUE`) portion of the data. The second argument is a string that specifies the fully qualified name of the column to be mined. The `MINE_VALUE` function always returns a string that can be converted back to the original datatype.

Syntax

```
DBMS_LOGMNR.MINE_VALUE (
     sql_redo_undo      IN  RAW,
     column_name        IN  VARCHAR2 default '') RETURN VARCHAR2;
```

Parameters

Table 116-10 MINE_VALUE Function Parameters

Parameter	Description
`sql_redo_undo`	Specifies either the `REDO_VALUE` or the `UNDO_VALUE` column in the `V$LOGMNR_CONTENTS` view from which to extract data values. See the Usage Notes for more information.
`column_name`	Specifies the fully qualified name (`schema.table.column`) of the column for which this function will return information. In a CDB, the column name is specified as follows: `container_name:schema.table.column`

Return Values

Table 116-11 Return Values for MINE_VALUE Function

Return	Description
`NULL`	The column is not contained within the self-describing record, or the column value is `NULL`. To distinguish between the two different null possibilities, use the `DBMS_LOGMNR.COLUMN_PRESENT` function.
`NON-NULL`	The column is contained within the self-describing record; the value is returned in string format.

Exceptions

Table 116-12 MINE_VALUE Function Exceptions

Exception	Description
`ORA-01323`	Invalid state. Currently, a LogMiner dictionary is not associated with the LogMiner session. You must specify a LogMiner dictionary for the LogMiner session.
`ORA-00904`	Invalid identifier. The value specified for the `column_name` parameter was not a fully qualified column name.

Usage Notes

- To use the MINE_VALUE function, you must have successfully started LogMiner.
- The MINE_VALUE function must be invoked in the context of a select operation from the V$LOGMNR_CONTENTS view.
- The MINE_VALUE function does not support LONG, LOB, ADT, or COLLECTION datatypes.
- The value for the sql_redo_undo parameter depends on the operation performed and the data of interest:
  - If an update operation was performed and you want to know what the value was prior to the update operation, specify UNDO_VALUE.
  - If an update operation was performed and you want to know what the value is after the update operation, specify REDO_VALUE.
  - If an insert operation was performed, typically you would specify REDO_VALUE (because the value of a column prior to an insert operation will always be null).
  - If a delete operation was performed, typically you would specify UNDO_VALUE (because the value of a column after a delete operation will always be null).
- If the DBMS_LOGMNR.MINE_VALUE function is used to get an NCHAR value that includes characters not found in the database character set, then those characters are returned as the replacement character (for example, an inverted question mark) of the database character set.

116.6.5 REMOVE_LOGFILE Procedure

This procedure removes a redo log file from an existing list of redo log files for LogMiner to process.

In a CDB, the REMOVE_LOGFILE procedure must be called from the root database. You must have the LOGMINING administrative privilege to use this procedure.

Syntax

```
DBMS_LOGMNR.REMOVE_LOGFILE (
   LogFileName     IN VARCHAR2);
```

Parameters

Table 116-13 REMOVE_LOGFILE Procedure Parameters

Parameter	Description
LogFileName	Specifies the name of the redo log file to be removed from the list of redo log files to be analyzed during this session.

Exceptions

Table 116-14 REMOVE_LOGFILE Procedure Exception

Exception	Description
ORA-01290	Cannot remove unlisted log file

Usage Notes

- Before querying the V$LOGMNR_CONTENTS view, you must make a successful call to the DBMS_LOGMNR.START_LOGMNR procedure (within the current LogMiner session).
- You can use this procedure to remove a redo log file from the list of redo log files for LogMiner to process if you know that redo log file does not contain any data of interest.
- Multiple redo log files can be removed by calling this procedure repeatedly.
- The redo log files do not need to be removed in any particular order.
- To start a new list of redo log files for analysis, use the END_LOGMNR procedure to end the current LogMiner session, and then build a new list using the ADD_LOGFILE procedure.
- Even if you remove all redo log files from the list, any subsequent calls you make to the ADD_LOGFILE procedure must match the database ID and RESETLOGS SCN of the removed redo log files. Therefore, to analyze the redo log files from a different database (or a database incarnation with a different database RESETLOGS SCN) than that with which the current list of redo log files is associated, use the END_LOGMNR procedure to end the current LogMiner session, and then build a new list using the ADD_LOGFILE procedure.

116.6.6 START_LOGMNR Procedure

This procedure starts LogMiner by loading the dictionary that LogMiner will use to translate internal schema object identifiers to names.

In a CDB, the START_LOGMNR procedure must be called from the root database. You must have the LOGMINING administrative privilege to use this procedure.

> **Note:**
>
> The ability to create flat file dictionary dumps of pluggable databases (PDBs) is desupported in Oracle Database 21c.
>
> In previous releases, using a flat file dictionary was one means of mining the redo logs for the changes associated with a specific PDB whose data dictionary was contained within the flat file. This feature is now desupported. Starting with Oracle Database 21c, Oracle recommends that you call DBMS_LOGMNR.START_LOGMNR, and supply the system change number (SCN) or time range that you want to mine. The SCN or time range options of START_LOGMNR are enhanced to support mining of individual PDBs.

Syntax

```
DBMS_LOGMNR.START_LOGMNR (
   startScn           IN NUMBER default 0,
   endScn             IN NUMBER default 0,
   startTime          IN DATE default '01-jan-1988',
   endTime            IN DATE default '31-dec-2110',
   DictFileName       IN VARCHAR2 default '',
   Options            IN BINARY_INTEGER default 0 );
```

Parameters

Table 116-15 START_LOGMNR Procedure Parameters

Parameter	Description
startScn	Directs LogMiner to return only redo records with an SCN greater than or equal to the startScn specified. This fails if there is no redo log file containing the specified startScn value. (You can query the FILENAME, LOW_SCN, and NEXT_SCN columns in the V$LOGMNR_LOGS view for each redo log file to determine the range of SCN values contained in each redo log file.)
endScn	Directs LogMiner to return only redo records with an SCN less than or equal to the endScn specified. If you specify an endScn value that is beyond the value in any redo log file, then LogMiner uses the greatest endScn value in the redo log file that contains the most recent changes. (You can query the FILENAME, LOW_SCN, and NEXT_SCN columns in the V$LOGMNR_LOGS view for each redo log file to determine the range of SCN values contained in each redo log file.)
startTime	Directs LogMiner to return only redo records with a timestamp greater than or equal to the startTime specified. This fails if there is no redo log file containing the specified startTime value. (You can query the FILENAME, LOW_TIME, and HIGH_TIME columns in the V$LOGMNR_LOGS view for each redo log file to determine the range of time covered in each redo log file.) This parameter is ignored if startScn is specified. See the Usage Notes for additional information.
endTime	Directs LogMiner to return only redo records with a timestamp less than or equal to the endTime specified. If you specify an endTime value that is beyond the value in any redo log file, then LogMiner will use the greatest endTime in the redo log file that contains the most recent changes. You can query the FILENAME, LOW_TIME, and HIGH_TIME columns in the V$LOGMNR_LOGS view for each redo log file to determine the range of time covered in each redo log file.) This parameter is ignored if endScn is specified. See the Usage Notes for additional information.
DictFileName	Specifies the flat file that contains the LogMiner dictionary. It is used to reconstruct SQL_REDO and SQL_UNDO columns in V$LOGMNR_CONTENTS, as well as to fully translate SEG_NAME, SEG_OWNER, SEG_TYPE_NAME, TABLE_NAME, and TABLE_SPACE columns. The fully qualified path name for the LogMiner dictionary file must be specified. (This file must have been created previously through the DBMS_LOGMNR_D.BUILD procedure.) You need to specify this parameter only if neither DICT_FROM_REDO_LOGS nor DICT_FROM_ONLINE_CATALOG is specified.
options	See Table 116-2.

Exceptions

Table 116-16 START_LOGMNR Procedure Exceptions

Exception	Description
ORA-01280	Internal error encountered.
ORA-01281	startScn or endScn parameter value is not a valid SCN, or endScn is less than startScn.

Table 116-16 (Cont.) START_LOGMNR Procedure Exceptions

Exception	Description
ORA-01282	value for the startTime parameter was greater than the value specified for the endTime parameter, or there was no redo log file that was compatible with the date range specified with the startTime and endTime parameters.
ORA-01283	Options parameter specified is invalid.
ORA-01284	LogMiner dictionary file specified in the DictFileName parameter has a full path length greater than 256 characters, or the file cannot be opened.
ORA-01285	Error reading specified file.
ORA-01291	Redo log files that are needed to satisfy the user's requested SCN or time range are missing.
ORA-01292	No log file has been specified for the current LogMiner session.
ORA-01293	Mounted database required for specified LogMiner options.
ORA-01294	Error occurred while processing information in the specified dictionary file, possible corruption.
ORA-01295	Specified LogMiner dictionary does not correspond to the database that produced the log files being analyzed.
ORA-01296	Character set mismatch between specified LogMiner dictionary and log files.
ORA-01297	Redo version mismatch between LogMiner dictionary and log files.
ORA-01299	Specified LogMiner dictionary corresponds to a different database incarnation.
ORA-01300	Writable database required for specified LogMiner options.

Usage Notes

> **Note:**
>
> The continuous_mine option for the dbms_logmnr.start_logmnr package is desupported in Oracle Database 19c (19.1), and is no longer available.

- The DBMS_LOGMNR.ADD_LOGFILE call from a PDB connection is not supported for adhoc users and returns error 65040.
- Users can specify the SCN range or the time range. The required logfiles will be added programmatically.
- Dumping Flat File dictionary and mining using Flat File dictionary are not supported from a PDB connection.
- LogMiner can use a dictionary that you previously extracted to the redo log files or to a flat file, or you can specify that LogMiner use the online catalog if LogMiner is mining data from the source system.
- After executing the START_LOGMNR procedure, you can query the following views:
  - V$LOGMNR_CONTENTS - contains history of information in redo log files

  - V$LOGMNR_DICTIONARY - contains current information about the LogMiner dictionary file extracted to a flat file
  - V$LOGMNR_PARAMETERS - contains information about the LogMiner session

  (You can query the V$LOGMNR_LOGS view after a redo log file list has been added to the list of files that LogMiner is to mine.)

- Parameters and options are not persistent across calls to DBMS_LOGMNR.START_LOGMNR. You must specify all desired parameters and options (including SCN and time ranges) each time you call DBMS_LOGMNR.START_LOGMNR
- Be aware that specifying redo log files using a timestamp is not precise.
- Keep the following in mind regarding starting and ending times or SCN ranges:
  - If you specify neither a startTime nor a startScn parameter, LogMiner will set the startScn parameter to use the lowest SCN value from the redo log file that contains the oldest changes.
  - If you specify both time and SCN values, LogMiner uses the SCN value or values and ignores the time values.
  - If you specify starting and ending time or SCN values and they are found in the LogMiner redo log file list, then LogMiner mines the logs indicated by those values.
  - If you specify starting and ending times or SCN values that are not in the LogMiner redo log file list, and you specify DBMS_LOGMNR.START_LOGMNR, and you specify:
    * 0 for the startTime or startScn value, then the lowest SCN in the LogMiner redo log file list will be used as the startScn
    * A nonzero number for the startTime or startScn value, then an error is returned
    * 0 or a nonzero number for the endTime or endScn value, then the highest SCN in the LogMiner redo log file list will be used as the endScn
  - If you specify starting and ending times or SCN values and they are not found in the LogMiner redo log file list, and you specify DBMS_LOGMNR.START_LOGMNR, and you specify:
    * 0 for the startTime or startScn value, then an error is returned.
    * A startTime or startScn value that is greater than any value in the database's archived redo log files, then LogMiner starts mining in the online redo log file. LogMiner will continue to process the online redo log file until it finds a change at, or beyond, the requested starting point before it returns rows from the V$LOGMNR_CONTENTS view.
    * An endTime or endScn parameter value that indicates a time or SCN in the future, then LogMiner includes the online redo log files when it mines. When you query the V$LOGMNR_CONTENTS view, rows will be returned from this view as changes are made to the database, and will not stop until LogMiner sees a change beyond the requested ending point.
    * 0 for the endTime or endScn parameter value, then LogMiner includes the online redo log files when it mines. When you query the V$LOGMNR_CONTENTS view, rows will be returned from this view as changes are made to the database, and will not stop until you enter CTL+C or you terminate the PL/SQL cursor.

117
DBMS_LOGSTDBY

The DBMS_LOGSTDBY package provides subprograms for configuring and managing the logical standby database environment.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_LOGSTDBY Subprograms

> **See Also:**
>
> *Oracle Data Guard Concepts and Administration*

117.1 DBMS_LOGSTDBY Overview

The DBMS_LOGSTDBY package helps you manage the SQL Apply (logical standby database) environment.

The subprograms in the DBMS_LOGSTDBY package help you to accomplish the following main objectives:

- Manage configuration parameters used by SQL Apply.

  For example, controlling how transactions are applied on the logical standby database, how much shared pool is used, and how many processes are used by SQL Apply to mine and apply the changes.
- Ensure an appropriate level of supplemental logging is enabled, and a LogMiner dictionary is built correctly for logical standby database creation.
- Provide a way to skip the application of changes to selected tables or entire schemas in the logical standby database, and specify ways to handle exceptions encountered by SQL Apply.
- Allow controlled access to tables in the logical standby database that may require maintenance.

117.2 DBMS_LOGSTDBY Security Model

You must have the DBA role to use the DBMS_LOGSTDBY package.

A prototype role, LOGSTDBY_ADMINISTRATOR, is created by default with RESOURCE and EXECUTE privileges on DBMS_LOGSTDBY. If you choose to use this role, consider granting ALTER DATABASE and ALTER SESSION privileges to the role so that the grantee can start and stop SQL Apply and can enable and disable the database guard.

The procedures associated with skipping transactions (SKIP and UNSKIP, SKIP_ERROR and UNSKIP_ERROR, and SKIP_TRANSACTION and UNSKIP_TRANSACTION) all require DBA privileges to execute because their scope may contain wildcard schemas. Oracle recommends that where SKIP procedures are specified, these be owned by a secure account with appropriate privileges on the schemas they act on (for example, SYS).

117.3 DBMS_LOGSTDBY Constants

The DBMS_LOGSTDBY package defines several enumerated constants for specifying parameter values. Enumerated constants must be prefixed with the package name, for example, DBMS_LOGSTDBY.SKIP_ACTION_SKIP.

The following table describes the constants for the proc_name parameter in the DBMS_LOGSTDBY.SKIP procedure.

Table 117-1 Constants for SKIP Options Flag

Constant	Description
MAX_EVENTS	Maximum number of events to log in the DBA_LOGSTDBY_EVENTS view. See the DBMS_LOGSTDBY APPLY_SET Procedure.
SKIP_ACTION_APPLY	Used inside the user-defined procedure registered with DBMS_LOGSTDBY.SKIP. Use this constant when setting the value of the SKIP_ACTION parameter with DBMS_LOGSTDBY_CONTEXT if you want SQL Apply to apply the DDL or PL/SQL statement.
SKIP_ACTION_ERROR	Used inside the user-defined procedure registered with DBMS_LOGSTDBY.SKIP. Use this constant when setting the value of the SKIP_ACTION parameter with DBMS_LOGSTDBY_CONTEXT if you want SQL Apply to error out.
SKIP_ACTION_REPLACE	Used inside the user-defined procedure registered with DBMS_LOGSTDBY.SKIP. Use this constant when setting the value of the SKIP_ACTION parameter with DBMS_LOGSTDBY_CONTEXT if you want SQL Apply to apply the replacement DDL. (This constant should not be used while handling an Oracle Supplied PL/SQL procedure call).
SKIP_ACTION_SKIP	Used inside the user-defined procedure registered with DBMS_LOGSTDBY.SKIP. Use this constant when setting the value of the SKIP_ACTION parameter with DBMS_LOGSTDBY_CONTEXT if you want SQL Apply to skip the associated DDL or Oracle supplied PL/SQL procedure call.

117.4 Summary of DBMS_LOGSTDBY Subprograms

This table describes each subprogram of the DBMS_LOGSTDBY procedure, including a reference to the section where each procedure is described in more detail.

In a multitenant container database (CDB), some subprograms must be called from the root. There may be other differences as well. See the individual subprogram descriptions for details.

Table 117-2 DBMS_LOGSTDBY Package Subprograms

Subprogram	Description
APPLY_SET Procedure	Sets the values of various parameters that configure and maintain SQL Apply.

Table 117-2 (Cont.) DBMS_LOGSTDBY Package Subprograms

Subprogram	Description
APPLY_UNSET Procedure	Restores the default values of various parameters that configure and maintain SQL Apply.
BUILD Procedure	Ensures supplemental logging is enabled properly and builds the LogMiner dictionary.
INSTANTIATE_TABLE Procedure	Creates and populates a table in the standby database from a corresponding table in the primary database.
IS_APPLY_SERVER Function	This function returns `TRUE` if it is executed from PL/SQL in the context of a logical standby apply server process. This function is used in conjunction with triggers that have the `fire_once` parameter in the `DBMS_DDL.SET_TRIGGER_FIRING_PROPERTY` subprogram set to `FALSE` (the default is `TRUE`). Such triggers are executed when the relevant target is updated by an apply process. This function can be used within the body of the trigger to ensure that the trigger takes different (or no) actions on the primary or on the standby.
MAP_PRIMARY_SCN Function	Maps an SCN relevant to the primary database to a corresponding SCN at the logical standby database. The mapped SCN is conservative in nature, and can thus be used to flash back the logical standby database to compensate for a flashback database operation performed at the primary database.
PREPARE_FOR_NEW_PRIMARY Procedure	Used after a failover, this procedure ensures a local logical standby database that was not involved in the failover has not processed more redo than the new primary database and reports the set of archive redo log files that must be replaced to ensure consistency
PURGE_SESSION Procedure	Identifies the archived redo log files that have been applied to the logical standby database and are no longer needed by SQL Apply
REBUILD Procedure	Records relevant metadata (including the LogMiner dictionary) in the redo stream in case a database that has recently changed its role to a primary database following a failover operation fails to do so during the failover process
SET_TABLESPACE Procedure	Moves metadata tables required by SQL Apply to the user-specified tablespace. By default, the metadata tables are created in the `SYSAUX` tablespace.
SKIP Procedure	Specifies rules that control database operations that should not be applied to the logical standby database
SKIP_ERROR Procedure	Specifies rules regarding what action to take upon encountering errors.
SKIP_TRANSACTION Procedure	Specifies transactions that should not be applied on the logical standby database. Be careful in using this procedure, because not applying specific transactions may cause data corruption at the logical standby database.
UNSKIP Procedure	Deletes rules specified by the `SKIP` procedure.
UNSKIP_ERROR Procedure	Deletes rules specified by the `SKIP_ERROR` procedure.
UNSKIP_TRANSACTION Procedure	Deletes rules specified by the `SKIP_TRANSACTION` procedure.

117.4.1 APPLY_SET Procedure

Use this procedure to set values of parameters that configure and manage SQL Apply in a logical standby database environment. All parameters, except for PRESERVE_COMMIT_ORDER, can be changed without having to stop SQL Apply.

In a CDB, the APPLY_SET procedure must be called from the root database.

Syntax

```
DBMS_LOGSTDBY.APPLY_SET (
     inname             IN VARCHAR,
     value              IN VARCHAR);
```

Parameters

Table 117-3 APPLY_SET Procedure Parameters

Parameter	Description
APPLY_SERVERS	Controls the number of APPLIER processes used to apply changes. The maximum number allowed is 1024, provided the MAX_SERVERS parameter is set to accommodate this.
EVENT_LOG_DEST	Controls where SQL Apply records the occurrence of an interesting event. It takes the following values: • DEST_ALL - All events will be recorded in the DBA_LOGSTDBY_EVENTS view and in the alert log. • DEST_EVENTS_TABLE - All events that contain information about user data will be recorded only in the DBA_LOGSTDBY_EVENTS view. This is the default value. For example, if SQL Apply receives an ORA-1403 error, the whole event is recorded in the DBA_LOGSTDBY_EVENTS view. Whereas, the alert log records only that SQL Apply stopped because of ORA-1403. No information regarding the user table or offending statement is logged in the alert log. However, if you stop the SQL Apply engine, it gets recorded in both the DBA_LOGSTDBY_EVENTS view and in the alert log. Note that this parameter affects the behavior of the following parameters: RECORD_APPLIED_DDL, RECORD_SKIP_DDL, RECORD_SKIP_ERRORS, and RECORD_UNSUPPORTED_OPERATIONS. For example, if RECORD_APPLIED_DDL is set to TRUE, but EVENT_LOG_DEST is set to DEST_EVENTS_TABLE, then the applied DDL string will only be recorded in the DBA_LOGSTDBY_EVENTS view.
LOG_AUTO_DEL_RETENTION_TARGET	This parameter setting is only meaningful if LOG_AUTO_DELETE has been set to TRUE. The value you supply for this parameter controls how long (in minutes) a remote archived log that is received from the primary database will be retained at the logical standby database once all redo records contained in the log have been applied at the logical standby database. The default value is 1440 minutes.

Table 117-3 (Cont.) APPLY_SET Procedure Parameters

Parameter	Description
LOG_AUTO_DELETE	Automatically deletes foreign archived redo log files as soon as they have been applied on the logical standby database. By default, a foreign archived redo log file is not deleted until 24 hours (the default value of LOG_AUTO_DEL_RETENTION_TARGET parameter) after it has been applied at the logical standby database. Set to TRUE to enable automatic deletion of archived redo log files. Set to FALSE to disable automatic deletion. The default value is TRUE.
MAX_EVENTS_RECORDED	Number of recent events that will be visible through the DBA_LOGSTDBY_EVENTS view. To record all events encountered by SQL Apply, use the DBMS_LOGSTDBY.MAX_EVENTS constant as the number value. The default value is 10,000.
MAX_SERVERS	Number of processes that SQL Apply uses to read and apply redo. The default value is 9. The maximum number allowed is 2048.
MAX_SGA	Number of megabytes from shared pool in System Global Area (SGA) that SQL Apply will use. The default value is 30 megabytes or one quarter of the value set for SHARED_POOL_SIZE, whichever is lower. The maximum size allowed is 4095 megabytes.
PREPARE_SERVERS	Controls the number of PREPARER processes used to prepare changes. The maximum number allowed is 1024, provided the MAX_SERVERS parameter is set to accommodate this.
PRESERVE_COMMIT_ORDER	TRUE: Transactions are applied to the logical standby database in the exact order in which they were committed on the primary database. This is the default parameter setting. FALSE: Transactions containing non-overlapping sets of rows may be committed in a different order than they were committed on the primary database. Regardless of the level chosen, modifications done to the same row are always applied in the same order as they happened on the primary database. See the Usage Notes for details and recommendations. You cannot modify this parameter while SQL Apply is running.
RECORD_APPLIED_DDL	Controls whether DDL statements that have been applied to the logical standby database are recorded in the location specified by the EVENT_LOG_DEST parameter. Specify one of the following values: TRUE: Indicates that DDL statements applied to the logical standby database are recorded in the DBA_LOGSTDBY_EVENTS table and the alert log. FALSE: Indicates that applied DDL statements are not recorded. This is the default parameter setting.
RECORD_SKIP_DDL	Controls whether skipped DDL statements are recorded in the location specified by the EVENT_LOG_DEST parameter. Specify one of the following values: TRUE: Skipped DDL statements are recorded in the DBA_LOGSTDBY_EVENTS table and the alert log. This is the default parameter setting. FALSE: Skipped DDL statements are not recorded in the DBA_LOGSTDBY_EVENTS table and the alert log.

Table 117-3 (Cont.) APPLY_SET Procedure Parameters

Parameter	Description
RECORD_SKIP_ERRORS	Controls whether skipped errors (as described by the SKIP_ERROR procedure) are recorded in the location specified by the EVENT_LOG_DEST parameter. Specify one of the following values: TRUE: Skipped errors are recorded in the DBA_LOGSTDBY_EVENTS table and the alert log. This is the default parameter setting. FALSE: Skipped errors are not recorded in the DBA_LOGSTDBY_EVENTS table and the alert log.
RECORD_UNSUPPORTED_OPERATIONS	Captures information about transactions running on the primary database that will not be supported by a logical standby database. This procedure records its information as events in the DBA_LOGSTDBY_EVENTS table. Specify one of the following values: TRUE: The information is captured and recorded as events in the DBA_LOGSTDBY_EVENTS table. FALSE: The information is not captured. This is the default.

If a parameter is changed while SQL Apply is running, the change will take effect at some point in the future. In such a case, an informational row is inserted into the DBA_LOGSTDBY_EVENTS view at the time the parameter change takes effect.

Additionally, if you are modifying a parameter while SQL Apply is running on an Oracle RAC configuration, you must be connected to the same instance where SQL Apply is running.

Exceptions

Table 117-4 APPLY_SET Procedure Exceptions

Exception	Description
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested
ORA-16236	Logical Standby metadata operation in progress

Usage Notes

- Use the APPLY_UNSET procedure to restore the default settings of a parameter.
- See *Oracle Data Guard Concepts and Administration* for help with tuning SQL Apply and for information about setting appropriate values for different parameters.

Examples

To record DDLs in the DBA_LOGSTDBY_EVENTS view and in the alert log, issue the following statement:

```
SQL> EXECUTE DBMS_LOGSTDBY.APPLY_SET('RECORD_APPLIED_DDL', TRUE);
```

117.4.2 APPLY_UNSET Procedure

Use the APPLY_UNSET procedure to restore the default values of the parameters that you changed with the APPLY_SET procedure.

In a CDB, the APPLY_UNSET procedure must be called from the root database.

Syntax

```
DBMS_LOGSTDBY.APPLY_UNSET (
     inname          IN VARCHAR);
```

Parameters

The parameter information for the APPLY_UNSET procedure is the same as that described for the APPLY_SET procedure. See Table 117-3 for complete parameter information.

Exceptions

Table 117-5 APPLY_UNSET Procedure Exceptions

Exception	Description
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested
ORA-16236	Logical Standby metadata operation in progress

Usage Notes

- Use the APPLY_SET procedure to specify a nondefault value for a parameter.

Examples

If you previously specified that applied DDLs show up in the DBA_LOGSTDBY_EVENTS view and the alert log, you can restore the default behavior of SQL Apply regarding applied DDL statements with the following statement:

```
SQL> EXECUTE DBMS_LOGSTDBY.APPLY_UNSET('RECORD_APPLIED_DDL');
```

117.4.3 BUILD Procedure

Use this procedure on the primary database to record relevant metadata (LogMiner dictionary) information in the redo log, which will subsequently be used by SQL Apply. This procedure will enable database-wide primary- and unique-key supplemental logging, if necessary.

In a CDB, the BUILD procedure must be called from the root database on the primary. Additionally, you cannot add or remove PDBs from a CDB while this procedure is executing.

> **Note:**
>
> In databases created using Oracle Database 11*g* release 2 (11.2) or later, supplemental logging information is automatically propagated to any existing physical standby databases. However, for databases in earlier releases, or if the database was created using an earlier release and then upgraded to 11.2, you must check whether supplemental logging is enabled at the physical standby(s) if it is also enabled at the primary database. If it is not enabled at the physical standby(s), then before performing a switchover or failover, you must enable supplemental logging on all existing physical standby databases. To do so, issue the following SQL command on each physical standby:
>
> ```
> SQL>  ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (PRIMARY KEY, UNIQUE INDEX)
> COLUMNS;
> ```
>
> If you do not do this, then any logical standby that is also in the same Data Guard configuration will be unusable if a switchover or failover is performed to one of the physical standby databases. If a switchover or failover has already occurred and supplemental logging was not enabled, then you must recreate all logical standby databases.

Syntax

```
DBMS_LOGSTDBY.BUILD;
```

Usage Notes

- Supplemental log information includes extra information in the redo logs that uniquely identifies a modified row in the logical standby database, and also includes information that helps efficient application of changes to the logical standby database.
- LogMiner dictionary information allows SQL Apply to interpret data in the redo logs.
- `DBMS_LOGSTDBY.BUILD` should be run only once for each logical standby database you want to create. You do not need to use DBMS_LOGSTDBY.BUILD for each Oracle RAC instance.
- `DBMS_LOGSTDBY.BUILD` waits for all transactions (including distributed transactions) that are active at the time of the procedure invocation to complete before returning. See *Oracle Database Administrator's Guide* for information about how to handle in-doubt transactions.

Examples

To build the LogMiner dictionary in the redo stream of the primary database and to record additional information so that a logical standby database can be instantiated, issue the following SQL statement at the primary database

```
SQL> EXECUTE DBMS_LOGSTDBY.BUILD;
```

117.4.4 INSTANTIATE_TABLE Procedure

This procedure creates and populates a table in the standby database from a corresponding table in the primary database.

The table requires the name of the database link (`dblink`) as an input parameter. If the table already exists in the logical standby database, it will be dropped and re-created based on the

table definition at the primary database. This procedure only brings over the data associated with the table, and not the associated indexes and constraints.

Use the INSTANTIATE_TABLE procedure to:

- Add a table to a standby database.
- Re-create a table in a standby database.

In a CDB, the INSTANTIATE_TABLE procedure must be called from within the container in which the table to be instantiated resides. Additionally, the database link that is provided to the primary database must point to the corresponding container on the primary.

Syntax

```
DBMS_LOGSTDBY.INSTANTIATE_TABLE (
     schema_name         IN VARCHAR2,
     table_name          IN VARCHAR2,
     dblink              IN VARCHAR2);
```

Parameters

Table 117-6 INSTANTIATE_TABLE Procedure Parameters

Parameter	Description
schema_name	Name of the schema
table_name	Name of the table to be created or re-created in the standby database
dblink	Name of the database link account that has privileges to read and lock the table in the primary database, as well as the SELECT_CATALOG_ROLE on the primary database

Exceptions

Table 117-7 INSTANTIATE_TABLE Procedure Exceptions

Exception	Description
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16236	Logical Standby metadata operation in progress
ORA-16276	Specified database link does not correspond to primary database
ORA-16277	Specified table is not supported by logical standby database
ORA-16278	Specified table has a multi-object skip rule defined

Usage Notes

- Use this procedure to create and populate a table in a way that keeps the data on the standby database transactionally consistent with the primary database.
- This table will not be synchronized with the rest of the tables being maintained by SQL Apply and SQL Apply will not start to maintain it until SQL Apply encounters redo that occurred after the table was instantiated from the primary. The SCN at which the table was instantiated from the primary database is available in the DBA_LOGSTDBY_EVENTS view.
- The specified table must be a table that is supported by logical standby (that is, it does not appear in the DBA_LOGSTDBY_UNSUPPORTED_TABLES view on the primary database).

- If there are any skip rules that specifically name this table (without any wildcards), those skip rules will be dropped as part of `INSTANTIATE_TABLE`, so that the table will be properly maintained by SQL Apply in the future. If there are skip rules that indirectly reference this table (match a skip rule with a wildcard in the `schema_name` or `table_name`, and have a TABLE, DML, or SCHEMA_DDL statement type), `INSTANTIATE_TABLE` will fail with an ORA-16278 error. Any multi-object skip rules that pertain to the table must be dropped or changed before re-attempting the `INSTANTIATE_TABLE` call.

Examples

```
SQL> EXECUTE DBMS_LOGSTDBY.INSTANTIATE_TABLE (-
     SCHEMA_NAME => 'HR', TABLE_NAME => 'EMPLOYEES', -
     DBLINK => 'INSTANTIATE_TBL_LINK');
```

117.4.5 IS_APPLY_SERVER Function

This function returns `TRUE` if it is executed from PL/SQL in the context of a logical standby apply server process.

This function is used in conjunction with triggers that have the `fire_once` parameter in the `DBMS_DDL.SET_TRIGGER_FIRING_PROPERTY` subprogram set to `FALSE` (the default is `TRUE`). Such triggers are executed when the relevant target is updated by an apply process. This function can be used within the body of the trigger to ensure that the trigger takes different (or no) actions on the primary or on the standby.

Syntax

```
DBMS_LOGSTDBY.IS_APPLY_SERVER
RETURN BOOLEAN;
```

Parameters

None

117.4.6 MAP_PRIMARY_SCN Function

This function returns an SCN on the standby that predates the supplied SCN from the primary database by at least 5 minutes.

It can be used to determine a safe SCN to use in a compensating flashback database operation at the logical standby database, following a flashback database operation or a point-in-time recovery operation at the primary database.

Syntax

```
DBMS_LOGSTDBY.MAP_PRIMARY_SCN(primary_scn NUMBER) RETURN NUMBER;
```

Exceptions

Table 117-8 MAP_PRIMARY_SCN Function Exceptions

Exception	Description
ORA-20001	Primary SCN is before mapped range
ORA-20002	SCN mapping requires `PRESERVE_COMMIT_ORDER` to be `TRUE`

Usage Notes

Use this function to get a conservative SCN at the logical standby database that corresponds to an SCN at the primary database. This function is useful in the context of doing compensating flashback database operations at the logical standby following a flashback database or a point-in-time recovery operation done at the primary database.

117.4.7 PREPARE_FOR_NEW_PRIMARY Procedure

The `PREPARE_FOR_NEW_PRIMARY` procedure must be invoked at a logical standby database following a failover, if that standby database was not the target of the failover operation.

Such a standby database must process the exact same set of redo logs processed at the new primary database. This routine ensures that the local logical standby database has not processed more redo than the new primary database and reports the set of archive logs that must be replaced to ensure consistency. The set of replacement logs will be reported in the alert.log. These logs must be copied to the logical standby and registered using the `ALTER DATABASE REGISTER LOGICAL LOGFILE` statement.

In a CDB, the `PREPARE_FOR_NEW_PRIMARY` procedure must be called from the root database.

Syntax

```
DBMS_LOGSTDBY.PREPARE_FOR_NEW_PRIMARY (
           FORMER_STANDBY_TYPE         IN VARCHAR2,
           DBLINK                      IN VARCHAR2);
```

Parameters

Table 117-9 PREPARE_FOR_NEW_PRIMARY Procedure Parameters

Parameter	Description
`FORMER_STANDBY_TYPE`	The type of standby database that was the target of the failover operation to become the new primary database. Valid values are '`PHYSICAL`' if the new primary was formerly a physical standby, and '`LOGICAL`' if the new primary database was formerly a logical standby database.
`DBLINK`	The name of a database link to the new primary database

Exceptions

Table 117-10 PREPARE_FOR_NEW_PRIMARY Procedure Exceptions

Exception	Description
`ORA-16104`	Invalid Logical Standby option.
`ORA-16109`	Failed to apply log data from previous primary.

Usage Notes

- This routine is intended only for logical standby systems.This routine will fail if the new primary database was formerly a logical standby database and the LogMiner dictionary build has not completed successfully.Log files displayed in the alert log will be referred to as *terminal logs*. Users should keep in mind that file paths are relative to the new primary

database and may not resolve locally.Upon manual registration of the terminal logs, users should complete the process by calling either `START LOGICAL STANDBY APPLY` if the new primary database was formerly a physical standby database or `START LOGICAL STANDBY APPLY NEW PRIMARY` if the new primary database was formerly a logical standby database.See the alert log for more details regarding the reasons for any exception.

Examples

```
SQL> EXECUTE DBMS_LOGSTDBY.PREPARE_FOR_NEW_PRIMARY (  -
                FORMER_STANDBY_TYPE => 'LOGICAL',    -
                DBLINK => 'dblink_to_newprimary');
```

117.4.8 PURGE_SESSION Procedure

`PURGE_SESSION` identifies all archived redo log files that have been applied to the logical standby database and are no longer needed by SQL Apply.

Once identified, you can issue operating system commands to delete some or all of the unnecessary archived redo log files.

In a CDB, the `PURGE_SESSION` procedure must be called from the root database.

Syntax

```
DBMS_LOGSTDBY.PURGE_SESSION;
```

Exceptions

Table 117-11 PURGE_SESSION Procedure Exceptions

Exception	Description
ORA-01309	Invalid session

Usage Notes

- This procedure does not delete the archived redo log files. You must issue operating system commands to delete unneeded files.
- This procedure updates the `DBA_LOGMNR_PURGED_LOG` view that displays the archived redo log files that have been applied to the logical standby database.
- In Oracle Database 10*g* Release 2, metadata related to the archived redo log files (and the actual archived redo log files) are purged automatically based on the default setting of the `LOG_AUTO_DELETE` parameter described in the `DBMS_LOGSTDBY.APPLY_SET` procedure described.

Example

To identify and remove unnecessary files:

1. Enter the following statement on the logical standby database:

   ```
   SQL> EXECUTE DBMS_LOGSTDBY.PURGE_SESSION;
   ```

2. Query the `DBA_LOGMNR_PURGED_LOG` view to list the archived redo log files that can be removed:

   ```
   SQL> SELECT * FROM DBA_LOGMNR_PURGED_LOG;
   ```

```
FILE_NAME
   ----------------------------------
   /boston/arc_dest/arc_1_40_509538672.log
   /boston/arc_dest/arc_1_41_509538672.log
   /boston/arc_dest/arc_1_42_509538672.log
   /boston/arc_dest/arc_1_43_509538672.log
   /boston/arc_dest/arc_1_44_509538672.log
   /boston/arc_dest/arc_1_45_509538672.log
   /boston/arc_dest/arc_1_46_509538672.log
   /boston/arc_dest/arc_1_47_509538672.log
```

3. Use operating system-specific commands to delete archived redo log files from the file system.

117.4.9 REBUILD Procedure

This procedure is used if a database that has recently changed its role to a primary database following a failover operation fails to record relevant metadata (including the LogMiner dictionary) in the redo stream required for other logical standby databases.

In a CDB, the `REBUILD` procedure must be called from the root database.

Syntax

```
DBMS_LOGSTDBY.REBUILD;
```

Usage Notes

- LogMiner dictionary information is logged in the redo log files.The standby redo log files (if present) are archived.

Examples

```
SQL> EXECUTE DBMS_LOGSTDBY.REBUILD;
```

117.4.10 SET_TABLESPACE Procedure

This procedure moves metadata tables required by SQL Apply to the user-specified tablespace.

By default, the metadata tables are created in the `SYSAUX` tablespace. SQL Apply cannot be running when you invoke this procedure.

In a CDB, the `SET_TABLESPACE` procedure must be called from the root database.

Syntax

```
DBMS_LOGSTDBY.SET_TABLESPACE(
          NEW_TABLESPACE IN VARCHAR2)
```

Parameters

Table 117-12 SET_TABLE SPACE Procedure Parameters

Parameter	Description
NEW_TABLESPACE	Name of the new tablespace where metadata tables will reside.

Exceptions

Table 117-13 SET_TABLESPACE Procedure Exceptions

Exception	Description
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16236	Logical Standby metadata operation in progress

Examples

To move metadata tables to a new tablespace named LOGSTDBY_TBS, issue the following statement:

```
SQL> EXECUTE DBMS_LOGSTDBY.SET_TABLESPACE (new_tablespace => 'LOGSTDBY_TBS');
```

117.4.11 SKIP Procedure

The SKIP procedure can be used to define rules that will be used by SQL Apply to skip the application of certain changes to the logical standby database.

For example, the SKIP procedure can be used to skip changes to a subset of tables in the logical standby database. It can also be used to specify DDL statements that should not be applied at the logical standby database or should be modified before they are applied in the logical standby database. One reason why a DDL statement may need to be modified is to accommodate a different directory structure on the logical standby database.

> **Note:**
>
> For information about skipping containers, see "Skipping Containers."

Syntax

```
DBMS_LOGSTDBY.SKIP (
     stmt                      IN VARCHAR2,
     schema_name               IN VARCHAR2 DEFAULT NULL,
     object_name               IN VARCHAR2 DEFAULT NULL,
     proc_name                 IN VARCHAR2 DEFAULT NULL,
     use_like                  IN BOOLEAN DEFAULT TRUE,
     esc                       IN CHAR1 DEFAULT NULL);
```

Parameters

Table 117-14 SKIP Procedure Parameters

Parameter	Description
stmt	Either a keyword that identifies a set of SQL statements or a specific SQL statement. The use of keywords simplifies configuration since keywords, generally defined by the database object, identify all SQL statements that operate on the specified object. Table 117-15 shows a list of keywords and the equivalent SQL statements, either of which is a valid value for this parameter. The keyword PL/SQL is used for the execution of Oracle-supplied packages which are supported for replication.
schema_name	The name of one or more schemas (wildcards are permitted) associated with the SQL statements identified by the stmt parameter. If not applicable, this value must be set to NULL.
object_name	The name of one or more objects (wildcards are permitted) associated with the SQL statements identified by the stmt. If not applicable, this value must be set to NULL.
proc_name	Name of a stored procedure to call when SQL Apply determines that a particular statement matches the filter defined by the stmt, schema_name, and object_name parameters. Specify the procedure in the following format: 'schema.package.procedure' This procedure returns a value that directs SQL Apply to perform one of the following: execute the statement, skip the statement, or execute a replacement statement. The procedures to be invoked in the case of DDL or PL/SQL take no arguments. You can access the various information needed inside the procedure by accessing the context associated with the namespace,LSBY_APPLY_CONTEXT. For a full list of parameters that are accessible in the context of the skip procedure, see the DBMS_LOGSTBDY_CONTEXT package. The parameters of interest in the case of DDLs are: STATEMENT, STATEMENT_TYPE, SCHEMA, NAME, CURENT_SCHEMA, XIDUSN, XIDSLT, XIDSQN and SKIP_ACTION. The parameters of interest in the case of PL/SQL are: STATEMENT, PACKAGE_SCHMEMA, PACKAGE_NAME, PROCEDURE_NAME, CURRENT_SCHEMA, XIDUSN, XIDSLT, XIDSQN, EXIT_STATUS, and SKIP_ACTION.. **Note 1**: The DBMS_LOGSTDBY.SKIP_ACTION_REPLACE constant is not supported for PL/SQL. **Note 2**: SQL Apply calls the skip handler when the procedure's exit is processed. **Note 3**: The use_like parameter must be set to FALSE for PL/SQL since wildcarding PL/SQL is not supported.

Table 117-14 (Cont.) SKIP Procedure Parameters

Parameter	Description
proc_name (cont.)	A sample conditional skip rule on DBMS_RLS.DROP_POLICY is as follows:

```
create or replace procedure sec.mgr.skip_drop_policy
is
l_stmt  CLOB;
l_pkgown varchar2(30);
l_pkgnam varchar2(30);
l_procnm varchar2(30);
l_cur_schema varchar2(30);
l_xidusn   number;
l_xidslt   number;
l_xidsqn   number;
l_exit_status number;
l_skip_action number;
Begin
-- read all relevant info
dbms_logstdby_context.get_context(name => 'STATEMENT', value => l_stmt);
dbms_logstdby_context.get_context(name => 'PACKAGE_SCHEMA', value =>
l_pkgown);
dbms_logstdby_context.get_context(name => 'PACKAGE_NAME', value =>
l_pkgnam);
dbms_logstdby_context.get_context(name => 'PROCEDURE_NAME', value =>
l_procnm);
dbms_logstdby_context.get_context(name => 'CURRENT_SCHEMA', value =>
l_cur_schema);
dbms_logstdby_context.get_context(name => 'XIDUSN', value => l_xidusn);
dbms_logstdby_context.get_context(name => 'XIDSLT', value => l_xidslt);
dbms_logstdby_context.get_context(name => 'XIDSQN', value => l_xidsqn);
dbms_logstdby_context.get_context(name => 'EXIT_STATUS', value =>
l_ext_status);

if 0 == l_ext_status then
Insert Into sec_mgr.logit Values
      ('Success: '||l_pkgown||'.'||l_pkgnm||'.'||l_procnm|| ' by '||
l_current_user);

If l_current_user != 'TESTSCHEMA' Then
      l_skip_action := DBMS_LOGSTDBY.SKIP_ACTION_APPLY;
    Else
      l_skip_action := DBMS_LOGSTDBY.SKIP_ACTION_SKIP;
    End If;
  End If;

dbms_logstdby_context.set_context(name=>'SKIP_ACTION',
value => l_skip_action);

End skip_drop_policy;

EXECUTE DBMS_LOGSTDBY.SKIP( -
     stmt => 'PL/SQL', -
     schema_name => 'SYS', -
     object_name => 'DBMS_RLS.DROP_POLICY', -
     proc_name => 'SEC_MGR.SKIP_DROP_POLICY' -
     use_like=> FALSE);
```

Table 117-14 (Cont.) SKIP Procedure Parameters

Parameter	Description
use_like	Allows pattern matching to isolate the tables that you want to skip on the logical standby database. The use_like parameter matches a portion of one character value to another by searching the first value for the pattern specified by the second, and calculates strings using characters as defined by the input character set. This parameter follows the same rules for pattern matching described in *Oracle Database SQL Language Reference*.
esc	Identifies an escape character (such as the character "/") that you can use for pattern matching. If the escape character appears in the pattern before the character "%" or "_" then Oracle interprets this character literally in the pattern, rather than as a special pattern matching character. See *Oracle Database SQL Language Reference* for more information about pattern matching.

Usage Notes

- This procedure requires DBA privileges to execute.
- You cannot associate a stored procedure to be invoked in the context of a DML statement. For example, the following statement returns the ORA-16104: invalid Logical Standby option requested error:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP(-
     stmt => 'DML', -
     schema_name => 'HR', -
     object_name => 'EMPLOYEES', -
     proc_name => 'DML_HANDLER');
```

  Also, if an event matches multiple rules either because of the use of wildcards while specifying the rule or because of a specification of overlapping rules. For example, if you specify a rule for the SCHEMA_DDL event for the HR.EMPLOYEES table, and a rule for the ALTER TABLE event for the HR.EMPLOYEES table, only one of the matching procedures will be invoked (alphabetically, by procedure). In the following code example, consider the following rules:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP( -
     stmt => 'SCHEMA_DDL', -
     schema_name => 'HR', -
     object_name => 'EMPLOYEES', -
     proc_name => 'SCHEMA_DDL_HANDLER');
SQL> EXECUTE DBMS_LOGSTDBY.SKIP( -
     stmt => 'ALTER TABLE', -
     schema_name => 'HR', -
     object_name => 'EMPLOYEES', -
     proc_name => 'TABLE_ALTER_HANDLER');
```

  On encountering an ALTER TABLE statement, the schema_ddl_handler procedure will be invoked because its name will be at the top of an alphabetically sorted list of procedures that are relevant to the statement.Collisions on a rule set because of a specification containing wildcard entries are resolved in a similar fashion. For example, the rules in the following example will result in the empddl_handler procedure being invoked upon encountering the ALTER TABLE HR.EMPLOYEES ADD COLUMN RATING NUMBER statement:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP(-
     stmt => 'ALTER TABLE', -
     schema_name => 'HR', -
     object_name => 'EMP%', -
```

```
     proc_name => 'EMPDDL_HANDLER');
SQL> EXECUTE DBMS_LOGSTDBY.SKIP( -
     stmt => 'ALTER TABLE', -
     schema_name => 'HR', -
     object_name => 'EMPLOYEES', -
     proc_name => 'EMPLOYEE_DDL_HANDLER');
```

- Use the SKIP procedure with caution, particularly when skipping DDL statements. If a CREATE TABLE statement is skipped, for example, you must also specify other DDL statements that refer to that table in the SKIP procedure. Otherwise, the statements will fail and cause an exception. When this happens, SQL Apply stops running.
- Before calling the SKIP procedure, SQL Apply must be halted. Do this by issuing an ALTER DATABASE STOP LOGICAL STANDBY APPLY statement. Once all desired filters have been specified, issue an ALTER DATABASE START LOGICAL STANDBY APPLY IMMEDIATE statement to start SQL Apply using the new filter settings.
- See the UNSKIP procedure for information about reversing (undoing) the settings of the SKIP procedure.
- For USER statements, the SCHEMA_NAME parameter will be the user and specify '%' for the OBJECT_NAME parameter.
- If the PROC_NAME parameter is supplied, it must already exist in DBA_PROCEDURES and it must execute with DEFINER rights. If the procedure is declared with INVOKER rights, the ORA-1031: insufficient privileges message will be returned.
- If the procedure returns a REPLACEMENT statement, the REPLACEMENT statement will be executed using the SYSTEM and OBJECT privileges of the owner of the procedure.
- The PL/SQL block of a SKIP procedure cannot contain transaction control statements (for example, COMMIT, ROLLBACK, SAVEPOINT, and SET CONSTRAINT) unless the block is declared to be an autonomous transaction.

Skip Statement Options

Table 117-15 lists the supported values for the stmt parameter of the SKIP procedure. The left column of the table lists the keywords that may be used to identify the set of SQL statements to the right of the keyword. In addition, any of the SQL statements listed in the sys.audit_actions table (shown in the right column of Table 117-15) are also valid values. Note that keywords are generally defined by database object.

Table 117-15 Supported Values for the stmt Parameter

Keyword	Associated SQL Statements
There is no keyword for this group of SQL statements.	GRANT REVOKE ANALYZE TABLE ANALYZE INDEX ANALYZE CLUSTER
CLUSTER	AUDIT CLUSTER CREATE CLUSTER DROP CLUSTER TRUNCATE CLUSTER
CONTAINER	See "Skipping Containers"

Table 117-15 (Cont.) Supported Values for the `stmt` Parameter

Keyword	Associated SQL Statements
CONTEXT	CREATE CONTEXT DROP CONTEXT
DATABASE LINK	CREATE DATABASE LINK CREATE PUBLIC DATABASE LINK DROP DATABASE LINK DROP PUBLIC DATABASE LINK
DIMENSION	ALTER DIMENSION CREATE DIMENSION DROP DIMENSION
DIRECTORY[1]	CREATE DIRECTORY DROP DIRECTORY
DML	Includes DML statements on a table (for example: INSERT, UPDATE, and DELETE)
INDEX	ALTER INDEX CREATE INDEX DROP INDEX
NON_SCHEMA_DDL	*All DDL that does not pertain to a particular schema* **Note:** SCHEMA_NAME and OBJECT_NAME must be null
PL/SQL[2]	Execute Oracle-supplied package.
PROCEDURE[3]	ALTER FUNCTION ALTER PACKAGE ALTER PACKAGE BODY ALTER PROCEDURE CREATE FUNCTION CREATE LIBRARY CREATE PACKAGE CREATE PACKAGE BODY CREATE PROCEDURE DROP FUNCTION DROP LIBRARY DROP PACKAGE DROP PACKAGE BODY DROP PROCEDURE
PROFILE	ALTER PROFILE CREATE PROFILE DROP PROFILE

Table 117-15 (Cont.) Supported Values for the `stmt` Parameter

Keyword	Associated SQL Statements
ROLE	ALTER ROLE CREATE ROLE DROP ROLE SET ROLE
ROLLBACK STATEMENT	ALTER ROLLBACK SEGMENT CREATE ROLLBACK SEGMENT DROP ROLLBACK SEGMENT
SCHEMA_DDL	*All DDL statements that create, modify, or drop schema objects (for example: tables, indexes, and columns)* **Note:** SCHEMA_NAME and OBJECT_NAME must *not* be null
SEQUENCE	ALTER SEQUENCE CREATE SEQUENCE DROP SEQUENCE
SYNONYM	CREATE PUBLIC SYNONYM CREATE SYNONYM DROP PUBLIC SYNONYM DROP SYNONYM
SYSTEM AUDIT	AUDIT *SQL_statements* NOAUDIT *SQL_statements*
TABLE	CREATE TABLE ALTER TABLE DROP TABLE TRUNCATE TABLE
TABLESPACE	CREATE TABLESPACE DROP TABLESPACE ALTER TABLESPACE
TRIGGER	ALTER TRIGGER CREATE TRIGGER DISABLE ALL TRIGGERS DISABLE TRIGGER DROP TRIGGER ENABLE ALL TRIGGERS ENABLE TRIGGER

Table 117-15 (Cont.) Supported Values for the `stmt` Parameter

Keyword	Associated SQL Statements
TYPE	ALTER TYPE ALTER TYPE BODY CREATE TYPE CREATE TYPE BODY DROP TYPE DROP TYPE BODY
USER	ALTER USER CREATE USER DROP USER
VIEW	CREATE VIEW DROP VIEW
VIEW	CREATE VIEW DROP VIEW

[1] All directory objects are owned by SYS, but for the purpose of filtering them with a skip directive the schema should be specified as '%'".
[2] See *Oracle Data Guard Concepts and Administration* for information about supported packages.
[3] Java schema objects (sources, classes, and resources) are considered the same as procedure for purposes of skipping (ignoring) SQL statements.

Exceptions

Table 117-16 DBMS_LOGSTDBY.SKIP Procedure Exceptions

Exception	Description
ORA-01031	Insufficient privileges: • Procedure used INVOKER rights • Procedure needs DBA privileges
ORA-16103	Logical standby apply must be stopped to allow this operation.
ORA-16104	Invalid logical standby option requested.
ORA-16203	"Unable to interpret SKIP procedure return values." Indicates that a SKIP procedure has either generated an exception or has returned ambiguous values. You can identify the offending procedure by examining the DBA_LOGSTDBY_EVENTS view.
ORA-16236	Logical standby metadata operation in progress.

Examples

Example 1. Skipping all DML and DDL changes made to a schema
The following example shows how to specify rules so that SQL Apply will skip both DDL and DML statements made to the HR schema.

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP(STMT => 'SCHEMA_DDL', -
     schema_name => 'HR', -
     object_name => '%', -
     proc_name => null);
SQL> EXECUTE DBMS_LOGSTDBY.SKIP(STMT => 'DML', -
     schema_name => 'HR', -
     object_name => '%', -
     proc_name => null);
```

Example 2. Creating a procedure to handle different file system organization
For example, if the file system organization in the logical standby database is different than that in the primary database, you can write a `SKIP` procedure to handle DDL statements with file specifications transparently.The following procedure can handle DDL statements as long as you follow a specific naming convention for the file specification string.

1. Create the `SKIP` procedure to handle tablespace DDL statements:

```
CREATE OR REPLACE PROCEDURE sys.handle_tbs_ddl
IS
l_old_stmt varchar2(4000);
l_stmt_typ varcahr2(40);
l_schema   varchar2(30);
l_name     varchar2(30);
l_xidusn   number;
l_xidslt   number;
l_xidsqn   number;
l_skip_action number;
l_new_stmt varchar2(4000);

-- read all information
dbms_logstdby_context.get_context(name=>'STATEMENT',value=>l_old_stmt);
  dbms_logstdby_context.get_context(name=>'STATEMENT_TYPE',value=>l_stmt_type);
  dbms_logstdby_context.get_context(name=>'OWNER',value=>l_schema);
  dbms_logstdby_context.get_context(name=>'NAME',value=>l_name);
  dbms_logstdby_context.get_context(name=>'XIDUSN',value=>l_xidusn);
  dbms_logstdby_context.get_context(name=>'XIDSLT',value=>l_xidslt);
  dbms_logstdby_context.get_context(name=>'XIDSQN',value=>l_xidsqn);
  dbms_logstdby_context.get_context(name=>'CONTAINER_NAME',value=>l_conname);

--
-- All primary file specification that contains a directory
-- /usr/orcl/primary/dbs
-- should go to /usr/orcl/stdby directory specification

  BEGIN
  l_new_stmt := replace (l_old_stmt, '/usr/orcl/primary/dbs','/usr/orcl/stdby');
l_skip_action := DBMS_LOGSTDBY.SKIP_ACTION_REPLACE;
  EXCEPTION
  WHEN OTHERS THEN
    l_skip_action := DBMS_LOGSTDBY.SKIP_ACTION_ERROR;
    l_new_stmt := NULL;
  END;

  dbms_logstdby_context.set_context(name=>new_statement, value => l_new_stmt);
dbms_logstdby_context.set_context(name=>'SKIP_ACTION', value => l_skip_action);
END handle_tbs_ddl;
```

2. Register the `SKIP` procedure with SQL Apply:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP (stmt => 'TABLESPACE', -
             proc_name => 'SYS.HANDLE_TBS_DDL');
```

Skipping Containers

To skip a container (either a PDB or the root), use the `CONTAINER` keyword. All SQL statements executed on the container, as well as any other actions taken on the container, are skipped.

You can skip a particular PDB within a CDB. For example, the following command skips the PDB named `PDB1`. The command must be executed at the root level:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP(stmt => 'CONTAINER', object_name => 'PDB1');
```

As shown in the following example, you could also skip only the root of the CDB, but not any of the PDBs that exist under the root. The command must be executed at the root level:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP(stmt => 'CONTAINER', object_name => 'CDB$ROOT');
```

> **Note:**
>
> To create other skip rules for a container, create the rules from within the container. The container to which the rules will apply is automatically derived from the container in which the rules are created.

117.4.12 SKIP_ERROR Procedure

The SKIP_ERROR procedure specifies the action to take when a logical standby database detects an error.

Upon encountering an error, the logical standby database uses the criteria contained in this procedure to determine a course of action. The default action when a match is found is to skip the error and continue with applying changes. However, if a procedure is supplied, then `SKIP_ERROR` can take other actions depending on the situation. It can do nothing, which causes SQL Apply to stop, or it can change the error message text and stop SQL Apply, or it can actually skip the error.

Syntax

```
DBMS_LOGSTDBY.SKIP_ERROR (
     stmt                      IN VARCHAR2,
     schema_name               IN VARCHAR2 DEFAULT NULL,
     object_name               IN VARCHAR2 DEFAULT NULL,
     proc_name                 IN VARCHAR2 DEFAULT NULL,
     use_like                  IN BOOLEAN  DEFAULT NULL,
     esc                       IN CHAR1    DEFAULT NULL);
```

Parameters

Table 117-17 SKIP_ERROR Procedure Parameters

Parameter	Description
`stmt`	Either a keyword that identifies a set of SQL statements or a specific SQL statement. The use of keywords simplifies configuration because keywords, generally defined by the database object, identify all SQL statements that operate on the specified object. Table 117-15 shows a list of keywords and the equivalent SQL statements, either of which is a valid value for this parameter.

Table 117-17 (Cont.) SKIP_ERROR Procedure Parameters

Parameter	Description
schema_name	The name of one or more schemas (wildcards are permitted) associated with the SQL statements identified by the stmt parameter. If not applicable, this value must be set to NULL.
object_name	The name of one or more objects (wildcards are permitted) associated with the SQL statements identified by the stmt. If not applicable, this value must be set to NULL.
proc_name	Name of a stored procedure to call when SQL Apply encounters an error and determines a particular statement matches the filter defined by the stmt, schema_name, and object_name parameters. Specify the procedure in the following format: '"schema"."package"."procedure"' This procedure returns an error message that directs SQL Apply to perform one of the following actions: • Silently skip the error and continue with SQL Apply • Replace the error message that would have been created with a custom one, and stop SQL Apply • Do nothing, causing SQL Apply to stop and the original error message to be logged The procedure registered with SQL Apply does not take any parameters. The context associated with LSBY_APPLY_CONTEXT can be used to retrieve all relevant information related to the error. See the DBMS_LOGSTDBY_CONTEXT package for a list of all parameters associated with LSBY_APPLY_CONTEXT. The parameters of interest for procedures registered with SKIP_ERROR are CONTAINER_NAME, STATEMENT, STATEMENT_TYPE, SCHEMA, NAME, XIDUSN, XIDSLT, XIDSQN, ERROR and NEW_ERROR.
use_like	Allows pattern matching to isolate the tables that you want to skip on the logical standby database. The use_like parameter matches a portion of one character value to another by searching the first value for the pattern specified by the second, and calculates strings using characters as defined by the input character set. This parameter follows the same rules for pattern matching described in *Oracle Database SQL Language Reference*.
esc	Identifies an escape character (such as the characters "%" or "_") that you can use for pattern matching. If the escape character appears in the pattern before the character "%" or "_" then Oracle interprets this character literally in the pattern, rather than as a special pattern matching character.

Usage Notes

- A stored procedure provided to the SKIP_ERROR procedure is called when SQL Apply encounters an error that could shut down the application of redo logs to the standby database.
- Running this stored procedure affects the error being written in the STATUS column of the DBA_LOGSTDBY_EVENTS table. The STATUS_CODE column remains unchanged. If the stored procedure is to have no effect, that is, apply will be stopped, then the NEW_ERROR is written to the events table. To truly have no effect, set NEW_ERROR to ERROR in the procedure.
- If the stored procedure requires that a shutdown be avoided, then you must set NEW_ERROR to NULL.
- This procedure requires DBA privileges to execute.

- For USER statements, the SCHEMA_NAME parameter will be the user and you should specify '%' for the OBJECT_NAME parameter.
- If the PROC_NAME parameter is specified, it must already exist in DBA_PROCEDURES and it must execute with DEFINERS rights. If the procedure is declared with INVOKERS rights, the ORA-1031: insufficient privileges message will be returned.
- The PL/SQL block of a SKIP_ERROR procedure cannot contain transaction control statements (for example: COMMIT, ROLLBACK, SAVEPOINT, and SET CONSTRAINT) unless the block is declared to be an autonomous transaction using the following syntax:

  ```
  PRAGMA AUTONOMOUS_TRANSACTION
  ```

Exceptions

Table 117-18 SKIP_ERROR Procedure Exceptions

Exception	Description
ORA-01031	Insufficient privileges: • Procedure used INVOKER rights • Procedure needs DBA privileges
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested
ORA-16236	Logical Standby metadata operation in progress

Example 1

The following example shows how to specify rules so that SQL Apply will skip any error raised from any GRANT DDL command.

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP_ERROR('GRANT')
```

Example 2

To skip errors on GRANT statements on SYS or HR schemas, define a procedure handle_error_ddl and register it. In the following example, assume that handle_error_ddl is a free-standing procedure in the SYS schema.

1. Create the error-handler procedure:

```
CREATE OR REPLACE PROCEDURE sys.handle_error_ddl
is
l_stmt    VARCHAR2(4000);
l_stmt_type VARCHAR2(40);
l_schema    VARCHAR2(30);
l_name      VARCHAR2(30);
l_xidusn      NUMBER;
l_xidslt      NUMBER;
l_xidsqn      NUMBER;
l_error     VARCHAR2(4000);
l_conname    VARCHAR2(30);
l_newerr   VARCHAR2(4000);

BEGIN
 dbms_logstdby_context.get_context(name=>'STATEMENT',value=>l_stmt);
  dbms_logstdby_context.get_context(name=>'STATEMENT_TYPE',value=>l_stmt_type);
  dbms_logstdby_context.get_context(name=>'SCHEMA',value=>l_schema);
  dbms_logstdby_context.get_context(name=>'NAME',value=>l_name);
```

```
  dbms_logstdby_context.get_context(name=>'XIDUSN',value=>l_xidusn);
  dbms_logstdby_context.get_context(name=>'XIDSLT',value=>l_xidslt);
  dbms_logstdby_context.get_context(name=>'XIDSQN',value=>l_xidsqn);
  dbms_logstdby_context.get_context(name=>'ERROR',value=>l_error);
  dbms_logstdby_context.get_context(name=>'CONTAINER_NAME',value=>l_conname);

-- default error to what we already have
l_new_error := l_error;

-- Ignore any GRANT errors on SYS or HR schemas

IF INSTR(UPPER(l_stmt), 'GRANT') > 0
THEN
IF l_schema is NULL
OR (l_schema is NOT NULL AND
(UPPER(l_schema) = 'SYS' OR
UPPER(l_schema) = 'HR')
THEN
l_new_error := NULL;
-- record the fact that we just skipped an error on 'SYS' or 'HR' schemas
-- code not shown here
END IF;
END IF;

dbms_logstdby_context.set_context(name => 'NEW_ERROR', value => l_new_error);

END handle_error_ddl;
/
```

2. Register the error handler with SQL Apply:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP_ERROR ( -
     statement => 'NON_SCHEMA_DDL', -
     schema_name => NULL, -
     object_name => NULL, -
     proc_name => 'SYS.HANDLE_ERROR_DDL');
```

117.4.13 SKIP_TRANSACTION Procedure

This procedure provides a way to skip (ignore) applying transactions to the logical standby database. You can skip specific transactions by specifying transaction identification information.

Syntax

```
DBMS_LOGSTDBY.SKIP_TRANSACTION (
     xidusn_p          IN NUMBER,
     xidslt_p          IN NUMBER,
     xidsqn_p          IN NUMBER,
     con_name_p      VARCHAR2     IN     DEFAULT  NULL
);
```

Parameters

Table 117-19 SKIP_TRANSACTION Procedure Parameters

Parameter	Description
xidusn_p NUMBER	Transaction ID undo segment number of the transaction being skipped
xidslt_p NUMBER	Transaction ID slot number of the transaction being skipped

Table 117-19 (Cont.) SKIP_TRANSACTION Procedure Parameters

Parameter	Description
xidsqn_p NUMBER	Transaction ID sequence number of the transaction being skipped
con_name_p	The transaction name.

Usage Notes

If SQL Apply stops due to a particular transaction (for example, a DDL transaction), you can specify that transaction ID and then continue to apply. You can call this procedure multiple times for as many transactions as you want SQL Apply to ignore.

> **WARNING:**
>
> **SKIP_TRANSACTION** is an inherently dangerous operation. Do not invoke this procedure unless you have examined the transaction in question through the **V$LOGMNR_CONTENTS** view and have taken compensating actions at the logical standby database. **SKIP_TRANSACTION** is not the appropriate procedure to invoke to skip DML changes to a table.
>
> To skip a DML failure, use a **SKIP** procedure, such as **SKIP('DML','MySchema','MyFailed Table').** Using the **SKIP_TRANSACTION** procedure for DML transactions may skip changes for other tables, thus logically corrupting them.

- This procedure requires DBA privileges to execute.
- Use the DBA_LOGSTDBY_SKIP_TRANSACTION view to list the transactions that are going to be skipped by SQL Apply.

Exceptions

Table 117-20 SKIP_TRANSACTION Procedure Exceptions

Exception	Description
ORA-01031	Need DBA privileges
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested

Examples

To skip a DDL transaction with (xidusn_p, xidslt_p, xidsqn_p) of (1.13.1726) you can register a rule as shown in the following example:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP_TRANSACTION (-
     xidusn_p => 1, xidslt_p => 13, xidsqn_p => 1726);
```

117.4.14 UNSKIP Procedure

Use the UNSKIP procedure to delete rules specified earlier with the SKIP procedure.

The parameters specified in the UNSKIP procedure must match exactly for it to delete an already-specified rule.

The container_name argument is valid only in a CDB.

Syntax

```
DBMS_LOGSTDBY.UNSKIP (
     stmt                      IN VARCHAR2,
     schema_name               IN VARCHAR2 DEFAULT NULL,
     object_name               IN VARCHAR2 DEFUALT NULL,
     container_name            IN VARCHAR2 DEFAULT NULL);
```

Parameters

The parameter information for the UNSKIP procedure is the same as that described for the SKIP procedure. See Table 117-14 for complete parameter information.

Exceptions

Table 117-21 UNSKIP Procedure Exceptions

Exception	Description
ORA-01031	need DBA privileges to execute this procedure
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested

Usage Notes

> **WARNING:**
>
> If DML changes for a table have been skipped and not compensated for, you must follow the call to the **UNSKIP** procedure with a call to the **INSTANTIATE_TABLE** procedure to synchronize this table with those maintained by SQL Apply.

- This procedure requires DBA privileges to execute.
- Wildcards passed in the schema_name or the object_name parameter are not expanded. The wildcard character is matched at the character level. Thus, you can delete only one specified rule by invoking the UNSKIP procedure, and you will need a distinct UNSKIP procedure call to delete each rule that was previously specified.

  For example, assume you have specified the following two rules to skip applying DML statements to the HR.EMPLOYEE and HR.EMPTEMP tables:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP (STMT => 'DML',-
     SCHEMA_NAME => 'HR', -
     OBJECT_NAME => 'EMPLOYEE', -
     PROC_NAME => null);
```

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP (STMT => 'DML',-
     SCHEMA_NAME => 'HR', -
     OBJECT_NAME => 'EMPTEMP', -
     PROC_NAME => null);
```

In the following example, the wildcard in the TABLE_NAME parameter cannot be used to delete the rules that were specified:

```
SQL> EXECUTE DBMS_LOGSTDBY.UNSKIP (STMT => 'DML',-
     SCHEMA_NAME => 'HR', -
     OBJECT_NAME => 'EMP%');
```

In fact, this UNSKIP procedure matches neither of the rules, because the wildcard character in the TABLE_NAME parameter is not expanded. Instead, the wildcard character will be used in an exact match to find the corresponding SKIP rule.

117.4.15 UNSKIP_ERROR Procedure

Use the UNSKIP_ERROR procedure to delete rules specified earlier with the SKIP_ERROR procedure.

The parameters specified in the UNSKIP_ERROR procedure must match exactly for the procedure to delete an already-specified rule.

The container_name argument is valid only in a CDB.

Syntax

```
DBMS_LOGSTDBY.UNSKIP_ERROR (
     stmt                      IN VARCHAR2,
     schema_name               IN VARCHAR2 DEFAULT NULL,
     object_name               IN VARCHAR2 DEFAULT NULL,
     container_name            IN VARCHAR2 DEFAULT NULL);
```

Parameters

The parameter information for the UNSKIP_ERROR procedure is the same as that described for the SKIP_ERROR procedure. See Table 117-17 for complete parameter information.

Exceptions

Table 117-22 UNSKIP_ERROR Procedure Exceptions

Exception	Description
ORA-01031	Need DBA privileges
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested

Usage Notes

- This procedure requires DBA privileges to execute.
- Wildcards passed in the schema_name or the object_name parameters are not expanded. Instead, the wildcard character is treated as any other character and an exact match is made. Thus, you can delete only one specified rule by invoking the UNSKIP_ERROR procedure, and you need a distinct UNSKIP_ERROR procedure call to delete each rule that you previously specified.

For example, assume you have specified the following two rules to handle the HR.EMPLOYEE and HR.EMPTEMP tables:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP_ERROR (STMT => 'DML',-
     SCHEMA_NAME => 'HR', -
     OBJECT_NAME => 'EMPLOYEE', -
     PROC_NAME => 'hr_employee_handler');
SQL> EXECUTE DBMS_LOGSTDBY.SKIP_ERROR (STMT => 'DML',-
     SCHEMA_NAME => 'HR', -
     OBJECT_NAME => 'EMPTEMP', -
     PROC_NAME => 'hr_tempemp_handler');
```

In this case, the following UNSKIP procedure cannot be used to delete the rules that you have specified:

```
SQL> EXECUTE DBMS_LOGSTDBY.UNSKIP_ERROR (STMT => 'DML',-
     SCHEMA_NAME => 'HR', -
     OBJECT_NAME => 'EMP%');
```

In fact, the UNSKIP procedure will match neither of the rules, because the wildcard character in the OBJECT_NAME parameter will not be expanded.

Example

To remove a handler that was previously registered with SQL Apply from getting called on encountering an error, you can issue the following statement:

```
DBMS_LOGSTDBY.UNSKIP_ERROR ( -
      statement => 'NON_SCHEMA_DDL', -
      schema_name => NULL, -
      object_name => NULL);
```

117.4.16 UNSKIP_TRANSACTION Procedure

Use the UNSKIP_TRANSACTION procedure to delete rules specified earlier with the SKIP_TRANSACTION procedure.

The parameters specified in the UNSKIP_TRANSACTION procedure must match exactly for the procedure to delete an already-specified rule.

Syntax

```
DBMS_LOGSTDBY.UNSKIP_TRANSACTION (
     xidusn_p         IN NUMBER,
     xidslt_p         IN NUMBER,
     xidsqn_p         IN NUMBER);
```

Parameters

Table 117-23 UNSKIP_TRANSACTION Procedure Parameters

Parameter	Description
XIDUSN	Transaction ID undo segment number of the transaction being skipped
XIDSLT	Transaction ID slot number of the transaction being skipped
XIDSQN	Transaction ID sequence number of the transaction being skipped

Exceptions

Table 117-24 UNSKIP_TRANSACTION Procedure Exceptions

Exception	Description
ORA-01031	need DBA privileges to execute this procedure
ORA-16103	Logical Standby apply must be stopped to allow this operation
ORA-16104	invalid Logical Standby option requested

Usage Notes

- This procedure requires DBA privileges to execute.
- Query the DBA_LOGSTDBY_SKIP_TRANSACTION view to list the transactions that are going to be skipped by SQL Apply.

Examples

To remove a rule that was originally specified to skip the application of a transaction with (XIDUSN, XIDSLT, XIDSQN) of (1.13.1726) issue the following statement:

```
SQL> DBMS_LOGSTDBY.UNSKIP_TRANSACTION (XIDUSN => 1, XIDSLT => 13, XIDSQN => 1726);
```

118
DBMS_LOGMNR_D

The `DBMS_LOGMNR_D` package, one of a set of LogMiner packages, contains two subprograms: the `BUILD` procedure and the `SET_TABLESPACE` procedure.

- The `BUILD` procedure extracts the LogMiner data dictionary to either the redo log files or to a flat file. This information is saved in preparation for future analysis of redo log files using the LogMiner tool.
- The `SET_TABLESPACE` procedure re-creates all LogMiner tables in an alternate tablespace.

  The **LogMiner data dictionary** consists of the memory data structures and the database tables that are used to store and retrieve information about objects and their versions. It is referred to as the **LogMiner dictionary** throughout the LogMiner documentation.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_LOGMNR_D Subprograms

> **See Also:**
>
> *Oracle Database Utilities* for information regarding LogMiner.

118.1 DBMS_LOGMNR_D Overview

LogMiner requires a dictionary to translate object IDs into object names when it returns redo data to you.

LogMiner gives you two options for supplying the dictionary:

- Using the online catalog
- Extracting a LogMiner dictionary to the redo log files
- Extracting a LogMiner dictionary to a flat file (for non-CDBs only)

> **Note:**
>
> In previous releases, using a flat file dictionary was one means of mining the redo logs for the changes associated with a specific PDB whose data dictionary was contained within the flat file. This feature is desupported for PDBs in Oracle Database 19c, and desupported in later releases. With Oracle Database 19c and later releases, Oracle recommends that you call `DBMS_LOGMNR.START_LOGMNR`, and supply the system change number (SCN) or time range that you want to mine. The SCN or time range options of `START_LOGMNR` are enhanced to support mining of individual PDBs.

Use the BUILD procedure to extract the LogMiner dictionary to the redo log files or a flat file. If you want to specify the online catalog as the dictionary source, you do so when you start LogMiner with the DBMS_LOGMNR.START_LOGMNR package.

Use the SET_TABLESPACE procedure if you want LogMiner tables to use a tablespace other than the default SYSAUX tablespace.

> **See Also:**
>
> DBMS_LOGMNR for information on the package subprograms used in running a LogMiner session.

118.2 DBMS_LOGMNR_D Security Model

You must have the EXECUTE_CATALOG_ROLE role to use the DBMS_LOGMNR_D package.

118.3 Summary of DBMS_LOGMNR_D Subprograms

This table lists and briefly describes the DBMS_LOGMNR_D subprograms.

In a multitenant container database (CDB), some subprograms must be called from the root. There can be other differences as well. See the individual subprogram descriptions for details.

Table 118-1 DBMS_LOGMNR_D Package Subprograms

Subprogram	Description
BUILD Procedure	Extracts the LogMiner dictionary to either a flat file or one or more redo log files
SET_TABLESPACE Procedure	Re-creates all LogMiner tables in an alternate tablespace

118.3.1 BUILD Procedure

This procedure extracts the LogMiner data dictionary to the redo log files.

The following considerations apply to a multitenant container database (CDB) environment.

- In a CDB environment, when you extract to the redo log files, the BUILD procedure must be called from the root database. The LogMiner data dictionary for the entire CDB is extracted to the redo log files.
- You cannot add or remove PDBs from a CDB while this procedure is running.

> **Note:**
>
> In previous releases, using a flat file dictionary was one means of mining the redo logs for the changes associated with a specific PDB whose data dictionary was contained within the flat file. This feature is desupported for PDBs in Oracle Database 19c, and desupported in later releases. With Oracle Database 19c and later releases, Oracle recommends that you call `DBMS_LOGMNR.START_LOGMNR`, and supply the system change number (SCN) or time range that you want to mine. The SCN or time range options of `START_LOGMNR` are enhanced to support mining of individual PDBs.

Syntax

```
DBMS_LOGMNR_D.BUILD (
     dictionary_filename  IN  VARCHAR2,
     dictionary_location  IN  VARCHAR2,
     options              IN  NUMBER);
```

Parameters

Table 118-2 BUILD Procedure Parameters

Parameter	Description
`dictionary_filename`	Specifies the name of the LogMiner dictionary file.
`dictionary_location`	Specifies the directory object for the LogMiner dictionary file.
`options`	Specifies that the LogMiner dictionary is written to the redo log files (`STORE_IN_REDO_LOGS`).

Exceptions

Table 118-3 BUILD Procedure Exceptions

Exception	Description
`ORA-01302`	Dictionary build options are missing or incorrect. This error is returned under the following conditions: • If the value of the `OPTIONS` parameter is not one of the supported values (`STORE_IN_REDO_LOGS`), or is not specified • If the `STORE_IN_REDO_LOGS` option is not specified and neither the `dictionary_filename` nor the `dictionary_location` parameter is specified • If the `STORE_IN_REDO_LOGS` option is specified and either the `dictionary_filename` or the `dictionary_location` parameter is specified

Table 118-3 (Cont.) BUILD Procedure Exceptions

Exception	Description
ORA-01308	Initialization parameter UTL_FILE_DIR is not set. **Note:** In earlier releases, you used the UTL_FILE_DIR initialization parameter to specify a directory location. However, as of Oracle Database 18c, the UTL_FILE_DIR initialization parameter is desupported. It is still supported for backward compatibility, but Oracle recommends that you instead use directory objects.
ORA-01336	Specified dictionary file cannot be opened. This error is returned when the dictionary file is read-only.
ORA-01308	Dictionary directory is not set. This error is returned under the following conditions: • The specified value for the dictionary_location is not a directory object. • The specified value for the dictionary_location is a directory object that is defined to be a file path that cannot be accessed.

Usage Notes

- To extract the LogMiner dictionary to the redo log files, specify only the STORE_IN_REDO_LOGS option. The size of the LogMiner dictionary may cause it to be contained in multiple redo log files.

  The combinations of parameters used result in the following behavior:

  – If you do not specify any parameters, an error is returned.

  – If you do not specify a filename and location, but do specify the STORE_IN_REDO_LOGS option, the LogMiner dictionary is extracted to the redo log files.

  – If you specify a filename and location, as well as the STORE_IN_REDO_LOGS option, an error is returned.

- Ideally, the LogMiner dictionary file is created after all database dictionary changes have been made, and before the creation of any redo log files that you want to analyze. You can use LogMiner to dump the LogMiner dictionary to the redo log files, perform DDL operations, and dynamically apply the DDL changes to the LogMiner dictionary.
- The database must be open when you run the DBMS_LOGMNR_D.BUILD procedure.
- To extract a LogMiner dictionary file to the redo log files, the following conditions must be met:

  – Archivelog mode must be enabled in order to generate usable redo log files.

  – The COMPATIBLE parameter in the initialization parameter file must be set to 9.2.0 or higher.

- The database to which LogMiner is attached must be Oracle9i or later.

In addition, supplemental logging (at least the minimum level) should be enabled to ensure that you can take advantage of all the features that LogMiner offers.

Examples

Example: Extracting the LogMiner Dictionary to the Redo Log Files

The following example extracts the LogMiner dictionary to the redo log files.

```
SQL> EXECUTE dbms_logmnr_d.build( -
     options => dbms_logmnr_d.store_in_redo_logs);
```

118.3.2 SET_TABLESPACE Procedure

This procedure moves LogMiner tables from the default `SYSAUX` tablespace to an alternate tablespace.

By default, all LogMiner tables are created to use the `SYSAUX` tablespace. However, it may be desirable to have LogMiner tables use an alternate tablespace. Use this procedure to move LogMiner tables to this alternate tablespace

In a CDB, only the LogMiner metadata in the local container is moved to the requested tablespace.

Syntax

```
DBMS_LOGMNR_D.SET_TABLESPACE (
     new_tablespace        IN VARCHAR2);
```

Parameters

Table 118-4 SET_TABLESPACE Parameter

Parameter	Description
`new_tablespace`	A string naming a preexisting tablespace. To move all LogMiner tables to employ this tablespace, supply this parameter.

Usage Notes

- Users upgrading from earlier versions of Oracle Database may find LogMiner tables in the `SYSTEM` tablespace. Oracle encourages such users to consider using the `SET_TABLESPACE` procedure to move the tables to the `SYSAUX` tablespace once they are confident that they will not be downgrading to an earlier version of Oracle Database.
- Users of this routine must supply an existing tablespace.

Example: Using the DBMS_LOGMNR_D.SET_TABLESPACE Procedure

The following example shows the creation of an alternate tablespace and execution of the `DBMS_LOGMNR_D.SET_TABLESPACE` procedure.

```
SQL> CREATE TABLESPACE  logmnrts$ datafile '/usr/oracle/dbs/logmnrts.f'
     SIZE 25 M REUSE AUTOEXTEND ON MAXSIZE UNLIMITED;

SQL> EXECUTE dbms_logmnr_d.set_tablespace('logmnrts$');
```

119
DBMS_LOGSTDBY_CONTEXT

As of Oracle Database 12*c* release 1 (12.1), SQL Apply processes have access to a context namespace called `LSBY_APPLY_CONTEXT`. You can use the procedures provided in the `DBMS_LOGSTDBY_CONTEXT` package to set and retrieve various parameters associated with `LSBY_APPLY_CONTEXT`. This is useful when writing skip procedures that are registered with SQL Apply using the `DBMS_LOGSTDBY.SKIP` and `DBMS_LOGSTDBY.SKIP_ERROR` procedures.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_LOGSTDBY_CONTEXT Subprograms

> **See Also:**
>
> *Oracle Data Guard Concepts and Administration* for more information about SQL Apply and logical standby databases

119.1 DBMS_LOGSTDBY_CONTEXT Overview

SQL Apply processes have access to a context namespace called `LSBY_APPLY_CONTEXT`.

The following table lists the predefined parameters associated with `LSBY_APPLY_CONTEXT` that you can set and retrieve by using the procedures provided in the `DBMS_LOGSTDBY_CONTEXT` package. The ability to set and retrieve the parameters in this way is useful when writing skip procedures that are registered with SQL Apply using the `DBMS_LOGSTDBY.SKIP` and `DBMS_LOGSTDBY.SKIP_ERROR` procedures.

Table 119-1 Predefined Parameters of Namespace LSBY_APPLY_CONTEXT

Parameter	Description
`STATEMENT`	First 4000 bytes of the statement that the Apply process is processing.
`STATEMENT_TYPE`	See Table 117-15 for a list of various statement types.
`PACKAGE_SCHEMA`	Schema that owns the PL/SQL package being processed (if applicable).
`PACKAGE_NAME`	Name of the PL/SQL package being processed (if applicable).
`PROCEDURE_NAME`	Name of the PL/SQL procedure being processed (if applicable).
`CURRENT_SCHEMA`	Current schema in effect when the DDL or PL/SQL procedure was executed at the primary.
`XIDUSN`	XIDUSN of the transaction at the primary database.
`XIDSLT`	XIDSLT of the transaction at the primary database.

Table 119-1 (Cont.) Predefined Parameters of Namespace LSBY_APPLY_CONTEXT

Parameter	Description
XIDSQN	XIDSQN of the transaction at the primary database.
SCHEMA	Schema that owns the object being processed by SQL Apply (in the case of DML or DDL operations).
NAME	Object name being processed by SQL Apply (in case of DML or DDL operations).
CONTAINER_NAME	Container where the target object or the PL/SQL procedure resides (in the case of a multitenant container database (CDB)).
ERROR	Text of the original error encountered by the SQL Apply process (if applicable).
NEW_ERROR	Text of the new error to be raised by the SQL Apply process (if applicable, See the DBMS_LOGSTDBY SKIP_ERROR Procedure).
NEW_STMT	Text of the replacement statement that SQL Apply must execute (If applicable, See the DBMS_LOGSTDBY SKIP Procedure).
SKIP_ACTION	The skip action to be performed by SQL Apply (See the DBMS_LOGSTDBY SKIP Procedure).

119.2 DBMS_LOGSTDBY_CONTEXT Security Model

The security model for the DBMS_LOGSTDBY_CONTEXT package is the same as for the DBMS_LOGSTDBY package.

> **See Also:**
>
> DBMS_LOGSTDBY

119.3 Summary of DBMS_LOGSTDBY_CONTEXT Subprograms

This table lists and briefly describes the DBMS_LOGSTDBY_CONTEXT package subprograms.

Table 119-2 DBMS_LOGSTDBY_CONTEXT Package Subprograms

Subprogram	Description
CLEAR_ALL_CONTEXT Procedure	Clears all parameters contained within namespace LSBY_APPLY_CONTEXT.
CLEAR_CONTEXT Procedure	Clears the specific parameter.
GET_CONTEXT Procedure	Retrieves the value for the specified parameter.
SET_CONTEXT Procedure	Sets the named parameter with the specified value.

119.3.1 CLEAR_ALL_CONTEXT Procedure

This procedure clears all parameters contained within namespace `LSBY_APPLY_CONTEXT`.

Syntax

```
DBMS_LOGSTDBY_CONTEXT.CLEAR_ALL_CONTEXT;
```

Parameters

None

119.3.2 CLEAR_CONTEXT Procedure

This procedure clears the specific parameter.

Syntax

```
DBMS_LOGSTDBY_CONTEXT.CLEAR_CONTEXT (
   name             IN VARCHAR2);
```

Parameters

Table 119-3 *CLEAR_CONTEXT Procedure Parameters*

Parameter	Description
`name`	Name of the parameter to be cleared.

119.3.3 GET_CONTEXT Procedure

This procedure retrieves the value for the specified parameter.

Syntax

```
DBMS_LOGSTDBY_CONTEXT.GET_CONTEXT (
   name             IN VARCHAR2,
   value            OUT VARCHAR2);
```

Parameters

Table 119-4 *GET_CONTEXT Procedure Parameters*

Parameter	Description
`name`	Name of the parameter.
value	The value retrieved for the parameter.

119.3.4 SET_CONTEXT Procedure

This procedure sets the named parameter with the specified value.

Syntax

```
DBMS_LOGSTDBY_CONTEXT.SET_CONTEXT (
   name             IN VARCHAR2
   value            IN VARCHAR2);
```

Parameters

Table 119-5 *SET_CONTEXT Procedure Parameters*

Parameter	Description
name	Name of the parameter to be set.
value	Value to be assigned to the parameter being set.

120
DBMS_MEMOPTIMIZE

The `DBMS_MEMOPTIMIZE` package provides the interface for managing the Memoptimized Rowstore data buffered in the large pool and the memoptimize pool.

This chapter contains the following topics:

- DBMS_MEMOPTIMIZE Overview
- Summary of DBMS_MEMOPTIMIZE Subprograms

120.1 DBMS_MEMOPTIMIZE Overview

The `DBMS_MEMOPTIMIZE` package provides the interface for managing Memoptimized Rowstore data buffered in the large pool and the memoptimize pool.

The Memoptimized Rowstore provides the following functionality:

- Fast Ingest

  Fast ingest optimizes the processing of high-frequency, single-row data inserts. Fast ingest uses the large pool in the SGA for buffering the inserts before writing them to disk.

- Fast Lookup

  Fast lookup enables fast retrieval of data for high-frequency queries. Fast lookup uses a separate memory area in the SGA called the *memoptimize pool* for buffering data queried from tables.

The `DBMS_MEMOPTIMIZE` package provides the following operations related to the Memoptimized Rowstore:

- Fast ingest operations:
  - Provide the low high-water mark (low HWM) of the sequence numbers of rows that have been successfully written from the large pool to disk across all the sessions.
  - Provide the high-water mark (HWM) sequence number of the row that has been written to the large pool for the current session.
  - Flush all the fast ingest data from the large pool to disk for the current session.
- Fast lookup operations:
  - Remove data for a table from the memoptimize pool.
  - Populate data for a table in the memoptimize pool.

120.2 Summary of DBMS_MEMOPTIMIZE Subprograms

This table lists the `DBMS_MEMOPTIMIZE` subprograms and briefly describes them.

Table 120-1 DBMS_MEMOPTIMIZE Package Subprograms Related to Fast Ingest

Procedure	Description
GET_APPLY_HWM_SEQID Function	Returns the low high-water mark (low HWM) of the sequence numbers of rows that have been successfully written to disk globally across all the sessions.
GET_WRITE_HWM_SEQID Function	Returns the high-water mark (HWM) sequence number of the row that has been written to the large pool for the current session.
WRITE_END Procedure	Flushes all the fast ingest data from the large pool to disk for the current session.

Table 120-2 DBMS_MEMOPTIMIZE Package Subprograms Related to Fast Lookup

Procedure	Description
DROP_OBJECT Procedure	Removes data for a table from the memoptimize pool.
POPULATE Procedure	Populates data for a table in the memoptimize pool.

120.2.1 DROP_OBJECT Procedure

This procedure removes data for a table from the memoptimize pool.

Syntax

```
DBMS_MEMOPTIMIZE.DROP_OBJECT (
   schema_name         IN VARCHAR2,
   table_name          IN VARCHAR2,
   partition_name      IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 120-3 DROP_OBJECT Procedure Parameters

Parameter	Description
schema_name	Name of the schema.
table_name	Name of the table for which the data needs to be removed from the memoptimize pool.
partition_name	Name of the table partition for which the data needs to be removed from the memoptimize pool. This is an optional parameter and its default value is NULL.

120.2.2 GET_APPLY_HWM_SEQID Function

This function returns the low high-water mark (low HWM) of sequence numbers of the records that have been successfully written to disk across all the sessions.

Syntax

```
DBMS_MEMOPTIMIZE.GET_APPLY_HWM_SEQID
  RETURN number;
```

Return Value

Returns the low high-water mark (low HWM) of sequence numbers of the records that have been successfully written to disk across all the sessions.

120.2.3 GET_WRITE_HWM_SEQID Function

This function returns the high-water mark (HWM) sequence number of the record that has been written to the large pool for the current session.

Syntax

```
DBMS_MEMOPTIMIZE.GET_WRITE_HWM_SEQID
  RETURN number;
```

Return Value

Returns the high-water mark (HWM) sequence number of the record that has been written to the large pool for the current session.

120.2.4 POPULATE Procedure

This procedure populates the data for a table in the memoptimize pool.

Syntax

```
DBMS_MEMOPTIMIZE.POPULATE (
   schema_name        IN VARCHAR2,
   table_name         IN VARCHAR2,
   partition_name     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 120-4 POPULATE Procedure Parameters

Parameter	Description
schema_name	Name of the schema.
table_name	Name of the table for which the data needs to be populated in the memoptimize pool.

Table 120-4 (Cont.) POPULATE Procedure Parameters

Parameter	Description
partition_name	Name of the table partition for which the data needs to be populated in the memoptimize pool. This is an optional parameter and its default value is NULL.

120.2.5 WRITE_END Procedure

This procedure flushes all the fast ingest data from the large pool to disk for the current session.

Syntax

```
DBMS_MEMOPTIMIZE.WRITE_END;
```

121
DBMS_MEMOPTIMIZE_ADMIN

The `DBMS_MEMOPTIMIZE_ADMIN` package provides the interface for flushing the Memoptimized Rowstore data buffered in the large pool to disk.

This chapter contains the following topics:

- DBMS_MEMOPTIMIZE_ADMIN Overview
- Summary of DBMS_MEMOPTIMIZE_ADMIN Subprograms

121.1 DBMS_MEMOPTIMIZE_ADMIN Overview

The `DBMS_MEMOPTIMIZE_ADMIN` package provides the interface for flushing the Memoptimized Rowstore data that is buffered in the large pool to disk.

The Memoptimized Rowstore provides the following functionality:

- Fast Ingest

  Fast ingest optimizes the processing of high-frequency, single-row data inserts. Fast ingest uses the large pool in the SGA for buffering the inserts before writing them to disk.

- Fast Lookup

  Fast lookup enables fast retrieval of data for high-frequency queries. Fast lookup uses a separate memory area in the SGA called the *memoptimize pool* for buffering the data queried from tables.

The `DBMS_MEMOPTIMIZE_ADMIN` package provides the following operation related to fast ingest of the Memoptimized Rowstore:

- Flush fast ingest data from the large pool to disk for all sessions.

121.2 Summary of DBMS_MEMOPTIMIZE_ADMIN Subprograms

This table lists the `DBMS_MEMOPTIMIZE_ADMIN` subprograms and briefly describes them.

Table 121-1 DBMS_MEMOPTIMIZE_ADMIN Package Subprograms

Procedure	Description
WRITES_FLUSH Procedure	Flushes all the fast ingest data of the Memoptimize Rowstore from the large pool to disk for all the sessions.

121.2.1 WRITES_FLUSH Procedure

This procedure flushes all the fast ingest data of the Memoptimized Rowstore from the large pool to disk for all the sessions.

Syntax

```
DBMS_MEMOPTIMIZE_ADMIN.WRITES_FLUSH;
```

122
DBMS_METADATA

The DBMS_METADATA package provides a way for you to retrieve metadata from the database dictionary as XML or creation DDL and to submit the XML to re-create the object.

> **See Also:**
>
> *Oracle Database Utilities* for more information and for examples of using the DBMS_METADATA package.

This chapter contains the following topics:

- Overview
- Security Model
- Rules and Limits
- Data Structures - Object and Table Types
- Subprogram Groupings
  - Subprograms for Retrieving Multiple Objects From the Database
  - Subprograms for Submitting XML to the Database
- Summary of All DBMS_METADATA Subprograms

122.1 DBMS_METADATA Overview

You can use the DBMS_METADATA package to retrieve metadata and to submit XML.

Retrieving Metadata

If you are retrieving metadata, you can specify:

- The kind of object to be retrieved. This can be either a particular object type (such as a table, index, or procedure) or a heterogeneous collection of object types that form a logical unit (such as a database export or schema export).
- Optional selection criteria, such as owner or name.
- Parse items (attributes of the returned objects to be parsed and returned separately).
- Optional transformations on the output, implemented by XSLT (Extensible Stylesheet Language Transformation) scripts. By default the output is represented in XML, but you can specify transformations (into SQL DDL, for example), which are implemented by XSLT stylesheets stored in the database or externally.

DBMS_METADATA provides the following retrieval interfaces:

- For programmatic use: OPEN, SET_FILTER, SET_COUNT, GET_QUERY, SET_PARSE_ITEM, ADD_TRANSFORM, SET_TRANSFORM_PARAM,SET_REMAP_PARAM, FETCH_xxx, and CLOSE retrieve multiple objects.

- For use in SQL queries and for browsing: `GET_XML`, `GET_DDL` and `GET_SXML` return metadata for a single named object. The `GET_DEPENDENT_XML`, `GET_DEPENDENT_DDL`, `GET_GRANTED_XML`, and `GET_GRANTED_DDL` interfaces return metadata for one or more dependent or granted objects. These procedures do not support heterogeneous object types.

Submitting XML

If you are submitting XML, you specify:

- The type of object
- Optional transform parameters to modify the object (for example, changing the object's owner)
- Parse items (attributes of the submitted objects to be parsed and submitted separately)
- Whether to execute the operation or simply return the generated DDL

`DBMS_METADATA` provides a programmatic interface for submission of XML. It is comprised of the following procedures: `OPENW`, `ADD_TRANSFORM`, `SET_TRANSFORM_PARAM`, `SET_REMAP_PARAM`, `SET_PARSE_ITEM`, `CONVERT`, `PUT`, and `CLOSE`.

122.2 DBMS_METADATA Security Model

The `DBMS_METADATA` package considers a privileged user to be one who is connected as user `SYS` or who has the `SELECT_CATALOG_ROLE` role.

The object views of the Oracle metadata model implement security as follows:

- Nonprivileged users can see the metadata of only their own objects.
- Nonprivileged users can also retrieve public synonyms, system privileges granted to them, and object privileges granted to them or by them to others. This also includes privileges granted to `PUBLIC`.
- If callers request objects they are not privileged to retrieve, no exception is raised; the object is simply not retrieved.
- If nonprivileged users are granted some form of access to an object in someone else's schema, they will be able to retrieve the grant specification through the Metadata API, but not the object's actual metadata.
- In stored procedures, functions, and definers-rights packages, roles (such as `SELECT_CATALOG_ROLE`) are disabled. Therefore, such a PL/SQL program can only fetch metadata for objects in its own schema. If you want to write a PL/SQL program that fetches metadata for objects in a different schema (based on the invoker's possession of `SELECT_CATALOG_ROLE`), you must make the program invokers-rights.
- For all objects that have passwords, except database links (for example, users and roles), the following rules apply:
  - A user who has the `SELECT_CATALOG_ROLE` can see all metadata for an object except the passwords for that object.
  - The `SYS` user, users who have the `EXP_FULL_DATABASE` role, and users who own an object can see all metadata for that object, including passwords.
- For database links the password is never displayed. For security reasons Oracle restricts visibility of the password value to `SYS` users who query the `link$.passwordx` column directly. Instead of the password, `DBMS_METADATA` returns the following invalid syntax:

```
IDENTIFIED BY VALUES ':1'
```

A user who knows the password of the database link can manually replace the `:1` with the password.

122.3 Rules and Limits

In an Oracle Shared Server (OSS) environment, the `DBMS_METADATA` package must disable session migration and connection pooling.

This results in any shared server process that is serving a session running the package to effectively become a default, dedicated server for the life of the session. You should ensure that sufficient shared servers are configured when the package is used and that the number of servers is not artificially limited by too small a value for the `MAX_SHARED_SERVERS` initialization parameter.

122.4 DBMS_METADATA Data Structures - Object and Table Types

The `DBMS_METADATA` package defines, in the `SYS` schema, the `OBJECT` and `TABLE` types shown in this code.

```
CREATE TYPE sys.ku$_parsed_item AS OBJECT (
  item            VARCHAR2(30),
  value           VARCHAR2(4000),
  object_row      NUMBER )
/

CREATE PUBLIC SYNONYM ku$_parsed_item FOR sys.ku$_parsed_item;

CREATE TYPE sys.ku$_parsed_items IS TABLE OF sys.ku$_parsed_item
/

CREATE PUBLIC SYNONYM ku$_parsed_items FOR sys.ku$_parsed_items;

CREATE TYPE sys.ku$_ddl AS OBJECT (
   ddlText        CLOB,
parsedItems sys.ku$_parsed_items )
/

CREATE PUBLIC SYNONYM ku$_ddl FOR sys.ku$_ddl;

CREATE TYPE sys.ku$_ddls IS TABLE OF sys.ku$_ddl
/

CREATE PUBLIC SYNONYM ku$_ddls FOR sys.ku$_ddls;

CREATE TYPE sys.ku$_multi_ddl AS OBJECT (
   object_row     NUMBER,
   ddls           sys.ku$_ddls )
/

CREATE OR REPLACE PUBLIC SYNONYM ku$_multi_ddl FOR sys.ku$_multi_ddl;

CREATE TYPE sys.ku$_multi_ddls IS TABLE OF sys.ku$_multi_ddl;
/

CREATE OR REPLACE PUBLIC SYNONYM ku$_multi_ddls FOR
```

```
                          sys.ku$_multi_ddls;

CREATE TYPE sys.ku$_ErrorLine IS OBJECT (
   errorNumber    NUMBER,
   errorText      VARCHAR2(2000) )
/

CREATE PUBLIC SYNONYM ku$_ErrorLine FOR sys.ku$_ErrorLine;

CREATE TYPE sys.ku$_ErrorLines IS TABLE OF sys.ku$_ErrorLine
/
CREATE PUBLIC SYNONYM ku$ErrorLines FOR sys.ku$_ErrorLines;

CREATE TYPE sys.ku$_SubmitResult AS OBJECT (
   ddl          sys.ku$_ddl,
   errorLines   sys.ku$_ErrorLines );
/

CREATE TYPE sys.ku$_SubmitResults IS TABLE OF sys.ku$_SubmitResult
/

CREATE PUBLIC SYNONYM ku$_SubmitResults FOR sys.ku$_SubmitResults;
```

> **Note:**
>
> The maximum size of the `VARCHAR2`, `NVARCHAR2`, and `RAW` datatypes has been increased to 32 KB when the `COMPATIBLE` initialization parameter is set to 12.0 and the `MAX_STRING_SIZE` initialization parameter is set to `EXTENDED`. The `DBMS_METADATA` package supports this increased size unless the version of the metadata is earlier than Oracle Database 12*c* Release 1 (12.1).

122.5 DBMS_METADATA Subprogram Groupings

The `DBMS_METADATA` subprograms retrieve objects from, and submit XML to, a database. Some subprograms are used for both activities, while others are used only for retrieval or only for submission.

- Table 122-1 provides a summary, in alphabetical order, of `DBMS_METADATA` subprograms used to retrieve multiple objects from a database.
- Table 122-2 provides a summary, in alphabetical order, of `DBMS_METADATA` subprograms used to submit XML metadata to a database.

122.5.1 DBMS_METADATA Subprograms for Retrieving Multiple Objects From the Database

`DBMS_METADATA` uses these subprograms used for retrieving multiple objects from the database.

Table 122-1 DBMS_METADATA Subprograms for Retrieving Multiple Objects

Subprogram	Description
ADD_TRANSFORM Function	Specifies a transform that `FETCH_xxx` applies to the XML representation of the retrieved objects

Table 122-1 (Cont.) DBMS_METADATA Subprograms for Retrieving Multiple Objects

Subprogram	Description
CLOSE ProcedureCLOSE Procedure	Invalidates the handle returned by `OPEN` and cleans up the associated state
FETCH_xxx Functions and Procedures	Returns metadata for objects meeting the criteria established by `OPEN`, `SET_FILTER`, `SET_COUNT`, `ADD_TRANSFORM`, and so on
GET_QUERY Function	Returns the text of the queries that are used by `FETCH_xxx`
GET_xxx Functions	Fetches the metadata for a specified object as XML, SXML, or DDL, using only a single call
OPEN Function	Specifies the type of object to be retrieved, the version of its metadata, and the object model
SET_COUNT Procedure	Specifies the maximum number of objects to be retrieved in a single `FETCH_xxx` call
SET_FILTER Procedure	Specifies restrictions on the objects to be retrieved, for example, the object name or schema
SET_PARSE_ITEM Procedure	Enables output parsing by specifying an object attribute to be parsed and returned
SET_TRANSFORM_PARAM and SET_REMAP_PARAM Procedures	Specifies parameters to the XSLT stylesheets identified by `transform_handle`

122.5.2 DBMS_METADATA Subprograms for Submitting XML to the Database

`DBMS_METADATA` uses these subprograms for submitting XML to the database.

Table 122-2 DBMS_METADATA Subprograms for Submitting XML

Subprogram	Description
ADD_TRANSFORM Function	Specifies a transform for the XML documents
CLOSE ProcedureCLOSE Procedure	Closes the context opened with `OPENW`
CONVERT Functions and Procedures	Converts an XML document to DDL
OPENW Function	Opens a write context
PUT Function	Submits an XML document to the database
SET_PARSE_ITEM Procedure	Specifies an object attribute to be parsed
SET_TRANSFORM_PARAM and SET_REMAP_PARAM Procedures	`SET_TRANSFORM_PARAM` specifies a parameter to a transform `SET_REMAP_PARAM` specifies a remapping for a transform

122.6 Summary of All DBMS_METADATA Subprograms

This table lists the `DBMS_METADATA` subprograms and briefly describes them.

Table 122-3 DBMS_METADATA Package Subprograms

Subprogram	Description
ADD_TRANSFORM Function	Specifies a transform that FETCH_xxx applies to the XML representation of the retrieved objects
CLOSE ProcedureCLOSE Procedure	Invalidates the handle returned by OPEN and cleans up the associated state
CONVERT Functions and Procedures	Converts an XML document to DDL
FETCH_xxx Functions and Procedures	Returns metadata for objects meeting the criteria established by OPEN, SET_FILTER, SET_COUNT, ADD_TRANSFORM, and so on
GET_xxx Functions	Fetches the metadata for a specified object as XML, SXML, or DDL, using only a single call
GET_QUERY Function	Returns the text of the queries that are used by FETCH_xxx
OPEN Function	Specifies the type of object to be retrieved, the version of its metadata, and the object model
OPENW Function	Opens a write context
PUT Function	Submits an XML document to the database
SET_COUNT Procedure	Specifies the maximum number of objects to be retrieved in a single FETCH_xxx call
SET_FILTER Procedure	Specifies restrictions on the objects to be retrieved, for example, the object name or schema
SET_PARSE_ITEM Procedure	Enables output parsing by specifying an object attribute to be parsed and returned
SET_TRANSFORM_PARAM and SET_REMAP_PARAM Procedures	Specifies parameters to the XSLT stylesheets identified by transform_handle

122.6.1 ADD_TRANSFORM Function

The DBMS_METADATA.ADD_TRANSFORM function is used for both retrieval and submission.

- When this procedure is used to retrieve objects, it specifies a transform that FETCH_xxx applies to the XML representation of the retrieved objects.
- When used to submit objects, it specifies a transform that CONVERT or PUT applies to the XML representation of the submitted objects. It is possible to add more than one transform.

> **See Also:**
>
> – Subprograms for Retrieving Multiple Objects From the Database
> – Subprograms for Submitting XML to the Database
> – "SET_TRANSFORM_PARAM and SET_REMAP_PARAM Procedures" for information about how to modify and customize transform output

Syntax

```
DBMS_METADATA.ADD_TRANSFORM (
   handle       IN NUMBER,
```

```
   name         IN VARCHAR2,
   encoding     IN VARCHAR2 DEFAULT NULL,
   object_type  IN VARCHAR2 DEFAULT NULL)
 RETURN NUMBER;
```

Parameters

Table 122-4 ADD_TRANSFORM Function Parameters

Parameters	Description
handle	The handle returned from OPEN when this transform is used to retrieve objects. Or the handle returned from OPENW when this transform is used in the submission of XML metadata.
name	The name of the transform. The name can be an internal keyword like DDL to use internally stored stylesheets. If the name contains a colon, it is interpreted as directory_object_name:file_name of a user-supplied Extensible Stylesheet Language Transformation (XSLT) script. Otherwise, name designates a transform implemented by DBMS_METADATA. See Table 122-5 for descriptions of available transforms.
encoding	The name of the Globalization Support character set in which the stylesheet pointed to by name is encoded. This is only valid if name is a URL. If left NULL and the URL is external to the database, UTF-8 encoding is assumed. If left NULL and the URL is internal to the database (that is, it begins with /oradb/), then the encoding is assumed to be the database character set.
object_type	The definition of this parameter depends upon whether you are retrieving objects or submitting XML metadata. 1. When you use ADD_TRANFORM to retrieve objects, the following definition of object_type applies: Designates the object type to which the transform applies. (Note that this is an object type name, not a path name.) By default the transform applies to the object type of the OPEN handle. When the OPEN handle designates a heterogeneous object type, the following behavior can occur: • if object_type is omitted, the transform applies to all object types within the heterogeneous collection • if object_type is specified, the transform only applies to that specific object type within the collection If you omit this parameter you can add the DDL transform to all objects in a heterogeneous collection with a single call. If you supply this parameter, you can add a transform for a specific object type. 2. When you use ADD_TRANSFORM in the submission of XML metadata, this parameter is the object type to which the transform applies. By default, it is the object type of the OPENW handle. Because the OPENW handle cannot designate a heterogeneous object type, the caller would normally leave this parameter NULL in the ADD_TRANSFORM calls.

The following table describes the transforms available on the ADD_TRANSFORM function.

Because new transforms are occasionally added, you might want to query the DBMS_METADATA_TRANSFORMS view to see all valid Oracle-supplied transforms for specific object types.

Table 122-5 Transforms Available on ADD_TRANSFORM Function

Object Type	Transform Name	Input Doc Type	Output Doc Type	Description
All	DDL	XML	DDL	Convert XML to SQL to create the object
All	MODIFY	XML	XML	Modify XML document according to transform parameters
Subset	SXML	XML	SXML	Convert XML to SXML
Subset	MODIFYSXML	SXML	SXML	Modify SXML document according to transform parameters
Subset	SXMLDDL	SXML	DDL	Convert SXML to DDL
Subset	ALTERXML	SXML difference document	ALTER_XML	Generate ALTER_XML from SXML difference document. (See the DBMS_METADATA_DIFF PL/SQL package for more information about SXML difference format.) The following parameters are valid for the ALTERXML transform: • XPATH - The XPATH of the object being altered • NAME - Name of the object being altered • ALTERABLE - Affirms that the object can be altered. If the object cannot be altered, a NOT_ALTERABLE element is inserted whose value indicates the reason. • CLAUSE_TYPE - The type of clause (for example, ADD_COLUMN) • COLUMN_ATTRIBUTE - The attribute being modified • CONSTRAINT_TYPE - The type of constraint (for example, UNIQUE or PRIMARY)
Subset	ALTERDDL	ALTER_XML	ALTER_DDL	Convert ALTER_XML to ALTER_DDL

Return Values

The opaque handle that is returned is used as input to SET_TRANSFORM_PARAM and SET_REMAP_PARAM. Note that this handle is different from the handle returned by OPEN or OPENW; it refers to the transform, not the set of objects to be retrieved.

Usage Notes

- With no transforms added, objects are returned by default as XML documents. You call ADD_TRANSFORM to specify the XSLT stylesheets to be used to transform the returned XML documents.

- You can call ADD_TRANSFORM more than once to apply multiple transforms to XML documents. The transforms are applied in the order in which they were specified, the output of the first transform being used as input to the second, and so on.
- The output of a DDL transform is *not* an XML document. Therefore, no transform should be added after the DDL transform.
- Each transform expects a certain format XML document as input. If the input document is unspecified, metadata XML format is assumed.
- When the ALTERXML transform is used, parse items are returned in a PARSE_LIST element of the ALTER_XML document. Each PARSE_LIST_ITEM element contains an ITEM and a VALUE. For example:

```
<PARSE_LIST>
 <PARSE_LIST_ITEM>
  <ITEM>XPATH</ITEM>
  <VALUE>/sxml:TABLE/sxml:RELATIONAL_TABLE/sxml:COL_LIST/sxml:COL_LIST_ITEM[14]</
VALUE>
 </PARSE_LIST_ITEM>
 <PARSE_LIST_ITEM>
  <ITEM>NAME</ITEM>
  <VALUE>Z1</VALUE>
 </PARSE_LIST_ITEM>
 <PARSE_LIST_ITEM>
  <ITEM>CLAUSE_TYPE</ITEM>
  <VALUE>ADD_COLUMN</VALUE>
 </PARSE_LIST_ITEM>
 <PARSE_LIST_ITEM>
  <ITEM>COLUMN_ATTRIBUTE</ITEM>
  <VALUE>NOT_NULL</VALUE>
 </PARSE_LIST_ITEM>
</PARSE_LIST>
```

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INVALID_OPERATION. ADD_TRANSFORM was called after the first call to FETCH_xxx for the OPEN context. After the first call to FETCH_xxx is made, no further calls to ADD_TRANSFORM for the current OPEN context are permitted.
- INCONSISTENT_ARGS. The arguments are inconsistent. Possible inconsistencies include the following:
  - encoding is specified even though name is not a URL.
  - object_type is not part of the collection designated by handle.

122.6.2 CLOSE Procedure

This procedure is used for both retrieval and submission. This procedure invalidates the handle returned by `OPEN` (or `OPENW`) and cleans up the associated state.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database
> - Subprograms for Submitting XML to the Database

Syntax

```
DBMS_METADATA.CLOSE (
   handle  IN NUMBER);
```

Parameters

Table 122-6 CLOSE Procedure Parameters

Parameter	Description
`handle`	The handle returned from `OPEN` (or `OPENW`).

Usage Notes

> **Note:**
>
> The following notes apply only to object retrieval

You can prematurely terminate the stream of objects established by `OPEN` or (`OPENW`).

- If a call to `FETCH_xxx` returns `NULL`, indicating no more objects, a call to `CLOSE` is made transparently. In this case, you can still call `CLOSE` on the handle and not get an exception. (The call to `CLOSE` is not required.)
- If you know that only one specific object will be returned, you should explicitly call `CLOSE` after the single `FETCH_xxx` call to free resources held by the handle.

Exceptions

- `INVALID_ARGVAL`. The value for the `handle` parameter is `NULL` or invalid.

122.6.3 CONVERT Functions and Procedures

The `DBMS_METADATA.CONVERT` functions and procedures transform input XML documents.

The `CONVERT` functions return creation DDL. The `CONVERT` procedures return either XML or DDL, depending on the specified transforms.

> **See Also:**
>
> - Subprograms for Submitting XML to the Database

Syntax

The `CONVERT` functions are as follows:

```
DBMS_METADATA.CONVERT (
   handle   IN NUMBER,
   document IN sys.XMLType)
 RETURN sys.ku$_multi_ddls;

DBMS_METADATA.CONVERT (
  handle   IN NUMBER,
  document IN CLOB)
 RETURN sys.ku$_multi_ddls;
```

The `CONVERT` procedures are as follows:

```
DBMS_METADATA.CONVERT (
  handle   IN NUMBER,
  document IN sys.XMLType,
  result   IN OUT NOCOPY CLOB);

DBMS_METADATA.CONVERT (
  handle   IN NUMBER,
  document IN CLOB,
  result   IN OUT NOCOPY CLOB);
```

Parameters

Table 122-7 CONVERT Subprogram Parameters

Parameter	Description
`handle`	The handle returned from `OPENW`
`document`	The XML document containing object metadata of the type of the `OPENW` handle
result	The converted document

Return Values

Either XML or DDL, depending on the specified transforms.

Usage Notes

You can think of `CONVERT` as the second half of `FETCH_xxx`, either `FETCH_DDL` (for the function variants) or `FETCH_CLOB` (for the procedure variants). There are two differences:

- `FETCH_xxx` gets its XML document from the database, but `CONVERT` gets its XML document from the caller
- `FETCH_DDL` returns its results in a `sys.ku$_ddls` nested table, but `CONVERT` returns a `sys.ku$_multi_ddls` nested table

The transforms specified with `ADD_TRANSFORM` are applied in turn, and the result is returned to the caller. For the function variants, the DDL transform must be specified. If parse items were specified, they are returned in the `parsedItems` column. Parse items are ignored by the procedure variants.

The encoding of the XML document is embedded in its CLOB or XMLType representation. The version of the metadata is embedded in the XML. The generated DDL is valid for the database version specified in `OPENW`.

Exceptions

- `INVALID_ARGVAL`. A `NULL` or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- `INCONSISTENT_OPERATION`. No transform was specified. The DDL transform was not specified (function variants only).
- `INCOMPATIBLE_DOCUMENT`. The version of the XML document is not compatible with this version of the software.

122.6.4 FETCH_xxx Functions and Procedures

These functions and procedures return metadata for objects meeting the criteria established by `OPEN`, `SET_FILTER`, `SET_COUNT`, `ADD_TRANSFORM`, and so on.

See "Usage Notes" for the variants.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database

Syntax

The `FETCH` functions are as follows:

```
DBMS_METADATA.FETCH_XML (
   handle  IN NUMBER)
RETURN sys.XMLType;

DBMS_METADATA.FETCH_DDL (
   handle  IN NUMBER)
RETURN sys.ku$_ddls;

DBMS_METADATA.FETCH_CLOB (
   handle       IN NUMBER,
   cache_lob    IN BOOLEAN DEFAULT TRUE,
   lob_duration IN PLS INTEGER DEFAULT DBMS_LOB.SESSION)
RETURN CLOB;
```

The `FETCH` procedures are as follows:

```
DBMS_METADATA.FETCH_CLOB (
   handle  IN NUMBER,
   doc     IN OUT NOCOPY CLOB);

DBMS_METADATA.FETCH_XML_CLOB (
```

```
handle  IN NUMBER,
doc     IN OUT NOCOPY CLOB,
parsed_items OUT sys.ku$_parsed_items,
object_type_path OUT VARCHAR2);
```

Parameters

Table 122-8 FETCH_xxx Function Parameters

Parameters	Description
handle	The handle returned from OPEN.
cache_lob	TRUE=read LOB into buffer cache
lob_duration	The duration for the temporary LOB created by FETCH_CLOB, either DBMS_LOB.SESSION (the default) or DBMS_LOB.CALL.
doc	The metadata for the objects, or NULL if all objects have been returned.
parsed_items	A nested table containing the items specified by SET_PARSE_ITEM. If SET_PARSE_ITEM was not called, a NULL is returned.
object_type_path	For heterogeneous object types, this is the full path name of the object type for the objects returned by the call to FETCH_XXX. If handle designates a homogeneous object type, a NULL is returned.

Return Values

The metadata for the objects or NULL if all objects have been returned.

Usage Notes

These functions and procedures return metadata for objects meeting the criteria established by the call to OPEN that returned the handle, and subsequent calls to SET_FILTER, SET_COUNT, ADD_TRANSFORM, and so on. Each call to FETCH_xxx returns the number of objects specified by SET_COUNT (or less, if fewer objects remain in the underlying cursor) until all objects have been returned. After the last object is returned, subsequent calls to FETCH_xxx return NULL and cause the stream created by OPEN to be transparently closed.

There are several different FETCH_xxx functions and procedures:

- The FETCH_XML function returns the XML metadata for an object as an XMLType. It assumes that if any transform has been specified, that transform will produce an XML document. In particular, it assumes that the DDL transform has not been specified.
- The FETCH_DDL function returns the DDL (to create the object) in a sys.ku$_ddls nested table. It assumes that the DDL transform has been specified. Each row of the sys.ku$_ddls nested table contains a single DDL statement in the ddlText column; if requested, parsed items for the DDL statement will be returned in the parsedItems column. Multiple DDL statements may be returned under the following circumstances:
  - When you call SET_COUNT to specify a count greater than 1
  - When an object is transformed into multiple DDL statements. For example, A TYPE object that has a DDL transform applied to it can be transformed into both CREATE TYPE and CREATE TYPE BODY statements. A TABLE object can be transformed into a CREATE TABLE, and one or more ALTER TABLE statements
- The FETCH_CLOB function simply returns the object, transformed or not, as a CLOB. By default, the CLOB is read into the buffer cache and has session duration, but these defaults can be overridden with the cache_lob and lob_duration parameters.

- The FETCH_CLOB procedure returns the objects by reference in an IN OUT NOCOPY parameter. This is faster than the function variant, which returns LOBs by value, a practice that involves an expensive LOB copy.
- The FETCH_XML_CLOB procedure returns the XML metadata for the objects as a CLOB in an IN OUT NOCOPY parameter. This helps to avoid LOB copies, which can consume a lot of resources. It also returns a nested table of parse items and the full path name of the object type of the returned objects.
- All LOBs returned by FETCH_xxx are temporary LOBs. You must free the LOB. If the LOB is supplied as an IN OUT NOCOPY parameter, you must also create the LOB.
- If SET_PARSE_ITEM was called, FETCH_DDL and FETCH_XML_CLOB return attributes of the object's metadata (or the DDL statement) in a sys.ku$_parsed_items nested table. For FETCH_XML_CLOB, the nested table is an OUT parameter. For FETCH_DDL, it is a column in the returned sys.ku$_ddls nested table. Each row of the nested table corresponds to an item specified by SET_PARSE_ITEM and contains the following columns:
  - item—the name of the attribute as specified in the name parameter to SET_PARSE_ITEM.
  - value—the attribute value, or NULL if the attribute is not present in the DDL statement.
  - object-row—a positive integer indicating the object to which the parse item applies. If multiple objects are returned by FETCH_xxx, (because SET_COUNT specified a count greater than 1) then object_row=1 for all items for the first object, 2 for the second, and so on.
- The rows of the sys.ku$_parsed_items nested table are ordered by ascending object_row, but otherwise the row order is undetermined. To find a particular parse item within an object row the caller must search the table for a match on item.
- In general there is no guarantee that a requested parse item will be returned. For example, the parse item may not apply to the object type or to the particular line of DDL, or the item's value may be NULL.
- If SET_PARSE_ITEM was not called, NULL is returned as the value of the parsed items nested table.
- It is expected that the same variant of FETCH_xxx will be called for all objects selected by OPEN. That is, programs will not intermix calls to FETCH_XML, FETCH_DDL, FETCH_CLOB, and so on using the same OPEN handle. The effect of calling different variants is undefined; it might do what you expect, but there are no guarantees.
- Every object fetched will be internally consistent with respect to on-going DDL (and the subsequent recursive DML) operations against the dictionary. In some cases, multiple queries may be issued, either because the object type is heterogeneous or for performance reasons (for example, one query for heap tables, one for index-organized tables). Consequently the FETCH_xxx calls may in fact be fetches from different underlying cursors (meaning that read consistency is not guaranteed).

Caution:

Do not run functions and procedures within a query, because they can fail when the handle is for an open context for a heterogeneous object.

Exceptions

Most exceptions raised during execution of the query are propagated to the caller. Also, the following exceptions may be raised:

- `INVALID_ARGVAL`. A `NULL` or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- `INCONSISTENT_OPERATION`. Either `FETCH_XML` was called when the DDL transform had been specified, or `FETCH_DDL` was called when the DDL transform had *not* been specified.

122.6.5 GET_xxx Functions

`GET_xxx` functions let you fetch metadata for objects with a single call.

Functions

The `GET_xxx` functions are:

- `GET_XML`
- `GET_DDL`
- `GET_SXML`
- `GET_DEPENDENT_XML`
- `GET_DEPENDENT_DDL`
- `GET_GRANTED_XML`
- `GET_GRANTED_DDL`

> **See Also:**
>
> Subprograms for Retrieving Multiple Objects From the Database for more information about related subprograms

Syntax

```
DBMS_METADATA.GET_XML (
object_type     IN VARCHAR2,
name            IN VARCHAR2,
schema          IN VARCHAR2 DEFAULT NULL,
version         IN VARCHAR2 DEFAULT 'COMPATIBLE',
model           IN VARCHAR2 DEFAULT 'ORACLE',
transform       IN VARCHAR2 DEFAULT NULL)
RETURN CLOB;

DBMS_METADATA.GET_DDL (
object_type     IN VARCHAR2,
name            IN VARCHAR2,
schema          IN VARCHAR2 DEFAULT NULL,
version         IN VARCHAR2 DEFAULT 'COMPATIBLE',
model           IN VARCHAR2 DEFAULT 'ORACLE',
transform       IN VARCHAR2 DEFAULT 'DDL')
```

```
RETURN CLOB;

DBMS_METADATA.GET_SXML (
object_type     IN VARCHAR2,
name            IN VARCHAR2 DEFAULT NULL,
schema          IN VARCHAR2 DEFAULT NULL,
version         IN VARCHAR2 DEFAULT 'COMPATIBLE',
model           IN VARCHAR2 DEFAULT 'ORACLE',
transform       IN VARCHAR2 DEFAULT 'SXML')
RETURN CLOB;

DBMS_METADATA.GET_DEPENDENT_XML (
object_type        IN VARCHAR2,
base_object_name   IN VARCHAR2,
base_object_schema IN VARCHAR2 DEFAULT NULL,
version            IN VARCHAR2 DEFAULT 'COMPATIBLE',
model              IN VARCHAR2 DEFAULT 'ORACLE',
transform          IN VARCHAR2 DEFAULT NULL,
object_count       IN NUMBER   DEFAULT 10000)
RETURN CLOB;

DBMS_METADATA.GET_DEPENDENT_DDL (
object_type         IN VARCHAR2,
base_object_name    IN VARCHAR2,
base_object_schema  IN VARCHAR2 DEFAULT NULL,
version             IN VARCHAR2 DEFAULT 'COMPATIBLE',
model               IN VARCHAR2 DEFAULT 'ORACLE',
transform           IN VARCHAR2 DEFAULT 'DDL',
object_count        IN NUMBER   DEFAULT 10000)
RETURN CLOB;

DBMS_METADATA.GET_GRANTED_XML (
object_type     IN VARCHAR2,
grantee         IN VARCHAR2 DEFAULT NULL,
version         IN VARCHAR2 DEFAULT 'COMPATIBLE',
model           IN VARCHAR2 DEFAULT 'ORACLE',
transform       IN VARCHAR2 DEFAULT NULL,
object_count    IN NUMBER   DEFAULT 10000)
RETURN CLOB;

DBMS_METADATA.GET_GRANTED_DDL (
object_type     IN VARCHAR2,
grantee         IN VARCHAR2 DEFAULT NULL,
version         IN VARCHAR2 DEFAULT 'COMPATIBLE',
model           IN VARCHAR2 DEFAULT 'ORACLE',
transform       IN VARCHAR2 DEFAULT 'DDL',
object_count    IN NUMBER   DEFAULT 10000)
RETURN CLOB;
```

Parameters

Table 122-9 GET_xxx Function Parameters

Parameter	Description
object_type	The type of object that you want to be retrieved. This parameter takes the same values as the OPEN object_type parameter, except that it cannot be a heterogeneous object type. The attributes of the object type must be appropriate to the function. That is, for GET_*xxx* it must be a named object.
name	The object name. This object name is used internally in a NAME filter. (If the name is longer than 30 characters, then it will be used in a LONGNAME filter.) If this parameter is NULL, then no NAME or LONGNAME filter is specified. See SET_FILTER Procedure for a list of filters.
schema	The object schema. This object schema is used internally in a SCHEMA filter. The default is the current user.
version	The version of metadata that you want to be extracted. This parameter takes the same values as the OPEN version parameter.
model	The object model that you want to use. This parameter takes the same values as the OPEN model parameter.
transform	The name of a transformation on the output. This parameter takes the same values as the ADD_TRANSFORM *name* parameter. For GET_XML this name must not be DDL.
base_object_name	The base object name. The base object name is used internally in a BASE_OBJECT_NAME filter.
base_object_schema	The base object schema. The base object schema is used internally in a BASE_OBJECT_SCHEMA filter. The default is the current user.
grantee	The grantee. The grantee is used internally in a GRANTEE filter. The default is the current user.
object_count	The maximum number of objects to return. See SET_COUNT Procedure.

Return Values

The metadata for the specified object as XML or DDL.

Usage Notes

- These functions enable you to fetch metadata for objects with a single call. They encapsulate calls to OPEN, SET_FILTER, and so on. The function that you use depends on the characteristics of the object type, and on whether you want XML, SXML, or DDL.
  - GET_*xxx* is used to fetch named objects, especially schema objects (tables, views).
  - GET_DEPENDENT_*xxx* is used to fetch dependent objects (audits, object grants).
  - GET_GRANTED_*xxx* is used to fetch granted objects (system grants, role grants).
- For some object types, you can use more than one function. For example, you can use GET_*xxx* to fetch an index by name, or GET_DEPENDENT_*xxx* to fetch the same index by specifying the table on which it is defined.
- GET_*xxx* only returns a single named object.

- For GET_DEPENDENT_*xxx* and GET_GRANTED_*xxx*, an arbitrary number of dependent or granted objects can match the input criteria. You can specify an object count when fetching these objects. (The default count of 10000 should be adequate in most cases.)
- If the DDL transform is specified, then session-level transform parameters are inherited.
- If you start these functions from SQL*Plus, then to obtain complete, uninterrupted output, you should set the PAGESIZE to 0, and set LONG to some large number.
- Currently, GET_DDL for a USER object does not return the CONTAINER clause.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- OBJECT_NOT_FOUND. The specified object was not found in the database.

Examples

Example: Fetch the XML Representation of SCOTT.EMP

To generate complete, uninterrupted output, set the PAGESIZE to 0 and set LONG to some large number, as shown, before executing your query.

```
SET LONG 2000000
SET PAGESIZE 0
SELECT DBMS_METADATA.GET_XML('TABLE','EMP','SCOTT')
FROM DUAL;
```

Example: Fetch the DDL for all Complete Tables in the Current Schema, Filter Out Nested Tables and Overflow Segments

This example fetches the DDL for all "complete" tables in the current schema, filtering out nested tables and overflow segments. The example uses SET_TRANSFORM_PARAM (with the handle value = DBMS_METADATA.SESSION_TRANSFORM meaning "for the current session") to specify that storage clauses are not to be returned in the SQL DDL. Afterwards, the example resets the session-level parameters to their defaults.

To generate complete, uninterrupted output, set the PAGESIZE to 0 and set LONG to some large number, as shown, before running your query.

```
SET LONG 2000000
SET PAGESIZE 0
EXECUTE
DBMS_METADATA.SET_TRANSFORM_PARAM(DBMS_METADATA.SESSION_TRANSFORM,'STORAGE',fa
lse);
SELECT DBMS_METADATA.GET_DDL('TABLE',u.table_name)
     FROM USER_ALL_TABLES u
     WHERE u.nested='NO'
     AND (u.iot_type is null or u.iot_type='IOT');
EXECUTE
DBMS_METADATA.SET_TRANSFORM_PARAM(DBMS_METADATA.SESSION_TRANSFORM,'DEFAULT');
```

Example: Fetch the DDL For All Object Grants On HR.EMPLOYEES

```
SELECT DBMS_METADATA.GET_DEPENDENT_DDL('OBJECT_GRANT','EMPLOYEES','HR')
    FROM DUAL;
```

Example: Fetch the DDL For System Grants Granted To `SCOTT`

```
SELECT DBMS_METADATA.GET_GRANTED_DDL('SYSTEM_GRANT','SCOTT')
    FROM DUAL;
```

> **Note:**
>
> If you use the package `DBMS_METADATA.GET_GRANTED_DDL` to grant system privileges to a nonprivileged user granted system privileges, and you then attempt to use that nonprivileged user to grant privileges to other users by using `GRANT INHERIT PRIVILEGES`, then you receive the error "ORA-31703: cannot grant `INHERIT PRIVILEGES` privilege on behalf of other users". The `DBMS_METADATA.GET_GRANTED_DDL` procedure returns only privileges directly granted to the user.

Example: Fetch the DDL For All User `SCOTT` Object Grants on `HR.EMPLOYEES`

```
SQL> SELECT DBMS_METADATA.GET_GRANTED_DDL('OBJECT_GRANT','SCOTT')
     FROM dual ;
DBMS_METADATA.GET_GRANTED_DDL('OBJECT_GRANT','SCOTT')
```

Example: Fetch the DDL for the Data Use Case Domain `DAY_OF_WEEK`

```
SQL> SET LONG 2000000
SELECT DBMS_METADATA.GET_DDL('SQL_DOMAIN', 'DAY_OF_WEEK');
```

122.6.6 GET_QUERY Function

This function returns the text of the queries that are used by `FETCH_xxx`. This function assists in debugging.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database

Syntax

```
DBMS_METADATA.GET_QUERY (
   handle  IN NUMBER)
 RETURN VARCHAR2;
```

Parameters

Table 122-10 GET_QUERY Function Parameters

Parameter	Description
handle	The handle returned from OPEN. It cannot be the handle for a heterogeneous object type.

Return Values

The text of the queries that will be used by FETCH_xxx.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for the handle parameter.

122.6.7 OPEN Function

The DBMS_METADATA.OPEN function specifies the type of object to be retrieved, the version of its metadata, and the object model.

The return value is an opaque context handle for the set of objects to be used in subsequent calls.

> **See Also:**
>
> For more information about related subprograms, see:
>
> Subprograms for Retrieving Multiple Objects From the Database

Syntax

```
DBMS_METADATA.OPEN (
   object_type  IN VARCHAR2,
   version      IN VARCHAR2 DEFAULT 'COMPATIBLE',
   model        IN VARCHAR2 DEFAULT 'ORACLE',
   network_link IN VARCHAR2 DEFAULT NULL)
 RETURN NUMBER;
```

Parameters

Table 122-11 Open Function Parameters

Parameter	Description
object_type	The type of object to be retrieved. The table "DBMS_METADATA: Object Types" lists the valid object type names and their meanings. These object types are supported for the ORACLE model of metadata (see model in this table). The Attributes column in the DBMS_METADATA: Object Types table specifies some object type attributes: • Schema objects, such as tables, belong to schemas. • Named objects have unique names (if they are schema objects, the name is unique to the schema). • Dependent objects, such as indexes, are defined with reference to a base schema object. • Granted objects are granted or assigned to a user or role and therefore have a named grantee. • Heterogeneous object types denote a collection of related objects of different types. See the table "Object Types Returned for the Heterogeneous Object Type" for a listing of object types returned for the heterogeneous object type. These attributes are relevant when choosing object selection criteria. See "SET_FILTER Procedure" for more information.
version	The version of metadata to be extracted. Database objects or attributes that are incompatible with the version will not be extracted. Legal values for this parameter are as follows: COMPATIBLE (default)—the version of the metadata corresponds to the database compatibility level. LATEST—the version of the metadata corresponds to the database version. A specific database version. The value cannot be lower than 9.2.0.
model	Specifies which view to use, because the API can support multiple views on the metadata. Only the ORACLE model is supported.
network_link	The name of a database link to the database whose metadata is to be retrieved. If NULL (the default), then metadata is retrieved from the database on which the caller is running

The table Table 122-12 provides the name, meaning, attributes, and notes for the DBMS_METADATA package object types. In the attributes column, S represents a schema object, N represents a named object, D represents a dependent object, G represents a granted object, and H represents a heterogeneous object.

Table 122-12 DBMS_METADATA: Object Types

Type Name	Meaning	Attributes	Notes
AQ_QUEUE	queues	SND	Dependent on table
AQ_QUEUE_TABLE	additional metadata for queue tables	ND	Dependent on table
AQ_TRANSFORM	transforms	SN	None
ASSOCIATION	associate statistics	D	None

Table 122-12 (Cont.) DBMS_METADATA: Object Types

Type Name	Meaning	Attributes	Notes
AUDIT	audits of SQL statements	DG	Modeled as dependent, granted object. The base object name is the statement audit option name (for example, ALTER SYSTEM). There is no base object schema. The grantee is the user or proxy whose statements are audited.
AUDIT_OBJ	audits of schema objects	D	None
CLUSTER	clusters	SN	None
COMMENT	comments	D	None
CONSTRAINT	constraints	SND	Constraints do not include the following: • Primary key constraint for IOT • Column NOT NULL constraints • Certain REF SCOPE and WITH ROWID constraints for tables with REF columns
CONTEXT	application contexts	N	None
DATABASE_EXPORT	all metadata objects in a database	H	Corresponds to a full database export
DB_LINK	database links	SN	Modeled as schema objects because they have owners. For public links, the owner is PUBLIC. For private links, the creator is the owner.
DEFAULT_ROLE	default roles	G	Granted to a user by ALTER USER
DIMENSION	dimensions	SN	None
DIRECTORY	directories	N	None
FGA_POLICY	fine-grained audit policies	D	Not modeled as named object because policy names are not unique.
FUNCTION	stored functions	SN	None
INDEX_STATISTICS	precomputed statistics on indexes	D	The base object is the index's table.
INDEX	indexes	SND	None
INDEXTYPE	indextypes	SN	None
JAVA_SOURCE	Java sources	SN	None
JOB	jobs	S	None
LIBRARY	external procedure libraries	SN	None
MATERIALIZED_VIEW	materialized views	SN	None
MATERIALIZED_VIEW_LOG	materialized view logs	D	None
OBJECT_GRANT	object grants	DG	None
ON_USER_GRANT	Grants	G	Modeled as user grants. Grants the privileges of one user to other user in the form GRANT ... ON USER .... The grantee is the user. Example: GRANT INHERIT PRIVILEGES ON USER "USER1" TO "USER2".
OPERATOR	operators	SN	None

Table 122-12 (Cont.) DBMS_METADATA: Object Types

Type Name	Meaning	Attributes	Notes
PACKAGE	stored packages	SN	By default, both package specification and package body are retrieved. See "SET_FILTER Procedure".
PACKAGE_SPEC	package specifications	SN	None
PACKAGE_BODY	package bodies	SN	None
PROCEDURE	stored procedures	SN	None
PROFILE	profiles	N	None
PROXY	proxy authentications	G	Granted to a user by ALTER USER
REF_CONSTRAINT	referential constraint	SND	None
REFRESH_GROUP	refresh groups	SN	None
RESOURCE_COST	resource cost info	H	None
RLS_CONTEXT	driving contexts for enforcement of fine-grained access-control policies	D	Corresponds to the DBMS_RLS.ADD_POLICY_CONTENT procedure
RLS_GROUP	fine-grained access-control policy groups	D	Corresponds to the DBMS_RLS.CREATE_GROUP procedure
RLS_POLICY	fine-grained access-control policies	D	Corresponds to DBMS_RLS.ADD_GROUPED_POLICY. Not modeled as named objects because policy names are not unique.
RMGR_CONSUMER_GROUP	resource consumer groups	SN	Oracle Data Pump does not use these object types. Instead, it exports resource manager objects as procedural objects.
RMGR_INTITIAL_CONSUMER_GROUP	assign initial consumer groups to users	G	None
RMGR_PLAN	resource plans	SN	None
RMGR_PLAN_DIRECTIVE	resource plan directives	D	Dependent on resource plan
ROLE	roles	N	None
ROLE_GRANT	role grants	G	None
ROLLBACK_SEGMENT	rollback segments	N	None
SCHEMA_EXPORT	all metadata objects in a schema	H	Corresponds to user-mode export.
SEQUENCE	sequences	SN	None
SQL_DOMAIN	Data Use Case Domains	SN	A **Data Use Case Domain** is a high-level dictionary object that belongs to a schema and encapsulates a set of properties and constraints. The attributes and constraints are defined and managed only for the domain, and are automatically applied to all the columns of the given domain.

Table 122-12 (Cont.) DBMS_METADATA: Object Types

Type Name	Meaning	Attributes	Notes
SYNONYM	synonyms	See notes	Private synonyms are schema objects. Public synonyms are not, but for the purposes of this API, their schema name is PUBLIC. The name of a synonym is considered to be the synonym itself. For example, in CREATE PUBLIC SYNONYM FOO FOR BAR, the resultant object is considered to have name FOO and schema PUBLIC.
SYSTEM_GRANT	system privilege grants	G	None
TABLE	tables	SN	None
TABLE_DATA	metadata describing row data for a table, nested table, or partition	SND	For partitions, the object name is the partition name. For nested tables, the object name is the storage table name. The base object is the top-level table to which the table data belongs. For nested tables and partitioning, this is the top-level table (*not* the parent table or partition). For nonpartitioned tables and non-nested tables this is the table itself.
TABLE_EXPORT	metadata for a table and its associated objects	H	Corresponds to table-mode export
TABLE_STATISTICS	precomputed statistics on tables	D	None
TABLESPACE	tablespaces	N	None
TABLESPACE_QUOTA	tablespace quotas	G	Granted with ALTER USER
TRANSPORTABLE_EXPORT	metadata for objects in a transportable tablespace set	H	Corresponds to transportable tablespace export
TRIGGER	triggers	SND	None
TRUSTED_DB_LINK	trusted links	N	None
TYPE	user-defined types	SN	By default, both type and type body are retrieved. See "SET_FILTER Procedure". .
TYPE_SPEC	type specifications	SN	None
TYPE_BODY	type bodies	SN	None
USER	users	N	None
VIEW	views	SN	None
XMLSCHEMA	XML schema	SN	The object's name is its URL (which can be longer than 30 characters). Its schema is the user who registered it.
XS_USER	Real Application Security (RAS) user	N	Corresponds to RAS users
XS_ROLE	Real Application Security (RAS) role	N	Corresponds to RAS roles
XS_ROLESET	Real Application Security (RAS) rolesets	N	Corresponds to RAS rolesets
XS_ROLE_GRANT	Real Application Security (RAS) role grants	N	Corresponds to RAS role grants

Table 122-12 (Cont.) DBMS_METADATA: Object Types

Type Name	Meaning	Attributes	Notes
XS_SECURITY_CLASS	Real Application Security (RAS) security class	SN	Corresponds to RAS security classes
XS_DATA_SECURITY	Real Application Security (RAS) data security policy	SN	Corresponds to RAS data security policies
XS_ACL	Real Application Security (RAS) ACL	SN	Corresponds to RAS access control lists (ACLs) and associated access control entries (ACEs)
XS_ACL_PARAM	Real Application Security (RAS) ACL parameter	N	Corresponds to RAS access control lists (ACL) parameters
XS_NAMESPACE	Real Application Security (RAS) namespace	N	Corresponds to RAS namespaces.

Table Table 122-13 lists the types of objects returned for the major heterogeneous object types. For SCHEMA_EXPORT, certain object types are only returned if the INCLUDE_USER filter is specified at TRUE. In the table, such object types are marked INCLUDE_USER.

Table 122-13 Object Types Returned for the Heterogeneous Object Type

Object Type	DATABASE_EXPORT	SCHEMA_EXPORT	TABLE_EXPORT	TRANSPORTABLE_EXPORT
ASSOCIATION	Yes	No	No	No
AUDIT	Yes	No	No	No
AUDIT_OBJ	Yes	Yes	Yes	Yes
CLUSTER	Yes	Yes	No	Yes
COMMENT	Yes	Yes	Yes	Yes
CONSTRAINT	Yes	Yes	Yes	Yes
CONTEXT	Yes	No	No	No
DB_LINK	Yes	Yes	No	No
DEFAULT_ROLE	Yes	INCLUDE_USER	No	No
DIMENSION	Yes	Yes	No	No
DIRECTORY	Yes	No	No	No
FGA_POLICY	Yes	No	No	Yes
FUNCTION	Yes	Yes	No	No
INDEX_STATISTICS	Yes	Yes	Yes	Yes
INDEX	Yes	Yes	Yes	Yes
INDEXTYPE	Yes	Yes	No	No
JAVA_SOURCE	Yes	Yes	No	No
JOB	Yes	Yes	No	No
LIBRARY	Yes	Yes	No	No
MATERIALIED_VIEW	Yes	Yes	No	No
MATERIALIZED_VIEW_LOG	Yes	Yes	No	No
OBJECT_GRANT	Yes	Yes	Yes	Yes

Table 122-13 (Cont.) Object Types Returned for the Heterogeneous Object Type

Object Type	DATABASE_EXPORT	SCHEMA_EXPORT	TABLE_EXPORT	TRANSPORTABLE_EXPORT
OPERATOR	Yes	Yes	No	No
PACKAGE	Yes	Yes	No	No
PACKAGE_SPEC	Yes	Yes	No	No
PACKAGE_BODY	Yes	Yes	No	No
PASSWORD_HISTORY	Yes	INCLUDE_USER	No	No
PASSWORD_VERIFY_FUNCTION	Yes	No	No	No
PROCEDURE	Yes	Yes	No	No
PROFILE	Yes	No	No	No
PROXY	Yes	No	No	No
REF_CONSTRAINT	Yes	Yes	Yes	Yes
REFRESH_GROUP	Yes	Yes	No	No
RESOURCE_COST	Yes	No	No	No
RLS_CONTEXT	Yes	No	No	Yes
RLS_GROUP	Yes	No	No	Yes
RLS_POLICY	Yes	Table data is retrieved according to policy	Table data is retrieved according to policy	Yes
ROLE	Yes	No	No	No
ROLE_GRANT	Yes	No	No	No
ROLLBACK_SEGMENT	Yes	No	No	No
SEQUENCE	Yes	Yes	No	No
SYNONYM	Yes	Yes	No	No
SYSTEM_GRANT	Yes	INCLUDE_USER	No	No
TABLE	Yes	Yes	Yes	Yes
TABLE_DATA	Yes	Yes	Yes	Yes
TABLE_STATISTICS	Yes	Yes	Yes	Yes
TABLESPACE	Yes	No	No	No
TABLESPACE_QUOTA	Yes	INCLUDE_USER	No	No
TRIGGER	Yes	Yes	Yes	Yes
TRUSTED_DB_LINK	Yes	No	No	No
TYPE	Yes	Yes	No	Yes, if the types are used by tables in the transportable set
TYPE_SPEC	Yes	Yes	No	Yes, if the types are used by tables in the transportable set

Table 122-13 (Cont.) Object Types Returned for the Heterogeneous Object Type

Object Type	DATABASE_EXPORT	SCHEMA_EXPORT	TABLE_EXPORT	TRANSPORTABLE_EXPORT
TYPE_BODY	Yes	Yes	No	Yes, if the types are used by tables in the transportable set
USER	Yes	INCLUDE_USER	No	No
VIEW	Yes	Yes	No	No
XMLSCHEMA	Yes	Yes	No	No

Return Values

An opaque handle to the class of objects. This handle is used as input to SET_FILTER, SET_COUNT, ADD_TRANSFORM, GET_QUERY, SET_PARSE_ITEM, FETCH_xxx, and CLOSE.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INVALID_OBJECT_PARAM. The version or model parameter was not valid for the object_type.

122.6.8 OPENW Function

This function specifies the type of object to be submitted and the object model. The return value is an opaque context handle.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Submitting XML to the Database

Syntax

```
DBMS_METADATA.OPENW
  (object_type  IN VARCHAR2,
  version       IN VARCHAR2 DEFAULT 'COMPATIBLE',
  model         IN VARCHAR2 DEFAULT 'ORACLE')
 RETURN NUMBER;
```

Parameters

Table 122-14 OPENW Function Parameters

Parameter	Description
object_type	The type of object to be submitted. Valid types names and their meanings are listed in Table 122-12. The type cannot be a heterogeneous object type.

Table 122-14 (Cont.) OPENW Function Parameters

Parameter	Description
version	The version of DDL to be generated by the CONVERT function. DDL clauses that are incompatible with the version will not be generated. The legal values for this parameter are as follows: • COMPATIBLE - This is the default. The version of the DDL corresponds to the database compatibility level. Database compatibility must be set to 9.2.0 or higher. • LATEST - The version of the DDL corresponds to the database version. • A specific database version. The value cannot be lower than 9.2.0.
model	Specifies which view to use. Only the Oracle proprietary (ORACLE) view is supported by DBMS_METADATA.

Return Values

An opaque handle to write context. This handle is used as input to the ADD_TRANSFORM, CONVERT, PUT, and CLOSE procedures.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INVALID_OBJECT_PARAM. The model parameter was not valid for the object_type.

122.6.9 PUT Function

This function submits an XML document containing object metadata to the database to create the object.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Submitting XML to the Database

Syntax

```
DBMS_METADATA.PUT (
   handle     IN             NUMBER,
   document   IN             sys.XMLType,
   flags      IN             NUMBER,
   results    IN OUT NOCOPY  sys.ku$_SubmitResults)
  RETURN BOOLEAN;

DBMS_METADATA.PUT (
   handle     IN             NUMBER,
   document   IN             CLOB,
   flags      IN             NUMBER,
   results    IN OUT NOCOPY  sys.ku$_SubmitResults)
 RETURN BOOLEAN;
```

Parameters

Table 122-15 PUT Function Parameters

Parameter	Description
handle	The handle returned from OPENW.
document	The XML document containing object metadata for the type of the OPENW handle.
flags	Reserved for future use
results	Detailed results of the operation.

Return Values

TRUE if all SQL operations succeeded; FALSE if there were any errors.

Usage Notes

The PUT function converts the XML document to DDL just as CONVERT does (applying the specified transforms in turn) and then submits each resultant DDL statement to the database. As with CONVERT, the DDL transform must be specified. The DDL statements and associated parse items are returned in the sys.ku$_SubmitResults nested table. With each DDL statement is a nested table of error lines containing any errors or exceptions raised by the statement.

The encoding of the XML document is embedded in its CLOB or XMLType representation. The version of the metadata is embedded in the XML. The generated DDL is valid for the database version specified in OPENW.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INCONSISTENT_OPERATION. The DDL transform was not specified.
- INCOMPATIBLE_DOCUMENT. The version of the XML document is not compatible with this version of the software.

122.6.10 SET_COUNT Procedure

This procedure specifies the maximum number of objects to be retrieved in a single FETCH_xxx call.

By default, each call to FETCH_xxx returns one object. You can use the SET_COUNT procedure to override this default. If FETCH_xxx is called from a client, specifying a count value greater than 1 can result in fewer server round trips and, therefore, improved performance.

For heterogeneous object types, a single FETCH_xxx operation only returns objects of a single object type.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database

Syntax

```
DBMS_METADATA.SET_COUNT (
   handle           IN NUMBER,
   value            IN NUMBER,
   object_type_path IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 122-16 SET_COUNT Procedure Parameters

Parameter	Description
`handle`	The handle returned from `OPEN`.
`value`	The maximum number of objects to retrieve.
`object_type_path`	A path name designating the object types to which the count value applies. By default, the count value applies to the object type of the `OPEN` handle. When the `OPEN` handle designates a heterogeneous object type, behavior can be either of the following: • If `object_type_path` is omitted, then the count applies to all object types within the heterogeneous collection. • If `object_type_path` is specified, then the count only applies to the specific node (or set of nodes) within the tree of object types forming the heterogeneous collection.

Exceptions

- `INVALID_ARGVAL`. A `NULL` or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- `INVALID_OPERATION`. `SET_COUNT` was called after the first call to `FETCH_xxx` for the `OPEN` context. After the first call to `FETCH_xxx` is made, no further calls to `SET_COUNT` for the current `OPEN` context are permitted.
- `INCONSISTENT_ARGS`. `object_type` parameter is not consistent with handle.

122.6.11 SET_FILTER Procedure

This procedure specifies restrictions on the objects to be retrieved, for example, the object name or schema.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database

Syntax

```
DBMS_METADATA.SET_FILTER (
   handle           IN NUMBER,
   name             IN VARCHAR2,
   value            IN VARCHAR2,
   object_type_path IN VARCHAR2 DEFAULT NULL);

DBMS_METADATA.SET_FILTER (
   handle            IN NUMBER,
   name              IN VARCHAR2,
   value             IN BOOLEAN DEFAULT TRUE,
   object_type_path  IN VARCHAR2 DEFAULT NULL);

DBMS_METADATA.SET_FILTER (
   handle            IN NUMBER,
   name              IN VARCHAR2,
   value             IN NUMBER,
   object_type_path  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 122-17 SET_FILTER Procedure Parameters

Parameter	Description
handle	The handle returned from OPEN.
name	The name of the filter. For each filter, Table 122-18 lists the object_type it applies to, its name, its datatype (text or Boolean), and its meaning or effect (including its default value, if any). The Datatype column of Table 122-18 also indicates whether a text filter is an expression filter. An expression filter is the right-hand side of a SQL comparison (that is, a SQL comparison operator (=, !=, and so on.)) and the value compared against. The value must contain parentheses and quotation marks where appropriate. Note that in PL/SQL and SQL*Plus, two single quotes (*not* a double quote) are needed to represent an apostrophe. For example, an example of a NAME_EXPR filter in PL/SQL is as follows: 'IN (''DEPT'',''EMP'')' The filter value is combined with a particular object attribute to produce a WHERE condition in the query that fetches the objects. In the preceding example, the filter is combined with the attribute corresponding to an object name; objects named 'DEPT' and 'EMP' are selected.
value	The value of the filter. Text, Boolean, and Numeric filters are supported.
object_type_path	A path name designating the object types to which the filter applies. By default, the filter applies to the object type of the OPEN handle. When the OPEN handle designates a heterogeneous object type, you can use this parameter to specify a filter for a specific node or set of nodes within the tree of object types that form the heterogeneous collection. See Table 122-19 for a listing of some of the values for this parameter.

Table 122-18 describes the object type, name, datatype, and meaning of the filters available with the SET_FILTER procedure.

Table 122-18 SET_FILTER: Filters

Object Type	Name	Datatype	Meaning
Named objects	NAME	Text	Objects with this exact name are selected.
Named objects	NAME_EXPR	Text expression	The filter value is combined with the object attribute corresponding to the object name to produce a WHERE condition in the query that fetches the objects. By default, all named objects of object_type are selected.
Named objects	EXCLUDE_NAME_EXPR	Text expression	The filter value is combined with the attribute corresponding to the object name to specify objects that are to be excluded from the set of objects fetched. By default, all named objects of the object type are selected.
TABLE_EXPORT	EXCLUDE_TABLES	Boolean	If TRUE, all paths associated with tables are excluded from the set of objects fetched. If FALSE (the default), all paths associated with tables are fetched.
Schema objects	SCHEMA	Text	Objects in this schema are selected. If the object type is SYNONYM, then specify PUBLIC to select public synonyms.
Schema objects	SCHEMA_EXPR	Text expression	The filter value is combined with the attribute corresponding to the object's schema. The default is determined as follows: - if BASE_OBJECT_SCHEMA is specified, then objects in that schema are selected; - otherwise, objects in the current schema are selected.
PACKAGE, TYPE	SPECIFICATION	Boolean	If TRUE, retrieve the package or type specification. Defaults to TRUE.
PACKAGE, TYPE	BODY	Boolean	If TRUE, retrieve the package or type body. Defaults to TRUE.
TABLE, CLUSTER, INDEX, TABLE_DATA, TABLE_EXPORT, TRANSPORTABLE_EXPORT	TABLESPACE	Text	Objects in this tablespace (or having a partition in this tablespace) are selected.
TABLE, CLUSTER, INDEX,TABLE_DATA, TABLE_EXPORT, TRANSPORTABLE_EXPORT	TABLESPACE_EXPR	Text expression	The filter value is combined with the attribute corresponding to the object's tablespace (or in the case of a partitioned table or index, the partition's tablespaces). By default, objects in all tablespaces are selected.
TABLE, objects dependent on tables	PRIMARY	Boolean	If TRUE, retrieve primary tables (that is, tables for which the secondary object bit in obj$ is clear. Defaults to TRUE.

Table 122-18 (Cont.) SET_FILTER: Filters

Object Type	Name	Datatype	Meaning
TABLE, objects dependent on tables	SECONDARY	Boolean	If TRUE, retrieve secondary tables (that is, tables for which the secondary object bit in obj$ is set). Defaults to TRUE.
Dependent Objects	BASE_OBJECT_NAME	Text	Objects are selected that are defined or granted on objects with this name. Specify SCHEMA for triggers on schemas. Specify DATABASE for database triggers. Column-level comments cannot be selected by column name; the base object name must be the name of the table, view, or materialized view containing the column.
Dependent Objects	BASE_OBJECT_SCHEMA	Text	Objects are selected that are defined or granted on objects in this schema. If BASE_OBJECT_NAME is specified with a value other than SCHEMA or DATABASE, this defaults to the current schema.
Dependent Objects	BASE_OBJECT_NAME_EXPR	Text expression	The filter value is combined with the attribute corresponding to the name of the base object. Not valid for schema and database triggers.
Dependent Objects	EXCLUDE_BASE_OBJECT_NAME_EXPR	Text expression	The filter value is combined with the attribute corresponding to the name of the base object to specify objects that are to be excluded from the set of objects fetched. Not valid for schema and database triggers.
Dependent Objects	BASE_OBJECT_SCHEMA_EXPR	Text expression	The filter value is combined with the attribute corresponding to the schema of the base object.
Dependent Objects	BASE_OBJECT_TYPE	Text	The object type of the base object.
Dependent Objects	BASE_OBJECT_TYPE_EXPR	Text expression	The filter value is combined with the attribute corresponding to the object type of the base object. By default no filtering is done on object type.
Dependent Objects	BASE_OBJECT_TABLESPACE	Text	The tablespace of the base object.
Dependent Objects	BASE_OBJECT_TABLESPACE_EXPR	Text expression	The filter value is combined with the attribute corresponding to the tablespaces of the base object. By default, no filtering is done on the tablespace.
INDEX, TRIGGER	SYSTEM_GENERATED	Boolean	If TRUE, select indexes or triggers even if they are system-generated. If FALSE, omit system-generated indexes or triggers. Defaults to TRUE.
Granted Objects	GRANTEE	Text	Objects are selected that are granted to this user or role. Specify PUBLIC for grants to PUBLIC.
Granted Objects	PRIVNAME	Text	The name of the privilege or role to be granted. For TABLESPACE_QUOTA, only UNLIMITED can be specified.
Granted Objects	GRANTEE_EXPR	Text expression	The filter value is combined with the attribute corresponding to the grantee name.

Table 122-18 (Cont.) SET_FILTER: Filters

Object Type	Name	Datatype	Meaning
Granted Objects	EXCLUDE_GRANTEE_EXPR	Text expression	The filter value is combined with the attribute corresponding to the grantee name to specify objects that are to be excluded from the set of objects fetched.
OBJECT_GRANT	GRANTOR	Text	Object grants are selected that are granted by this user.
SYNONYM, JAVA_SOURCE, XMLSCHEMA	LONGNAME	Text	A name longer than 30 characters. Objects with this exact name are selected. If the object name is 30 characters or less, the NAME filter must be used.
SYNONYM, JAVA_SOURCE, XMLSCHEMA	LONGNAME_EXPR	Text	The filter value is combined with the attribute corresponding to the object's long name. By default, no filtering is done on the long name of an object.
All objects	CUSTOM_FILTER	Text	The text of a WHERE condition. The condition is appended to the query that fetches the objects. By default, no custom filter is used. The other filters are intended to meet the needs of the majority of users. Use CUSTOM_FILTER when no defined filters exists for your purpose. Of necessity such a filter depends on the detailed structure of the UDTs and views used in the query. Because filters may change from version to version, upward compatibility is not guaranteed.
All objects	EDITION	Text	The edition filter is accepted for any object type, but affects only objects that support editions. The filter is only accepted for local objects (that is, the network_link parameter is not specified in the OPEN call). The edition name must be a valid edition name. If an edition is not specified, the edition of the active session is used.
SCHEMA_EXPORT	SCHEMA	Text	The schema whose objects are selected.
SCHEMA_EXPORT	SCHEMA_EXPR	Text expression	The filter value is either: combined with the attribute corresponding to a schema name to produce a WHERE condition in the query that fetches schema objects, combined with the attribute corresponding to a base schema name to produce a WHERE condition in the query that fetches dependent objects. By default the current user's objects are selected.
SCHEMA_EXPORT	INCLUDE_USER	Boolean	If TRUE, retrieve objects containing privileged information about the user. For example, USER, PASSWORD_HISTORY, TABLESPACE_QUOTA. Defaults to FALSE.
TABLE_EXPORT	SCHEMA	Text	Objects (tables and their dependent objects) in this schema are selected.

Table 122-18 (Cont.) SET_FILTER: Filters

Object Type	Name	Datatype	Meaning
TABLE_EXPORT	SCHEMA_EXPR	Text expression	The filter value is either: combined with the attribute corresponding to a schema name to produce a WHERE condition in the query that fetches the tables, combined with the attribute corresponding to a base schema name to produce a WHERE condition in the query that fetches the tables' dependent objects. By default the current user's objects are selected.
TABLE_EXPORT	NAME	Text	The table with this exact name is selected along with its dependent objects.
TABLE_EXPORT	NAME_EXPR	Text expression	The filter value is combined with the attribute corresponding to a table name in the queries that fetch tables and their dependent objects. By default all tables in the selected schemas are selected, along with their dependent objects.
Heterogeneous objects	BEGIN_WITH	Text	The fully qualified path name of the first object type in the heterogeneous collection to be retrieved. Objects normally fetched prior to this object type will not be retrieved.
Heterogeneous objects	BEGIN_AFTER	Text	The fully qualified path name of an object type after which the heterogeneous retrieval should begin. Objects of this type will not be retrieved, nor will objects normally fetched prior to this object type.
Heterogeneous objects	END_BEFORE	Text	The fully qualified path name of an object type where the heterogeneous retrieval should end. Objects of this type will not be retrieved, nor will objects normally fetched after this object type.
Heterogeneous objects	END_WITH	Text	The fully qualified path name of the last object type in the heterogeneous collection to be retrieved. Objects normally fetched after this object type will not be retrieved.

Table 122-18 (Cont.) SET_FILTER: Filters

Object Type	Name	Datatype	Meaning
Heterogeneous objects	INCLUDE_PATH_EXPR, EXCLUDE_PATH_EXPR	Text expression	For these two filters, the filter value is combined with the attribute corresponding to an object type path name to produce a WHERE condition in the query that fetches the object types belonging to the heterogeneous collection. Objects of types satisfying this condition are included (INCLUDE_PATH_EXPR) or excluded (EXCLUDE_PATH_EXPR) from the set of object types fetched. Path names in the filter value do not have to be fully qualified. See Table 122-19 for valid path names that can be used with these filters. BEGIN_WITH, BEGIN_AFTER, END_BEFORE, END_WITH, INCLUDE_PATH_EXPR, and EXCLUDE_PATH_EXPR all restrict the set of object types in the heterogeneous collection. By default, objects of all object types in the heterogeneous collection are retrieved.

Usage Notes

- Each call to SET_FILTER causes a WHERE condition to be added to the underlying query that fetches the set of objects. The WHERE conditions are concatenated with the AND keyword so that you can use multiple SET_FILTER calls to refine the set of objects to be returned. For example to specify that you want the object named EMP in schema SCOTT, do the following:

```
SET_FILTER(handle,'SCHEMA','SCOTT');
SET_FILTER(handle,'NAME','EMP');
```

- You can use the same text expression filter multiple times with different values. All the filter conditions will be applied to the query. For example, to get objects with names between Felix and Oscar, do the following:

```
SET_FILTER(handle,'NAME_EXPR','>=''FELIX''');
SET_FILTER(handle,'NAME_EXPR','<=''OSCAR''');
```

- With SET_FILTER, you can specify the schema of objects to be retrieved, but security considerations may override this specification. If the caller is SYS or has the SELECT_CATALOG_ROLE role, then any object can be retrieved; otherwise, only the following can be retrieved:
  - Schema objects owned by the current user
  - Public synonyms
  - System privileges granted to the current user or to PUBLIC
  - Grants on objects for which the current user is owner, grantor, or grantee (either explicitly or as PUBLIC).
  - SCHEMA_EXPORT where the name is the current user
  - TABLE_EXPORT where SCHEMA is the current user

  If you request objects that you are not privileged to retrieve, no exception is raised; the object is not retrieved, as if it did not exist.

In stored procedures, functions, and definers-rights packages, roles (such as SELECT_CATALOG_ROLE) are disabled. Therefore, such a PL/SQL program can only fetch metadata for objects in its own schema. If you want to write a PL/SQL program that fetches metadata for objects in a different schema (based on the invoker's possession of SELECT_CATALOG_ROLE), you must make the program invokers-rights.

- For heterogeneous object types, the BEGIN_WITH and BEGIN_AFTER filters allow restart on an object type boundary. Appropriate filter values are returned by the FETCH_XML_CLOB procedure.

  Filters on heterogeneous objects provide default values for filters on object types within the collection. You can override this default for a particular object type by specifying the appropriate filter for the specific object type path. For example, for SCHEMA_EXPORT the NAME filter specifies the schema to be fetched including all the tables in the schema, but you can further restrict this set of tables by supplying a NAME_EXPR filter explicitly for the TABLE object type path. Table 122-19 lists valid object type path names for the major heterogeneous object types along with an explanation of the scope of each path name. (The same information is available in the following catalog views: DATABASE_EXPORT_OBJECTS, SCHEMA_EXPORT_OBJECTS, and TABLE_EXPORT_OBJECTS.) See Table 122-18 for filters defined for each path name. These path names are valid in the INCLUDE_PATH_EXPR and EXCLUDE_PATH_EXPR filters. Path names marked with an asterisk (*) are *only* valid in those filters; they cannot be used as values of the SET_FILTER object_type_path parameter.

Table 122-19 Object Type Path Names for Heterogeneous Object Types

Heterogeneous Type	Path Name (*=valid only in xxx_PATH_EXPR)	Scope
TABLE_EXPORT	AUDIT_OBJ	Object audits on the selected tables
TABLE_EXPORT	COMMENT	Table and column comments for the selected tables
TABLE_EXPORT	CONSTRAINT	Constraints (including referential constraints) on the selected tables
TABLE_EXPORT	*GRANT	Object grants on the selected tables
TABLE_EXPORT	INDEX	Indexes (including domain indexes) on the selected tables
TABLE_EXPORT	OBJECT_GRANT	Object grants on the selected tables
TABLE_EXPORT	REF_CONSTRAINT	Referential (foreign key) constraints on the selected tables
TABLE_EXPORT	STATISTICS	Statistics on the selected tables
TABLE_EXPORT	TABLE_DATA	Row data for the selected tables
TABLE_EXPORT	TRIGGER	Triggers on the selected tables
SCHEMA_EXPORT	ASSOCIATION	Statistics type associations for objects in the selected schemas
SCHEMA_EXPORT	AUDIT_OBJ	Audits on all objects in the selected schemas
SCHEMA_EXPORT	CLUSTER	Clusters in the selected schemas and their indexes
SCHEMA_EXPORT	COMMENT	Comments on all objects in the selected schemas
SCHEMA_EXPORT	CONSTRAINT	Constraints (including referential constraints) on all objects in the selected schemas
SCHEMA_EXPORT	DB_LINK	Private database links in the selected schemas
SCHEMA_EXPORT	DEFAULT_ROLE	Default roles granted to users associated with the selected schemas
SCHEMA_EXPORT	DIMENSION	Dimensions in the selected schemas

Table 122-19 (Cont.) Object Type Path Names for Heterogeneous Object Types

Heterogeneous Type	Path Name (*=valid only in xxx_PATH_EXPR)	Scope
SCHEMA_EXPORT	FUNCTION	Functions in the selected schemas and their dependent grants and audits
SCHEMA_EXPORT	*GRANT	Grants on objects in the selected schemas
SCHEMA_EXPORT	INDEX	Indexes (including domain indexes) on tables and clusters in the selected schemas
SCHEMA_EXPORT	INDEXTYPE	Indextypes in the selected schemas and their dependent grants and audits
SCHEMA_EXPORT	JAVA_SOURCE	Java sources in the selected schemas and their dependent grants and audits
SCHEMA_EXPORT	JOB	Jobs in the selected schemas
SCHEMA_EXPORT	LIBRARY	External procedure libraries in the selected schemas
SCHEMA_EXPORT	MATERIALIZED_VIEW	Materialized views in the selected schemas
SCHEMA_EXPORT	MATERIALIZED_VIEW_LOG	Materialized view logs on tables in the selected schemas
SCHEMA_EXPORT	OBJECT_GRANT	Grants on objects in the selected schemas
SCHEMA_EXPORT	OPERATOR	Operators in the selected schemas and their dependent grants and audits
SCHEMA_EXPORT	PACKAGE	Packages (both specification and body) in the selected schemas, and their dependent grants and audits
SCHEMA_EXPORT	PACKAGE_BODY	Package bodies in the selected schemas
SCHEMA_EXPORT	PACKAGE_SPEC	Package specifications in the selected schemas
SCHEMA_EXPORT	PASSWORD_HISTORY	The password history for users associated with the selected schemas
SCHEMA_EXPORT	PROCEDURE	Procedures in the selected schemas and their dependent grants and audits
SCHEMA_EXPORT	REF_CONSTRAINT	Referential (foreign key) constraints on tables in the selected schemas
SCHEMA_EXPORT	REFRESH_GROUP	Refresh groups in the selected schemas
SCHEMA_EXPORT	SEQUENCE	Sequences in the selected schemas and their dependent grants and audits
SCHEMA_EXPORT	STATISTICS	Statistics on tables and indexes in the selected schemas
SCHEMA_EXPORT	SYNONYM	Private synonyms in the selected schemas
SCHEMA_EXPORT	TABLE	Tables in the selected schemas and their dependent objects (indexes, constraints, triggers, grants, audits, comments, table data, and so on)
SCHEMA_EXPORT	TABLE_DATA	Row data for tables in the selected schemas
SCHEMA_EXPORT	TABLESPACE_QUOTA	Tablespace quota granted to users associated with the selected schemas
SCHEMA_EXPORT	TRIGGER	Triggers on tables in the selected schemas
SCHEMA_EXPORT	XS_SECURITY_CLASS	Oracle Real Application Security (RAS) security classes
SCHEMA_EXPORT	XS_DATA_SECURITY	Oracle Real Application Security (RAS) data security policies

Table 122-19 (Cont.) Object Type Path Names for Heterogeneous Object Types

Heterogeneous Type	Path Name (*=valid only in xxx_PATH_EXPR)	Scope
SCHEMA_EXPORT	XS_ACL	Oracle Real Application Security (RAS) access control lists (ACLs)
SCHEMA_EXPORT	TYPE	Types (both specification and body) in the selected schemas, and their dependent grants and audits
SCHEMA_EXPORT	TYPE_BODY	Type bodies in the selected schemas
SCHEMA_EXPORT	TYPE_SPEC	Type specifications in the selected schemas
SCHEMA_EXPORT	USER	User definitions for users associated with the selected schemas
SCHEMA_EXPORT	VIEW	Views in the selected schemas and their dependent objects (grants, constraints, comments, audits)
DATABASE_EXPORT	ASSOCIATION	Statistics type associations for objects in the database
DATABASE_EXPORT	AUDIT	Audits of SQL statements
DATABASE_EXPORT	AUDIT_OBJ	Audits on all objects in the database
DATABASE_EXPORT	CLUSTER	Clusters and their indexes
DATABASE_EXPORT	COMMENT	Comments on all objects
DATABASE_EXPORT	CONSTRAINT	Constraints (including referential constraints)
DATABASE_EXPORT	CONTEXT	Application contexts
DATABASE_EXPORT	DB_LINK	Private and public database links
DATABASE_EXPORT	DEFAULT_ROLE	Default roles granted to users in the database
DATABASE_EXPORT	DIMENSION	Dimensions in the database
DATABASE_EXPORT	DIRECTORY	Directory objects in the database
DATABASE_EXPORT	FGA_POLICY	Fine-grained audit policies
DATABASE_EXPORT	FUNCTION	Functions
DATABASE_EXPORT	* GRANT	Object and system grants
DATABASE_EXPORT	INDEX	Indexes (including domain indexes) on tables and clusters
DATABASE_EXPORT	INDEXTYPE	Indextypes and their dependent grants and audits
DATABASE_EXPORT	JAVA_SOURCE	Java sources and their dependent grants and audits
DATABASE_EXPORT	JOB	Jobs
DATABASE_EXPORT	LIBRARY	External procedure libraries
DATABASE_EXPORT	MATERIALIZED_VIEW	Materialized views
DATABASE_EXPORT	MATERIALIZED_VIEW_LOG	Materialized view logs
DATABASE_EXPORT	OBJECT_GRANT	All object grants in the database
DATABASE_EXPORT	OPERATOR	Operators and their dependent grants and audits
DATABASE_EXPORT	PACKAGE	Packages (both specification and body) and their dependent grants and audits
DATABASE_EXPORT	PACKAGE_BODY	Package bodies
DATABASE_EXPORT	PACKAGE_SPEC	Package specifications
DATABASE_EXPORT	PASSWORD_HISTORY	Password histories for database users

Table 122-19 (Cont.) Object Type Path Names for Heterogeneous Object Types

Heterogeneous Type	Path Name (*=valid only in xxx_PATH_EXPR)	Scope
DATABASE_EXPORT	*PASSWORD_VERIFY_FUNCTION	The password complexity verification function
DATABASE_EXPORT	PROCEDURE	Procedures and their dependent grants and objects
DATABASE_EXPORT	PROFILE	Profiles
DATABASE_EXPORT	PROXY	Proxy authentications
DATABASE_EXPORT	REF_CONSTRAINT	Referential (foreign key) constraints on tables in the database
DATABASE_EXPORT	REFRESH_GROUP	Refresh groups
DATABASE_EXPORT	*RESOURCE_ COST	Resource cost information
DATABASE_EXPORT	RLS_CONTEXT	Fine-grained access-control driving contexts
DATABASE_EXPORT	RLS_GROUP	Fine-grained access-control policy groups
DATABASE_EXPORT	RLS_POLICY	Fine-grained access-control policies
DATABASE_EXPORT	ROLE	Roles
DATABASE_EXPORT	ROLE_GRANT	Role grants to users in the database
DATABASE_EXPORT	ROLLBACK_SEGMENT	Rollback segments
DATABASE_EXPORT	*SCHEMA (named object)	Database schemas including for each schema all related and dependent objects: user definitions and their attributes (default roles, role grants, tablespace quotas, and so on), objects in the schema (tables, view, packages, types, and so on), and their dependent objects (grants, audits, indexes, constraints, and so on). The NAME and NAME_EXPR filters can be used with this object type path name to designate the database schemas to be fetched.
DATABASE_EXPORT	SEQUENCE	Sequences
DATABASE_EXPORT	STATISTICS	Statistics on tables and indexes
DATABASE_EXPORT	SYNONYM	Public and private synonyms
DATABASE_EXPORT	SYSTEM_GRANT	System privilege grants
DATABASE_EXPORT	TABLE	Tables and their dependent objects (indexes, constraints, triggers, grants, audits, comments, table data, and so on)
DATABASE_EXPORT	TABLE_DATA	Row data for all tables
DATABASE_EXPORT	TABLESPACE	Tablespace definitions
DATABASE_EXPORT	TABLESPACE_QUOTA	Tablespace quota granted to users in the database
DATABASE_EXPORT	TRIGGER	Triggers on the database, on schemas, and on schema objects
DATABASE_EXPORT	XS_USER	Oracle Real Application Security (RAS) users
DATABASE_EXPORT	XS_ROLE	Oracle Real Application Security (RAS) roles
DATABASE_EXPORT	XS_SECURITY_CLASS	Oracle Real Application Security (RAS) security classes
DATABASE_EXPORT	XS_DATA_SECURITY	Oracle Real Application Security (RAS) data security policies
DATABASE_EXPORT	XS_ACL	Oracle Real Application Security (RAS) access control lists (ACLs)
DATABASE_EXPORT	XS_NAMESPACE	Oracle Real Application Security (RAS) namespaces
DATABASE_EXPORT	TRUSTED_DB_LINK	Trusted links

Table 122-19 (Cont.) Object Type Path Names for Heterogeneous Object Types

Heterogeneous Type	Path Name (*=valid only in xxx_PATH_EXPR)	Scope
DATABASE_EXPORT	TYPE	Types (both specification and body) and their dependent grants and audits
DATABASE_EXPORT	TYPE_BODY	Type bodies
DATABASE_EXPORT	TYPE_SPEC	Type specifications
DATABASE_EXPORT	USER	User definitions
DATABASE_EXPORT	VIEW	Views

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INVALID_OPERATION. SET_FILTER was called after the first call to FETCH_xxx for the OPEN context. After the first call to FETCH_xxx is made, no further calls to SET_FILTER are permitted.
- INCONSISTENT_ARGS. The arguments are inconsistent. Possible inconsistencies include the following:
  - The filter name is not valid for the object type associated with the OPEN context.
  - The filter name is not valid for the object_type_path.
  - The object_type_path is not part of the collection designated by handle .
  - The filter value is the wrong datatype.

122.6.12 SET_PARSE_ITEM Procedure

This procedure is used for both retrieval and submission. This procedure enables output parsing and specifies an object attribute to be parsed and returned.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database
> - Subprograms for Submitting XML to the Database

Syntax

The following syntax applies when SET_PARSE_ITEM is used for object retrieval:

```
DBMS_METADATA.SET_PARSE_ITEM (
   handle       IN  NUMBER,
   name         IN  VARCHAR2,
   object_type  IN  VARCHAR2 DEFAULT NULL);
```

The following syntax applies when SET_PARSE_ITEM is used for XML submission:

```
DBMS_METADATA.SET_PARSE_ITEM (
   handle     IN NUMBER,
   name        IN VARCHAR2);
```

Parameters

Table 122-20 SET_PARSE_ITEM Procedure Parameters

Parameter	Description
handle	The handle returned from OPEN (or OPENW).
name	The name of the object attribute to be parsed and returned. See Table 122-21 for the attribute object type, name, and meaning.
object_type	Designates the object type to which the parse item applies (this is an object type name, not a path name). By default, the parse item applies to the object type of the OPEN handle. When the OPEN handle designates a heterogeneous object type, behavior can be either of the following: • If object_type is omitted, then the parse item applies to all object types within the heterogeneous collection . • If object_type is specified, then the parse item only applies to that specific object type within the collection . This parameter only applies when SET_PARSE_ITEM is used for object retrieval.

Table 122-21 describes the object type, name, and meaning of the items available in the SET_PARSE_ITEM procedure.

Because new items are occasionally added, you can query the DBMS_METADATA_PARSE_ITEMS view to see a complete list of valid parse items or to find valid parse items for a specific object type.

Table 122-21 SET_PARSE_ITEM: Parse Items

Object Type	Name	Meaning
All objects	VERB	If FETCH_XML_CLOB is called, no value is returned. If FETCH_DDL is called, then for every row in the sys.ku$_ddls nested table returned by FETCH_DDL the verb in the corresponding ddlText is returned. If the ddlText is a SQL DDL statement, then the SQL verb (for example, CREATE, GRANT, AUDIT) is returned. If the ddlText is a procedure call (for example, DBMS_AQADM.CREATE_QUEUE_TABLE()) then the package.procedure-name is returned.
All objects	OBJECT_TYPE	If FETCH_XML_CLOB is called, an object type name from Table 122-12 is returned. If FETCH_DDL is called and the ddlText is a SQL DDL statement whose verb is CREATE or ALTER, the object type as used in the DDL statement is returned (for example, TABLE, PACKAGE_BODY, and so on). Otherwise, an object type name from Table 122-12 is returned.

Table 122-21 (Cont.) SET_PARSE_ITEM: Parse Items

Object Type	Name	Meaning
Schema objects	SCHEMA	The object schema is returned. If the object is not a schema object, no value is returned.
Named objects	NAME	The object name is returned. If the object is not a named object, no value is returned.
TABLE, TABLE_DATA, INDEX	TABLESPACE	The name of the object's tablespace or, if the object is a partitioned table, the default tablespace is returned. For a TABLE_DATA object, this is always the tablespace where the rows are stored.
TRIGGER	ENABLE	If the trigger is enabled, ENABLE is returned. If the trigger is disabled, DISABLE is returned.
OBJECT_GRANT, TABLESPACE_QUOTA	GRANTOR	The grantor is returned.
Dependent objects (including domain index secondary tables)	BASE_OBJECT_NAME	The name of the base object is returned. If the object is not a dependent object, no value is returned.
Dependent objects (including domain index secondary tables)	BASE_OBJECT_SCHEMA	The schema of the base object is returned. If the object is not a dependent object, no value is returned.
Dependent objects (including domain index secondary tables)	BASE_OBJECT_TYPE	The object type of the base object is returned. If the object is not a dependent object, no value is returned.
Granted objects	GRANTEE	The grantee is returned. If the object is not a granted object, no value is returned.

Usage Notes

These notes apply when using SET_PARSE_ITEM to retrieve objects.

By default, the FETCH_xxx routines return an object's metadata as XML or creation DDL. By calling SET_PARSE_ITEM you can request that individual attributes of the object be returned as well.

You can call SET_PARSE_ITEM multiple times to ask for multiple items to be parsed and returned. Parsed items are returned in the sys.ku$_parsed_items nested table.

For TABLE_DATA objects, the following parse item return values are of interest:

If Object Is	NAME, SCHEMA	BASE_OBJECT_NAME, BASE_OBJECT_SCHEMA
nonpartitioned table	table name, schema	table name, schema
table partition	partition name, schema	table name, schema
nested table	storage table name, schema	name and schema of top-level table (*not* the parent nested table)

Tables are not usually thought of as dependent objects. However, secondary tables for domain indexes are dependent on the domain indexes. Consequently, the BASE_OBJECT_NAME,

BASE_OBJECT_SCHEMA and BASE_OBJECT_TYPE parse items for secondary TABLE objects return the name, schema, and type of the domain index.

> **See Also:**
>
> - "FETCH_xxx Functions and Procedures"
> - *Oracle Database Utilities* for more information about using the metadata APIs.

By default, the CONVERT and PUT procedures simply transform an object's XML metadata to DDL. By calling SET_PARSE_ITEM you can request that individual attributes of the object be returned as well.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INVALID_OPERATION. SET_PARSE_ITEM was called after the first call to FETCH_xxx for the OPEN context. After the first call to FETCH_xxx is made, no further calls to SET_PARSE_ITEM are permitted.
- INCONSISTENT_ARGS. The attribute name is not valid for the object type associated with the OPEN context.

122.6.13 SET_TRANSFORM_PARAM and SET_REMAP_PARAM Procedures

These procedures are used for both retrieval and submission. SET_TRANSFORM_PARAM and SET_REMAP_PARAM specify parameters to the XSLT stylesheet identified by transform_handle.

Purpose

Use these parameters to modify or customize the output of the transform.

> **See Also:**
>
> For more information about related subprograms:
>
> - Subprograms for Retrieving Multiple Objects From the Database
> - Subprograms for Submitting XML to the Database

Syntax

```
DBMS_METADATA.SET_TRANSFORM_PARAM (
   transform_handle    IN NUMBER,
   name                IN VARCHAR2,
   value               IN VARCHAR2,
   object_type         IN VARCHAR2 DEFAULT NULL);

DBMS_METADATA.SET_TRANSFORM_PARAM (
   transform_handle    IN NUMBER,
```

```
   name               IN VARCHAR2,
   value              IN BOOLEAN DEFAULT TRUE,
   object_type        IN VARCHAR2 DEFAULT NULL);

DBMS_METADATA.SET_TRANSFORM_PARAM (
   transform_handle   IN NUMBER,
   name               IN VARCHAR2,
   value              IN NUMBER,
   object_type        IN VARCHAR2 DEFAULT NULL);

DBMS_METADATA.SET_REMAP_PARAM (
   transform_handle   IN NUMBER,
   name               IN VARCHAR2,
   old_value          IN VARCHAR2,
   new_value          IN VARCHAR2,
   object_type        IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 122-22 describes the parameters for the SET_TRANSFORM_PARAM and SET_REMAP_PARAM procedures.

Because new parameters are occasionally added, Oracle suggests that you query the DBMS_METADATA_TRANSFORM_PARAMS view to see all the valid transform parameters for each transform or to find valid transform parameters for specific object types.

Table 122-22 SET_TRANSFORM_PARAM and SET_REMAP_PARAM Parameters

Parameters	Description
name	The name of the transform parameter. For descriptions of the parameters available for each transform on the SET_TRANSFORM_PARAM procedure, see the following: Table 122-23 - DDL transform Table 122-24 - MODIFY transform Table 122-26 - SXML transform Table 122-27 - MODIFYSXML transform Table 122-28 - SXMLDDL transform For descriptions of the parameters available for the MODIFY transform on the SET_REMAP_PARAM procedure, see Table 122-25. For descriptions of the parameters available for the ALTERXML transform, see Table 122-4.
new_value	The new value for the remapping. This parameter is valid only for SET_REMAP_PARAM.
object_type	Designates the object type to which the transform or remap parameter applies. By default, it applies to the same object type as the transform. In cases where the transform applies to all object types within a heterogeneous collection, the following apply: • If object_type is omitted, the parameter applies to all applicable object types within the heterogeneous collection. • If object_type is specified, the parameter only applies to that object type. This allows a caller who has added a transform to a heterogeneous collection to specify different transform parameters for different object types within the collection.

Table 122-22 (Cont.) SET_TRANSFORM_PARAM and SET_REMAP_PARAM Parameters

Parameters	Description
old_value	The old value for the remapping. This parameter is valid only for SET_REMAP_PARAM.
transform_handle	Either (1) the handle returned from ADD_TRANSFORM, or (2) the enumerated constant SESSION_TRANSFORM that designates the DDL transform for the whole session. Note that the handle returned by OPEN is not a valid transform handle. For SET_REMAP_PARAM, the transform handle must designate the MODIFY transform.
value	The value of the transform. This parameter is valid only for SET_TRANSFORM_PARAM.

Table 122-23 describes the object type, name, datatype, and meaning of the parameters for the DDL transform in the SET_TRANSFORM_PARAM procedure.

Table 122-23 SET_TRANSFORM_PARAM: Transform Parameters for the DDL Transform

Object Type	Name	Datatype	Meaning
USER, TABLE, CLUSTER, VIEW, MATERIALIZED_VIEW PROCEDURE, FUNCTION, PACKAGE, TYPE, TRIGGER	COLLATION_CLAUSE	Text	There are three possible values: • NEVER — Collation clauses are never generated. • ALWAYS — Collation clauses are always generated. • NON_DEFAULT — Collation clauses other than USING_NLS_COMP are generated.
TABLE	OMIT_ENCRYPTION_CLAUSE	BOOLEAN	If set to Y, directs Data Pump to suppress column encryption clauses. Columns encrypted in the source database are not encrypted in imported tables. Defaults to N. If set to N, directs Data Pump to create column encryption clauses, as in the source database.
TABLE	DWCS_CVT_IOTS	BOOLEAN	If set to Y, directs Data Pump to transform Index Organized tables to heap organized tables by suppressing the ORGANIZATION INDEX clause when creating the table. Defaults to N. If set to N, the generated DDL retains the table characteristics of the source object.
TABLE, CONSTRAINT	DWCS_CVT_CONSTRAINTS	BOOLEAN	If set to Y, directs Data Pump to create pk/fk/uk constraints as disabled. Defaults to N. If set to N, directs Data Pump to create pk/fk/uk constraints based on the source database status.

Table 122-23 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the DDL Transform

Object Type	Name	Datatype	Meaning
TABLE, CONSTRAINT	CONSTRAINT_USE_DEFAULT_INDEX	BOOLEAN	This transform parameter affects the generation of index relating the pk/uk constraint. If set to Y, forces the name of an index automatically created to enforce the constraint to be identical to the constraint name. Defaults to N. If set to N, the index is created as it was named on the source database.
TABLE, CONSTRAINT	CONSTRAINT_NAME_FROM_INDEX	BOOLEAN	This transform parameter affects the generation of pk/uk constraints which reference user created indexes. If set to Y, forces the name of the constraint to match the name of the index. Defaults to N. If set to N, the constraint is created as it was named on the source database.
TABLE	INCLUDE_SHARDING_CLAUSES	BOOLEAN	This transform parameter enables sharding keywords for the transform to facilitate creating sharded tables, sequences, tablespaces and tablespace set. Options: [TRUE | FALSE] Default: FALSE. When set to TRUE, sharding syntax keywords are available: SHARDED keyword DUPLICATED keyword PARTITION BY keyword, with the keyword options CONSISTENT HASH and PARTITIONS AUTO When the INCLUDE_SHARDING_CLAUSES parameter is set to FALSE in CREATE TABLE, the DDL will contain PARTITION BY RANGE and not include the PARTITIONS AUTO clause.
All objects	PRETTY	BOOLEAN	If TRUE, format the output with indentation and line feeds. Defaults to TRUE.
All objects	SQLTERMINATOR	BOOLEAN	If TRUE, append a SQL terminator (; or /) to each DDL statement. Defaults to FALSE.
TABLE	CONSTRAINTS	BOOLEAN	If TRUE, include all non-referential table constraints in the DDL. If FALSE, omit them. Defaults to TRUE.
TABLE	REF_CONSTRAINTS	BOOLEAN	If TRUE, include all referential constraints (foreign keys) in the DDL. If FALSE, omit them. Defaults to TRUE.

Table 122-23 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the DDL Transform

Object Type	Name	Datatype	Meaning
TABLE	CONSTRAINTS_AS_ALTER	BOOLEAN	If TRUE, include table constraints as separate ALTER TABLE (and, if necessary, CREATE INDEX) statements. If FALSE, specify table constraints as part of the CREATE TABLE statement. Defaults to FALSE. Requires that CONSTRAINTS be TRUE. **Note**: The CONSTRAINTS_AS_ALTER parameter has no effect when an index and a constraint are created for a table in two separate DDL statements.
TABLE, TYPE	OID	BOOLEAN	If TRUE, include the Object ID (OID) clause in the DDL. If FALSE, omit it. Defaults to FALSE.
TABLE	SIZE_BYTE_KEYWORD	BOOLEAN	If TRUE, include the BYTE keyword as part of the size specification of CHAR and VARCHAR2 columns that use byte semantics. If FALSE, omit the keyword. Defaults to FALSE.
TABLE, INDEX	PARTITIONING	BOOLEAN	If TRUE, include partitioning clauses in the DDL. If FALSE, omit them. Defaults to TRUE.
INDEX, CONSTRAINT, ROLLBACK_SEGMENT, CLUSTER, TABLE, TABLESPACE	SEGMENT_ATTRIBUTES	BOOLEAN	If TRUE, include segment attributes clauses (physical attributes, storage attributes, tablespace, logging) in the DDL. If FALSE, omit them. Defaults to TRUE.
INDEX, CONSTRAINT, ROLLBACK_SEGMENT, CLUSTER,TABLE	STORAGE	BOOLEAN	If TRUE, include storage clauses in the DDL. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
INDEX, CONSTRAINT, ROLLBACK_SEGMENT, CLUSTER,TABLE	TABLESPACE	BOOLEAN	If TRUE, include tablespace clauses in the DDL. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
TYPE, PACKAGE	SPECIFICATION	BOOLEAN	If TRUE, include the type or package specification in the DDL. If FALSE, omit it. Defaults to TRUE.
TYPE, PACKAGE	BODY	BOOLEAN	If TRUE, include the type body or package body in the DDL. If FALSE, omit it. Defaults to TRUE.
VIEW	FORCE	BOOLEAN	If TRUE, use the FORCE keyword in the CREATE VIEW statement. If FALSE, do not use the FORCE keyword in the CREATE VIEW statement. Defaults to TRUE.
OUTLINE	INSERT	BOOLEAN	If TRUE, include the INSERT statements into the OL$ dictionary tables that will create the outline and its hints. If FALSE, omit a CREATE OUTLINE statement. Defaults to FALSE. Note: This object type is being deprecated.

Table 122-23 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the DDL Transform

Object Type	Name	Datatype	Meaning
All objects	DEFAULT	BOOLEAN	Calling SET_TRANSFORM_PARAM with this parameter set to TRUE has the effect of resetting all parameters for the transform to their default values. Setting this FALSE has no effect. There is no default.
All objects	INHERIT	BOOLEAN	If TRUE, inherits session-level parameters. Defaults to FALSE. If an application calls ADD_TRANSFORM to add the DDL transform, then by default the only transform parameters that apply are those explicitly set for that transform handle. This has no effect if the transform handle is the session transform handle.
ROLE	REVOKE_FROM	Text	The name of a user from whom the role must be revoked. If this is a non-null string and if the CREATE ROLE statement grants you the role, a REVOKE statement is included in the DDL after the CREATE ROLE statement. Note: When you issue a CREATE ROLE statement, Oracle may grant you the role. You can use this transform parameter to undo the grant. Defaults to null string.
TABLESPACE	REUSE	BOOLEAN	If TRUE, include the REUSE parameter for data files in a tablespace to indicate that existing files can be reused. If FALSE, omit the REUSE parameter. Defaults to FALSE.
CLUSTER, INDEX, ROLLBACK_SEGMENT, TABLE, TABLESPACE	PCTSPACE	NUMBER	A number representing the percentage by which space allocation for the object type is to be modified. The value is the number of one-hundredths of the current allocation. For example, 100 means 100%. If the object type is TABLESPACE, the following size values are affected: - in file specifications, the value of SIZE - MINIMUM EXTENT - EXTENT MANAGEMENT LOCAL UNIFORM SIZE For other object types, INITIAL and NEXT are affected.

Table 122-23 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the DDL Transform

Object Type	Name	Datatype	Meaning
TABLE	LOB_STORAGE	Text	Specifies the storage type to use for LOB segments. The options are as follows: • SECUREFILE - LOB storage is returned as SECUREFILE • BASICFILE - LOB storage is returned as BASICFILE • DEFAULT - The keyword (SECUREFILE or BASICFILE) is omitted in the LOB STORE AS clause. • NO_CHANGE - LOB segments are created with the same storage they had in the source database. This is the default. Specifying this transform changes the LOB storage for all tables in the job, including tables that provide storage for materialized views.
TABLE	TABLE_COMPRESSION_CLAUSE	Text	Specifies a table compression clause (for example, COMPRESS BASIC) to use when the table is created. Specify NONE to omit the table compression clause. The table will have the default compression for the tablespace. Specifying this transform changes the compression type for all tables in the job, including tables that provide storage for materialized views.

Table 122-24 describes the object type, name, datatype, and meaning of the parameters for the MODIFY transform in the SET_TRANSFORM_PARAM procedure.

Table 122-24 SET_TRANSFORM_PARAM: Transform Parameters for the MODIFY Transform

Object Type	Name	Datatype	Meaning
All objects	OBJECT_ROW	NUMBER	A number designating the object row for an object. The object in the document that corresponds to this number will be copied to the output document. This parameter is additive. By default, all objects are copied to the output document.

Table 122-25 describes the object type, name, datatype, and meaning of the parameters for the MODIFY transform in the SET_REMAP_PARAM procedure.

Table 122-25 SET_REMAP_PARAM: Transform Parameters for the MODIFY Transform

Object Type	Name	Datatype	Meaning
LIBRARY, TABLESPACE, DIRECTORY	REMAP_DATAFILE	Text	Objects in the document will have their filespecs renamed as follows: any filespec matching old_value will be changed to new_value. Filespecs should *not* be enclosed in quotes. This parameter is additive. By default, filespecs are not renamed.
Named objects and all objects dependent on named objects	REMAP_NAME	Text	Any named object in the document whose name matches old_value will have its name changed to new_value. Any dependent object whose base object name matches old_value will have its base schema name changed to new_value. This parameter is additive. By default, names are not remapped.
Schema Objects, Dependent Objects, Granted Objects, USER	REMAP_SCHEMA	Text	Any schema object in the document whose name matches old_value will have its schema name changed to new_value. This parameter is additive. By default, schemas are not remapped. NOTE: The mapping may not be 100 percent complete because there are certain schema references that Import is not capable of finding. For example, Import will not find schema references embedded within the body of definitions of triggers, types, views, procedures, and packages.
TABLE, CLUSTER, CONSTRAINT, INDEX, ROLLBACK_SEGMENT, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG, TABLESPACE_QUOTA	REMAP_TABLESPACE	Text	Objects in the document will have their tablespaces renamed as follows: any tablespace name matching old_value will be changed to new_value. This parameter is additive. By default, tablespaces are not remapped. Starting with Oracle Database 21c, you can use the % wildcard in place of the source tablespaces. The wildcard is applied to DDL for the object types USER, TABLE, MVIEW, MVIEW_LOG, MVIEW_ZONEMAP and CLUSTERS. TRANSFORM=TABLESPACE:N is mutually exclusive with REMAP_TABLESPACE.

Table 122-26 SET_TRANSFORM_PARAM: Transform Parameters for the SXML Transform

Object type	Name	Datatype	Meaning
USER, TABLE, CLUSTER, VIEW, MATERIALIZED_VIEW PROCEDURE, FUNCTION, PACKAGE, TYPE, TRIGGER	COLLATION_CLAUSE	Text	There are three possible values: • NEVER — Collation clauses are never generated. • ALWAYS — Collation clauses are always generated. • NON_DEFAULT — Collation clauses other than USING_NLS_COMP are generated.

Table 122-26 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the SXML Transform

Object type	Name	Datatype	Meaning
TABLE, TYPE	OID	Boolean	If TRUE, include the Oracle Internet Directory (OID) clause in the SXML. If FALSE, omit it. Defaults to FALSE.
TABLE, INDEX, CLUSTER, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG.	STORAGE	Boolean	If TRUE, include storage clauses in the SXML. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
TABLE, INDEX, CLUSTER, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG.	TABLESPACE	Boolean	If TRUE, include tablespace clauses in the SXML. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
TABLE	REF_CONSTRAINTS	Boolean	If TRUE, include all referential constraints (foreign keys) in the SXML. If FALSE, omit them. Defaults to TRUE.
TABLE, INDEX, MATERIALIZED_VIEW	PHYSICAL_PROPERTIES	Boolean	If TRUE, include segment attributes clauses (physical attributes, storage attributes, tablespace, logging) in the SXML. If FALSE, omit them. Defaults to TRUE.
INDEX, CONSTRAINT, ROLLBACK_SEGMENT, CLUSTER, TABLE, TABLESPACE	SEGMENT_ATTRIBUTES	Boolean	If TRUE, include segment attributes clauses (physical attributes, storage attributes, tablespace, logging) in the SXML. If FALSE, omit them. Defaults to TRUE.
TABLE, INDEX	PARTITIONING	Boolean	If TRUE, include partitioning clauses in the SXML. If FALSE, omit them. Defaults to TRUE.
TABLE	CONSTRAINTS	Boolean	If TRUE, include all non-referential table constraints in the SXML. If FALSE, omit them. Defaults to TRUE.

Table 122-27 SET_TRANSFORM_PARAM: Transform Parameters for the MODIFYSXML Transform

Object type	Name	Datatype	Meaning
TABLE, INDEX, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	STORAGE	Boolean	If TRUE, include storage clauses in the output SXML. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
TABLE, INDEX, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	TABLESPACE	Boolean	If TRUE, include tablespace clauses in the output SXML. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.

Table 122-27 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the MODIFYSXML Transform

Object type	Name	Datatype	Meaning
TABLE	REF_CONSTRAINTS	Boolean	If TRUE, include all referential constraints (foreign keys) in the output SXML. If FALSE, omit them. Defaults to TRUE.
TABLE, INDEX, VIEW, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	REMAP_NAME	Text	Any NAME element in the document that matches old_value will be changed to new_value. This does not apply to column names. (See REMAP_COLUMN_NAME).
TABLE, INDEX, VIEW, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	REMAP_SCHEMA	Text	Any SCHEMA element in the document matching old_value will be changed to new_value.
TABLE, INDEX, VIEW, MATERIALIZED_VIEW	REMAP_COLUMN_NAME	Text	Any column in the document whose name matches old_value will have its name changed to new_value.
TABLE, INDEX, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	SEGMENT_ATTRIBUTES	Boolean	If TRUE, include segment attributes clauses (physical attributes, storage attributes, tablespace, logging) in the output SXML. If FALSE, omit them. Defaults to TRUE.
TABLE	CONSTRAINTS	Boolean	If TRUE, include all non-referential table constraints in the output SXML. If FALSE, omit them. Defaults to TRUE.

Table 122-28 SET_TRANSFORM_PARAM: Transform Parameters for the SXMLDDL Transform

Object type	Name	Datatype	Meaning
USER, TABLE, CLUSTER, VIEW, MATERIALIZED_VIEW PROCEDURE, FUNCTION, PACKAGE, TYPE, TRIGGER	COLLATION_CLAUSE	Text	There are three possible values: • NEVER — Collation clauses are never generated. • ALWAYS — Collation clauses are always generated. • NON_DEFAULT — Collation clauses other than USING_NLS_COMP are generated.
TABLE	OID	Boolean	If TRUE, include OIDs in the DDL. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
TABLE, INDEX, CLUSTER, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG.	TABLESPACE	Boolean	If TRUE, include tablespace clauses in the DDL. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.

Table 122-28 (Cont.) SET_TRANSFORM_PARAM: Transform Parameters for the SXMLDDL Transform

Object type	Name	Datatype	Meaning
TABLE, INDEX, CLUSTER, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	STORAGE	Boolean	If TRUE, include storage clauses in the DDL. If FALSE, omit them. (Ignored if SEGMENT_ATTRIBUTES is FALSE.) Defaults to TRUE.
TABLE	REF_CONSTRAINTS	Boolean	If TRUE, include all referential constraints (foreign keys) in the DDL. If FALSE, omit them. Defaults to TRUE.
INDEX	PRESERVE_LOCAL	Boolean	If PARTITIONING is FALSE and PRESERVE_LOCAL is TRUE and the index is a locally partitioned index, include the LOCAL keyword in the DDL. Defaults to FALSE.
TABLE, INDEX, CLUSTER, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG	SEGMENT_ATTRIBUTES	Boolean	If TRUE, include segment attributes clauses (physical attributes, storage attributes, tablespace, logging) in the DDL. If FALSE, omit them. Defaults to TRUE.
TABLESPACE	REUSE	Boolean	If TRUE, include the REUSE parameter for data files in a tablespace to indicate that existing files can be reused. If FALSE, omit the REUSE parameter. Defaults to FALSE.
TABLE, INDEX	PARTITIONING	Boolean	If TRUE, include partitioning clauses in the DDL. If FALSE, omit them. Defaults to TRUE.
TABLE	CONSTRAINTS	Boolean	If TRUE, include all non-referential table constraints in the output SXML. If FALSE, omit them. Defaults to TRUE.

Exceptions

- INVALID_ARGVAL. A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.
- INVALID_OPERATION. Either SET_TRANSFORM_PARAM or SET_REMAP_PARAM was called after the first call to FETCH_xxx for the OPEN context. After the first call to FETCH_xxx is made, no further calls to SET_TRANSFORM_PARAM or SET_REMAP_PARAM are permitted.
- INCONSISTENT_ARGS. The arguments are inconsistent. This can mean the following:
  - The transform parameter name is not valid for the object type associated with the OPEN context or for the transform associated with the transform handle.
  - The transform applies to all object types in a heterogeneous collection, but object_type is not part of the collection.
- TRANSFORM=TABLESPACE:N is mutually exclusive with REMAP_TABLESPACE.
- The following restrictions apply to use of REMAP_TABLESPACE transform parameter:
  - The target tablespace must be a permanent tablespace, and it must exist before the import.
  - The target tablespace cannot be a temporary tablespace.
  - The % wildcard cannot be used with multiple REMAP_TABLESPACE parameters.

- The REMAP_TABLESPACE parameter and TRANSFORM=TABLESPACE:N transform parameter are mutually exclusive.

Usage Notes

XSLT allows parameters to be passed to stylesheets. You call SET_TRANSFORM_PARAM or SET_REMAP_PARAM to specify the value of a parameter that you want to be passed to the stylesheet, identified by transform_handle.

Normally, if you call SET_TRANSFORM_PARAMETER multiple times for the same parameter name, each call overrides the prior call. For example, the following sequence sets the STORAGE transform parameter to TRUE, because TRUE is the last value specified for SET_TRANSFORM_PARAM:

```
SET_TRANSFORM_PARAM(tr_handle,'STORAGE',false);
SET_TRANSFORM_PARAM(tr_handle,'STORAGE',true);
```

However, some transform parameters are additive. This means that all specified parameter values are applied to the document, not just the last one. For example, the OBJECT_ROW parameter to the MODIFY transform is additive. If you specify the following, then both specified rows are copied to the output document:

```
SET_TRANSFORM_PARAM(tr_handle,'OBJECT_ROW',5);
SET_TRANSFORM_PARAM(tr_handle,'OBJECT_ROW',8);
```

The REMAP_TABLESPACE parameter is also additive. If you specify the following, then tablespaces TBS1 and TBS3 are changed to TBS2 and TBS4, respectively:

```
SET_REMAP_PARAM(tr_handle,'REMAP_TABLESPACE','TBS1','TBS2');
SET_REMAP_PARAM(tr_handle,'REMAP_TABLESPACE','TBS3','TBS4');
```

The order in which the transformations are performed is undefined. For example, if you specify the following, then the result is undefined:

```
SET_REMAP_PARAM(tr_handle,'REMAP_TABLESPACE','TBS1','TBS2');
SET_REMAP_PARAM(tr_handle,'REMAP_TABLESPACE','TBS2','TBS3');
```

> **Note:**
>
> The number of remap parameters that can be specified for a MODIFY transform is limited to ten. That is, you can specify up to ten REMAP_DATAFILE parameters, up to ten REMAP_SCHEMA parameters and so on. Additional instances are ignored. To work around this limit, you can perform another DBMS_METADATA.ADD_TRANSFORM and specify additional remap parameters.

The GET_DDL, GET_DEPENDENT_DDL, and GET_GRANTED_DDL functions allow the casual browser to extract the creation DDL for an object. So that you can specify transform parameters, this package defines an enumerated constant SESSION_TRANSFORM as the handle of the DDL transform at the session level. You can call SET_TRANSFORM_PARAM using DBMS_METADATA.SESSION_TRANSFORM as the transform handle to set transform parameters for the whole session. GET_DDL, GET_DEPENDENT_DDL, and GET GRANTED_DDL inherit these parameters when they invoke the DDL transform.

> **Note:**
>
> The enumerated constant must be prefixed with the package name `DBMS_METADATA.SESSION_TRANSFORM`.

123
DBMS_METADATA_DIFF

The DBMS_METADATA_DIFF package contains the interfaces for comparing two metadata documents in SXML format.

This chapter contains the following topics:

- Overview
- Security Model
- Browsing APIs for Fetching and Comparing Objects
- Summary of DBMS_METADATA_DIFF Subprograms

> **See Also:**
>
> *Oracle Database Utilities* for more information and for examples of using the DBMS_METADATA_DIFF package.

123.1 DBMS_METADATA_DIFF Overview

You can use the interfaces contained in the DBMS_METADATA_DIFF package to compare two metadata documents in SXML format.

The result of the comparison is an SXML difference document. This document can be converted to other formats using the DBMS_METADATA submit interface and the CONVERT API.

123.2 DBMS_METADATA_DIFF Security Model

The browsing interface of the DBMS_METADATA_DIFF package actually uses the DBMS_METADATA package to fetch the metadata to be compared. Therefore, the security model used for DBMS_METADATA also applies to DBMS_METADATA_DIFF. (Note, however, that DBMS_METADATA_DIFF does not support all object types.)

> **See Also:**
>
> DBMS_METADATA for information about the DBMS_METADATA security model

123.3 DBMS_METADATA_DIFF — Browsing APIs for Fetching and Comparing Objects

These functions allow you to compare the metadata for two objects with a single call.

Syntax

```
DBMS_METADATA_DIFF.COMPARE_SXML(
object_type   IN VARCHAR2,
name1         IN VARCHAR2,
name2         IN VARCHAR2,
schema1       IN VARCHAR2 DEFAULT NULL,
schema2       IN VARCHAR2 DEFAULT NULL,
network_link1 IN VARCHAR2 DEFAULT NULL,
network_link2 IN VARCHAR2 DEFAULT NULL)
RETURN CLOB;

DBMS_METADATA_DIFF.COMPARE_ALTER(
object_type   IN VARCHAR2,
name1         IN VARCHAR2,
name2         IN VARCHAR2,
schema1       IN VARCHAR2 DEFAULT NULL,
schema2       IN VARCHAR2 DEFAULT NULL,
network_link1 IN VARCHAR2 DEFAULT NULL,
network_link2 IN VARCHAR2 DEFAULT NULL)
RETURN CLOB;

DBMS_METADATA_DIFF.COMPARE_ALTER_XML(
object_type   IN VARCHAR2,
name1         IN VARCHAR2,
name2         IN VARCHAR2,
schema1       IN VARCHAR2 DEFAULT NULL,
schema2       IN VARCHAR2 DEFAULT NULL,
network_link1 IN VARCHAR2 DEFAULT NULL,
network_link2 IN VARCHAR2 DEFAULT NULL)
RETURN CLOB;
```

Parameters

Table 123-1 COMPARE_xxx Function Parameters

Parameters	Description
object_type	The type of object to be compared. Valid type names are CLUSTER, CONTEXT, DB_LINK, FGA_POLICY, INDEX, MATERIALIZED_VIEW, MATERIALIZED_VIEW_LOG, QUEUE, QUEUE_TABLE, RLS_CONTEXT, RLS_GROUP, RLS_POLICY, ROLE, SEQUENCE, SYNONYM, TABLE, TABLESPACE, TRIGGER, TYPE, TYPE_SPEC, TYPE_BODY, USER, and VIEW.
name1	The name of the first object in the comparison.
name2	The name of the second object in the comparison.
schema1	The schema of the first object in the comparison. The default is the current user.
schema2	The schema of the second object in the comparison. The default is the value of schema1.

Table 123-1 (Cont.) COMPARE_xxx Function Parameters

Parameters	Description
network_link1	The name of a database link to the database on which the first object resides. If NULL (the default), then the object is assumed to be in the database on which the caller is running.
network_link2	The name of a database link to the database on which the second object resides. The default is the value of network_link1.

Return Values

DBMS_METADATA_DIFF.COMPARE_xxx returns the differences between two objects.

Exceptions

- INVALID_ARGVAL

  A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.

- OBJECT_NOT_FOUND

  The specified object was not found in the database.

Usage Notes

These functions encapsulate calls to both DBMS_METADATA and DBMS_METADATA_DIFF functions and procedures to fetch the metadata for each of the two objects and compare them.

Which function you use depends on the comparison format you want:

- COMPARE_SXML returns an SXML difference document.
- COMPARE_ALTER returns a set of ALTER statements for making the first object like the second object.
- COMPARE_ALTER_XML returns an ALTER_XML document.

123.4 Summary of DBMS_METADATA_DIFF Subprograms

The DBMS_METADATA_DIFF subprograms provide comparison functionality for different object types.

These subprograms are used to:

- Specify the type of objects to be compared
- Specify the SXML documents to be compared
- Show the differences between the compared documents
- Clean up after the comparison

Table 123-2 provides a summary of DBMS_METADATA_DIFF subprograms.

Table 123-2 DBMS_METADATA_DIFF Package Subprograms

Subprogram	Description
OPENC Function	Specifies the type of objects to be compared
ADD_DOCUMENT Procedure	Specifies an SXML document to be compared
FETCH_CLOB Functions and Procedures	Returns a CLOB showing the differences between the two documents specified by `ADD_DOCUMENT`
CLOSE Procedure	Invalidates the handle returned by `OPENC` and cleans up associated state

123.4.1 OPENC Function

This function specifies the type of objects to be compared. The return value is an opaque context handle.

Syntax

```
DBMS_METADATA_DIFF.OPENC (
object_type  IN VARCHAR2)
RETURN NUMBER;
```

Parameters

Table 123-3 OPENC Function Parameters

Parameters	Description
`object_type`	The type of object to be compared. Valid type names are `CLUSTER`, `CONTEXT`, `DB_LINK`, `FGA_POLICY`, `INDEX`, `MATERIALIZED_VIEW`, `MATERIALIZED_VIEW_LOG`, `QUEUE`, `QUEUE_TABLE`, `RLS_CONTEXT`, `RLS_GROUP`, `RLS_POLICY`, `ROLE`, `SEQUENCE`, `SYNONYM`, `TABLE`, `TABLESPACE`, `TRIGGER`, `TYPE`, `TYPE_SPEC`, `TYPE_BODY`, `USER`, and `VIEW`.

Return Values

The opaque handle that is returned is used as input to `ADD_DOCUMENT`, `FETCH_xxx` and `CLOSE`.

Exceptions

- `INVALID_ARGVAL`

  A `NULL` or invalid value was supplied for an input parameter. The error message text identifies the parameter.

123.4.2 ADD_DOCUMENT Procedure

This procedure specifies an SXML document that is to be compared.

Syntax

```
DBMS_METADATA_DIFF.ADD_DOCUMENT(
handle IN NUMBER, document IN sys.XMLType);
```

```
DBMS_METADATA_DIFF.ADD_DOCUMENT(
handle IN NUMBER, document IN CLOB);
```

Parameters

Table 123-4 ADD_DOCUMENT Procedure Parameters

Parameter	Description
handle	The handle returned from OPENC
document	A document to be compared. The document must be of the type specified in OPENC.

Usage Notes

Because the comparison interface allows you to compare exactly two SXML documents, a program must call ADD_DOCUMENT exactly twice for each OPENC handle. In the comparison result, the document specified by the first call is document 1, and the document specified by the second call is document 2.

Exceptions

- INVALID_ARGVAL

  A NULL or invalid value was supplied for an input parameter. The error message text identifies the parameter.

123.4.3 FETCH_CLOB Functions and Procedures

The FETCH_CLOB functions and procedures return a CLOB showing the differences between the two documents specified by ADD_DOCUMENT.

Syntax

```
DBMS_METADATA_DIFF.FETCH_CLOB(
handle IN NUMBER)
RETURN CLOB;

DBMS_METADATA_DIFF.FETCH_CLOB(

handle IN  NUMBER,
doc    IN  OUT NOCOPY CLOB);

DBMS_METADATA_DIFF.FETCH_CLOB(
handle IN  NUMBER,
doc    IN  OUT NOCOPY CLOB
diffs  OUT BOOLEAN);
```

Parameters

Table 123-5 FETCH_CLOB Subprogram Parameters

Parameter	Description
handle	The handle returned from OPENC.
doc	A CLOB containing the differences between documents 1 and 2.
diffs	TRUE if the documents are different or FALSE if they are identical.

Return Values

The differences between documents 1 and 2.

Exceptions

- `INVALID_ARGVAL`

  A `NULL` or invalid value was supplied for an input parameter. The error message text identifies the parameter.

123.4.4 CLOSE Procedure

This procedure invalidates the handle returned by `OPENC` and cleans up associated state.

Syntax

```
DBMS_METADATA_DIFF.CLOSE(
handle IN NUMBER);
```

Parameters

Table 123-6 CLOSE Function Parameters

Parameters	Description
`handle`	The handle returned from `OPENC`

Exceptions

- `INVALID_ARGVAL`

  A `NULL` or invalid value was supplied for an input parameter. The error message text identifies the parameter.

124
DBMS_MGD_ID_UTL

The DBMS_MGD_ID_UTL package contains various utility functions and procedures.

These consist of the following utility subprograms:

- A logging utility that sets and gets Java and PL/SQL logging levels.
- A proxy utility consisting of two procedures used to set and unset the host and port of the proxy server.
- A metadata utility consisting of functions and procedures used for managing metadata.

> **See Also:**
>
> *Oracle Database Development Guide* for more information.

This chapter describes each of these utility subprograms and contains the following topics:

- Security Model
- Constants
- Exceptions
- Summary of DBMS_MGD_ID_UTL Subprograms

The examples in this chapter assume that the user has run the following set of commands before running the contents of each script:

```
SQL> connect / as sysdba;
Connected.
SQL> create user mgduser identified by password;
SQL> grant connect, resource to mgduser;
SQL> connect mgduser
Enter password: mgduserpassword
Connected.
SQL> set serveroutput on;
```

124.1 DBMS_MGD_ID_UTL Security Model

You must run the catmgd.sql script to load the DBMS_MGD_ID_UTL package and Identity Code Package schema objects in the MGDSYS schema.

DBMS_MGD_ID_UTL is a MGDSYS-owned package. Any DBMS_MGD_ID_UTL subprogram called from an anonymous PL/SQL block is run using the privileges of the current user.

A user must be granted connect and resource roles to use the DBMS_MGD_ID_UTL package and its subprograms.

EXECUTE privilege is granted to PUBLIC for these ADTs: MGD_ID, MGD_ID_COMPONENT, MGD_ID_COMPONENT_VARRAY, and for this package DBMS_MGD_ID_UTL.

SELECT or READ privilege is granted to PUBLIC for these read-only views: MGD_ID_CATEGORY and MGD_ID_SCHEME and for these metadata views: USER_MGD_ID_CATEGORY and USER_MGD_ID_SCHEME, and for table MGD_ID_XML_VALIDATOR, and for sequence MGD$SEQUENCE_CATEGORY.

INSERT, UPDATE and DELETE privilege is granted to PUBLIC for these metadata views: USER_MGD_ID_CATEGORY and USER_MGD_ID_SCHEME.

Public synonyms, by the same name, are created for these ADTs: MGD_ID, MGD_ID_COMPONENT, MGD_ID_COMPONENT_VARRAY and for this package DBMS_MGD_ID_UTL, as well as for these read-only views: MGD_ID_CATEGORY and MGD_ID_SCHEME and for these metadata views: USER_MGD_ID_CATEGORY and USER_MGD_ID_SCHEME, and for table MGD_ID_XML_VALIDATOR.

124.2 DBMS_MGD_ID_UTL Constants

The DBMS_MGD_ID_UTL package defines several constants for specifying parameter values.

These constants are shown in the following tables.

Table 124-1 DBMS_MGD_ID_UTL Constants — Installed Category IDs and Names

Name	Value
EPC_ENCODING_CATEGORY_ID	1
EPC_ENCODING_CATEGORY_NAME	EPC

Table 124-2 DBMS_MGD_ID_UTL Constants — Logging Levels

Name	Value
LOGGING_LEVEL_OFF	0
LOGGING_LEVEL_SEVERE	1
LOGGING_LEVEL_WARNING	2
LOGGING_LEVEL_INFO	3
LOGGING_LEVEL_FINE	4
LOGGING_LEVEL_FINER	5
LOGGING_LEVEL_FINEST	6
LOGGING_LEVEL_ALL	7

124.3 DBMS_MGD_ID_UTL Exceptions

The table in this topic lists the DBMS_MGD_ID_UTL exceptions.

Table 124-3 Exceptions Raised by DBMS_MGD_ID_UTL Package

Name	Error Code	Description
TDTJavaException	-55200	During the tag data translation, a Java exception was raised.
TDTCategoryNotFound	-55201	The specified category was not found.
TDTSchemeNotFound	-55202	During the tag data translation, the specified scheme was not found.
TDTLevelNotFound	-55203	During the tag data translation, the specified level was not found.
TDTOptionNotFound	-55204	During the tag data translation, the specified option was not found.
TDTFieldValidationException	-55205	During the tag data translation, the validation operation failed on a field.
TDTUndefinedField	-55206	During the tag data translation, an undefined field was detected.
TDTRuleEvaluationFailed	-55207	During the tag data translation, the rule evaluation operation failed.
TDTTooManyMatchingLevels	-55208	During the tag data translation, too many matching levels were found.

124.4 Summary of DBMS_MGD_ID_UTL Subprograms

This table describes the utility subprograms in the DBMS_MGD_ID_UTL package.

All the values and names passed to the procedures defined in the DBMS_MGD_ID_UTL package are case insensitive unless otherwise mentioned. To preserve the case, enclose the values with double quotation marks.

Table 124-4 DBMS_MGD_ID_UTL Package Subprograms

Subprogram	Description
ADD_SCHEME Procedure	Adds a tag data translation scheme to an existing category
CREATE_CATEGORY Function	Creates a new category or a new version of a category
EPC_TO_ORACLE_SCHEME Function	Converts the EPCglobal tag data translation (TDT) XML to Oracle tag data translation XML
GET_CATEGORY_ID Function	Returns the category ID given the category name and the category version
GET_COMPONENTS Function	Returns all relevant separated component names separated by semicolon (';') for the specified scheme
GET_ENCODINGS Function	Returns a list of semicolon (';') separated encodings (formats) for the specified scheme
GET_JAVA_LOGGING_LEVEL Function	Returns an integer representing the current Java trace logging level
GET_PLSQL_LOGGING_LEVEL Function	Returns an integer representing the current PL/SQL trace logging level
GET_SCHEME_NAMES Function	Returns a list of semicolon (';') separated scheme names for the specified category
GET_TDT_XML Function	Returns the Oracle tag data translation XML for the specified scheme

Table 124-4 (Cont.) DBMS_MGD_ID_UTL Package Subprograms

Subprogram	Description
GET_VALIDATOR Function	Returns the Oracle Database tag data translation schema
REFRESH_CATEGORY Function	Refreshes the metadata information on the Java stack for the specified category
REMOVE_CATEGORY Procedure	Removes a category including all the related TDT XML if the value of `category_version` parameter is NULL
REMOVE_PROXY Procedure	Unsets the host and port of the proxy server
REMOVE_SCHEME Procedure	Removes a tag data translation scheme from a category
SET_JAVA_LOGGING_LEVEL Procedure	Sets the Java logging level
SET_PLSQL_LOGGING_LEVEL Procedure	Sets the PL/SQL tracing logging level
SET_PROXY Procedure	Sets the host and port of the proxy server for Internet access
VALIDATE_SCHEME Function	Validates the input tag data translation XML against the Oracle tag data translation schema

124.4.1 ADD_SCHEME Procedure

This procedure adds a tag data translation scheme to an existing category.

Syntax

```
DBMS_MGD_ID_UTL.ADD_SCHEME (
   category_id IN  VARCHAR2,
   tdt_xml     IN  CLOB);
```

Parameters

Table 124-5 ADD_SCHEME Procedure Parameters

Parameter	Description
`category_id`	Category ID
`tdt_xml`	Tag data translation XML

Examples

This example performs the following actions:

1. Creates a category.
2. Adds a contractor scheme and an employee scheme to the `MGD_SAMPLE_CATEGORY` category.
3. Validates the `MGD_SAMPLE_CATEGORY` scheme.
4. Tests the tag translation of the contractor scheme and the employee scheme.
5. Removes the contractor scheme.
6. Tests the tag translation of the contractor scheme and this returns the expected exception for the removed contractor scheme.
7. Tests the tag translation of the employee scheme and this returns the expected values.

8. Removes the `MGD_SAMPLE_CATEGORY` category.

```
--contents of add_scheme2.sql
SET LINESIZE 160
-----------------------------------------------------------------------
---CREATE CATEGORY, ADD_SCHEME, REMOVE_SCHEME, REMOVE_CATEGORY-------
-----------------------------------------------------------------------
DECLARE
  amt          NUMBER;
  buf          VARCHAR2(32767);
  pos          NUMBER;
  tdt_xml      CLOB;
  validate_tdtxml VARCHAR2(1042);
  category_id  VARCHAR2(256);
BEGIN
  -- remove the testing category if already existed
  DBMS_MGD_ID_UTL.remove_category('MGD_SAMPLE_CATEGORY', '1.0');
  -- Step 1. Create the testing category 'MGD_SAMPLE_CATEGORY', version 1.0.
  category_id := DBMS_MGD_ID_UTL.CREATE_CATEGORY('MGD_SAMPLE_CATEGORY', '1.0', 'Oracle',
'http://www.example.com/mgd/sample');
  -- Step 2. Add contractor scheme to the category.
  DBMS_LOB.CREATETEMPORARY(tdt_xml, true);
  DBMS_LOB.OPEN(tdt_xml, DBMS_LOB.LOB_READWRITE);

  buf := '<?xml version="1.0" encoding="UTF-8"?>
<TagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                    xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                    xmlns="oracle.mgd.idcode">
 <scheme name="CONTRACTOR_TAG" optionKey="1" xmlns="">
  <level type="URI" prefixMatch="example.contractor.">
   <option optionKey="1" pattern="example.contractor.([0-9]*).([0-9]*)"
           grammar="''example.contractor.'' contractorID ''.'' divisionID">
    <field seq="1" characterSet="[0-9]*" name="contractorID"/>
    <field seq="2" characterSet="[0-9]*" name="divisionID"/>
   </option>
  </level>
  <level type="BINARY" prefixMatch="11">
   <option optionKey="1" pattern="11([01]{7})([01]{6})"
           grammar="''11'' contractorID divisionID ">
    <field seq="1" characterSet="[01]*" name="contractorID"/>
    <field seq="2" characterSet="[01]*" name="divisionID"/>
   </option>
  </level>
 </scheme>
</TagDataTranslation>';

  amt := length(buf);
  pos := 1;
  DBMS_LOB.WRITE(tdt_xml, amt, pos, buf);
  DBMS_LOB.CLOSE(tdt_xml);

  DBMS_MGD_ID_UTL.ADD_SCHEME(category_id, tdt_xml);

  -- Add the employee scheme to the category.
  DBMS_LOB.CREATETEMPORARY(tdt_xml, true);
  DBMS_LOB.OPEN(tdt_xml, DBMS_LOB.LOB_READWRITE);

  buf := '<?xml version="1.0" encoding="UTF-8"?>
<TagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                    xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                    xmlns="oracle.mgd.idcode">
 <scheme name="EMPLOYEE_TAG" optionKey="1" xmlns="">
```

```
  <level type="URI" prefixMatch="example.employee.">
   <option optionKey="1" pattern="example.employee.([0-9]*).([0-9]*)"
           grammar="''example.employee.'' employeeID ''.'' divisionID">
    <field seq="1" characterSet="[0-9]*" name="employeeID"/>
    <field seq="2" characterSet="[0-9]*" name="divisionID"/>
   </option>
  </level>
  <level type="BINARY" prefixMatch="01">
   <option optionKey="1" pattern="01([01]{7})([01]{6})"
           grammar="''01'' employeeID divisionID ">
    <field seq="1" characterSet="[01]*" name="employeeID"/>
    <field seq="2" characterSet="[01]*" name="divisionID"/>
   </option>
  </level>
 </scheme>
</TagDataTranslation>';

  amt := length(buf);
  pos := 1;
  DBMS_LOB.WRITE(tdt_xml, amt, pos, buf);
  DBMS_LOB.CLOSE(tdt_xml);
  DBMS_MGD_ID_UTL.ADD_SCHEME(category_id, tdt_xml);

  -- Step 3. Validate the scheme.
  dbms_output.put_line('Validate the MGD_SAMPLE_CATEGORY Scheme');
  validate_tdtxml := DBMS_MGD_ID_UTL.validate_scheme(tdt_xml);
  dbms_output.put_line(validate_tdtxml);
  dbms_output.put_line('Length of scheme xml is: '||DBMS_LOB.GETLENGTH(tdt_xml));

  -- Step 4. Test tag translation of contractor scheme.
  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     'example.contractor.123.45',
                     NULL, 'BINARY'));

  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     '111111011101101',
                     NULL, 'URI'));

  -- Test tag translation of employee scheme.
  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     'example.employee.123.45',
                     NULL, 'BINARY'));

  dbms_output.put_line(
    mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                     '0111110111101101',
                     NULL, 'URI'));

  DBMS_MGD_ID_UTL.REMOVE_SCHEME(category_id, 'CONTRACTOR_TAG');

  -- Step 6. Test tag translation of contractor scheme. Doesn't work any more.
  BEGIN
    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                       'example.contractor.123.45',
                       NULL, 'BINARY'));

    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
```

```
                    '111111011101101',
                    NULL, 'URI'));
  EXCEPTION
    WHEN others THEN
      dbms_output.put_line('Contractor tag translation failed: '||SQLERRM);
  END;

  -- Step 7. Test tag translation of employee scheme. Still works.
  BEGIN
    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                    'example.employee.123.45',
                    NULL, 'BINARY'));
    dbms_output.put_line(
      mgd_id.translate('MGD_SAMPLE_CATEGORY', NULL,
                    '011111011101101',
                    NULL, 'URI'));
  EXCEPTION
    WHEN others THEN
      dbms_output.put_line('Employee tag translation failed: '||SQLERRM);
  END;

  -- Step 8. Remove the testing category, which also removes all the associated schemes
  DBMS_MGD_ID_UTL.remove_category('MGD_SAMPLE_CATEGORY', '1.0');
END;
/
SHOW ERRORS;

SQL> @add_scheme3.sql
.
.
.
Validate the MGD_SAMPLE_CATEGORY Scheme
EMPLOYEE_TAG;URI,BINARY;divisionID,employeeID
Length of scheme xml is: 933
111111011101101
example.contractor.123.45
011111011101101
example.employee.123.45
Contractor tag translation failed: ORA-55203: Tag data translation level not found
ORA-06512: at "MGDSYS.DBMS_MGD_ID_UTL", line 54
ORA-06512: at "MGDSYS.MGD_ID", line 242
ORA-29532: Java call terminated by uncaught Java
exception: oracle.mgd.idcode.exceptions.TDTLevelNotFound: Matching level not
found for any configured scheme
011111011101101
example.employee.123.45
.
.
.
```

124.4.2 CREATE_CATEGORY Function

This function creates a new category or a new version of a category.

Syntax

```
DBMS_MGD_ID_UTL.CREATE_CATEGORY (
   category_name    IN  VARCHAR2,
   category_version IN  VARCHAR2,
   agency           IN  VARCHAR2,
```

```
  URI              IN  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 124-6 CREATE_CATEGORY Function Parameters

Parameter	Description
category_name	Name of category
category_version	Category version
agency	Organization that owns the category. For example, EPCglobal owns the category EPC.
URI	URI that provides additional information about the category

Usage Notes

The return value is the category ID.

Examples

See the ADD_SCHEME Procedure for an example of creating the MGD_SAMPLE_CATEGORY category.

124.4.3 EPC_TO_ORACLE_SCHEME Function

This function converts the EPCglobal tag data translation (TDT) XML to Oracle Database tag data translation XML.

Syntax

```
DBMS_MGD_ID_UTL.EPC_TO_ORACLE_SCHEME (
   xml_scheme IN CLOB)
 RETURN CLOB;
```

Parameters

Table 124-7 EPC_TO_ORACLE_SCHEME Function Parameters

Parameter	Description
xml_scheme	Name of EPC tag scheme to be converted

Usage Notes

The return value is the contents of the CLOB containing the Oracle Datanase tag data translation XML.

Examples

The following example converts standard EPCglobal Tag Data Translation (TDT) files into Oracle Database TDT files:

```
--Contents of MGD_ID_DOC2.sql
----------------------------
-- EPC_TO_ORACLE_SCHEME   --
----------------------------
```

```
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');

BEGIN
 DBMS_JAVA.set_output(1000000);
 DBMS_OUTPUT.ENABLE(1000000);
 DBMS_MGD_ID_UTL.set_java_logging_level(DBMS_MGD_ID_UTL.LOGGING_LEVEL_SEVERE);
END;
/

DECLARE
  epcScheme                 CLOB;
  oracleScheme              CLOB;
  amt                       NUMBER;
  buf                       VARCHAR2(32767);
  pos                       NUMBER;
  seq                       BINARY_INTEGER;
  validate_epcscheme        VARCHAR2(256);
  validate_oraclescheme     VARCHAR2(256);
BEGIN

  DBMS_LOB.CREATETEMPORARY(epcScheme, true);
  DBMS_LOB.OPEN(epcScheme, DBMS_LOB.LOB_READWRITE);

  buf := '<?xml version="1.0" encoding="UTF-8"?>
<epcTagDataTranslation version="0.04" date="2005-04-18T16:05:00Z"
                       epcTDSVersion="1.1r1.27"
                       xmlns:xsi="http://www.w3.org/2001/XMLSchema"
                       xsi:noNamespaceSchemaLocation="EpcTagDataTranslation.xsd">
  <scheme name="GID-96" optionKey="1" tagLength="96">
    <level type="BINARY" prefixMatch="00110101"
           requiredFormattingParameters="taglength">
      <option optionKey="1" pattern="00110101([01]{28})([01]{24})([01]{36})"
              grammar="''00110101'' generalmanager objectclass serial">
        <field seq="1" decimalMinimum="0" decimalMaximum="268435455"
               characterSet="[01]*" bitLength="28" name="generalmanager"/>
        <field seq="2" decimalMinimum="0" decimalMaximum="16777215"
               characterSet="[01]*" bitLength="24" name="objectclass"/>
        <field seq="3" decimalMinimum="0" decimalMaximum="68719476735"
               characterSet="[01]*" bitLength="36" name="serial"/>
      </option>
    </level>
    <level type="TAG_ENCODING" prefixMatch="urn:epc:tag:gid-96"
           requiredFormattingParameters="taglength">
      <option optionKey="1"
              pattern="urn:epc:tag:gid-96:([0-9]*)\.([0-9]*)\.([0-9]*)"
              grammar="''urn:epc:tag:gid-96:'' generalmanager ''.'' objectclass ''.'' serial">
        <field seq="1" decimalMinimum="0" decimalMaximum="268435455"
               characterSet="[0-9]*" name="generalmanager"/>
        <field seq="2" decimalMinimum="0" decimalMaximum="16777215"
               characterSet="[0-9]*" name="objectclass"/>
        <field seq="3" decimalMinimum="0" decimalMaximum="68719476735"
               characterSet="[0-9]*" name="serial"/>
      </option>
    </level>
    <level type="PURE_IDENTITY" prefixMatch="urn:epc:id:gid">
      <option optionKey="1"
              pattern="urn:epc:id:gid:([0-9]*)\.([0-9]*)\.([0-9]*)"
              grammar="''urn:epc:id:gid:'' generalmanager ''.'' objectclass ''.'' serial">
        <field seq="1" decimalMinimum="0" decimalMaximum="268435455"
               characterSet="[0-9]*" name="generalmanager"/>
        <field seq="2" decimalMinimum="0" decimalMaximum="16777215"
               characterSet="[0-9]*" name="objectclass"/>
```

```
        <field seq="3" decimalMinimum="0" decimalMaximum="68719476735"
               characterSet="[0-9]*" name="serial"/>
      </option>
    </level>
    <level type="LEGACY" prefixMatch="generalmanager=">
      <option optionKey="1"
              pattern="generalmanager=([0-9]*);objectclass=([0-9]*);serial=([0-9]*)"
              grammar="''generalmanager=''generalmanager'';objectclass=''objectclass '';serial=''
serial">
        <field seq="1" decimalMinimum="0" decimalMaximum="268435455"
               characterSet="[0-9]*" name="generalmanager"/>
        <field seq="2" decimalMinimum="0" decimalMaximum="16777215"
               characterSet="[0-9]*" name="objectclass"/>
        <field seq="3" decimalMinimum="0" decimalMaximum="68719476735"
               characterSet="[0-9]*" name="serial"/>
      </option>
    </level>
  </scheme>
</epcTagDataTranslation>';
  amt := length(buf);
  pos := 1;
  DBMS_LOB.WRITE(epcScheme, amt, pos, buf);
  DBMS_LOB.CLOSE(epcScheme);
    oracleScheme := DBMS_MGD_ID_UTL.epc_to_oracle_scheme(epcScheme);
  dbms_output.put_line('Length of oracle scheme xml is: '||DBMS_LOB.GETLENGTH(oracleScheme));
  dbms_output.put_line(DBMS_LOB.SUBSTR(oracleScheme, DBMS_LOB.GETLENGTH(oracleScheme), 1));
  dbms_output.put_line(' ');
  dbms_output.put_line('Validate the Oracle Scheme');
  validate_oraclescheme := DBMS_MGD_ID_UTL.validate_scheme(oracleScheme);
  dbms_output.put_line('Validation result: '||validate_oraclescheme);
END;
/
SHOW ERRORS;

SQL> @mgd_id_doc2.sql
PL/SQL procedure successfully completed.

Length of oracle scheme xml is: 2475
<?xml version = '1.0' encoding = 'UTF-8'?>
<TagDataTranslation version="0.04"
date="2005-04-18T16:05:00Z" xmlns:xsi="http://www.w3.org/2001/XMLSchema"
xmlns="oracle.mgd.idcode"><scheme name="GID-96" optionKey="1" xmlns=""><level
type="BINARY" prefixMatch="00110101" requiredFormattingParameters=""><option
optionKey="1" pattern="00110101([01]{28})([01]{24})([01]{36})"
grammar="'00110101' generalmanager objectclass serial"><field seq="1"
decimalMinimum="0" decimalMaximum="268435455" characterSet="[01]*"
bitLength="28" name="generalmanager"/><field seq="2" decimalMinimum="0"
decimalMaximum="16777215" characterSet="[01]*" bitLength="24"
name="objectclass"/><field seq="3" decimalMinimum="0"
decimalMaximum="68719476735" characterSet="[01]*" bitLength="36"
name="serial"/></option></level><level type="TAG_ENCODING"
prefixMatch="urn:epc:tag:gid-96" requiredFormattingParameters=""><option
optionKey="1" pattern="urn:epc:tag:gid-96:([0-9]*)\.([0-9]*)\.([0-9]*)"
grammar="'urn:epc:tag:gid-96:' generalmanager '.' objectclass '.' serial"><field
seq="1" decimalMinimum="0" decimalMaximum="268435455" characterSet="[0-9]*"
name="generalmanager"/><field seq="2" decimalMinimum="0"
decimalMaximum="16777215" characterSet="[0-9]*" name="objectclass"/><field
seq="3" decimalMinimum="0" decimalMaximum="68719476735" characterSet="[0-9]*"
name="serial"/></option></level><level type="PURE_IDENTITY"
prefixMatch="urn:epc:id:gid"><option optionKey="1"
pattern="urn:epc:id:gid:([0-9]*)\.([0-9]*)\.([0-9]*)" grammar="'urn:epc:id:gid:'
generalmanager '.' objectclass '.' serial"><field seq="1" decimalMinimum="0"
```

```
decimalMaximum="268435455" characterSet="[0-9]*" name="generalmanager"/><field
seq="2" decimalMinimum="0" decimalMaximum="16777215" characterSet="[0-9]*"
name="objectclass"/><field seq="3" decimalMinimum="0"
decimalMaximum="68719476735" characterSet="[0-9]*"
name="serial"/></option></level><level type="LEGACY"
prefixMatch="generalmanager="><option optionKey="1"
pattern="generalmanager=([0-9]*);objectclass=([0-9]*);serial=([0-9]*)"
grammar="'generalmanager='generalmanager';objectclass='objectclass ';serial='
serial"><field seq="1" decimalMinimum="0" decimalMaximum="268435455"
characterSet="[0-9]*" name="generalmanager"/><field seq="2" decimalMinimum="0"
decimalMaximum="16777215" characterSet="[0-9]*" name="objectclass"/><field
seq="3" decimalMinimum="0" decimalMaximum="68719476735" characterSet="[0-9]*"
name="serial"/></option></level></scheme></TagDataTranslation>
Validate the Oracle Scheme
Validation result:
GID-96;LEGACY,TAG_ENCODING,PURE_IDENTITY,BINARY;objectclass,generalmanager,serial,

PL/SQL procedure successfully completed.
.
.
.
```

124.4.4 GET_CATEGORY_ID Function

This function returns the category ID for a given category name and category version.

Syntax

```
DBMS_MGD_ID_UTL.GET_CATEGORY_ID (
   category_name     IN  VARCHAR2,
   category_version  IN  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 124-8 GET_CATEGORY_ID Function Parameters

Parameter	Description
`category_name`	Name of category
`category_version`	Category version

Usage Notes

- If the value of `category_version` is NULL, then the ID of the latest version of the specified category is returned.
- The return value is the category ID for the specified category name.

Examples

The following example returns a category ID given a category name and its version:

```
-- Contents of get_category1.sql file
SELECT DBMS_MGD_ID_UTL.get_category_id('EPC', NULL) FROM DUAL;

SQL> @get_category1.sql
.
.
.
```

```
DBMS_MGD_ID_UTL.GET_CATEGORY_ID('EPC',NULL)-----------------------------------------------------------------
-----------------1
.
.
.
```

124.4.5 GET_COMPONENTS Function

This function returns all relevant separated component names separated by semicolon (;) for the specified scheme.

Syntax

```
DBMS_MGD_ID_UTL.GET_COMPONENTS (
   category_id IN VARCHAR2,
   scheme_name IN VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 124-9 GET_COMPONENTS Function Parameters

Parameter	Description
category_id	Category ID
scheme_name	Name of scheme

Usage Notes

The return value contains the component names separated by a semicolon (;) for the specified scheme.

Examples

The following example gets the components:

```
--Contents of get_components.sql
DECLARE
  id             mgd_id;
  getcomps       VARCHAR2(1000);
  getencodings   VARCHAR2(1000);
  getschemenames VARCHAR2(1000);
BEGIN
 DBMS_MGD_ID_UTL.set_java_logging_level(DBMS_MGD_ID_UTL.LOGGING_LEVEL_OFF);
 DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
 getcomps := DBMS_MGD_ID_UTL.get_components(1,'SGTIN-64');
 dbms_output.put_line('Component names are: ' || getcomps);
 getencodings := DBMS_MGD_ID_UTL.get_encodings(1,'SGTIN-64');
 dbms_output.put_line('Encodings are: ' || getencodings);
 getschemenames := DBMS_MGD_ID_UTL.get_scheme_names(1);
 dbms_output.put_line('Scheme names are: ' || getschemenames);
END;
/
SHOW ERRORS;

SQL> @get_components.sql
.
.
.
Component names are:
```

```
filter,gtin,companyprefixlength,companyprefix,companyprefixindex,itemref,serial
Encodings are: ONS_HOSTNAME,LEGACY,TAG_ENCODING,PURE_IDENTITY,BINARY
Scheme names are:
GIAI-64,GIAI-96,GID-96,GRAI-64,GRAI-96,SGLN-64,SGLN-96,SGTIN-64,SGTIN-96,SSCC-64
,SSCC-96,USDOD-64,USDOD-96
PL/SQL procedure successfully completed.
.
.
.
```

124.4.6 GET_ENCODINGS Function

This function returns a list of semicolon (;) separated encodings (formats) for the specified scheme.

Syntax

```
DBMS_MGD_ID_UTL.GET_ENCODINGS (
   category_id IN VARCHAR2,
   scheme_name IN VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 124-10 GET_ENCODINGS Function Parameters

Parameter	Description
category_id	Category ID
scheme_name	Name of scheme

Usage Notes

The return value contains the encodings separated by a semicolon (;) for the specified scheme.

Examples

See the GET_COMPONENTS Function for an example.

124.4.7 GET_JAVA_LOGGING_LEVEL Function

This function returns an integer representing the current trace logging level.

Syntax

```
DBMS_MGD_ID_UTL.GET_JAVA_LOGGING_LEVEL
 RETURN INTEGER;
```

Usage Notes

The return value is the integer value denoting the current Java logging level.

Examples

The following example gets the Java logging level.

```
--Contents of getjavaloggінglevel.sql
DECLARE
```

```
     loglevel   NUMBER;
BEGIN
  DBMS_MGD_ID_UTL.set_java_logging_level(DBMS_MGD_ID_UTL.LOGGING_LEVEL_OFF);
  loglevel := DBMS_MGD_ID_UTL.get_java_logging_level();
  dbms_output.put_line('Java logging level = ' ||loglevel);
END;
/
SHOW ERRORS;

SQL> @getjavalogginglevel.sql
.
.
.
Java logging level = 0
PL/SQL procedure successfully completed.
.
.
.
```

124.4.8 GET_PLSQL_LOGGING_LEVEL Function

This function returns an integer representing the current PL/SQL trace logging level.

Syntax

```
DBMS_MGD_ID_UTL.GET_PLSQL_LOGGING_LEVEL
 RETURN INTEGER;

PRAGMA restrict_references(get_plsql_logging_level, WNDS);
```

Usage Notes

The return value is the integer value denoting the current PL/SQL logging level.

Examples

The following example gets the PL/SQL logging level.

```
--Contents of getplsqllogginglevel.sql
DECLARE
     loglevel   NUMBER;
BEGIN
  DBMS_MGD_ID_UTL.set_plsql_logging_level(0);
  loglevel := DBMS_MGD_ID_UTL.get_plsql_logging_level();
  dbms_output.put_line('PL/SQL logging level = ' ||loglevel);
END;
/
SHOW ERRORS;

SQL> @getplsqllogginglevel.sql
.
.
.
PL/SQL logging level = 0
PL/SQL procedure successfully completed.
.
.
.
```

124.4.9 GET_SCHEME_NAMES Function

This function returns a list of semicolon (;) separated scheme names for the specified category.

Syntax

```
DBMS_MGD_ID_UTL.GET_SCHEME_NAMES (
   category_id IN VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 124-11 GET_SCHEME_NAMES Function Parameters

Parameter	Description
category_id	Category ID

Usage Notes

The return value contains the scheme names for the specified category ID.

Examples

See the GET_COMPONENTS Function for an example.

124.4.10 GET_TDT_XML Function

This function returns the Oracle Database tag data translation XML for the specified scheme.

Syntax

```
DBMS_MGD_ID_UTL.GET_TDT_XML (
   category_id  IN VARCHAR2,
   scheme_name  IN VARCHAR2)
 RETURN CLOB;
```

Parameters

Table 124-12 GET_TDT_XML Function Parameters

Parameter	Description
category_id	Category ID
scheme_name	Name of scheme

Usage Notes

The return value contains the Oracle Database tag data translation XML for the specified scheme.

Examples

The following example gets the Oracle Database TDT XML for the specified scheme:

```
--Contents of get_tdtxml.sql
DECLARE
  gettdtxml      CLOB;

BEGIN
  gettdtxml := DBMS_MGD_ID_UTL.get_tdt_xml(1,'SGTIN-64');
  dbms_output.put_line('Length of tdt XML is '||DBMS_LOB.GETLENGTH(gettdtxml));
  dbms_output.put_line(DBMS_LOB.SUBSTR(gettdtxml, DBMS_LOB.GETLENGTH(gettdtxml), 1));
END;
/
SHOW ERRORS;

SQL> @get_tdtxml.sql
.
.
.
Length of tdt XML is 22884
<?xml version = '1.0' encoding = "UTF-8"?>
<TagDataTranslation version="0.04"
date="2005-04-18T16:05:00Z" xmlns:xsi="http://www.w3.org/2001/XMLSchema"
xmlns="oracle.mgd.idcode"><scheme name="SGTIN-64"
optionKey="companyprefixlength" xmlns="">
    <level type="BINARY"
prefixMatch="10" requiredFormattingParameters="filter">
      <option
optionKey="12" pattern="10([01]{3})([01]{14})([01]{20})([01]{25})" grammar="'10'
filter companyprefixindex itemref serial">
        <field seq="1"
decimalMinimum="0" decimalMaximum="7" characterSet="[01]*" bitLength="3"
length="1" padChar="0" padDir="LEFT" name="filter"/>
        <field seq="2"
decimalMinimum="0" decimalMaximum="16383" characterSet="[01]*" bitLength="14"
name="companyprefixindex"/>
        <field seq="3" decimalMinimum="0"
decimalMaximum="9" characterSet="[01]*" bitLength="20" length="1" padChar="0"
padDir="LEFT" name="itemref"/>
        <field seq="4" decimalMinimum="0"
decimalMaximum="33554431" characterSet="[01]*" bitLength="25" name="serial"/>
.
.
.
        <field seq="1" decimalMinimum="0" decimalMaximum="9999999" characterSet="[0-9]*"
length="7" padChar="0" padDir="LEFT" name="itemref"/>
        <field seq="2" decimalMinimum="0" decimalMaximum="999999" characterSet="[0-9]*" length="6"
padChar="0" padDir="LEFT" name="companyprefix"/>
      </option>
    </level>

</scheme></TagDataTranslation>
PL/SQL procedure successfully completed.
.
.
.
```

124.4.11 GET_VALIDATOR Function

This function returns the Oracle Database tag data translation schema.

Syntax

```
DBMS_MGD_ID_UTL.GET_VALIDATOR
 RETURN CLOB;
```

Usage Notes

The return value contains the Oracle Database tag data translation schema.

Examples

This example returns the Oracle Database TDT schema.

```
--Contents of get_validator.sql
DECLARE
  getvalidator              CLOB;
BEGIN
  getvalidator := DBMS_MGD_ID_UTL.get_validator;
  dbms_output.put_line('Length of validated oracle scheme xml is '||DBMS_LOB.GETLENGTH(getvalidator));
  dbms_output.put_line(DBMS_LOB.SUBSTR(getvalidator, DBMS_LOB.GETLENGTH(getvalidator), 1));
END;
/
SHOW ERRORS;

SQL> @get_validator.sql
.
.
.
Length of validated oracle scheme xml is 5780
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema
targetNamespace="oracle.mgd.idcode"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"

xmlns:tdt="oracle.mgd.idcode" elementFormDefault="unqualified"

attributeFormDefault="unqualified" version="1.0">
  <xsd:annotation>

<xsd:documentation>
      <![CDATA[
<epcglobal:copyright>Copyright ?2004
Epcglobal Inc., All
Rights
Reserved.</epcglobal:copyright>
<epcglobal:disclaimer>EPCglobal Inc., its
members, officers, directors,
employees, or agents shall not be liable for any
injury, loss, damages,
financial or otherwise, arising from, related to, or
caused by the use of this
document. The use of said document shall constitute
your express consent to
the foregoing
exculpation.</epcglobal:disclaimer>
<epcglobal:specification>Tag Data
Translation (TDT) version
```

```
1.0</epcglobal:specification>
]>

</xsd:documentation>
  </xsd:annotation>
  <xsd:simpleType
name="LevelTypeList">
    <xsd:restriction base="xsd:string">

</xsd:restriction>
  </xsd:simpleType>
  <xsd:simpleType name="TagLengthList"

<xsd:restriction base="xsd:string">
    </xsd:restriction>
  </xsd:simpleType>

<xsd:simpleType name="SchemeNameList">
    <xsd:restriction base="xsd:string">

</xsd:restriction>
  </xsd:simpleType>
  <xsd:simpleType
name="InputFormatList">
    <xsd:restriction base="xsd:string">

<xsd:enumeration value="BINARY"/>
      <xsd:enumeration value="STRING"/>

</xsd:restriction>
  </xsd:simpleType>
  <xsd:simpleType name="ModeList">

<xsd:restriction base="xsd:string">
      <xsd:enumeration value="EXTRACT"/>

<xsd:enumeration value="FORMAT"/>
    </xsd:restriction>
  </xsd:simpleType>

<xsd:simpleType name="CompactionMethodList">
    <xsd:restriction
base="xsd:string">
      <xsd:enumeration value="32-bit"/>
<xsd:enumeration value="16-bit"/>
      <xsd:enumeration value="8-bit"/>

<xsd:enumeration value="7-bit"/>      <xsd:enumeration value="6-bit"/>

<xsd:enumeration value="5-bit"/>
    </xsd:restriction>
  </xsd:simpleType>

<xsd:simpleType name="PadDirectionList">
    <xsd:restriction
base="xsd:string">
      <xsd:enumeration value="LEFT"/>
      <xsd:enumeration
value="RIGHT"/>
    </xsd:restriction>
  </xsd:simpleType>
  <xsd:complexType
```

```
name="Field">
    <xsd:attribute name="seq" type="xsd:integer" use="required"/>

<xsd:attribute name="name" type="xsd:string" use="required"/>
    <xsd:attribute
name="bitLength" type="xsd:integer"/>
    <xsd:attribute name="characterSet"
type="xsd:string" use="required"/>
    <xsd:attribute name="compaction"
type="tdt:CompactionMethodList"/>
    <xsd:attribute name="compression"
type="xsd:string"/>
    <xsd:attribute name="padChar" type="xsd:string"/>

<xsd:attribute name="padDir" type="tdt:PadDirectionList"/>
    <xsd:attribute
name="decimalMinimum" type="xsd:long"/>
    <xsd:attribute name="decimalMaximum"
type="xsd:long"/>
    <xsd:attribute name="length" type="xsd:integer"/>

</xsd:complexType>
  <xsd:complexType name="Option">
    <xsd:sequence>

<xsd:element name="field" type="tdt:Field" maxOccurs="unbounded"/>

</xsd:sequence>
    <xsd:attribute name="optionKey" type="xsd:string"
use="required"/>
    <xsd:attribute name="pattern" type="xsd:string"/>

<xsd:attribute name="grammar" type="xsd:string" use="required"/>

</xsd:complexType>
  <xsd:complexType name="Rule">
    <xsd:attribute
name="type" type="tdt:ModeList" use="required"/>
    <xsd:attribute
name="inputFormat" type="tdt:InputFormatList"

use="required"/>
    <xsd:attribute name="seq" type="xsd:integer"
use="required"/>
    <xsd:attribute name="newFieldName" type="xsd:string"
use="required"/>
    <xsd:attribute name="characterSet" type="xsd:string"
use="required"/>
    <xsd:attribute name="padChar" type="xsd:string"/>

<xsd:attribute name="padDir" type="tdt:PadDirectionList"/>
    <xsd:attribute
name="decimalMinimum" type="xsd:long"/>
    <xsd:attribute name="decimalMaximum"
type="xsd:long"/>
    <xsd:attribute name="length" type="xsd:string"/>

<xsd:attribute name="function" type="xsd:string" use="required"/>

<xsd:attribute name="tableURI" type="xsd:string"/>
    <xsd:attribute
name="tableParams" type="xsd:string"/>
```

```
    <xsd:attribute name="tableXPath"
type="xsd:string"/>
    <xsd:attribute name="tableSQL" type="xsd:string"/>

</xsd:complexType>
  <xsd:complexType name="Level">
    <xsd:sequence>
<xsd:element name="option" type="tdt:Option" minOccurs="1"

maxOccurs="unbounded"/>
      <xsd:element name="rule" type="tdt:Rule"
minOccurs="0"
                   maxOccurs="unbounded"/>
    </xsd:sequence>
<xsd:attribute name="type" type="tdt:LevelTypeList" use="required"/>
<xsd:attribute name="prefixMatch" type="xsd:string" use="optional"/>
<xsd:attribute name="requiredParsingParameters" type="xsd:string"/>
<xsd:attribute name="requiredFormattingParameters" type="xsd:string"/>

</xsd:complexType>
  <xsd:complexType name="Scheme">
    <xsd:sequence>
<xsd:element name="level" type="tdt:Level" minOccurs="1" maxOccurs="5"/>

</xsd:sequence>
    <xsd:attribute name="name" type="tdt:SchemeNameList"
use="required"/>
    <xsd:attribute name="optionKey" type="xsd:string"
use="required"/>
    <xsd:attribute name="tagLength" type="tdt:TagLengthList"
use="optional"/>
  </xsd:complexType>
  <xsd:complexType
name="TagDataTranslation">
    <xsd:sequence>
      <xsd:element name="scheme"
type="tdt:Scheme" maxOccurs="unbounded"/>
    </xsd:sequence>
    <xsd:attribute
name="version" type="xsd:string" use="required"/>
    <xsd:attribute name="date"
type="xsd:dateTime" use="required"/>
  </xsd:complexType>
  <xsd:element
name="TagDataTranslation" type="tdt:TagDataTranslation"/>
</xsd:schema>

PL/SQL procedure successfully completed.
.
.
.
```

124.4.12 REFRESH_CATEGORY Function

This function refreshes the metadata information on the Java stack for the specified category.

This function must be called before using `MGD_ID` functions.

Syntax

```
DBMS_MGD_ID_UTL.REFRESH_CATEGORY (
   category_id  IN VARCHAR2);
```

Parameters

Table 124-13 REFRESH_CATEGORY Function Parameters

Parameter	Description
category_id	Category ID

Examples

The following example refreshes the metadata information for the EPC category ID.

```
--Contents of tostring3.sql
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
DECLARE
id          MGD_ID;
BEGIN
 DBMS_MGD_ID_UTL.set_java_logging_level(DBMS_MGD_ID_UTL.LOGGING_LEVEL_OFF);
 DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
 dbms_output.put_line('..Testing to_string');
 DBMS_OUTPUT.PUT_LINE('test to_string');
   id := mgd_id('EPC', NULL, 'urn:epc:id:gid:0037000.30241.1041970', 'scheme=GID-96');
      DBMS_OUTPUT.PUT_LINE('mgd_id object as a string');
      DBMS_OUTPUT.PUT_LINE(id.to_string);
END;
/
SHOW ERRORS;
call DBMS_MGD_ID_UTL.remove_proxy();

SQL> @tostring3.sql
..Testing to_string
test to_string
mgd_id object as a string
category_id =1;schemes = GID-96;objectclass = 30241;generalmanager =
0037000;scheme = GID-96;1 = 1;serial = 1041970

PL/SQL procedure successfully completed.
```

124.4.13 REMOVE_CATEGORY Procedure

This procedure removes a category including all the related TDT XML.

This procedure is overloaded. The different functionality of each form of syntax is presented along with the definitions.

Syntax

Removes a category based on the specified category ID.

```
DBMS_MGD_ID_UTL.REMOVE_CATEGORY (
   category_id  IN VARCHAR2);
```

Removes a category based on the specified category name and category version.

```
DBMS_MGD_ID_UTL.REMOVE_CATEGORY (
   category_name    IN VARCHAR2,
   category_version IN VARCHAR2);
```

Parameters

Table 124-14 REMOVE_CATEGORY Procedure Parameters

Parameter	Description
category_id	Category ID
category_name	Name of category
category_version	Category version

Usage Notes

If the value of category_version is NULL, all versions for the specified category will be removed.

Examples

See the ADD_SCHEME Procedure for an example of removing a category.

124.4.14 REMOVE_PROXY Procedure

This procedure unsets the host and port of the proxy server.

Syntax

```
DBMS_MGD_ID_UTL.REMOVE_PROXY;
```

Examples

See the REFRESH_CATEGORY Function for an example.

124.4.15 REMOVE_SCHEME Procedure

This procedure removes a tag data translation scheme from a category.

Syntax

```
DBMS_MGD_ID_UTL.REMOVE_SCHEME (
   category_id  IN  VARCHAR2,
   scheme_name  IN  VARCHAR2);
```

Parameters

Table 124-15 REMOVE_SCHEME Procedure Parameters

Parameter	Description
category_id	Category ID
scheme_name	Name of scheme

Examples

See the ADD_SCHEME Procedure for an example of removing a scheme.

124.4.16 SET_JAVA_LOGGING_LEVEL Procedure

This procedure sets the Java trace logging level.

Syntax

```
DBMS_MGD_ID_UTL.SET_JAVA_LOGGING_LEVEL (
   logginglevel IN  INTEGER);
```

Parameters

Table 124-16 SET_JAVA_LOGGING_LEVEL Procedure Parameters

Parameter	Description
`logginglevel`	Logging level. The Java logging level can be one of the following values in descending order: • `LOGGING_LEVEL_OFF CONSTANT INTEGER := 0` • `LOGGING_LEVEL_SEVERE CONSTANT INTEGER := 1` • `LOGGING_LEVEL_WARNING CONSTANT INTEGER := 2` • `LOGGING_LEVEL_INFO CONSTANT INTEGER := 3` • `LOGGING_LEVEL_FINE CONSTANT INTEGER := 4` • `LOGGING_LEVEL_FINER CONSTANT INTEGER := 5` • `LOGGING_LEVEL_FINEST CONSTANT INTEGER := 6` • `LOGGING_LEVEL_ALL CONSTANT INTEGER := 7`

Examples

See the GET_JAVA_LOGGING_LEVEL Function for an example.

124.4.17 SET_PLSQL_LOGGING_LEVEL Procedure

This procedure sets the PL/SQL trace logging level.

Syntax

```
DBMS_MGD_ID_UTL.SET_PLSQL_LOGGING_LEVEL (
   level IN INTEGER);

PRAGMA restrict_references(set_plsql_logging_level, WNDS);
```

Parameters

Table 124-17 SET_PLSQL_LOGGING_LEVEL Procedure Parameters

Parameter	Description
level	Logging level. The PL/SQL logging level can be one of the following values in descending order: • LOGGING_LEVEL_OFF CONSTANT INTEGER := 0 • LOGGING_LEVEL_SEVERE CONSTANT INTEGER := 1 • LOGGING_LEVEL_WARNING CONSTANT INTEGER := 2 • LOGGING_LEVEL_INFO CONSTANT INTEGER := 3 • LOGGING_LEVEL_FINE CONSTANT INTEGER := 4 • LOGGING_LEVEL_FINER CONSTANT INTEGER := 5 • LOGGING_LEVEL_FINEST CONSTANT INTEGER := 6 • LOGGING_LEVEL_ALL CONSTANT INTEGER := 7

Examples

See the GET_PLSQL_LOGGING_LEVEL Function for an example.

124.4.18 SET_PROXY Procedure

This procedure sets the host and port of the proxy server for Internet access.

This procedure must be called if the database server accesses the Internet using a proxy server. Internet access is necessary because some rules need to look up the Object Naming Service (ONS) table to get the company prefix index.

You do not need to call this procedure does if you are only using schemes that do not contain any rules requiring Internet access.

Syntax

```
DBMS_MGD_ID_UTL.SET_PROXY (
   proxt_host   IN  VARCHAR2,
   proxy_port   IN  VARCHAR2);
```

Parameters

Table 124-18 SET_PROXY Procedure Parameters

Parameter	Description
proxy_host	Name of host
proxy_port	Host port number

Examples

See the REFRESH_CATEGORY Function for an example.

124.4.19 VALIDATE_SCHEME Function

This function validates the input tag data translation XML against the Oracle Database tag data translation schema.

Syntax

```
DBMS_MGD_ID_UTL.VALIDATE_SCHEME (
   xml_scheme IN CLOB)
 RETURN VARCHAR2;
```

Parameters

Table 124-19 VALIDATE_SCHEME Function Parameters

Parameter	Description
xml_scheme	Scheme to be validated.

Usage Notes

The return value contains the components names for the specified scheme.

Examples

See the ADD_SCHEME Procedure or the EPC_TO_ORACLE_SCHEME Function for an example.

125
DBMS_MGWADM

`DBMS_MGWADM` defines the Messaging Gateway administrative interface. The package and object types are owned by `SYS`.

> **Note:**
>
> - `DBMS_MGWADM` is deprecated in 23ai and will be desupported in a future release.
> - You must run the `catmgw.sql` script located under the `$ORACLE_HOME/mgw/admin` directory to load the Messaging Gateway packages and types into the database.

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* contains information on loading database objects and using `DBMS_MGWADM`

This chapter contains the following topics:

- Security Model
- Deprecated Subprograms
- Constants
- Data Structures
- Summary of DBMS_MGWADM Subprograms

125.1 DBMS_MGWADM Security Model

A user with administrative privilege can run all procedures in `DBMS_MGWADM`.

125.2 DBMS_MGWADM Deprecated Subprograms

Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

The following subprograms are deprecated with Oracle Database 11*g* Release 1 having been superseded by improved technology:

- ADD_SUBSCRIBER Procedure - use instead CREATE_JOB Procedure
- ALTER_PROPAGATION_SCHEDULE Procedure - use instead ALTER_JOB Procedure
- ALTER_SUBSCRIBER Procedure - use instead ALTER_JOB Procedure
- DB_CONNECT_INFO Procedure - use instead ALTER_AGENT Procedures

- DISABLE_PROPAGATION_SCHEDULE Procedure - use instead DISABLE_JOB Procedure
- ENABLE_PROPAGATION_SCHEDULE Procedure - use instead ENABLE_JOB Procedure
- REMOVE_SUBSCRIBER Procedure - use instead REMOVE_JOB Procedure
- RESET_SUBSCRIBER Procedure - use instead RESET_JOB Procedure
- SCHEDULE_PROPAGATION Procedure - use instead CREATE_JOB Procedure
- UNSCHEDULE_PROPAGATION Procedure - use instead REMOVE_JOB Procedure

125.3 DBMS_MGWADM Constants

The `DBMS_MGWADM` package defines various constants for specifying parameter values.

- Table 125-1
- Table 125-2
- Table 125-3
- Table 125-4
- Table 125-5
- Table 125-6
- Table 125-7
- Table 125-8
- Table 125-9
- Table 125-10
- Table 125-11

Table 125-1 DBMS_MGWADM Constants—Cleanup Actions

Name	Type	Description
CLEAN_STARTUP_STATE	CONSTANT BINARY_INTEGER	Sets the Messaging Gateway agent to a known state so that it can be started
CLEAN_LOG_QUEUES	CONSTANT BINARY_INTEGER	Messaging Gateway agent will clean log queues for all configured messaging system links
RESET_SUB_MISSING_LOG_REC	CONSTANT BINARY_INTEGER	Messaging Gateway agent recovers a Messaging Gateway subscriber that has failed due to a missing log record
RESET_SUB_MISSING_MESSAGE	CONSTANT BINARY_INTEGER	Messaging Gateway agent recovers a Messaging Gateway subscriber that has failed due to a missing persistent source message

Table 125-2 DBMS_MGWADM Constants—Force Values

Name	Type	Description
FORCE	CONSTANT BINARY_INTEGER	Represents a forced action
NO_FORCE	CONSTANT BINARY_INTEGER	Represents a normal, nonforced action

Table 125-3 DBMS_MGWADM Constants—Logging Levels

Name	Type	Description
BASIC_LOGGING	CONSTANT BINARY_INTEGER	The standard (the least) information written to the log file
TRACE_DEBUG_LOGGING	CONSTANT BINARY_INTEGER	The greatest information written to the log file
TRACE_HIGH_LOGGING	CONSTANT BINARY_INTEGER	The third level of detail of logging information written to the log file
TRACE_LITE_LOGGING	CONSTANT BINARY_INTEGER	The second level detail of logging information written to the log file

Table 125-4 DBMS_MGWADM Constants—Named Property Constants

Name	Type	Description
MGWPROP_PREFIX	CONSTANT VARCHAR2	A constant (MGWPROP$_) for the reserved property name prefix
MGWPROP_REMOVE	CONSTANT VARCHAR2	A constant (MGWPROP$_REMOVE) for the reserved property name used to remove an existing property
MGWPROP_REMOVE_ALL	CONSTANT VARCHAR2	A constant (MGWPROP$_REMOVE_ALL) for the reserved property name used to remove all properties

Table 125-5 DBMS_MGWADM Constants—Other Constants

Name	Type	Description
JMS_CONNECTION	CONSTANT BINARY_INTEGER	Used to indicate that JMS connections will be used to access JMS destinations in a domain-independent manner that supports a unified messaging model
JMS_QUEUE_CONNECTION	CONSTANT BINARY_INTEGER	Used to indicate that JMS queue connections will be used to access JMS destinations
JMS_TOPIC_CONNECTION	CONSTANT BINARY_INTEGER	Used to indicate that JMS topic connections will be used to access JMS destinations
NO_CHANGE	CONSTANT VARCHAR2	Indicates that an existing value should be preserved (not changed). This is used for certain APIs where the desire is to change one or more parameters but leave others unchanged.
DEFAULT_AGENT	CONSTANT VARCHAR2	Name of the Messaging Gateway default agent

Table 125-6 DBMS_MGWADM Constants—Propagation Types

Name	Type	Description
INBOUND_PROPAGATION	CONSTANT BINARY_INTEGER	Represents the propagation type for non-Oracle to Oracle Database Advanced Queuing propagation. The propagation source is a queue in a foreign (non-Oracle) messaging system and the destination is a local Oracle Database Advanced Queuing queue.
OUTBOUND_PROPAGATION	CONSTANT BINARY_INTEGER	Represents the propagation type for Oracle Database Advanced Queuing to non-Oracle propagation. The propagation source is a local Oracle Database Advanced Queuing queue and the destination is a queue in a foreign (non-Oracle) messaging system.

Table 125-7 DBMS_MGWADM Constants—Queue Domain Types

Name	Type	Description
DOMAIN_QUEUE	CONSTANT BINARY_INTEGER	Represents a queue destination. A JMS queue (point-to-point model) is classified as a queue.
DOMAIN_TOPIC	CONSTANT BINARY_INTEGER	Represents a topic destination. A JMS topic (publish-subscribe model) is classified as a topic.

Table 125-8 DBMS_MGWADM Constants—Shutdown Modes

Name	Type	Description
SHUTDOWN_IMMEDIATE	CONSTANT BINARY_INTEGER	Represents the immediate shutdown mode
SHUTDOWN_NORMAL	CONSTANT BINARY_INTEGER	Represents the normal shutdown mode

Table 125-9 DBMS_MGWADM Constants—WebSphere MQ Interface Types

Name	Type	Description
MQSERIES_BASE_JAVA_INTERFACE	CONSTANT BINARY_INTEGER	Represents the Base Java interface for the WebSphere MQ messaging system

Table 125-10 DBMS_MGWADM Constants—target_type Argument of SET_OPTION and REMOVE_OPTION Procedures

Name	Type	Description
AGENT_JAVA_PROP	CONSTANT PLS_INTEGER	Used for an agent option used to set a Java System property
MSGLINK_OPTION	CONSTANT PLS_INTEGER	Used for a messaging system link option
JOB_OPTION	CONSTANT PLS_INTEGER	Used for a propagation job option

Table 125-11 DBMS_MGWADM Constants—conntype Argument of CREATE_AGENT and ALTER_AGENT Procedures

Name	Type	Description
JDBC_OCI	CONSTANT VARCHAR2	Used to specify the JDBC OCI driver
JDBC_THIN	CONSTANT VARCHAR2	Used to specify the JDBC Thin driver

125.4 DBMS_MGWADM Data Structures

The DBMS_MGWADM package defines several OBJECT types.

DBMS_MGWADM Object Types

- SYS.MGW_MQSERIES_PROPERTIES Object Type
- SYS.MGW_PROPERTIES Object Type
- SYS.MGW_PROPERTY Object Type
- SYS.MGW_TIBRV_PROPERTIES Object Type

125.4.1 DBMS_MGWADM SYS.MGW_MQSERIES_PROPERTIES Object Type

This type specifies basic properties for a WebSphere MQ messaging system link.

Syntax

```
TYPE SYS.MGW_MQSERIES_PROPERTIES IS OBJECT (
   queue_manager       VARCHAR2(64),
   hostname            VARCHAR2(64),
   port                INTEGER,
   channel             VARCHAR2(64),
   interface_type      INTEGER,
   username            VARCHAR2(64),
   password            VARCHAR2(64),
   inbound_log_queue   VARCHAR2(64),
   outbound_log_queue  VARCHAR2(64),

-- Methods
STATIC FUNCTION construct
RETURN SYS.MGW_MQSERIES_PROPERTIES,

STATIC FUNCTION alter_construct
RETURN SYS.MGW_MQSERIES_PROPERTIES );
```

Attributes

Table 125-12 SYS.MGW_MQSERIES_PROPERTIES Attributes

Attribute	Description
queue_manager	The name of the WebSphere MQ queue manager

Table 125-12 (Cont.) SYS.MGW_MQSERIES_PROPERTIES Attributes

Attribute	Description
hostname	The host on which the WebSphere MQ messaging system resides. If hostname is NULL, then a WebSphere MQ bindings connection is used. If not NULL, then a client connection is used and requires that a port and channel be specified.
port	The port number. This is used only for client connections; that is, when hostname is not NULL.
channel	The channel used when establishing a connection to the queue manager. This is used only for client connections; that is, when hostname is not NULL.
interface_type	The type of messaging interface to use. Values: • DBMS_MGWADM.MQSERIES_BASE_JAVA_INTERFACE if the WebSphere MQ Base Java interface should be used. • DBMS_MGWADM.JMS_CONNECTION if the link is to be used to access JMS destinations in a unified, domain-independent manner. • DBMS_MGWADM.JMS_QUEUE_CONNECTION if the link is to be used for accessing JMS queues • DBMS_MGWADM.JMS_TOPIC_CONNECTION if the link is to be used for accessing JMS topics.
username	The username used for authentication to the WebSphere MQ messaging system
password	The password used for authentication to the WebSphere MQ messaging system
inbound_log_queue	The name of the WebSphere MQ queue used for propagation recovery purposes when this messaging link is used for inbound propagation; that is, when queues associated with this link serve as a propagation source: • For MQSERIES_BASE_JAVA_INTERFACE, this is the name of a physical WebSphere MQ queue created using WebSphere MQ administration tools. • For the JMS_CONNECTION interface and the JMS_QUEUE_CONNECTION interface, this is the name of a physical WebSphere MQ queue created using WebSphere MQ administration tools. • For JMS_TOPIC_CONNECTION interface, this specifies the name of a WebSphere MQ JMS topic. The physical WebSphere MQ queue used by subscribers of that topic must be created using WebSphere MQ administration tools. By default, the physical queue used is SYSTEM.JMS.D.SUBSCRIBER.QUEUE.

Table 125-12 (Cont.) SYS.MGW_MQSERIES_PROPERTIES Attributes

Attribute	Description
outbound_log_queue	The name of the WebSphere MQ queue used for propagation recovery purposes when this messaging link is used for outbound propagation; that is, when queues associated with this link serve as a propagation destination: • For MQSERIES_BASE_JAVA_INTERFACE, this is the name of a physical WebSphere MQ queue created using WebSphere MQ administration tools. • For the JMS_CONNECTION interface and the JMS_QUEUE_CONNECTION interface, this is the name of a physical WebSphere MQ queue created using WebSphere MQ administration tools. • For JMS_TOPIC_CONNECTION interface, this specifies the name of a WebSphere MQ JMS topic. The physical WebSphere MQ queue used by subscribers of that topic must be created using WebSphere MQ administration tools. By default, the physical queue used is SYSTEM.JMS.D.SUBSCRIBER.QUEUE.

Methods

Table 125-13 SYS.MGW_MQSERIES_PROPERTIES Methods

Method	Description
construct	Constructs a new SYS.MGW_MQSERIES_PROPERTIES instance. All attributes are assigned a value of NULL
alter_construct	Constructs a new SYS.MGW_MQSERIES_PROPERTIES instance for altering the properties of an existing messaging link. All attributes having a VARCHAR2 datatype are assigned a value of DBMS_MGWADM.NO_CHANGE. Attributes of other datatypes are assigned a value of NULL.

125.4.2 DBMS_MGWADM SYS.MGW_PROPERTIES Object Type

This type specifies an array of properties.

Syntax

```
TYPE SYS.MGW_PROPERTIES AS VARRAY (2000) OF SYS.MGW_PROPERTY;
```

Attributes

Table 125-14 SYS.MGW_PROPERTIES Attributes

Attribute	Description
name	Property name
value	Property value

Usage Notes

Unless noted otherwise, Messaging Gateway uses named properties as follows:

- Names with the `MGWPROP$_` prefix are reserved. They are used for special purposes and are invalid when used as a normal property name.
- A property name can exist only once in a property list; that is, a list can contain only one value for a given name. The name is case-insensitive.
- In general, a property list is order-independent, and the property names may appear in any order. An alter property list is an exception.
- You can use a new property list to alter an existing property list. Each new property modifies the original list in one of the following ways: adds a new property, modifies a property, removes a property, or removes all properties.

The alter list is processed in order, from the first element to the last element. Thus the order in which the elements appear in the alter list is meaningful, especially when the alter list is used to remove properties from an existing list.

The property name and value are used to determine how that element affects the original list. The following rules apply:

- Add or modify property

  ```
  MGW_PROPERTY.NAME = property_name
  MGW_PROPERTY.VALUE = property_value
  ```

  If a property of the given name already exists, then the current value is replaced with the new value; otherwise the new property is added to the end of the list.

- Remove property

  ```
  MGW_PROPERTY.NAME = 'MGWPROP$_REMOVE'
  MGW_PROPERTY.VALUE = name_of_property_to_remove
  ```

  No action is taken if the property name does not exist in the original list.

- Remove all properties

  ```
  MGW_PROPERTY.NAME = 'MGWPROP$_REMOVE_ALL'
  MGW_PROPERTY.VALUE = not used
  ```

> **See Also:**
>
> "The `DBMS_MGWADM` package defines constants to represent the reserved property names on Table 125-4

125.4.3 DBMS_MGWADM SYS.MGW_PROPERTY Object Type

This type specifies a named property which is used to specify optional properties for messaging links, foreign queues, and subscribers.

Syntax

```
TYPE SYS.MGW_PROPERTY IS OBJECT(
   name  VARCHAR2(500),
   value VARCHAR2(4000),

-- Methods
STATIC FUNCTION construct    --- (1)
RETURN SYS.MGW_PROPERTY,
```

```
STATIC FUNCTION construct(   --- (2)
   p_name   IN VARCHAR2,
   p_value  IN VARCHAR2)
RETURN SYS.MGW_PROPERTY );
```

Attributes

Table 125-15 SYS.MGW_PROPERTY Attributes

Attribute	Description
`name`	Property name
`value`	Property value

Methods

Table 125-16 SYS.MGW_PROPERTY Methods

Method	Description
`construct --- (1)`	Constructs a new `MGW_PROPERTY` instance. All attributes are assigned a value of `NULL`
`construct --- (2)`	Constructs a new `MGW_PROPERTY` instance initialized using the given parameters

> **See Also:**
>
> "The `DBMS_MGWADM` package defines constants to represent the reserved property names on Table 125-4

125.4.4 SYS.MGW_TIBRV_PROPERTIES Object Type

A type that specifies basic properties for a TIB/Rendezvous messaging system link. The Messaging Gateway agent creates a TIB/Rendezvous transport of type `TibrvRvdTransport` for each Messaging Gateway link.

Syntax

```
TYPE SYS.MGW_TIBRV_PROPERTIES IS OBJECT(
   service   VARCHAR2(128),
   daemon    VARCHAR2(128),
   network   VARCHAR2(256),
   cm_name   VARCHAR2(256),
   cm_ledger VARCHAR2(256),

-- Methods
STATIC FUNCTION construct
RETURN SYS.MGW_TIBRV_PROPERTIES,

STATIC FUNCTION alter_construct
RETURN SYS.MGW_TIBRV_PROPERTIES );
```

Attributes

Table 125-17 SYS.MGW_TIBRV_PROPERTIES Attributes

Attribute	Description
service	The service parameter for the rvd transport
daemon	The daemon parameter for the rvd transport
network	The network parameter for the rvd transport
cm_name	The CM correspondent name. Reserved for future use.
cm_ledger	The CM ledger file name. Reserved for future use.

Methods

Table 125-18 SYS.MGW_TIBRV_PROPERTIES Methods

Method	Description
construct	Constructs a new SYS.MGW_TIBRV_PROPERTIES instance. All attributes will be assigned a value of NULL.
alter_construct	Constructs a new SYS.MGW_TIBRV_PROPERTIES instance. This function is useful for altering the properties of an existing messaging link. All attributes having a VARCHAR2 datatype will be assigned a value of DBMS_MGWADM.NO_CHANGE. Attributes of other datatypes will be assigned a value of NULL.

125.5 Summary of DBMS_MGWADM Subprograms

This table lists the DBMS_MGWADM subprograms and briefly describes them.

Table 125-19 DBMS_MGWADM Package Subprograms

Subprogram	Description
ADD_SUBSCRIBER Procedure	Adds a subscriber used to consume messages from a source queue for propagation to a destination
ALTER_AGENT Procedures	Alters Messaging Gateway agent parameters
ALTER_JOB Procedure	Alters the properties of a propagation job
ALTER_MSGSYSTEM_LINK Procedure for TIB/Rendezvous	Alters the properties of a TIB/Rendezvous messaging system link
ALTER_MSGSYSTEM_LINK Procedure for WebSphere MQ	Alters the properties of a WebSphere MQ messaging system link
ALTER_PROPAGATION_SCHEDULE Procedure	Alters a propagation schedule
ALTER_SUBSCRIBER Procedure	Alters the parameters of a subscriber used to consume messages from a source queue for propagation to a destination
CLEANUP_GATEWAY Procedures	Cleans up Messaging Gateway
CREATE_AGENT Procedure	Creates a Messaging Gateway agent that will be used to process propagation jobs

Table 125-19 (Cont.) DBMS_MGWADM Package Subprograms

Subprogram	Description
CREATE_JOB Procedure	Creates a job used to propagate message from a source to a destination
CREATE_MSGSYSTEM_LINK Procedures for TIB/Rendezvous	Creates a messaging system link to a TIB/Rendezvous messaging system
CREATE_MSGSYSTEM_LINK Procedures for WebSphere MQ	Creates a messaging system link to a WebSphere MQ messaging system
DB_CONNECT_INFO Procedure	Configures connection information used by the Messaging Gateway agent for connections to Oracle Database
DISABLE_JOB Procedure	Disables a propagation job
DISABLE_PROPAGATION_SCHEDULE Procedure	Disables a propagation schedule
ENABLE_JOB Procedure	Enables a propagation job
ENABLE_PROPAGATION_SCHEDULE Procedure	Enables a propagation schedule
REGISTER_FOREIGN_QUEUE Procedure	Registers a non-Oracle queue entity in Messaging Gateway
REMOVE_AGENT Procedure	Removes a Messaging Gateway agent
REMOVE_JOB Procedure	Removes a propagation job
REMOVE_MSGSYSTEM_LINK Procedure	Removes a messaging system link for a non-Oracle messaging system
REMOVE_OPTION Procedure	Removes a Messaging Gateway configuration option
REMOVE_SUBSCRIBER Procedure	Removes a subscriber used to consume messages from a source queue for propagation to a destination
RESET_JOB Procedure	Resets the propagation error state for a propagation job
RESET_SUBSCRIBER Procedure	Resets the propagation error state for a subscriber
SCHEDULE_PROPAGATION Procedure	Schedules message propagation from a source to a destination
SET_LOG_LEVEL Procedures	Dynamically alters the Messaging Gateway agent logging level
SET_OPTION Procedure	Sets a Messaging Gateway configuration option
SHUTDOWN Procedures	Shuts down the Messaging Gateway agent
STARTUP Procedures	Starts the Messaging Gateway agent
UNREGISTER_FOREIGN_QUEUE Procedure	Removes a non-Oracle queue entity in Messaging Gateway
UNSCHEDULE_PROPAGATION Procedure	Removes a propagation schedule

125.5.1 ADD_SUBSCRIBER Procedure

This procedure adds a subscriber used to consume messages from a source queue for propagation to a destination.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see CREATE_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.ADD_SUBSCRIBER(
   subscriber_id     IN VARCHAR2,
   propagation_type  IN BINARY_INTEGER,
   queue_name        IN VARCHAR2,
   destination       IN VARCHAR2,
   rule              IN VARCHAR2 DEFAULT NULL,
   transformation    IN VARCHAR2 DEFAULT NULL,
   exception_queue   IN VARCHAR2 DEFAULT NULL
   options           IN SYS.MGW_PROPERTIES DEFAULT NULL);
```

Parameters

Table 125-20 ADD_SUBSCRIBER Procedure Parameters

Parameter	Description
subscriber_id	Specifies a user-defined name that identifies this subscriber
propagation_type	Specifies the type of message propagation. DBMS_MGWADM.OUTBOUND_PROPAGATION is for Oracle Database Advanced Queuing to non-Oracle propagation. DBMS_MGWADM.INBOUND_PROPAGATION is for non-Oracle to Oracle Database Advanced Queuing propagation
queue_name	Specifies the source queue to which this subscriber is being added. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
destination	Specifies the destination queue to which messages consumed by this subscriber are propagated. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
rule	Specifies an optional subscription rule used by the subscriber to dequeue messages from the source queue. This is NULL if no rule is needed. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
transformation	Specifies the transformation needed to convert between the Oracle Database Advanced Queuing payload and an ADT defined by Messaging Gateway. The type of transformation needed depends on the value specified for propagation_type. If NULL, then the Oracle Database Advanced Queuing payload type must be supported by Messaging Gateway.

Table 125-20 (Cont.) ADD_SUBSCRIBER Procedure Parameters

Parameter	Description
exception_queue	Specifies a queue used for exception message logging purposes. This queue must be on the same messaging system as the propagation source. If NULL, then an exception queue is not used and propagation stops if a problem occurs. The syntax and interpretation of this parameter depend on the value specified for propagation_type. The source queue and exception queue cannot be the same queue.
options	Optional subscriber properties. NULL if there are none. Typically these are lesser used configuration properties supported by the messaging system.

Usage Notes

> **See Also:**
>
> "Handling Arbitrary Payload Types Using Message Transformations", in *Oracle Database Advanced Queuing User's Guide* for more information regarding message conversion and transformation

If the non-Oracle messaging link being accessed for the subscriber uses a JMS interface, then the Messaging Gateway agent will use the Oracle JMS interface to access the Oracle Database Advanced Queuing queues. Otherwise the native Oracle Database Advanced Queuing interface will be used. Parameters are interpreted differently when the Messaging Gateway agent uses Oracle JMS for JMS connections.

Transformations are not currently supported if the Oracle JMS interface is used for propagation. The transformation parameter must be NULL.

> **See Also:**
>
> For additional information regarding subscriber options
>
> - "WebSphere MQ System Properties" in *Oracle Database Advanced Queuing User's Guide*
> - "TIB/Rendezvous System Properties" in *Oracle Database Advanced Queuing User's Guide*

OUTBOUND_PROPAGATION Subscribers

The parameters for a subscriber used for outbound propagation are interpreted as follows:

- queue_name specifies the local Oracle Database Advanced Queuing queue that is the propagation source. This must have a syntax of schema.queue.
- destination specifies the foreign queue to which messages are propagated. This must have a syntax of registered_queue@message_link.

- `rule` specifies an optional Oracle Database Advanced Queuing subscriber rule if the native Oracle Database Advanced Queuing interface is used, or a JMS selector if the Oracle JMS interface is used. If `NULL`, then no rule or selector is used.
- `transformation` specifies the transformation used to convert the Oracle Database Advanced Queuing payload to an ADT defined by Messaging Gateway.

  Messaging Gateway propagation dequeues messages from the Oracle Database Advanced Queuing queue using the transformation to convert the Oracle Database Advanced Queuing payload to a known ADT defined by Messaging Gateway. The message is then enqueued in the foreign messaging system based on the Messaging Gateway ADT.
- `exception_queue` specifies the name of a local Oracle Database Advanced Queuing queue to which messages are moved if an exception occurs. This must have a syntax of `schema.queue`.

If the native Oracle Database Advanced Queuing interface is used, then a subscriber will be added to the Oracle Database Advanced Queuing queue when this procedure is called, whether or not Messaging Gateway is running. The local subscriber will be of the form `sys.aq$_agent('MGW_subscriber_id', NULL, NULL)`.

If the Oracle JMS interface is used, then the Messaging Gateway agent will create a JMS durable subscriber with the name of `MGW_subscriber_id`. If the agent is not running when this procedure is called, then the durable subscriber will be created the next time the agent starts.

The exception queue has the following caveats:

- The user is responsible for creating the Oracle Database Advanced Queuing queue to be used as the exception queue.
- The payload type of the source and exception queue must match.
- The exception queue must be created as a queue type of `DBMS_AQADM.NORMAL_QUEUE` rather than `DBMS_AQADM.EXCEPTION_QUEUE`. Enqueue restrictions prevent Messaging Gateway propagation from using an Oracle Database Advanced Queuing queue of type `EXCEPTION_QUEUE` as a Messaging Gateway exception queue.

INBOUND_PROPAGATION Subscribers

The parameters for a subscriber used for inbound propagation are interpreted as follows:

- `queue_name` specifies the foreign queue that is the propagation source. This must have a syntax of `registered_queue@message_link`.
- `destination` specifies the local Oracle Database Advanced Queuing queue to which messages are propagated. This must have a syntax of `schema.queue`.
- `rule` specifies an optional subscriber rule that is valid for the foreign messaging system. This is `NULL` if no rule is needed.
- `transformation` specifies the transformation used to convert an ADT defined by Messaging Gateway to the Oracle Database Advanced Queuing payload type.

  Messaging Gateway propagation dequeues messages from the foreign messaging system and converts the message body to a known ADT defined by Messaging Gateway. The transformation is used to convert the Messaging Gateway ADT to an Oracle Database Advanced Queuing payload type when the message is enqueued to the Oracle Database Advanced Queuing queue.
- `exception_queue` specifies the name of a foreign queue to which messages are moved if an exception occurs. This must have a syntax of `registered_queue@message_link`.

Whether or not a subscriber is needed depends on the requirements of the non-Oracle messaging system. If a durable subscriber is necessary, then it will be created by the Messaging Gateway agent. If the agent is not running at the time this procedure is called, then the creation of the subscriber on the non-Oracle messaging system will occur when the agent next starts.

The exception queue has the following caveats:

- The exception queue must be a registered non-Oracle queue.
- The source and exception queues must use the same messaging system link.

125.5.2 ALTER_AGENT Procedures

This procedure configures Messaging Gateway agent parameters.

Syntax

```
DBMS_MGWADM.ALTER_AGENT (
   max_memory       IN  BINARY_INTEGER DEFAULT NULL,
   max_threads      IN  BINARY_INTEGER DEFAULT NULL,
   service          IN  VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE );

DBMS_MGWADM.ALTER_AGENT (
   agent_name      IN   VARCHAR2,
   username        IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   password        IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   database        IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   conntype        IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   max_memory      IN   PLS_INTEGER DEFAULT NULL,
   max_threads     IN   PLS_INTEGER DEFAULT NULL,
   service         IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   initfile        IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   comment         IN   VARCHAR2  DEFAULT DBMS_MGWADM.NO_CHANGE );
```

Parameters

Table 125-21 ALTER_AGENT Procedure Parameters

Parameter	Description
max_memory	The maximum heap size, in MB, used by the Messaging Gateway agent. If it is NULL, then the current value is unchanged.
max_threads	The number of messaging threads that the Messaging Gateway agent creates. If it is NULL, then the current value is unchanged. The maximum value of max_threads is limited to 128.
service	Specifies the database service that the Oracle Scheduler job class used by this agent will have affinity to. In an Oracle RAC environment, this means that the Messaging Gateway agent will run on only those database instances that are assigned to the service. If NULL, the job class used by this agent will be altered to belong to the default service which is mapped to every instance. If DBMS_MGWADM.NO_CHANGE, the current value is unchanged.
agent_name	Identifies the Messaging Gateway agent. DBMS_MGWADM.DEFAULT_AGENT specifies the default agent.

Table 125-21 (Cont.) ALTER_AGENT Procedure Parameters

Parameter	Description
username	Specifies the username used for connections to the Oracle Database. NULL is not allowed. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged. If a username is specified then a password must also be specified.
password	Specifies the password used for connections to the Oracle Database. NULL is not allowed. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged. A password must be specified if a username is specified.
database	Specifies the database connect string used for connections to the Oracle Database. NULL indicates that a local connection should be used. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged.Oracle strongly recommends that a connect string, rather than NULL, be specified. Usually it will be a net service name from tnsnames.ora.
conntype	Specifies the type of connection to the Oracle Database, DBMS_MGWADM.JDBC_OCI or DBMS_MGWADM.JDBC_THIN. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged
initfile	Specifies a Messaging Gateway initialization file used by this agent. NULL indicates that the default initialization file is used. If a value is specified, it should be the full path name of the file. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged.
comment	Optional comments for this agent. NULL if a comment is not desired. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged.

Usage Notes

- Default values for these configuration parameters are set when the Messaging Gateway agent is installed.
- Changes to the max_memory and max_threads parameters take effect the next time the Messaging Gateway agent is active. If the Messaging Gateway agent is currently active, then it must be shut down and restarted for the changes to take effect.
- The service parameter is used to set an Oracle Scheduler job class attribute. The job class is used to create a Scheduler job that starts the Messaging Gateway agent. An Oracle administrator must create the database service. If the value is NULL, the job class will belong to an internal service that is mapped to all instances.
- The username, password, and database parameters specify connection information used by the Messaging Gateway agent for connections to the Oracle Database. An Oracle administrator should create the user and grant it the role MGW_AGENT_ROLE.

125.5.3 ALTER_JOB Procedure

This procedure alters the properties of a propagation job.

Syntax

```
DBMS_MGWADM.ALTER_JOB (
   job_name          IN   VARCHAR2,
   rule              IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   transformation    IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
```

```
  exception_queue    IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
  poll_interval      IN   PLS_INTEGER DEFAULT 0,
  options            IN   SYS.MGW_PROPERTIES DEFAULT NULL,
  comments           IN   VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE );
```

Parameters

Table 125-22 ALTER_JOB Procedure Parameters

Parameter	Description
job_name	Identifies the propagation job
rule	Specifies an optional subscription rule used to dequeue messages from the propagation source. The syntax and interpretation of this parameter depend on the propagation type. A NULL value indicates that no subscription rule is needed. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged.
transformation	Specifies the transformation needed to convert between the Oracle Streams AQ payload and an ADT defined by Messaging Gateway. The type of transformation needed depends on the value specified for propagation_type. A NULL value indicates that no transformation is needed. If DBMS_MGWADM.NO_CHANGE, the current value is unchanged.
exception_queue	Specifies a queue used for exception message logging purposes. This queue must be on the same messaging system as the propagation source. In cases in which no exception queue is associated with the job, propagation stops if a problem occurs. The syntax and interpretation of this parameter depend on the propagation type. A NULL value indicates that no exception queue is used. If DBMS_MGWADM.NO_CHANGE, the current value is unchanged.
poll_interval	Specifies the polling interval, in seconds, used by the Messaging Gateway agent when checking for messages in the source queue. If no messages are available the agent will not poll again until the polling interval has passed. Once the agent detects a message it will continue propagating messages as long as any are available. Values: NULL, 0, or value > 0: • If zero (default), the current value will not be changed. • If NULL, the current value will be reset and the Messaging Gateway default polling interval will be used. The default polling interval is 5 seconds and can be overridden by the Messaging Gateway initialization file.
options	Optional job properties. If NULL, no options will be changed. If not NULL, then the properties specified in this list are combined with the current optional properties to form a new set of job options.
comments	An optional comment for this agent, or NULL if one is not desired. If DBMS_MGWADM.NO_CHANGE, the current value will not be changed.

Usage Notes

- If the non-Oracle messaging link being accessed for the propagation job uses a JMS interface, then the Messaging Gateway agent will use the Oracle JMS interface to access the Oracle Streams AQ queues. Otherwise the native Oracle Streams AQ interface will be used. Parameters are interpreted differently when the Messaging Gateway agent uses Oracle JMS for JMS connections.

- The subscriber rule cannot be altered when propagating from a JMS source. Instead, the propagation job must be dropped and re-created with the new rule. For JMS, changing the message selector on a durable subscription is equivalent to deleting and re-creating the subscription.
- Transformations are not currently supported if the Oracle JMS interface is used for propagation. The transformation parameter must be `DBMS_MGWADM.NO_CHANGE` (the default value).
- The options parameter specifies a set of properties used to alter the current optional properties. Each property affects the current property list in a particular manner; add a new property, replace an existing property, remove an existing property or remove all properties.

> **Note:**
>
> - SYS.MGW_PROPERTY Object Type for more information about the options parameter
> - OUTBOUND_PROPAGATION Jobs for outbound propagation parameter interpretation
> - INBOUND_PROPAGATION Jobs for inbound propagation parameter interpretation

125.5.4 ALTER_MSGSYSTEM_LINK Procedure for TIB/Rendezvous

This procedure alters the properties of a TIB/Rendezvous messaging system link.

Syntax

```
DBMS_MGWADM.ALTER_MSGSYSTEM_LINK (
   linkname      IN  VARCHAR2,
   properties    IN  SYS.MGW_TIBRV_PROPERTIES,
   options       IN  SYS.MGW_PROPERTIES DEFAULT NULL,
   comment       IN  VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE );
```

Parameters

Table 125-23 ALTER_MSGSYSTEM_LINK Procedure Parameters for TIB/Rendezvous

Parameters	Description
`linkname`	The messaging system link name
`properties`	Basic properties for a TIB/Rendezvous messaging system link. If `NULL`, then no link properties will be changed.
`options`	Optional link properties. If `NULL`, then no options will be changed. If not `NULL`, then the properties specified in this list are combined with the current options properties to form a new set of link options.
`comment`	A user-specified description, or `NULL` if one is not desired. If `DBMS_MGWADM.NO_CHANGE`, then the current value will not be changed.

Usage Notes

To retain an existing value for a messaging link property with a `VARCHAR2` datatype, specify `DBMS_MGWADM.NO_CHANGE` for that particular property. To preserve an existing value for a property of another datatype, specify `NULL` for that property.

The `options` parameter specifies a set of properties used to alter the current optional properties. Each property affects the current property list in a particular manner: add a new property, replace an existing property, remove an existing property, or remove all properties.

> **See Also:**
>
> SYS.MGW_PROPERTIES Object Type

Some properties cannot be modified, and this procedure will fail if an attempt is made to alter such a property. For properties and options that can be changed, a few are dynamic, and Messaging Gateway uses the new values immediately. Others require the Messaging Gateway agent to be shut down and restarted before they take effect.

> **See Also:**
>
> "TIB/Rendezvous System Properties" in *Oracle Database Advanced Queuing User's Guide* for more information about the messaging system properties and options

125.5.5 ALTER_MSGSYSTEM_LINK Procedure for WebSphere MQ

This procedure alters the properties of a WebSphere MQ messaging system link.

Syntax

```
DBMS_MGWADM.ALTER_MSGSYSTEM_LINK (
   linkname   IN  VARCHAR2,
   properties IN  SYS.MGW_MQSERIES_PROPERTIES,
   options    IN  SYS.MGW_PROPERTIES DEFAULT NULL,
   comment    IN  VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE);
```

Parameters

Table 125-24 ALTER_MSGSYSTEM_LINK Procedure Parameters for WebSphere MQ

Parameters	Description
`linkname`	The messaging system link name
`properties`	Basic properties for a WebSphere MQ messaging system link. If it is `NULL`, then no link properties are changed.
`options`	Optional link properties. `NULL` if no options are changed. If not `NULL`, then the properties specified in this list are combined with the current options properties to form a new set of link options.

Table 125-24 (Cont.) ALTER_MSGSYSTEM_LINK Procedure Parameters for WebSphere MQ

Parameters	Description
comment	An optional description or NULL if not desired. If DBMS_MGWADM.NO_CHANGE is specified, then the current value is not changed.

Usage Notes

To retain an existing value for a messaging link property with a VARCHAR2 datatype, specify DBMS_MGWADM.NO_CHANGE for that particular property. To preserve an existing value for a property of another datatype, specify NULL for that property.

The options parameter specifies a set of properties used to alter the current optional properties. Each property affects the current property list in a particular manner: add a new property, replace an existing property, remove an existing property, or remove all properties.

> **See Also:**
>
> SYS.MGW_PROPERTIES Object Type

Some properties cannot be modified, and this procedure will fail if an attempt is made to alter such a property. For properties and options that can be changed, a few are dynamic, and Messaging Gateway uses the new values immediately. Others require the Messaging Gateway agent to be shut down and restarted before they take effect.

> **See Also:**
>
> "WebSphere MQ System Properties" in *Oracle Database Advanced Queuing User's Guide* for more information about the messaging system properties and options

125.5.6 ALTER_PROPAGATION_SCHEDULE Procedure

This procedure alters a propagation schedule.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see ALTER_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.ALTER_PROPAGATION_SCHEDULE (
   schedule_id  IN VARCHAR2,
   duration     IN NUMBER DEFAULT NULL,
```

```
   next_time     IN VARCHAR2 DEFAULT NULL,
   latency       IN NUMBER DEFAULT NULL);
```

Parameters

Table 125-25 ALTER_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
schedule_id	Identifies the propagation schedule to be altered
duration	Reserved for future use
next_time	Reserved for future use
latency	Specifies the polling interval, in seconds, used by the Messaging Gateway agent when checking for messages in the source queue. If no messages are available in the source queue, then the agent will not poll again until the polling interval has passed. Once the agent detects a message it will continue propagating messages as long as any are available. Values: NULL or value > 0. If latency is NULL, then the Messaging Gateway agent default polling interval will be used. The default polling interval is 5 seconds, but it can be overridden by the Messaging Gateway initialization file.

Usage Notes

This procedure always overwrites the existing value for each parameter. If a given parameter is not specified, then the existing values are overwritten with the default value.

125.5.7 ALTER_SUBSCRIBER Procedure

This procedure alters the parameters of a subscriber used to consume messages from a source queue for propagation to a destination.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see ALTER_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.ALTER_SUBSCRIBER (
   subscriber_id    IN VARCHAR2,
   rule             IN VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   transformation   IN VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   exception_queue  IN VARCHAR2 DEFAULT DBMS_MGWADM.NO_CHANGE,
   options          IN SYS.MGW_PROPERTIES DEFAULT NULL );
```

Parameters

Table 125-26 ALTER_SUBSCRIBER Procedure Parameters

Parameter	Description
subscriber_id	Identifies the subscriber to be altered

Table 125-26 (Cont.) ALTER_SUBSCRIBER Procedure Parameters

Parameter	Description
rule	Specifies an optional subscription rule used by the subscriber to dequeue messages from the source queue. The syntax and interpretation of this parameter depend on the subscriber propagation type. A NULL value indicates that no subscription rule is needed. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged.
transformation	Specifies the transformation needed to convert between the Oracle Database Advanced Queuing payload and an ADT defined by Messaging Gateway. The type of transformation needed depends on the subscriber propagation type. A NULL value indicates that no transformation is needed. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged.
exception_queue	Specifies a queue used for exception message logging. This queue must be on the same messaging system as the propagation source. If no exception queue is associated with the subscriber, then propagation stops if a problem occurs. The syntax and interpretation of this parameter depend on the subscriber propagation type. A NULL value indicates that no exception queue is used. If DBMS_MGWADM.NO_CHANGE, then the current value is unchanged. The source queue and exception queue cannot be the same queue.
options	Optional subscriber properties. If NULL, then no options will be changed. If not NULL, then the properties specified in this list are combined with the current optional properties to form a new set of subscriber options.

Usage Notes

If the non-Oracle messaging link being accessed for the subscriber uses a JMS interface, then the Messaging Gateway agent will use the Oracle JMS interface to access the Oracle Database Advanced Queuing queues. Otherwise the native Oracle Database Advanced Queuing interface will be used. Parameters are interpreted differently when the Messaging Gateway agent uses Oracle JMS for JMS connections.

When propagating from a JMS source, the subscriber rule cannot be altered. Instead, the subscriber must be removed and added with the new rule. For JMS, changing the message selector on a durable subscription is equivalent to deleting and re-creating the subscription.

Transformations are not currently supported if the Oracle JMS interface is used for propagation. The transformation parameter must be DBMS_MGWADM.NO_CHANGE (the default value).

The options parameter specifies a set of properties used to alter the current optional properties. Each property affects the current property list in a particular manner: add a new property, replace an existing property, remove an existing property, or remove all properties.

See Also:

- SYS.MGW_PROPERTIES Object Type for more information on the options parameter
- "WebSphere MQ System Properties" in *Oracle Database Advanced Queuing User's Guide* for more information about WebSphere MQ subscriber options
- "TIB/Rendezvous System Properties" in *Oracle Database Advanced Queuing User's Guide* for more information about TIB/Rendezvous subscriber options
- "OUTBOUND_PROPAGATION Subscribers for outbound propagation parameter interpretation
- "INBOUND_PROPAGATION Subscribers for inbound propagation parameter interpretation

125.5.8 CLEANUP_GATEWAY Procedures

This procedure cleans up Messaging Gateway. The procedure performs cleanup or recovery actions that may be needed when Messaging Gateway is left in some abnormal or unexpected condition. The `MGW_GATEWAY` view lists Messaging Gateway status and configuration information that pertains to the cleanup actions.

Syntax

```
DBMS_MGWADM.CLEANUP_GATEWAY(
   action       IN   BINARY_INTEGER,
   sarg         IN   VARCHAR2 DEFAULT NULL);

DBMS_MGWADM.CLEANUP_GATEWAY(
   agent_name   IN   VARCHAR2,   action       IN   BINARY_INTEGER,   sarg         IN
VARCHAR2 DEFAULT NULL );
```

Parameters

Table 125-27 CLEANUP_GATEWAY Procedure Parameters

Parameter	Description
`action`	The cleanup action to be performed. Values: • `DBMS_MGWADM.CLEAN_STARTUP_STATE` for Messaging Gateway start up state recovery • `DBMS_MGWADM.CLEAN_LOG_QUEUES` for log queue cleanup • `DBMS_MGWADM.RESET_SUB_MISSING_LOG_REC` for propagation job recovery due to missing log record • `DBMS_MGWADM.RESET_SUB_MISSING_MESSAGE` for propagation job recovery due to missing message
`sarg`	Optional argument whose meaning depends on the value specified for `action`. This should be `NULL` if it is not used for the specified action.
`agent_name`	Identifies the Messaging Gateway agent. `DBMS_MGWADM.DEFAULT_AGENT` specifies the default agent.

Usage Notes

CLEAN_STARTUP_STATE

`sarg` is not used and must be `NULL`.

The `CLEAN_STARTUP_STATE` action recovers Messaging Gateway to a known state when the Messaging Gateway agent has crashed or some other abnormal event occurs, and Messaging Gateway cannot be restarted. This should be done only when the Messaging Gateway agent has been started but appears to have crashed or has been nonresponsive for an extended period of time.

The `CLEAN_STARTUP_STATE` action may be needed when the `MGW_GATEWAY` view shows that the `AGENT_STATUS` value is something other than `NOT_STARTED` or `START_SCHEDULED`, and the `AGENT_PING` value is `UNREACHABLE` for an extended period of time.

If the `AGENT_STATUS` value is `BROKEN`, then the Messaging Gateway agent cannot be started until the problem has been resolved and the `CLEAN_STARTUP_STATE` action used to reset the agent status. A `BROKEN` status can indicate that the Messaging Gateway start job detected a Messaging Gateway agent already running. This condition that should never occur under normal use.

Cleanup tasks include:

- Removing the Scheduler job used to start the external Messaging Gateway agent process.
- Setting certain configuration information to a known state. For example, setting the agent status to `NOT_STARTED`.

Execution of this command fails if:

- The agent status is `NOT_STARTED` or `START_SCHEDULED`.
- No shutdown attempt has been made prior to calling this procedure, except if the agent status is `STARTING`.
- The Messaging Gateway agent is successfully contacted.

  The assumption is that the agent is active, and this procedure fails. If the agent does not respond after several attempts have been made, then the cleanup tasks are performed. This procedure takes at least several seconds and possibly up to one minute. This is expected behavior under conditions where this particular cleanup action is appropriate and necessary.

> **Note:**
>
> Terminate any Messaging Gateway agent process that may still be running after a `CLEAN_STARTUP_STATE` action has been successfully performed. This should be done before calling `DBMS_MGWADM.STARTUP` to start Messaging Gateway. The process is usually named `extprocmgwextproc`.

CLEAN_LOG_QUEUES

`sarg` is not used and must be `NULL`.

The Messaging Gateway agent will clean log queues for all configured messaging system links. The agent will temporarily stop all propagation activity and then remove all obsolete and

bad log records from the log queues for all links. The procedure will fail if the Messaging Gateway agent is not running.

This cleanup action is automatically performed each time the Messaging Gateway agent is started.

> **Note:**
>
> The `CLEAN_LOG_QUEUES` action is performed only on agent startup. If this procedure is called when the agent is running, then the Messaging Gateway agent ignores it.

RESET_SUB_MISSING_LOG_REC

`sarg` specifies a Messaging Gateway job name (or subscriber ID) to be reset. It must not be `NULL`.

The Messaging Gateway agent recovers a Messaging Gateway propagation job that has failed due to a missing log record. The agent will reset the source and destination log records. The procedure will fail if the Messaging Gateway agent is not running.

> **Note:**
>
> If the messages in the source queue had already been propagated to the destination queue, then this action may result in duplicate messages.

RESET_SUB_MISSING_MESSAGE

`sarg` specifies a Messaging Gateway job name (or subscriber ID) to be reset. It must not be `NULL`.

The Messaging Gateway agent recovers a Messaging Gateway propagation job that has failed due to a missing persistent source message. The agent will treat the message as a non-persistent message and continue processing that propagation job. The procedure will fail if the Messaging Gateway agent is not running.

125.5.9 CREATE_AGENT Procedure

This procedure creates a Messaging Gateway agent that will be used to process propagation jobs.

Syntax

```
DBMS_MGWADM.CREATE_AGENT (
   agent_name     IN   VARCHAR2,
   username       IN   VARCHAR2 DEFAULT NULL,
   password       IN   VARCHAR2 DEFAULT NULL,
   database       IN   VARCHAR2 DEFAULT NULL,
   conntype       IN   VARCHAR2 DEFAULT DBMS_MGWADM.JDBC_OCI,
   max_memory     IN   PLS_INTEGER DEFAULT 64,
   max_threads    IN   PLS_INTEGER DEFAULT 1,
   service        IN   VARCHAR2 DEFAULT NULL,
   initfile       IN   VARCHAR2 DEFAULT NULL,
   comment        IN   VARCHAR2  DEFAULT NULL );
```

Parameters

Table 125-28 CREATE_AGENT Procedure Parameters

Parameter	Description
agent_name	A name used to identify the agent
username	Specifies the username used for connections to the Oracle Database
password	Specifies the password used for connections to the Oracle Database. A password must be specified if a username is specified.
database	Specifies the database connect string used for connections to the Oracle Database. NULL indicates that a local connection should be used. A value can be specified only if username is specified. Oracle strong recommends that a connect string, rather than NULL be specified. Usually it will be a net service name from tnsnames.ora.
conntype	Specifies the type of connection to the Oracle Database.Values: DBMS_MGWADM.JDBC_OCI, DBMS_MGWADM.JDBC_THIN
max_memory	Specifies the maximum heap size, in MB, used by the Messaging Gateway agent
max_threads	Specifies the number of messaging threads that the Messaging Gateway agent creates. This determines the number of propagation jobs that the agent can concurrently process. The maximum value of max_threads is limited to 128.
service	Specifies the database service that the Oracle Scheduler job class used by this agent will have affinity to. In an Oracle RAC environment, this means that the Messaging Gateway agent will only run on those database instances that are assigned to the service. If NULL, then the job class will belong to the default service which is mapped to every instance.
initfile	Specifies a Messaging Gateway initialization file used by this agent. NULL indicates that the default initialization file is used. If a value is specified, it should be the full path name of the file.
comment	An optional comment for this agent. NULL if one is not desired.

Usage Notes

- The Messaging Gateway automatically configures a default agent when Messaging Gateway is installed. The name of the default agent is DEFAULT_AGENT. This procedure can be used to create additional agents.
- The username, password, and database parameters specify connection information used by the Messaging Gateway agent for connections to the Oracle Database. An Oracle administrator should create the database user and grant it the role MGW_AGENT_ROLE. It is not mandatory that the connection information be specified when this procedure is called but it must be set before the agent can be started.
- The service parameter is used to create an Oracle Scheduler job class. The job class is used to create a Scheduler job that starts the Messaging Gateway agent. An Oracle administrator must create the database service. If the value is NULL, the job class will belong to an internal service that is mapped to all instances.

125.5.10 CREATE_JOB Procedure

This procedure creates a job used to propagate message from a source to a destination.

Syntax

```
DBMS_MGWADM.CREATE_JOB (
   job_name           IN   VARCHAR2,
   propagation_type   IN   PLS_INTEGER,
   source             IN   VARCHAR2,
   destination        IN   VARCHAR2,
   rule               IN   VARCHAR2 DEFAULT NULL,
   transformation     IN   VARCHAR2 DEFAULT NULL,
   exception_queue    IN   VARCHAR2 DEFAULT NULL,
   poll_interval      IN   PLS_INTEGER DEFAULT NULL,
   options            IN   SYS.MGW_PROPERTIES DEFAULT NULL,
   enabled            IN   BOOLEAN DEFAULT TRUE,
   comments           IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 125-29 CREATE_JOB Procedure Parameters

Parameter	Description
job_name	A user defined name to identify the propagation job
propagation_type	Specifies the type of message propagation. • DBMS_MGWADM.OUTBOUND_PROPAGATION for Oracle Streams AQ to non-Oracle propagation. • DBMS_MGWADM.INBOUND_PROPAGATION for non-Oracle to Oracle Streams AQ propagation.
source	Specifies the source queue whose messages are to be propagated. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
destination	Specifies the destination queue to which messages are propagated. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
rule	Specifies an optional subscription rule used to dequeue messages from the source queue. This should be NULL if no rule is needed. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
transformation	Specifies the transformation needed to convert between the Oracle Streams AQ payload and an ADT defined by Messaging Gateway. The type of transformation needed depends on the value specified for propagation_type. If no transformation is specified the Oracle Streams AQ payload type must be supported by Messaging Gateway.
exception_queue	Specifies a queue used for exception message logging purposes. This queue must be on the same messaging system as the propagation source. If NULL, an exception queue will not be used and propagation will stop if a problem occurs. The syntax and interpretation of this parameter depend on the value specified for propagation_type. The source queue and exception queue cannot be the same queue.

Table 125-29 (Cont.) CREATE_JOB Procedure Parameters

Parameter	Description
poll_interval	Specifies the polling interval, in seconds, used by the Messaging Gateway agent when checking for messages in the source queue. If no messages are available the agent will not poll again until the polling interval has passed. Once the agent detects a message it will continue propagating messages as long as any are available. Values: NULL or value > 0. If NULL, then the Messaging Gateway default polling interval will be used. The default polling interval is 5 seconds and can be overridden by the Messaging Gateway initialization file.
options	Optional job properties, NULL if there are none. Typically these are lesser used configuration properties supported by the messaging system.
enabled	Specifies whether this propagation job is enabled after creation. Values: TRUE, FALSE. • If TRUE (default), the job will be enabled after it is created. • If FALSE, the job will be disabled after it is created. A propagation job must be enabled and the Messaging Gateway agent running before messages can be propagated.
comments	An optional comment for this job. NULL if one is not desired.

Usage Notes

- The job must be enabled and Messaging Gateway agent started in order for messages to be propagated.
- If the non-Oracle messaging link being accessed for the propagation job uses a JMS interface, then the Messaging Gateway agent will use the Oracle JMS interface to access the Oracle Streams AQ queues. Otherwise the native Oracle Streams AQ interface will be used. Parameters are interpreted differently when the Messaging Gateway agent uses Oracle JMS for JMS connections.
- Transformations are not currently supported if the Oracle JMS interface is used for propagation. The transformation parameter must be NULL.

OUTBOUND_PROPAGATION Jobs

The parameters for an outbound propagation job are interpreted as follows:

- source specifies the local Oracle Streams AQ queue that is the propagation source. This must have syntax of schema.queue. This can be either a multiple consumer queue or a single consumer queue.
- destination specifies the non-Oracle queue to which messages are propagated. This must have syntax of registered_queue@message_link.
- rule specifies an optional Oracle Streams AQ subscriber rule if the native Oracle Stream AQ interface is used, or a JMS selector if the Oracle JMS interface is used. If NULL, then no rule or selector is used. This parameter must be NULL if the native Oracle Stream AQ interface is used and the propagation source is a single consumer queue.
- transformation specifies the transformation used to convert the Oracle Streams AQ payload to an ADT defined by Messaging Gateway. The full transformation name (schema.name) should be used if one is specified.

Messaging Gateway propagation dequeues messages from the Oracle Streams AQ queue using the transformation to convert the Oracle Streams AQ payload to a known ADT defined by Messaging Gateway. The message is then enqueued in the non-Oracle messaging system based on the Messaging Gateway ADT.

- `exception_queue` specifies the name of a local Oracle Streams AQ queue to which messages are moved if an exception occurs. The syntax must be `schema.queue`.

  If the native Oracle Streams AQ interface is used and the source is a multiple consumer queue, then a subscriber will be added to the Oracle Streams AQ queue when this procedure is called, whether or not the Messaging Gateway agent is running. The local subscriber will be of the form `sys.aq$_agent('MGW_job_name', NULL, NULL)`.

  If the Oracle JMS interface is used, then the Messaging Gateway agent will create a JMS durable subscriber with the name of `MGW_job_name`. If the agent is not running when this procedure is called, then the durable subscriber will be created the next time the agent starts.

  The exception queue has the following conditions:

  – The user is responsible for creating the Oracle Streams AQ queue to be used as the exception queue.

  – The payload type of the source queue and exception queue must match.

  – The exception queue must be created as a queue type of `DBMS_AQADM.NORMAL_QUEUE`. Enqueue restrictions prevent Messaging Gateway from using an Oracle Streams AQ queue of type `DBMS_AQADM.EXCEPTION_QUEUE` as a Messaging Gateway exception queue.

INBOUND_PROPAGATION Jobs

The parameters for an inbound propagation job are interpreted as follows:

- `source` specifies the non-Oracle queue that is the propagation source. The syntax must be `registered_queue@message_link`.
- `destination` specifies the local Oracle Streams AQ queue to which messages are propagated. The syntax must be `schema.queue`.
- `rule` specifies an optional subscriber rule that is valid for the non-Oracle messaging system. This should be `NULL` if no rule is needed.
- `transformation` specifies the transformation used to convert an ADT defined by Messaging Gateway to the Oracle Streams AQ payload type. The full transformation name (`schema.name`) should be used if one is specified

  Messaging Gateway propagation dequeues messages from the non-Oracle messaging system and converts the message body to a known ADT defined by Messaging Gateway. The transformation is used to convert the Messaging Gateway ADT to an Oracle Streams AQ payload type when the message is enqueued to the Oracle Streams AQ queue.

- `exception_queue` specifies the name of a registered non-Oracle queue to which messages are moved if an exception occurs. The syntax must be `registered_queue@message_link`.

  Whether or not a subscriber is needed for the source queue depends on the requirements of the non-Oracle messaging system. If a durable subscriber is necessary, then the Messaging Gateway agent will create it. If the agent is not running when this procedure is called, then the subscriber will be created on the non-Oracle messaging system the next time the agent starts.

  The exception queue has the following conditions:

  – The exception queue must be a registered non-Oracle queue.

– The source queue and exception queue must use the same messaging system link.

125.5.11 CREATE_MSGSYSTEM_LINK Procedures for TIB/Rendezvous

This procedure creates a link to a TIB/Rendezvous messaging system.

Syntax

```
DBMS_MGWADM.CREATE_MSGSYSTEM_LINK (
   linkname     IN  VARCHAR2,
   properties   IN  SYS.MGW_TIBRV_PROPERTIES,
   options      IN  SYS.MGW_PROPERTIES DEFAULT NULL,
   comment      IN  VARCHAR2 DEFAULT NULL );

DBMS_MGWADM.CREATE_MSGSYSTEM_LINK (
   linkname     IN  VARCHAR2,
   agent_name   IN  VARCHAR2,
   properties   IN  SYS.MGW_TIBRV_PROPERTIES,
   options      IN  SYS.MGW_PROPERTIES DEFAULT NULL,
   comment      IN  VARCHAR2 DEFAULT NULL );
```

Parameters

Table 125-30 CREATE_MSGSYSTEM_LINK Procedure Parameters for TIB/Rendezvous

Parameter	Description
`linkname`	A user-defined name to identify this messaging system link
`properties`	Basic properties of a TIB/Rendezvous messaging system link.
`options`	Optional link properties. `NULL` if there are none. These are less frequently used configuration properties supported by the messaging system
`comment`	A user-specified description. `NULL` if one is not desired.
`agent_name`	Specifies the Messaging Gateway agent that will be used to process all propagation jobs associated with this link. `DBMS_MGWADM.DEFAULT_AGENT` specifies the default agent.

Usage Notes

The Messaging Gateway default agent will process the propagation jobs associated with this link if an agent name is not specified.

> **See Also:**
>
> "TIB/Rendezvous System Properties" in *Oracle Database Advanced Queuing User's Guide* for more information about the messaging system properties and options

125.5.12 CREATE_MSGSYSTEM_LINK Procedures for WebSphere MQ

This procedure creates a messaging system link to a WebSphere MQ messaging system.

Syntax

```
DBMS_MGWADM.CREATE_MSGSYSTEM_LINK(
   linkname    IN VARCHAR2,
   properties  IN SYS.MGW_MQSERIES_PROPERTIES,
   options     IN SYS.MGW_PROPERTIES DEFAULT NULL,
   comment     IN VARCHAR2 DEFAULT NULL);

DBMS_MGWADM.CREATE_MSGSYSTEM_LINK(
   linkname    IN VARCHAR2,
   agent_name  IN  VARCHAR2,
   properties  IN SYS.MGW_MQSERIES_PROPERTIES,
   options     IN SYS.MGW_PROPERTIES DEFAULT NULL,
   comment     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 125-31 CREATE_MSGSYSTEM_LINK Procedure Parameters for WebSphere MQ

Parameter	Description
`linkname`	A user-defined name to identify the messaging system link
`properties`	Basic properties of a WebSphere MQ messaging system link
`options`	Optional link properties. `NULL` if there are none. These are less frequently used configuration properties supported by the messaging system.
`comment`	A user-specified description. `NULL` if one is not desired
`agent_name`	Specifies the Messaging Gateway agent that will be used to process all propagation jobs associated with this link. `DBMS_MGWADM.DEFAULT_AGENT` specifies the default agent.

Usage Notes

The Messaging Gateway default agent will process the propagation jobs associated with this link if an agent name is not specified.

> **See Also:**
>
> "WebSphere MQ System Properties" in *Oracle Database Advanced Queuing User's Guide* for more information about the messaging system properties and options

125.5.13 DB_CONNECT_INFO Procedure

This deprecated procedure configures connection information used by the Messaging Gateway default agent for connections to Oracle Database.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see ALTER_AGENT Procedures), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.DB_CONNECT_INFO (
   username      IN VARCHAR2,
   password      IN VARCHAR2,
   database      IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 125-32 DB_CONNECT_INFO Procedure Parameters

Parameter	Description
`username`	The username used for connections to Oracle Database. `NULL` is not allowed
`password`	The password used for connections to Oracle Database. `NULL` is not allowed
`database`	The database connect string used by the Messaging Gateway agent. `NULL` indicates that a local connection should be used. Oracle strongly recommends that a not `NULL` value be specified. Usually it will be a net service name from `tnsnames.ora`.

Usage Notes

The Messaging Gateway agent connects to Oracle Database as the user configured by this procedure. An Oracle administrator should create the user, grant it the role `MGW_AGENT_ROLE`, and then call this procedure to configure Messaging Gateway. Role `MGW_AGENT_ROLE` is used to grant this user special privileges needed to access Messaging Gateway configuration information stored in the database, enqueue or dequeue messages to and from Oracle Database Advanced Queuing queues, and perform certain Oracle Database Advanced Queuing administration tasks.

125.5.14 DISABLE_JOB Procedure

This procedure disables a propagation job.

Syntax

```
DBMS_MGWADM.DISABLE_JOB (
   job_name  IN VARCHAR2);
```

Parameters

Table 125-33 DISABLE_JOB Procedure Parameters

Parameter	Description
job_name	Identifies the propagation job

125.5.15 DISABLE_PROPAGATION_SCHEDULE Procedure

This **deprecated procedure** disables a propagation schedule.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see DISABLE_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.DISABLE_PROPAGATION_SCHEDULE (
   schedule_id  IN VARCHAR2);
```

Parameters

Table 125-34 DISABLE_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
schedule_id	Identifies the propagation schedule to be disabled

125.5.16 ENABLE_JOB Procedure

This procedure enables a propagation job.

Syntax

```
DBMS_MGWADM.ENABLE_JOB (
   job_name  IN VARCHAR2 );
```

Parameters

Table 125-35 ENABLE_JOB Procedure Parameters

Parameter	Description
job_name	Identifies the propagation job

125.5.17 ENABLE_PROPAGATION_SCHEDULE Procedure

This deprecated procedure enables a propagation schedule.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see ENABLE_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.ENABLE_PROPAGATION_SCHEDULE (
   schedule_id  IN VARCHAR2 );
```

Parameters

Table 125-36 ENABLE_PROPAGATION_SCHEDULE Procedure Parameters

Parameter	Description
schedule_id	Identifies the propagation schedule to be enabled

125.5.18 REGISTER_FOREIGN_QUEUE Procedure

This procedure registers a non-Oracle queue entity in Messaging Gateway.

Syntax

```
DBMS_MGWADM.REGISTER_FOREIGN_QUEUE(
   name            IN VARCHAR2,
   linkname        IN VARCHAR2,
   provider_queue  IN VARCHAR2 DEFAULT NULL,
   domain          IN INTEGER DEFAULT NULL,
   options         IN SYS.MGW_PROPERTIES DEFAULT NULL,
   comment         IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 125-37 REGISTER_FOREIGN_QUEUE Procedure Parameters

Parameters	Description
name	The registered queue name. This name identifies the foreign queue within Messaging Gateway and need not match the name of the queue in the foreign messaging system.
linkname	The link name for the messaging system on which this queue exists
provider_queue	The message provider (native) queue name. If NULL, then the value provided for the name parameter is used as the provider queue name.

Table 125-37 (Cont.) REGISTER_FOREIGN_QUEUE Procedure Parameters

Parameters	Description
`domain`	The domain type of the queue. `NULL` means the domain type is automatically determined based on the messaging system of the queue. `DBMS_MGWADM.DOMAIN_QUEUE` is for a queue (point-to-point model). `DBMS_MGWADM.DOMAIN_TOPIC` is for a topic (publish-subscribe model).
`options`	Optional queue properties
`comment`	A user-specified description. Can be `NULL`.

Usage Notes

This procedure does not create the physical queue in the non-Oracle messaging system. The non-Oracle queue must be created using the administration tools for that messaging system.

> **See Also:**
>
> For more information when registering queues for the WebSphere MQ messaging system or the TIB/Rendezvous messaging system, specifically "Optional Foreign Queue Configuration Properties" in *Oracle Database Advanced Queuing User's Guide*.

125.5.19 REMOVE_AGENT Procedure

This procedure removes a Messaging Gateway agent.

Syntax

```
DBMS_MGWADM.REMOVE_AGENT(
   agent_name    IN    VARCHAR2 );
```

Parameters

Table 125-38 REMOVE_AGENT Procedure Parameters

Parameters	Description
`agent_name`	Identifies the Messaging Gateway agent

Usage Notes

All messaging system links associated with this Messaging Gateway agent must be removed and the agent must be stopped before it can be removed. The Messaging Gateway default agent cannot be removed.

125.5.20 REMOVE_JOB Procedure

This procedure removes a propagation job.

Syntax

```
DBMS_MGWADM.REMOVE_JOB(
   job_name   IN   VARCHAR2,   force      IN   PLS_INTEGER DEFAULT DBMS_MGWADM.NO_FORCE);
```

Parameters

Table 125-39 REMOVE_JOB Procedure Parameters

Parameters	Description
job_name	Identifies the propagation job
force	Specifies whether the procedure should succeed even if Messaging Gateway is not able to perform all cleanup actions pertaining to this propagation job. Values: DBMS_MGWADM.NO_FORCE, DBMS_MGWADM.FORCE • NO_FORCE (default) means the job is not removed if Messaging Gateway is unable to clean up successfully • FORCE means the job is removed even though all cleanup actions may not be done

Usage Notes

- The Messaging Gateway agent uses various resources of the Oracle Database and the non-Oracle messaging system for its propagation work. These resources need to be released when the job is removed. For example, Messaging Gateway may create a durable subscriber on the source queue that should be removed when the job is removed. Therefore, this procedure should normally be called when the Messaging Gateway agent is running and able to access the non-Oracle messaging system associated with this job.
- For outbound propagation, a local subscriber is removed from the Oracle Streams AQ queue when the propagation source is a multiple consumer queue.

125.5.21 REMOVE_MSGSYSTEM_LINK Procedure

This procedure removes a messaging system link for a non-Oracle messaging system.

Syntax

```
DBMS_MGWADM.REMOVE_MSGSYSTEM_LINK(
   linkname  IN VARCHAR2);
```

Parameters

Table 125-40 REMOVE_MSGSYSTEM_LINK Procedure Parameters

Parameters	Description
linkname	The messaging system link name

Usage Notes

All registered queues associated with this link must be removed before the messaging system link can be removed. This procedure fails if there is a registered foreign (non-Oracle) queue that references this link.

125.5.22 REMOVE_OPTION Procedure

This procedure removes a Messaging Gateway configuration option. It can be used to remove an agent option, a messaging link option, or a propagation job option.

Syntax

```
DBMS_MGWADM.REMOVE_OPTION (
   target_type    IN   PLS_INTEGER,
   target_name    IN   VARCHAR2,
   option_name    IN   VARCHAR2);
```

Parameters

Table 125-41 REMOVE_OPTION Procedure Parameters

Parameter	Description
target_type	Specifies the target type of the Messaging Gateway entity: • DBMS_MGWADM.AGENT_JAVA_PROP to remove a Java System property for a Messaging Gateway agent • DBMS_MGWADM.MSGLINK_OPTION to remove a messaging link option • DBMS_MGWADM.JOB_OPTION to remove a propagation job option
target_name	Name or identifier of the target. The value for this parameter depends on the value specified for target_type parameter. This must not be NULL.
option_name	Option name. This must not be NULL.

> **See Also:**
>
> Table 125-10 regarding options for the option_type parameter

Usage Notes

DBMS_MGWADM.AGENT_JAVA_PROP Target

The procedure removes an agent option used to set a Java System property when the Messaging Gateway agent is started. The agent must be restarted for the change to take effect.

The parameters are interpreted as follows:

- target_name specifies the name of the Messaging Gateway agent. DBMS_MGWADM.DEFAULT_AGENT can be used for the default agent.
- option_name specifies the Java System property

- `encrypted` can be either `TRUE` or `FALSE`

DBMS_MGWADM.MSGLINK_OPTION Target

The procedure removes a single option for a Messaging Gateway messaging system link. This is equivalent to calling `DBMS_MGWADM.ALTER_MSGSYSTEM_LINK` and using the options parameter to remove an option.

The parameters are interpreted as follows:

- `target_name` specifies the name of the message system link
- `option_name` specifies the option to set
- `encrypted` must be `FALSE`

DBMS_MGWADM.JOB_OPTION Target

The procedure removes a single option for a Messaging Gateway propagation job. This is equivalent to calling `DBMS_MGWADM.ALTER_JOB` and using the options parameter to remove an option.

The parameters are interpreted as follows:

- `target_name` specifies the name of the propagation job
- `option_name` specifies the option to set
- `encrypted` must be `FALSE`

125.5.23 REMOVE_SUBSCRIBER Procedure

This procedure removes a subscriber used to consume messages from a source queue for propagation to a destination.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see REMOVE_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.REMOVE_SUBSCRIBER (
   subscriber_id  IN VARCHAR2,
   force          IN BINARY_INTEGER DEFAULT DBMS_MGWADM.NO_FORCE );
```

Parameters

Table 125-42 REMOVE_SUBSCRIBER Procedure Parameters

Parameter	Description
`subscriber_id`	Identifies the subscriber to be removed

Table 125-42 (Cont.) REMOVE_SUBSCRIBER Procedure Parameters

Parameter	Description
force	Specifies whether this procedure should succeed even if Messaging Gateway is not able to perform all cleanup actions pertaining to this subscriber. Values: DBMS_MGWADM.NO_FORCE, DBMS_MGWADM.FORCE • NO_FORCE means the subscriber is not removed if Messaging Gateway is unable to clean up successfully (default) • FORCE means the subscriber is removed even though all cleanup actions may not be done

Usage Notes

- The Messaging Gateway agent uses various resources of Oracle Database and the non-Oracle messaging system for its propagation work. These resources are typically associated with each subscriber and need to be released when the subscriber is no longer needed. Therefore, this procedure should only be called when the Messaging Gateway agent is running and able to access the non-Oracle messaging system associated with this subscriber.
- For outbound propagation, a local subscriber is removed from the Oracle Database Advanced Queuing queue.

125.5.24 RESET_JOB Procedure

This procedure resets the propagation error state for a propagation job.

Syntax

```
DBMS_MGWADM.RESET_JOB (
   job_name    IN    VARCHAR2);
```

Parameters

Table 125-43 RESET_JOB Procedure Parameters

Parameter	Description
job_name	Identifies the propagation job

Usage Notes

This procedure can be used to reset a propagation job that has been set to a failed state and propagation activities have been stopped. The administrator should correct the problem and then call this procedure to allow the agent to retry the propagation job. The STATUS field of the MGW_JOBS view indicates the job status.

125.5.25 RESET_SUBSCRIBER Procedure

This procedure resets the propagation error state for a subscriber.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see RESET_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.RESET_SUBSCRIBER (
   subscriber_id  IN VARCHAR2 );
```

Parameters

Table 125-44 RESET_SUBSCRIBER Procedure Parameters

Parameter	Description
subscriber_id	Identifies the subscriber

125.5.26 SCHEDULE_PROPAGATION Procedure

This procedure schedules message propagation from a source to a destination.

The schedule must be enabled and Messaging Gateway started in order for messages to be propagated.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see CREATE_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.SCHEDULE_PROPAGATION (
   schedule_id       IN VARCHAR2,
   propagation_type  IN BINARY_INTEGER,
   source            IN VARCHAR2,
   destination       IN VARCHAR2,
   start_time        IN DATE DEFAULT SYSDATE,
   duration          IN NUMBER DEFAULT NULL,
   next_time         IN VARCHAR2 DEFAULT NULL,
   latency           IN NUMBER DEFAULT NULL);
```

Parameters

Table 125-45 SCHEDULE_PROPAGATION Procedure Parameters

Parameter	Description
schedule_id	Specifies a user-defined name that identifies the schedule
propagation_type	Specifies the type of message propagation. DBMS_MGWADM.OUTBOUND_PROPAGATION is for Oracle Database Advanced Queuing to non-Oracle propagation. DBMS_MGWADM.INBOUND_PROPAGATION is for non-Oracle to Oracle Database Advanced Queuing propagation.
source	Specifies the source queue whose messages are to be propagated. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
destination	Specifies the destination queue to which messages are propagated. The syntax and interpretation of this parameter depend on the value specified for propagation_type.
start_time	Reserved for future use
duration	Reserved for future use
next_time	Reserved for future use
latency	Specifies the polling interval, in seconds, used by the Messaging Gateway agent when checking for messages in the source queue. If no messages are available in the source queue, then the agent will not poll again until the polling interval has passed. Once the agent detects a message it will continue propagating messages as long as any are available. Values: NULL or value > 0. If latency is NULL, then the Messaging Gateway agent default polling interval will be used. The default polling interval is 5 seconds but it can be overridden by the Messaging Gateway initialization file.

Usage Notes

For outbound propagation, parameters are interpreted as follows:

- source specifies the local Oracle Database Advanced Queuing queue from which messages are propagated. This must have a syntax of schema.queue.
- destination specifies the foreign queue to which messages are propagated. This must have a syntax of registered_queue@message_link.

For inbound propagation, parameters are interpreted as follows:

- source specifies the foreign queue from which messages are propagated. This must have a syntax of registered_queue@message_link.
- destination specifies the local Oracle Database Advanced Queuing queue to which messages are propagated. This must have a syntax of schema.queue.

The schedule is set to an enabled state when it is created.

125.5.27 SET_LOG_LEVEL Procedures

This procedure dynamically alters the Messaging Gateway agent logging level. The Messaging Gateway agent must be running.

Syntax

```
DBMS_MGWADM.SET_LOG_LEVEL (
   log_level    IN   BINARY_INTEGER);

DBMS_MGWADM.SET_LOG_LEVEL (
   agent_name   IN   VARCHAR2,
   log_level    IN   BINARY_INTEGER);
```

Parameters

Table 125-46 SET_LOG_LEVEL Procedure Parameters

Parameter	Description
`log_level`	Level at which the Messaging Gateway agent logs information. `DBMS_MGWADM.BASIC_LOGGING` generates the least information while `DBMS_MGWADM.TRACE_DEBUG_LOGGING` generates the most information.
`agent_name`	Identifies the Messaging Gateway agent. `DBMS_MGWADM.DEFAULT_AGENT` specifies the default agent.

> **See Also:**
>
> Table 125-3 for details on the `log_level` parameter

125.5.28 SET_OPTION Procedure

This procedure sets a Messaging Gateway configuration option. It can be used to set an agent option, a messaging link option, or a propagation job option.

Syntax

```
DBMS_MGWADM.SET_OPTION (
   target_type   IN   PLS_INTEGER,
   target_name   IN   VARCHAR2,
   option_name   IN   VARCHAR2,
   option_value  IN   VARCHAR2,
   encrypted     IN   BOOLEAN DEFAULT FALSE );
```

Parameters

Table 125-47 SET_OPTION Procedure Parameters

Parameter	Description
target_type	Specifies the target type of the Messaging Gateway entity: • DBMS_MGWADM.AGENT_JAVA_PROP to set a Java System property for a Messaging Gateway agent • DBMS_MGWADM.MSGLINK_OPTION to set a messaging link option • DBMS_MGWADM.JOB_OPTION to set a propagation job option
target_name	Name or identifier of the target. The value for this parameter depends on the value specified for target_type parameter. This must not be NULL.
option_name	Option name. This must not be NULL.
option_value	Option value
encrypted	Indicates whether the value should be stored as encrypted: • TRUE if the value should be stored in an encrypted form • FALSE if the value should be stored in a cleartext form

> **See Also:**
>
> Table 125-10 regarding options for the option_type parameter

Usage Notes

DBMS_MGWADM.AGENT_JAVA_PROP Target

The procedure will store an agent option used to set a Java System property when the Messaging Gateway agent is started. The agent must be restarted for the change to take effect.

The parameters are interpreted as follows:

- target_name specifies the name of the Messaging Gateway agent. DBMS_MGWADM.DEFAULT_AGENT can be used for the default agent.
- option_name specifies the Java System property
- encrypted can be either TRUE or FALSE

DBMS_MGWADM.MSGLINK_OPTION Target

The procedure will set or alter a single option for a Messaging Gateway messaging system link. This is equivalent to calling DBMS_MGWADM.ALTER_MSGSYSTEM_LINK and using the options parameter to set an option.

The parameters are interpreted as follows:

- target_name specifies the name of the message system link
- option_name specifies the option to set
- encrypted must be FALSE

DBMS_MGWADM.JOB_OPTION Target

The procedure will set or alter a single option for a Messaging Gateway propagation job. This is equivalent to calling `DBMS_MGWADM.ALTER_JOB` and using the options parameter to set an option.

The parameters are interpreted as follows:

- `target_name` specifies the name of the propagation job
- `option_name` specifies the option to set
- `encrypted` must be `FALSE`

125.5.29 SHUTDOWN Procedures

This procedure shuts down the Messaging Gateway agent. No propagation activity occurs until Messaging Gateway is restarted.

Syntax

```
DBMS_MGWADM.SHUTDOWN (
   sdmode      IN BINARY_INTEGER DEFAULT DBMS_MGWADM.SHUTDOWN_NORMAL);

DBMS_MGWADM.SHUTDOWN (
   agent_name  IN VARCHAR2);
```

Parameters

Table 125-48 SHUTDOWN Procedure Parameters

Parameter	Description
`sdmode`	The shutdown mode. The only value currently supported is `DBMS_MGWADM.SHUTDOWN_NORMAL` for normal shutdown. The Messaging Gateway agent may attempt to complete any propagation work currently in progress.
`agent_name`	Identifies the Messaging Gateway agent. `DBMS_MGWADM.DEFAULT_AGENT` specifies the default agent.

Usage Notes

The Messaging Gateway default agent is shut down if no agent name is specified.

125.5.30 STARTUP Procedures

This procedure starts the Messaging Gateway agent. It must be called before any propagation activity can take place.

Syntax

```
DBMS_MGWADM.STARTUP(
   instance       IN  BINARY_INTEGER DEFAULT 0,
   force          IN  BINARY_INTEGER DEFAULT DBMS_MGWADM.NO_FORCE);

DBMS_MGWADM.STARTUP(
   agent_name     IN  VARCHAR2);
```

Parameters

Table 125-49 STARTUP Procedure Parameters

Parameter	Description
instance	Specifies which instance can run the job queue job used to start the Messaging Gateway agent. If this is zero, then the job can be run by any instance. **Caution: This parameter has been deprecated**.
force	If this is DBMS_MGWADM.FORCE, then any positive integer is acceptable as the job instance. If this is DBMS_MGWADM.NO_FORCE (the default), then the specified instance must be running; otherwise the routine raises an exception. **Caution: This parameter has been deprecated**.
agent_name	Identifies the Messaging Gateway agent. DBMS_MGWADM.DEFAULT_AGENT specifies the default agent.

Usage Notes

- The Messaging Gateway default agent will be started if an agent name is not specified.
- The force and instance parameters are no longer used and will be ignored. If the instance affinity parameters were being used to start the default agent on a specific instance, the administrator will need to create a database service and then assign that service to the default agent using the DBMS_MGWADM.ALTER_AGENT procedure.
- The Messaging Gateway agent cannot be started until an agent user has been configured by the DBMS_MGWADM.CREATE_AGENT or DBMS_MGWADM.ALTER_AGENT subprograms.

125.5.31 UNREGISTER_FOREIGN_QUEUE Procedure

This procedure removes a non-Oracle queue entity in Messaging Gateway.

Syntax

```
DBMS_MGWADM.UNREGISTER_FOREIGN_QUEUE(
   name          IN VARCHAR2,
   linkname      IN VARCHAR2);
```

Parameters

Table 125-50 UNREGISTER_FOREIGN_QUEUE Procedure Parameters

Parameter	Description
name	The queue name
linkname	The link name for the messaging system on which the queue exists

Usage Notes

- This procedure does not remove the physical queue in the non-Oracle messaging system.
- All propagation jobs, subscribers and schedules referencing this queue must be removed before it can be unregistered. This procedure fails if a propagation job, subscriber, or propagation schedule references the non-Oracle queue.

125.5.32 UNSCHEDULE_PROPAGATION Procedure

This deprecated procedure removes a propagation schedule.

> **Note:**
>
> This subprogram has been deprecated as a result of improved technology (see REMOVE_JOB Procedure), and is retained only for reasons of backward compatibility.

Syntax

```
DBMS_MGWADM.UNSCHEDULE_PROPAGATION (
   schedule_id   IN VARCHAR2 );
```

Parameters

Table 125-51 UNSCHEDULE_PROPAGATION Procedure Parameters

Parameter	Description
schedule_id	Identifies the propagation schedule to be removed

126
DBMS_MGWMSG

DBMS_MGWMSG provides object types used by the canonical message types to convert message bodies, and methods, constants, and subprograms for working with Messaging Gateway message types.

This chapter contains the following topics:

- Security Model
- Constants
- Types
- Summary of DBMS_MGWMSG Subprograms

> **See Also:**
>
> DBMS_MGWADM which describes the Messaging Gateway administrative interface, DBMS_MGWADM

126.1 DBMS_MGWMSG Security Model

The EXECUTE privilege is granted to PUBLIC on all types defined in the DBMS_MGWMSG package as well as the canonical types. The DBMS_MGWMSG packages and object types are owned by SYS.

> **Note:**
>
> You must run the catmgw.sql script to load the Messaging Gateway packages and object types into the database. Refer to the *Oracle Database Advanced Queuing User's Guide* for information on loading database objects and using DBMS_MGWMSG.

126.2 DBMS_MGWMSG Constants

The DBMS_MGWMSG package defines constants for specifying various parameter values.

Table 126-1 DBMS_MGWMSG Constants: Value Types and Constants Representing the Type of Value for a SYS.MGW_NAME_VALUE_T Object

Value	Constant
TEXT_VALUE	CONSTANT BINARY_INTEGER := 1
RAW_VALUE	CONSTANT BINARY_INTEGER := 2
BOOLEAN_VALUE	CONSTANT BINARY_INTEGER := 3

Table 126-1 (Cont.) DBMS_MGWMSG Constants: Value Types and Constants Representing the Type of Value for a SYS.MGW_NAME_VALUE_T Object

Value	Constant
BYTE_VALUE	CONSTANT BINARY_INTEGER := 4
SHORT_VALUE	CONSTANT BINARY_INTEGER := 5
INTEGER_VALUE	CONSTANT BINARY_INTEGER := 6
LONG_VALUE	CONSTANT BINARY_INTEGER := 7
FLOAT_VALUE	CONSTANT BINARY_INTEGER := 8
DOUBLE_VALUE	CONSTANT BINARY_INTEGER := 9
DATE_VALUE	CONSTANT BINARY_INTEGER := 10

Table 126-2 DBMS_MGWMSG Constants: Boolean Values—Constants Representing a Boolean as a Numeric Value

Value	Constant
BOOLEAN_FALSE	CONSTANT BINARY_INTEGER := 0
BOOLEAN_TRUE	CONSTANT BINARY_INTEGER := 1

Table 126-3 DBMS_MGWMSG Constants: Case Comparisons

Value	Constant
CASE_SENSITIVE	CONSTANT BINARY_INTEGER := 0
CASE_INSENSITIVE	CONSTANT BINARY_INTEGER := 1

Table 126-4 Constants for the TIB/Rendezvous field type

Value	Constant
TIBRVMSG_BOOL	CONSTANT INTEGER := 1
TIBRVMSG_F32	CONSTANT INTEGER := 2
TIBRVMSG_F64	CONSTANT INTEGER := 3
TIBRVMSG_I8	CONSTANT INTEGER := 4
TIBRVMSG_I16	CONSTANT INTEGER := 5
TIBRVMSG_I32	CONSTANT INTEGER := 6
TIBRVMSG_I64	CONSTANT INTEGER := 7
TIBRVMSG_IPADDR32	CONSTANT INTEGER := 8
TIBRVMSG_IPPORT16	CONSTANT INTEGER := 9
TIBRVMSG_DATETIME	CONSTANT INTEGER := 10
TIBRVMSG_F32ARRAY	CONSTANT INTEGER := 11
TIBRVMSG_F64ARRAY	CONSTANT INTEGER := 12
TIBRVMSG_I8ARRAY	CONSTANT INTEGER := 13

Table 126-4 (Cont.) Constants for the TIB/Rendezvous field type

Value	Constant
TIBRVMSG_I16ARRAY	CONSTANT INTEGER := 14
TIBRVMSG_I32ARRAY	CONSTANT INTEGER := 15
TIBRVMSG_I64ARRAY	CONSTANT INTEGER := 16
TIBRVMSG_OPAQUE	CONSTANT INTEGER := 17
TIBRVMSG_STRING	CONSTANT INTEGER := 18
TIBRVMSG_XML	CONSTANT INTEGER := 19

126.3 DBMS_MGWMSG Types

DBMS_MGWMSG uses types to specify certain information.

- SYS.MGW_NAME_VALUE_T Type
- SYS.MGW_NAME_VALUE_T Type-Attribute Mapping
- SYS.MGW_NAME_TYPE_ARRAY_T Type
- SYS.MGW_TEXT_VALUE_T Type
- SYS.MGW_RAW_VALUE_T Type
- SYS.MGW_BASIC_MSG_T Type
- SYS.MGW_NUMBER_ARRAY_T Type
- SYS.MGW_TIBRV_FIELD_T Type
- SYS.MGW_TIBRV_MSG_T Type

126.3.1 DBMS_MGWMSG SYS.MGW_NAME_VALUE_T Type

This type specifies a named value.

The name attribute, type attribute, and one of the <>_value attributes are typically not NULL.

Syntax

```
TYPE SYS.MGW_NAME_VALUE_T IS OBJECT(
   name               VARCHAR2(250),
   type               INTEGER,
   integer_value      INTEGER,
   number_value       NUMBER,
   text_value         VARCHAR2(4000),
   raw_value          RAW(2000),
   date_value         DATE,

-- Methods
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_BOOLEAN (
   name   IN VARCHAR2,
   value  IN INTEGER )
RETURN SYS.MGW_NAME_VALUE_T,
```

```
STATIC FUNCTION CONSTRUCT_BYTE (
   name   IN VARCHAR2,
   value  IN INTEGER )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_SHORT (
   name   IN VARCHAR2,
   value  IN INTEGER )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_INTEGER (
   name   IN VARCHAR2,
   value  IN INTEGER )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_LONG (
   name   IN VARCHAR2,
   value  IN NUMBER )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_FLOAT (
   name   IN VARCHAR2,
   value  IN NUMBER )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_DOUBLE (
   name   IN VARCHAR2,
   value  IN NUMBER )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_TEXT (
   name   IN VARCHAR2,
   value  IN VARCHAR2 )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_RAW (
   name   IN VARCHAR2,
   value  IN RAW )
RETURN SYS.MGW_NAME_VALUE_T,

STATIC FUNCTION CONSTRUCT_DATE (
   name   IN VARCHAR2,
   value  IN DATE )
RETURN SYS.MGW_NAME_VALUE_T );
```

Attributes

Table 126-5 SYS.MGW_NAME_VALUE_T Attributes

Attribute	Description
name	Name associated with the value
type	Value type. Refer to the DBMS_MGWMSG.<>_VALUE constants in Table 126-1. This indicates which Java datatype and class are associated with the value. It also indicates which attribute stores the value.
integer_value	Stores a numeric integer value
number_value	Stores a numeric float or large integer value

Table 126-5 (Cont.) SYS.MGW_NAME_VALUE_T Attributes

Attribute	Description
text_value	Stores a text value
raw_value	Stores a RAW (bytes) value
date_value	Stores a date value

SYS.MGW_NAME_VALUE_T Type-Attribute Mapping

Table 126-6 shows the mapping between the value type and the attribute used to store the value.

Table 126-6 SYS.MGW_NAME_VALUE_T Type Attribute Mapping

Type	Value Stored in Attribute
DBMS_MGWMSG.TEXT_VALUE	text_value
DBMS_MGWMSG.RAW_VALUE	raw_value
DBMS_MGWMSG.BOOLEAN_VALUE	integer_value
DBMS_MGWMSG.BYTE_VALUE	integer_value
DBMS_MGWMSG.SHORT_VALUE	integer_value
DBMS_MGWMSG.INTEGER_VALUE	integer_value
DBMS_MGWMSG.LONG_VALUE	number_value
DBMS_MGWMSG.FLOAT_VALUE	number_value
DBMS_MGWMSG.DOUBLE_VALUE	number_value
DBMS_MGWMSG.DATE_VALUE	date_value

CONSTRUCT Method

This method constructs a new SYS.MGW_NAME_VALUE_T instance. All attributes are assigned a value of NULL.

Syntax

```
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_NAME_VALUE_T;
```

CONSTRUCT_*TYPE* Methods

These methods construct a new SYS.MGW_NAME_VALUE_T instance initialized with the value of a specific type. Each method sets the name and type attributes and one of the <>_value attributes, as shown in the mappings in Table 126-6.

Syntax

```
STATIC FUNCTION CONSTRUCT_<> (
   name   IN VARCHAR2,
   value  IN datatype )
RETURN SYS.MGW_NAME_VALUE_T;
```

Usage Notes

The `construct_boolean` method sets the value to either `DBMS_MGWMSG.BOOLEAN_TRUE` or `DBMS_MGWMSG.BOOLEAN_FALSE`.

126.3.2 DBMS_MGWMSG SYS.MGW_NAME_TYPE_ARRAY_T Type

This type specifies an array of name-value pairs. An object of `SYS.MGW_NAME_VALUE_ARRAY_T` type can have up to 1024 elements.

Syntax

```
TYPE SYS.MGW_NAME_VALUE_ARRAY_T
  AS VARRAY (1024) OF SYS.MGW_NAME_VALUE_T;
```

126.3.3 DBMS_MGWMSG SYS.MGW_TEXT_VALUE_T Type

This type specifies a `TEXT` value. It can store a large value as a `CLOB` or a smaller value (size <= 4000) as `VARCHAR2`. Only one of the `< >_ value` attributes should be set.

Syntax

```
TYPE SYS.MGW_TEXT_VALUE_T IS OBJECT(
   small_value VARCHAR2(4000),
   large_value CLOB,

-- Methods
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_TEXT_VALUE_T);
```

Attributes

Table 126-7 SYS.MGW_TEXT_VALUE_T Attributes

Attribute	Description
`small_value`	Small `TEXT` value. Used for values <= 4000.
`large_value`	Large `TEXT` value. Used when the value is too large for the `small_value` attribute.

CONSTRUCT Method

This method constructs a new `SYS.MGW_TEXT_VALUE_T` instance. All attributes are assigned a value of `NULL`.

Syntax

```
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_TEXT_VALUE_T;
```

126.3.4 DBMS_MGWMSG SYS.MGW_RAW_VALUE_T Type

This type specifies a `RAW` value. This type can store a large value as a `BLOB` or a smaller value (size <= 2000) as `RAW`. You must set no more than one of the `< >_value` attributes.

Syntax

```
TYPE SYS.MGW_RAW_VALUE_T IS OBJECT(
   small_value RAW(2000),
```

```
   large_value BLOB,

--Methods
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_RAW_VALUE_T);
```

Attributes

Table 126-8 SYS.MGW_RAW_VALUE_T Attributes

Attribute	Description
small_value	Small RAW (bytes) value <= 2000
large_value	Large RAW value. Used when the value is too large for the small_value attribute.

CONSTRUCT Method

This method constructs a new SYS.MGW_RAW_VALUE_T instance. All attributes are assigned a value of NULL.

Syntax

```
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_RAW_VALUE_T;
```

126.3.5 DBMS_MGWMSG SYS.MGW_BASIC_MSG_T Type

This is a canonical type for a basic TEXT or RAW message. Only a single TEXT or RAW value is typically set. An object of this type must not have both TEXT and RAW set to a not NULL value at the same time.

Syntax

```
TYPE SYS.MGW_BASIC_MSG_T IS OBJECT(
   header     SYS.MGW_NAME_VALUE_ARRAY_T,
   text_body  SYS.MGW_TEXT_VALUE_T,
   raw_body   SYS.MGW_RAW_VALUE_T,

--Methods
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_BASIC_MSG_T);
```

Attributes

Table 126-9 SYS.MGW_BASIC_MSG_T Attributes

Attribute	Description
header	Message header information as an array of name-value pairs
text_body	Message body for a TEXT message
raw_body	Message body for a RAW (bytes) message

CONSTRUCT Method

This method constructs a new SYS.MGW_BASIC_MSG_T instance. All attributes are assigned a value of NULL.

Syntax

```
STATIC FUNCTION CONSTRUCT
RETURN SYS.MGW_BASIC_MSG_T;
```

126.3.6 DBMS_MGWMSG SYS.MGW_NUMBER_ARRAY_T Type

A type that specifies an array of numbers.

Syntax

```
TYPE SYS.MGW_NUMBER_ARRAY_T AS VARRAY(1024) OF NUMBER;
```

126.3.7 DBMS_MGWMSG SYS.MGW_TIBRV_FIELD_T Type

This type represents a TIB/Rendezvous message field, typically used in a read-only fashion to retrieve field information from a `SYS.MGW_TIBRV_MSG_T` instance.

Syntax

```
TYPE SYS.MGW_TIBRV_FIELD_T IS OBJECT(
   field_name       VARCHAR2(256),
   field_id         INTEGER,
   field_type       INTEGER,
   number_value     NUMBER,
   number_array_value  SYS.MGW_NUMBER_ARRAY_T,
   text_value       VARCHAR2(4000),
   raw_value        RAW(2000),
   date_value       DATE,
   clob_value       CLOB,
   blob_value       BLOB);
```

Attributes

Table 126-10 SYS.MGW_TIBRV_FIELD_T Attributes

Attribute	Description
`field_name`	Field name. This will be `NULL` if the field has no name.
`field_id`	Field identifier. If the field identifier is zero (0), then that field is considered not to have a field identifier. Otherwise the field identifier is a nonzero value that is unique for all fields of that message.
`field_type`	Field wire format datatype. The `DBMS_MGWMSG.TIBRVMSG_<>` constants represent valid values for this attribute. The value of this field discriminates which value attribute is used to store the field data.
`number_value`	Used to store a numeric value
`number_array_value`	Used to store a numeric array value
`text_value`	Used to store a small text value
`raw_value`	Used to store a small raw value
`date_value`	Used to store a date value
`clob_value`	Used to store a large text value. This is used when the text data will not fit in `text_value`, that is, when size is larger than 4000.
`blob_value`	Used to store a large raw value. This is used when the raw data will not fit in `raw_value`; that is, when size is larger than 2000.

SYS.MGW_TIBRV_FIELD_T Type and Attribute Mapping

Table 126-11 describes the mapping in type `SYS.MGW_TIBRV_FIELD_T` between the field type and attribute used to store the value.

Table 126-11 SYS.MGW_TIBRV_FIELD_T Type and Attribute Mapping

Field Type (DBMS_MGWMSG constant)	Value Stored in Attribute
TIBRVMSG_BOOL	number_value
TIBRVMSG_F32	number_value
TIBRVMSG_F64	number_value
TIBRVMSG_I8	number_value
TIBRVMSG_I16	number_value
TIBRVMSG_I32	number_value
TIBRVMSG_I64	number_value
TIBRVMSG_IPADDR32	text_value
TIBRVMSG_IPPORT16	number_value
TIBRVMSG_DATETIME	date_value
TIBRVMSG_F32ARRAY	number_array_value
TIBRVMSG_F64ARRAY	number_array_value
TIBRVMSG_I8ARRAY	number_array_value
TIBRVMSG_I16ARRAY	number_array_value
TIBRVMSG_I32ARRAY	number_array_value
TIBRVMSG_I64ARRAY	number_array_value
TIBRVMSG_OPAQUE	raw_value or blob_value
TIBRVMSG_STRING	text_value or clob_value
TIBRVMSG_XML	raw_value or blob_value

126.3.8 DBMS_MGWMSG SYS.MGW_TIBRV_MSG_T Type

This type represents a TIB/Rendezvous message. You must never directly reference the attributes of this type. Instead use the type methods.

Syntax

```
TYPE SYS.MGW_TIBRV_MSG_T IS OBJECT(
   send_subject    VARCHAR2(256),
   reply_subject   VARCHAR2(256),
   cm_time_limit   NUMBER,
   cm_sender_name  VARCHAR2(256),
   cm_sequence_num NUMBER,
   fields          SYS.MGW_TIBRV_IFIELDS_T,
   clob_data1      CLOB,
   clob_data2      CLOB,
   clob_data3      CLOB,
   blob_data1      BLOB,
   blob_data2      BLOB,
```

```
   blob_data3      BLOB,

STATIC FUNCTION construct
RETURN SYS.MGW_TIBRV_MSG_T,

MEMBER PROCEDURE add_bool (
    name  IN   VARCHAR2,
    id    IN   INTEGER,
    value IN   INTEGER ),

MEMBER PROCEDURE add_f32 (
    name  IN   VARCHAR2,
    id    IN   INTEGER,
    value IN   FLOAT ),

MEMBER PROCEDURE add_f64 (
    name  IN   VARCHAR2,
    id    IN   INTEGER,
    value IN   DOUBLE ),

MEMBER PROCEDURE add_i8 (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  INTEGER ),

MEMBER PROCEDURE add_i16 (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  INTEGER ),

MEMBER PROCEDURE add_i32 (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  INTEGER ),

MEMBER PROCEDURE add_i64 (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  NUMBER ),

MEMBER PROCEDURE add_ipaddr32 (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  VARCHAR2 ),

MEMBER PROCEDURE add_ipport16 (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  INTEGER ),

MEMBER PROCEDURE add_datetime (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  DATE ),

MEMBER PROCEDURE add_f32array (
     name  IN  VARCHAR2,
     id    IN  INTEGER,
     value IN  SYS.MGW_NUMBER_ARRAY_T ),

MEMBER PROCEDURE add_f64array (
     name  IN  VARCHAR2,
```

```
    id    IN  INTEGER,
    value IN  SYS.MGW_NUMBER_ARRAY_T ),

MEMBER PROCEDURE add_i8array (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  SYS.MGW_NUMBER_ARRAY_T ),

MEMBER PROCEDURE add_i16array (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  SYS.MGW_NUMBER_ARRAY_T ),

MEMBER PROCEDURE add_i32array (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  SYS.MGW_NUMBER_ARRAY_T ),

MEMBER PROCEDURE add_i64array (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  SYS.MGW_NUMBER_ARRAY_T ),

MEMBER PROCEDURE add_string (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  VARCHAR2 ),

MEMBER PROCEDURE add_string (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  CLOB ),

MEMBER PROCEDURE add_opaque (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  RAW ),

MEMBER PROCEDURE add_opaque (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  BLOB ),

MEMBER PROCEDURE add_xml (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  RAW ),

MEMBER PROCEDURE add_xml (
    name  IN  VARCHAR2,
    id    IN  INTEGER,
    value IN  BLOB ),

MEMBER PROCEDURE set_send_subject (
    value  IN  VARCHAR2 ),

MEMBER PROCEDURE set_reply_subject (
    value  IN  VARCHAR2 ),

MEMBER PROCEDURE set_cm_time_limit (
    value  IN  NUMBER ),
```

```
MEMBER PROCEDURE set_cm_sender_name (
     value IN  VARCHAR2 ),

MEMBER PROCEDURE set_cm_sequence_num (
     value IN NUMBER ),

MEMBER FUNCTION get_send_subject
RETURN VARCHAR2,

MEMBER FUNCTION get_reply_subject
RETURN VARCHAR2,

MEMBER FUNCTION get_cm_time_limit
RETURN NUMBER,

MEMBER FUNCTION get_cm_sender_name
RETURN VARCHAR2,

MEMBER FUNCTION get_cm_sequence_num
RETURN NUMBER,

MEMBER FUNCTION get_field_count
RETURN INTEGER,

MEMBER FUNCTION get_field (
     idx    IN INTEGER )
RETURN SYS.MGW_TIBRV_FIELD_T,

MEMBER FUNCTION get_field_by_name (
     name   IN  VARCHAR2 )
RETURN SYS.MGW_TIBRV_FIELD_T,

MEMBER FUNCTION get_field_by_id (
     id     IN  INTEGER )
RETURN SYS.MGW_TIBRV_FIELD_T,

MEMBER FUNCTION find_field_name (
     name   IN  VARCHAR2,
     start_idx IN INTEGER )
RETURN INTEGER,

MEMBER FUNCTION find_field_id (
     id     IN INTEGER,
     start_idx IN INTEGER )
RETURN INTEGER
);
```

Attributes

Table 126-12 SYS.MGW_TIBRV_MSG_T Type Attributes

Attribute	Description
send_subject	Send subject name
reply_subject	Reply subject name
cm_time_limit	Time limit for a certified message
cm_sender_name	Sender name of a certified message
cm_sequence_num	Sequence number of a certified message

Table 126-12 (Cont.) SYS.MGW_TIBRV_MSG_T Type Attributes

Attribute	Description
fields	Collection of message fields
clob_data1	Used to store a large text value
clob_data2	Used to store a large text value
clob_data3	Used to store a large text value
blob_data1	Used to store a large raw value
blob_data2	Used to store a large raw value
blob_data3	Used to store a large raw value

Construct Method

Constructs a new SYS.MGW_TIBRV_MSG_T instance. All attributes are set to NULL.

Syntax

```
STATIC FUNCTION construct
RETURN SYS.MGW_TIBRV_MSG_T;
```

ADD_<> Methods

Adds a new field to the message.

Syntax

```
MEMBER PROCEDURE ADD_<> (
   name  IN  VARCHAR2,
   id    IN  INTEGER,
   value IN  datatype );
```

Parameters

Table 126-13 SYS.MGW_TIBRV_MSG_T ADD_<> Method Parameters

Parameter	Description
name	Field name
id	Field identifier
value	Field data

Table 126-14 shows, for each add method, the field type that will be assigned and valid values for the field data.

Table 126-14 MGW_TIBRV_MSG_T Add Method Field Types

Method Name	Field Type Assigned	Comment
add_bool	TIBRVMSG_BOOL	Valid values: 0 (false), 1 (true)
add_f32	TIBRVMSG_F32	n/a
add_f64	TIBRVMSG_F64	n/a
add_i8	TIBRVMSG_I8	Valid range: -128...127

Table 126-14 (Cont.) MGW_TIBRV_MSG_T Add Method Field Types

Method Name	Field Type Assigned	Comment
add_i16	TIBRVMSG_I16	Valid range: -32768...32767
add_i32	TIBRVMSG_I32	Valid range: -2147483648... 2147483647
add_i64	TIBRVMSG_I64	n/a
add_ipaddr32	TIBRVMSG_IPADDR32	n/a
add_ipport16	TIBRVMSG_IPPORT16	n/a
add_datetime	TIBRVMSG_DATETIME	n/a
add_f32array	TIBRVMSG_F32ARRAY	n/a
add_f64array	TIBRVMSG_F64ARRAY	n/a
add_i8array	TIBRVMSG_I8ARRAY	Valid range: -128...127
add_i16array	TIBRVMSG_I16ARRAY	Valid range: -32768...32767
add_i32array	TIBRVMSG_I32ARRAY	Valid range: -2147483648... 2147483647
add_i64array	TIBRVMSG_I64ARRAY	n/a
add_opaque	TIBRVMSG_OPAQUE	Value stored as RAW if size < 2000; otherwise value stored in BLOB
add_string	TIBRVMSG_STRING	Value stored as VARCHAR2 if size < 4000; otherwise value stored in CLOB
add_xml	TIBRVMSG_XML	Value stored as RAW if size < 2000; otherwise value stored in BLOB

SET_<> Methods

Accessor methods to set an instance attribute to a specific value.

Syntax

```
MEMBER PROCEDURE SET_<> (
  value IN datatype );
```

Parameters

Table 126-15 SYS.MGW_TIBRV_MSG_T SET_<> Method Parameters

Parameter	Description
value	Value to be assigned

GET_<> Methods

Accessor methods to retrieve the value for an instance attribute.

Syntax

```
MEMBER PROCEDURE GET_<>
RETURN datatype;
```

Parameters

None

Return Values

Returns the attribute value.

GET_FIELD_COUNT Function

Gets the number of message fields.

Syntax

```
MEMBER PROCEDURE get_field_count
RETURN INTEGER;
```

Parameters

None

Return Values

Returns the number of fields, or zero (0) if there are none.

GET_FIELD Function

Retrieves field information for the field having a given field collection index. This method should only be called if the `GET_FIELD_COUNT Function` returns a nonzero value and `idx` must specify a valid collection index; that is, `1<=idx<=get_field_count()`.

Syntax

```
MEMBER PROCEDURE get_field (
    idx    IN    INTEGER )
RETURN SYS.MGW_TIBRV_FIELD_T;
```

Parameters

Table 126-16 SYS.MGW_TIBRV_MSG_T GET_FIELD Function Parameters

Parameter	Description
`idx`	Specifies the 1-based field collection index of the field to retrieve

> **Note:**
>
> A 1-based index begins at one (1) instead of zero (0).

Return Values

Returns the field information.

GET_FIELD_BY_NAME Function

Retrieves field information for the first field that has a given field name. The name comparison is case-sensitive.

Syntax

```
MEMBER PROCEDURE get_field_by_name (
   name  IN   VARCHAR2 )
RETURN SYS.MGW_TIBRV_FIELD_T;
```

Parameters

Table 126-17 SYS.MGW_TIBRV_MSG_T GET_FIELD_BY_NAME Function Parameters

Parameter	Description
name	Specifies the field name to search for. This can be NULL to find the first field that does not have a field name.

Return Values

Returns the field information, or NULL if no match was found.

GET_FIELD_BY_ID Function

Retrieves field information for the first field that has a given field identifier.

A field can have either a unique identifier or no identifier. If the field identifier value is zero (0) or NULL, then the field is considered to have no identifier. Otherwise, the identifier is a nonzero value that is unique for all the fields of this message.

Syntax

```
MEMBER PROCEDURE get_field_by_id (
   id   IN  INTEGER )
RETURN SYS.MGW_TIBRV_FIELD_T;
```

Parameters

Table 126-18 SYS.MGW_TIBRV_MSG_T GET_FIELD_BY_ID Function Parameters

Parameter	Description
id	Specifies the field identifier to search for. This can be zero (0) or NULL to find the first field that does not have an identifier.

Return Values

Returns the field information, or NULL if no match was found.

FIND_FIELD_NAME Function

Searches for a field with a given field name, starting from a given index of the field collection. It returns the index of that field. The name comparison is case-sensitive. This function is useful for finding all the fields that have the same name.

Syntax

```
MEMBER PROCEDURE find_field_name (
   name      IN  VARCHAR2,
   start_idx IN  INTEGER )
RETURN INTEGER;
```

Parameters

Table 126-19 SYS.MGW_TIBRV_MSG_T FIND_FIELD_NAME Function Parameters

Parameter	Description
name	Specifies the field name to search for. This can be NULL to search for a field that does not have a field name.
start_idx	Specifies the 1-based field collection index from which the search should start.

Return Values

Returns the field index (> 0) if a match was found, or zero (0) if no match was found.

FIND_FIELD_ID Function

Searches for a field with a given field identifier, starting from a given index of the field collection. It returns the index of that field.

Syntax

```
MEMBER PROCEDURE find_field_id (
     id        IN  INTEGER,
     start_idx IN INTEGER )
RETURN    INTEGER;
```

Parameters

Table 126-20 SYS.MGW_TIBRV_MSG_T FIND_FIELD_ID Function Parameters

Parameter	Description
id	Specifies the field identifier to search for. This can be zero (0) or NULL to find a field that does not have an identifier.
start_idx	Specifies the 1-based field collection index from which the search should start.

Return Values

Returns the field index (> 0) if a match was found, or zero (0) if no match was found.

126.4 Summary of DBMS_MGWMSG Subprograms

This table lists the DDBMS_MGWMSG subprograms and briefly describes them.

Table 126-21 DBMS_MGWMSG Package Subprograms

Subprogram	Description
LCR_TO_XML Function	Converts a SYS.ANYDATA object encapsulating a row LCR (LCR$_ROW_RECORD) or a DDL LCR (LCR$_DDL_RECORD) to a SYS.XMLTYPE object
NVARRAY_ADD Procedure	Appends a name-value element to the end of a name-value array
NVARRAY_FIND_NAME Function	Searches a name-value array for the element with the name you specify in p_name

Table 126-21 (Cont.) DBMS_MGWMSG Package Subprograms

Subprogram	Description
NVARRAY_FIND_NAME_TYPE Function	Searches a name-value array for an element with the name and value type you specify
NVARRAY_GET Function	Gets the name-value element of the name you specify in `p_name` from a name-value array
NVARRAY_GET_BOOLEAN Function	Gets the value of the name-value array element that you specify in `p_name` and with the `BOOLEAN_VALUE` value type
NVARRAY_GET_BYTE Function	Gets the value of the name-value array element that you specify in `p_name` and with the `BYTE_VALUE` value type
NVARRAY_GET_DATE Function	Gets the value of the name-value array element that you specify in `p_name` and with the `DATE_VALUE` value type
NVARRAY_GET_DOUBLE Function	Gets the value of the name-value array element that you specify in `p_name` and with the `DOUBLE_VALUE` value type
NVARRAY_GET_FLOAT Function	Gets the value of the name-value array element that you specify in `p_name` and with the `FLOAT_VALUE` value type
NVARRAY_GET_INTEGER Function	Gets the value of the name-value array element that you specify in `p_name` and with the `INTEGER_VALUE` value type
NVARRAY_GET_LONG Function	Gets the value of the name-value array element that you specify in `p_name` and with the `LONG_VALUE` value type
NVARRAY_GET_RAW Function	Gets the value of the name-value array element that you specify in `p_name` and with the `RAW_VALUE` value type
NVARRAY_GET_SHORT Function	Gets the value of the name-value array element that you specify in `p_name` and with the `SHORT_VALUE` value type
NVARRAY_GET_TEXT Function	Gets the value of the name-value array element that you specify in `p_name` and with the `TEXT_VALUE` value type
XML_TO_LCR Function	Converts a `SYS.XMLTYPE` object to a `SYS.ANYDATA` object encapsulating a row LCR (`LCR$_ROW_RECORD`) or a DDL LCR (`LCR$_DDL_RECORD`)

126.4.1 LCR_TO_XML Function

This function converts a `SYS.ANYDATA` object encapsulating a row LCR (Logical Change Record, in this case a `LCR$_ROW_RECORD`) or a DDL LCR (`LCR$_DDL_RECORD`) to a `SYS.XMLTYPE` object.

> **See Also:**
>
> XML_TO_LCR Function

Syntax

```
DBMS_MGWMSG.LCR_TO_XML (
   p_anydata IN SYS.ANYDATA )
 RETURN SYS.XMLTYPE;
```

Parameters

Table 126-22 LCR_TO_XML Function Parameters

Parameter	Description
p_anydata	An ANYDATA object to be converted

Return Values

Returns a SYS.XMLTYPE object.

Usage Notes

An exception is raised if the encapsulated type p_anydata is not an LCR.

126.4.2 NVARRAY_ADD Procedure

This procedure appends a name-value element to the end of a name-value array.

Syntax

```
DBMS_MGWMSG.NVARRAY_ADD (
   p_array  IN OUT SYS.MGW_NAME_VALUE_ARRAY_T,
   p_value  IN     SYS.MGW_NAME_VALUE_T );
```

Parameters

Table 126-23 NVARRAY_ADD Procedure Parameters

Parameter	Description
p_array	On input, the name-value array instance to modify. If NULL, then a new array is created. On output, the modified name-value array instance.
p_value	The value to add. If NULL, then p_array is not changed.

126.4.3 NVARRAY_FIND_NAME Function

This function searches a name-value array for the element with the name you specify in p_name.

Syntax

```
DBMS_MGWMSG.NVARRAY_FIND_NAME (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN BINARY_INTEGER;
```

Parameters

Table 126-24 NVARRAY_FIND_NAME Function Parameters

Parameters	Description
p_array	The name-value array to search
p_name	The name to find
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns a positive integer that is the array index of the matching element or zero (0) if the specified name is not found.

126.4.4 NVARRAY_FIND_NAME_TYPE Function

This function searches a name-value array for an element with the name and value type you specify.

Syntax

```
DBMS_MGWMSG.NVARRAY_FIND_NAME_TYPE (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_type     IN BINARY_INTEGER
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN BINARY_INTEGER;
```

Parameters

Table 126-25 NVARRAY_FIND_NAME_TYPE Function Parameters

Parameter	Description
p_array	The name-value array to search
p_name	The name to find
p_type	The value type. Refer to the value type constants in Table 126-1 .
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns a positive integer that is the array index of the matching element, zero (0) if the specified name is not found, or negative one (-1) if the specified name is found but a type mismatch exists.

126.4.5 NVARRAY_GET Function

This function gets the name-value element of the name you specify in `p_name` from a name-value array.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN SYS.MGW_NAME_VALUE_T;
```

Parameters

Table 126-26 NVARRAY_GET Function Parameters

Parameter	Description
`p_array`	The name-value array
`p_name`	The value name
`p_compare`	Name comparison method. Values are `CASE_SENSITIVE` and `CASE_INSENSITIVE`.

Return Values

Returns the matching element, or `NULL` if the specified name is not found.

126.4.6 NVARRAY_GET_BOOLEAN Function

This function gets the value of the name-value array element that you specify in `p_name` and with the `BOOLEAN_VALUE` value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_BOOLEAN (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN INTEGER;
```

Parameters

Table 126-27 NVARRAY_GET_BOOLEAN Function Parameters

Parameter	Description
`p_array`	The name-value array
`p_name`	The value name
`p_compare`	Name comparison method. Values are `CASE_SENSITIVE` and `CASE_INSENSITIVE`.

Return Values

Returns the value, or `NULL` if either the specified name is not found or a type mismatch exists.

126.4.7 NVARRAY_GET_BYTE Function

This function gets the value of the name-value array element that you specify in `p_name` and with the `BYTE_VALUE` value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_BYTE (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN INTEGER;
```

Parameters

Table 126-28 NVARRAY_GET_BYTE Function Parameters

Parameter	Description
`p_array`	The name-value array
`p_name`	The value name
`p_compare`	Name comparison method. Values are `CASE_SENSITIVE` and `CASE_INSENSITIVE`.

Return Values

Returns the value, or `NULL` if either the specified name is not found or a type mismatch exists.

126.4.8 NVARRAY_GET_DATE Function

This function gets the value of the name-value array element that you specify in `p_name` and with the `DATE_VALUE` value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_DATE (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN DATE;
```

Parameters

Table 126-29 NVARRAY_GET_DATE Function Parameters

Parameters	Description
`p_array`	The name-value array
`p_name`	The value name

Table 126-29 (Cont.) NVARRAY_GET_DATE Function Parameters

Parameters	Description
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.9 NVARRAY_GET_DOUBLE Function

This function gets the value of the name-value array element that you specify in p_name and with the DOUBLE_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_DOUBLE (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN NUMBER;
```

Parameters

Table 126-30 NVARRAY_GET_DOUBLE Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.10 NVARRAY_GET_FLOAT Function

This function gets the value of the name-value array element that you specify in p_name and with the FLOAT_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_FLOAT (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN NUMBER;
```

Parameters

Table 126-31 NVARRAY_GET_FLOAT Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.11 NVARRAY_GET_INTEGER Function

This function gets the value of the name-value array element that you specify in p_name and with the INTEGER_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_INTEGER (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN INTEGER;
```

Parameters

Table 126-32 NVARRAY_GET_INTEGER Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.12 NVARRAY_GET_LONG Function

This function gets the value of the name-value array element that you specify in p_name and with the LONG_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_LONG (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
```

```
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN NUMBER;
```

Parameters

Table 126-33 NVARRAY_GET_LONG Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.13 NVARRAY_GET_RAW Function

This function gets the value of the name-value array element that you specify in p_name and with the RAW_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_RAW (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN RAW;
```

Parameters

Table 126-34 NVARRAY_GET_RAW Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.14 NVARRAY_GET_SHORT Function

This function gets the value of the name-value array element that you specify in p_name and with the SHORT_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_SHORT (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
```

```
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN INTEGER;
```

Parameters

Table 126-35 NVARRAY_GET_SHORT Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.15 NVARRAY_GET_TEXT Function

This function gets the value of the name-value array element that you specify in p_name and with the TEXT_VALUE value type.

Syntax

```
DBMS_MGWMSG.NVARRAY_GET_TEXT (
   p_array    IN SYS.MGW_NAME_VALUE_ARRAY_T,
   p_name     IN VARCHAR2,
   p_compare  IN BINARY_INTEGER DEFAULT CASE_SENSITIVE )
RETURN VARCHAR2;
```

Parameters

Table 126-36 NVARRAY_GET_TEXT Function Parameters

Parameter	Description
p_array	The name-value array
p_name	The value name
p_compare	Name comparison method. Values are CASE_SENSITIVE and CASE_INSENSITIVE.

Return Values

Returns the value, or NULL if either the specified name is not found or a type mismatch exists.

126.4.16 XML_TO_LCR Function

This function converts a `SYS.XMLTYPE` object to a `SYS.ANYDATA` object encapsulating a row LCR (`LCR$_ROW_RECORD`) or a DDL LCR (`LCR$_DDL_RECORD`).

> **See Also:**
>
> LCR_TO_XML Function

Syntax

```
DBMS_MGWMSG.XML_TO_LCR (
   p_xmldata IN SYS.XMLTYPE )
 RETURN SYS.ANYDATA;
```

Parameters

Table 126-37 XML_TO_LCR Function Parameters

Parameter	Description
`p_xmldata`	An `XMLTYPE` object representing an LCR

Return Values

Returns a `SYS.ANYDATA` object.

Usage Notes

An exception is raised if `p_xmldata` cannot be converted to an LCR.

127
DBMS_MLE

The `DBMS_MLE` package allows users to execute JavaScript code inside the Oracle Database and exchange data seamlessly between PL/SQL and JavaScript. The JavaScript code itself can execute PL/SQL and SQL through built-in JavaScript modules. JavaScript data types are automatically mapped to Oracle Database data types and vice versa.

With the `DBMS_MLE` package, developers can write their data processing logic in JavaScript.

This chapter contains the following topics:

- DBMS_MLE Overview
- DBMS_MLE Security Model
- Summary of DBMS_MLE Subprograms

127.1 DBMS_MLE Overview

Runtime state for MLE execution is encapsulated in execution contexts that users can explicitly create and destroy. Users can also export values from PL/SQL to MLE, and import values from MLE back into PL/SQL.

The following example captures the typical workflow for MLE execution:

```
set serveroutput on;
declare
  ctx dbms_mle.context_handle_t;
  source clob;
  greeting varchar2(100);
begin
  ctx := dbms_mle.create_context(); -- Create execution context for MLE execution
  dbms_mle.export_to_mle(ctx, 'person', 'World'); -- Export value from PL/SQL
  source := q'~
    var bindings = require("mle-js-bindings");
    var person = bindings.importValue("person"); // Import value previously exported
from PL/SQL
    var greeting = "Hello, " + person + "!";
    bindings.exportValue("greeting", greeting); // Export value to PL/SQL
  ~';
  dbms_mle.eval(ctx, 'JAVASCRIPT', source); -- Evaluate the source code snippet in the
execution context
  dbms_mle.import_from_mle(ctx, 'greeting', greeting); -- Import value previously
exported from MLE
  dbms_output.put_line('Greetings from MLE: ' || greeting);
  dbms_mle.drop_context(ctx); -- Drop the execution context once no longer required
end;
/
```

Executing the above code block produces the following output:

```
Greetings from MLE: Hello, World!
```

127.2 DBMS_MLE Security Model

Access to MLE features is protected by database privileges. The user must have the `EXECUTE DYNAMIC MLE` privilege while calling any of its functions and procedures that pertain to MLE execution. In addition to this, the user must also have the `EXECUTE ON JAVASCRIPT` privilege to execute JavaScript code. An `ORA-01031` error is raised if the user calling any of the `DBMS_MLE` subprograms does not have the appropriate privileges. See the summary of `DBMS_MLE` subprograms for the privileges required to call each `DBMS_MLE` subprogram.

127.3 Summary of DBMS_MLE Subprograms

This table lists the `DBMS_MLE` subprograms and briefly describes them.

Table 127-1 DBMS_MLE Package Subprograms

Subprogram	Description
CREATE_CONTEXT Function	Creates an MLE context for executing snippets in JavaScript.
DISABLE_CTX_STDERR Procedure	This procedure disables stderr stream of the specified MLE context, so that future writes to stderr are discarded.
DISABLE_CTX_STDOUT Procedure	This procedure disables stdout stream of the specified MLE context, so that future writes to `stdout` are discarded.
DISABLE_DEBUGGING Procedure	This procedure disables any currently enabled debugpoints for the current session.
DISABLE_ICS_STDERR Procedure	This procedure disables the `stderr` stream of the inlined MLE call specification context, so that future writes to `stderr` are discarded for the calling user in the current session.
DISABLE_ICS_STDOUT Procedure	This procedure disables the `stdout` stream of an inlined call specification context, so that future writes to `stdout` are discarded for the calling user in the current session.
DISABLE_STDERR Procedure	This procedure disables the stderr stream of MLE contexts, so that future writes to `stderr` are discarded.
DISABLE_STDOUT Procedure	This procedure disables the stdout stream of MLE contexts, so that future writes to stdout are discarded.
DROP_CONTEXT Procedure	This procedure is used to drop an MLE context that was previously created using the `CREATE_CONTEXT` procedure. After the context is dropped, the context handle is no longer valid and cannot be used anymore.
ENABLE_DEBUGGING Procedure	This procedure enables a set of debugpoints for the current session.
EVAL Procedure	This procedure executes the given JavaScript code within the context identified by the context handle.
EXPORT_TO_MLE Procedure	This procedure allows you to assign the given value, with appropriate conversion, to the named property in the MLE context. The property is created if it is not already present.
GET_AVAILABLE_LANGUAGES Function	This function returns the set of available MLE languages.
GET_CTX_ERROR_STACK Function	This function retrieves the JavaScript stack trace for the most recent application error in the given execution context.
GET_ERROR_STACK Function	This function retrieves the JavaScript stack trace for the most recent application error in the given module (and optional environment) call.

Table 127-1 (Cont.) DBMS_MLE Package Subprograms

Subprogram	Description
IMPORT_FROM_MLE Procedure	This procedure retrieves the value of the named property from the MLE context and converts it to the requested PL/SQL type.
PARSE_DEBUG_OUTPUT Function	Given a BLOB containing MLE debug output in the Java Heap Dump format, this function returns a textual representation of debug output.
SET_CTX_STDERR Procedure	This procedure redirects the stderr stream of the MLE context to the given `CLOB`.
SET_CTX_STDERR_TO_DBMS_OUTPUT Procedure	This procedure redirects the stderr stream of the MLE context to `DBMS_OUTPUT`.
SET_CTX_STDOUT Procedure	This procedure redirects the stdout stream of the MLE context to the given `CLOB`.
SET_CTX_STDOUT_TO_DBMS_OUTPUT Procedure	This procedure redirects the stdout stream of the MLE context to `DBMS_OUTPUT`.
SET_ICS_STDERR Procedure	This procedure redirects the `stderr` stream of the inlined MLE call specification context to the given `CLOB` for the calling user in the current session.
SET_ICS_STDERR_TO_DBMS_OUTPUT Procedure	This procedure redirects the `stderr` stream of the inlined MLE call specification context to `DBMS_OUTPUT` for the calling user in the current session.
SET_ICS_STDOUT Procedure	This procedure redirects the `stdout` stream of an inlined MLE call specification context in the current session to the given `CLOB`.
SET_ICS_STDOUT_TO_DBMS_OUTPUT Procedure	This procedure redirects the `stdout` stream of the inlined call specification context to `DBMS_OUTPUT` for the calling user in the current session.
SET_STDERR Procedure	This procedure redirects the stderr stream of MLE contexts to the given `CLOB`.
SET_STDERR_TO_DBMS_OUTPUT Procedure	This procedure redirects the stderr stream of MLE contexts to `DBMS_OUTPUT`.
SET_STDOUT Procedure	This procedure redirects the stdout stream of MLE contexts to the given `CLOB`.
SET_STDOUT_TO_DBMS_OUTPUT Procedure	This procedure redirects the stdout stream of MLE contexts to `DBMS_OUTPUT`.

127.3.1 CREATE_CONTEXT Function

Creates an MLE context for executing snippets in JavaScript. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.CREATE_CONTEXT
   RETURN CONTEXT_HANDLE_T;

DBMS_MLE.CREATE_CONTEXT(
   environment      IN      VARCHAR2)
RETURN CONTEXT_HANDLE_T;
```

Parameters

Table 127-2 CREATE_CONTEXT Function Parameters

Parameter	Description
`environment`	The case-sensitive schema name of the MLE environment that configures the context. Optional.

Return Values

This function returns a handle that uniquely identifies a context within a session for use in subsequent operations, such as `EXPORT_TO_MLE` and `EVAL`.

Usage Notes

A context has a lifetime limited to the session in which it was created. When a client session is terminated, all its contexts are dropped. All MLE contexts created in a session are also dropped when the session state is reset, for example, by calling `DBMS_SESSION.RESET_PACKAGE`. JavaScript code is evaluated in the context using the user, roles, and schema that are in effect at the time of context creation.

The function may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04105`: if the environment does not exist.

> **See Also:**
>
> *Oracle Database JavaScript Developer's Guide* for more details about privileges required to execute JavaScript code

127.3.2 DISABLE_CTX_STDERR Procedure

This procedure disables `stderr` stream of the specified MLE context, so that future writes to `stderr` are discarded. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.DISABLE_CTX_STDERR(
      context_handle            IN            context_handle_t);
```

Parameters

Table 127-3 DISABLE_CTX_STDERR Procedure Parameters

Parameter	Description
`context_handle`	The handle to an MLE context in the current session.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges

127.3.3 DISABLE_CTX_STDOUT Procedure

This procedure disables `stdout` stream of the specified MLE context, so that future writes to `stdout` are discarded. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.DISABLE_CTX_STDOUT(
      context_handle            IN            context_handle_t);
```

Parameters

Table 127-4 DISABLE_CTX_STDOUT Procedure Parameters

Parameter	Description
`context_handle`	The handle to an MLE context in the current session.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges

127.3.4 DISABLE_DEBUGGING Procedure

This procedure disables any currently enabled debugpoints for the current session. Post-execution debugging allows you to collect runtime state to be used for analysis after the program has been run. Post-execution debugging can only be applied to MLE code deployed in modules as opposed to code deployed using dynamic execution.

Syntax

```
DBMS_MLE.DISABLE_DEBUGGING();
```

Usage Notes

This procedure has no effect if no debugpoints are currently enabled. Debugging can be enabled again with a subsequent call to `DBMS_MLE.ENABLE_DEBUGGING`.

> **See Also:**
>
> *Oracle Database JavaScript Developer's Guide* for more information about post-execution debugging with MLE

127.3.5 DISABLE_ICS_STDERR Procedure

This procedure disables the `stderr` stream of the inlined MLE call specification context, so that future writes to `stderr` are discarded for the calling user in the current session. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.DISABLE_ICS_STDERR(
      name       IN       VARCHAR2);
```

Parameters

Table 127-5 DISABLE_ICS_STDERR Procedure Parameters

Parameter	Description
`name`	The name of the inlined MLE call specification.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04043`: if the call specification does not exist.

127.3.6 DISABLE_ICS_STDOUT Procedure

This procedure disables the `stdout` stream of an inlined call specification context, so that future writes to `stdout` are discarded for the calling user in the current session. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.DISABLE_ICS_STDOUT(
      name       IN       VARCHAR2);
```

Parameters

Table 127-6 DISABLE_ICS_STDOUT Procedure Parameters

Parameter	Description
`name`	The name of the inlined MLE call specification.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.

- ORA-04043: if the call specification does not exist.

127.3.7 DISABLE_STDERR Procedure

This procedure disables the stderr stream of MLE contexts, so that future writes to stderr are discarded. You need the EXECUTE DYNAMIC MLE privilege to execute this function.

Syntax

```
DBMS_MLE.DISABLE_STDERR();

DBMS_MLE.DISABLE_STDERR (
      module_name       IN      VARCHAR2);

DBMS_MLE.DISABLE_STDERR (
      module_name       IN      VARCHAR2,
      env_name          IN      VARCHAR2);
```

Parameters

Table 127-7 DISABLE_STDERR Function Parameters

Parameter	Description
module_name	The name of the MLE module.
env_name	The name of the MLE environment.

Usage Notes

When called without parameters, this procedure applies to all existing contexts and contexts created in the future. Otherwise, only the context associated with the given module (or module and environment combination) is affected.

If no environment is specified, the context defined by the given module and the built-in environment is used.

Any output that was buffered so far gets flushed to the pre-existing sink.

The procedure may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges
- ORA-04103: if the module does not exist
- ORA-04105: if the environment does not exist

127.3.8 DISABLE_STDOUT Procedure

This procedure disables the stdout stream of MLE contexts, so that future writes to stdout are discarded. You need the EXECUTE DYNAMIC MLE privilege to execute this function.

Syntax

```
DBMS_MLE.DISABLE_STDOUT();

DBMS_MLE.DISABLE_STDOUT(
```

```
      module_name       IN      VARCHAR2);

DBMS_MLE.DISABLE_STDOUT(
      module_name       IN      VARCHAR2,
      env_name          IN      VARCHAR2);
```

Parameters

Table 127-8 DISABLE_STDOUT Function Parameters

Parameter	Description
module_name	The name of the MLE module.
env_name	The name of the MLE environment.

Usage Notes

When called without parameters, this procedure applies to all existing contexts and contexts created in the future. Otherwise, only the context associated with the given module (or module and environment combination) is affected.

If no environment is specified, the context defined by the given module and the built-in environment is used.

Any output that was buffered so far gets flushed to the pre-existing sink.

The procedure may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges
- ORA-04103: if the module does not exist
- ORA-04105: if the environment does not exist

127.3.9 DROP_CONTEXT Procedure

This procedure is used to drop an MLE context that was previously created using the CREATE_CONTEXT function. After the context is dropped, the context handle is no longer valid and cannot be used anymore. You need the EXECUTE DYNAMIC MLE privilege to execute this procedure.

Syntax

```
DBMS_MLE.DROP_CONTEXT(
   context_handle    IN      context_handle_t);
```

Parameters

Table 127-9 DROP_CONTEXT Procedure Parameters

Parameter	Description
context_handle	The handle to an MLE context in the current session.

Usage Notes

The procedure may raise ORA-01031 error if the caller does not have sufficient privileges.

127.3.10 ENABLE_DEBUGGING Procedure

This procedure enables a set of debugpoints for the current session. Post-execution debugging allows you to collect runtime state to be used for analysis after the program has been run. Post-execution debugging can only be applied to MLE code deployed in modules as opposed to code deployed using dynamic execution.

Syntax

```
DBMS_MLE.ENABLE_DEBUGGING(
   debugspec          IN         JSON,
   sink               OUT        NOCOPY BLOB);
```

Parameters

Table 127-10 ENABLE_DEBUGGING Procedure Parameters

Parameter	Description
`debugspec`	The debug specification as a JSON document that identifies the debugging information to be collected.
`sink`	The sink to log debug output to.

Usage Notes

Calling this procedure multiple times in the same session replaces any existing set of debugpoints.

All enabled debugpoints are automatically disabled once the session ends.

The procedure may raise the following errors:

- `ORA-04103`: if the module does not exist.
- `ORA-04162`: if an attempt is made to debug an MLE built-in module.
- `ORA-04164`: if the caller is missing the `COLLECT DEBUG INFO` privilege on the module.
- `ORA-04165`: if the provided debug specification is invalid.

> **See Also:**
>
> *Oracle Database JavaScript Developer's Guide* for more information about post-execution debugging with MLE

127.3.11 EVAL Procedure

This procedure executes the given JavaScript code within the context identified by the context handle.

The evaluated code has access to all previous modifications to the state of the context, including variables defined by code previously evaluated in the context and values exported through `EXPORT_TO_MLE()`. The evaluated code can also import MLE built-in modules such as the MLE SQL driver.

You need the `EXECUTE DYNAMIC MLE` privilege to execute this procedure. It also requires the `EXECUTE ON JAVASCRIPT` privilege.

Syntax

```
DBMS_MLE.EVAL(
   context_handle   IN         context_handle_t,
   language_id      IN         language_t,
   source           IN         CLOB,
   result           IN OUT     NOCOPY CLOB CHARACTER SET ANY_CS,
   options          IN         VARCHAR2         DEFAULT NULL,
   source_name      IN         VARCHAR2         DEFAULT NULL);

DBMS_MLE.EVAL(
   context_handle   IN         context_handle_t,
   language_id      IN         language_t,
   source           IN         VARCHAR2,
   result           IN OUT     NOCOPY CLOB CHARACTER SET ANY_CS,
   options          IN         VARCHAR2         DEFAULT NULL,
   source_name      IN         VARCHAR2         DEFAULT NULL);
```

Parameters

Table 127-11 EVAL Procedure Parameters

Parameter	Description
`context_handle`	The handle to an MLE context.
`language_id`	The language of the provided source code. The value `'JAVASCRIPT'` must be provided as of Oracle 23ai.
`source`	The source code to be executed.
`result`	A buffer to which the result of the evaluation of the source code is appended. Optional.
`options`	Reserved for future use. Optional.
`source_name`	A name for the provided source code that is used to identify the snippet in stack traces. Optional.

Usage Notes

When specifying the optional `source_name` parameter, the `options` parameter must be defined as either `NULL` or `""`.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04108`: if the current container, the current user, or the currently enabled roles are different from those in effect at the time of context creation.
- `ORA-04153`: if the context handle is invalid.
- `ORA-04161`: if the source code being evaluated throws an exception.

127.3.12 EXPORT_TO_MLE Procedure

This procedure allows you to assign the given value, with appropriate conversion, to the named property in the MLE context. The property is created if it is not already present. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      BINARY_INTEGER);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      BINARY_DOUBLE);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      BINARY_FLOAT);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      BLOB);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      BOOLEAN);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      CLOB CHARACTER SET ANY_CS);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      DATE);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      DSINTERVAL_UNCONSTRAINED);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      JSON);

DBMS_MLE.EXPORT_TO_MLE (
```

```
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      NUMBER);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      TIMESTAMP_TZ_UNCONSTRAINED);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      TIMESTAMP_UNCONSTRAINED);

DBMS_MLE.EXPORT_UROWID (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      UROWID);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      VARCHAR2 CHARACTER SET ANY_CS);

DBMS_MLE.EXPORT_TO_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      YMINTERVAL_UNCONSTRAINED);

DBMS_MLE.EXPORT_CHAR (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      CHAR CHARACTER SET ANY_CS);

DBMS_MLE.EXPORT_RAW (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   property_value   IN      RAW);
```

Parameters

Table 127-12 EXPORT_TO_MLE Procedure Parameters

Parameter	Description
`context_handle`	The handle to an MLE context in the current session.
`property_name`	The name of the variable to be set. If the `property_name` value is `NULL` or an empty string, `ORA-04157` error is thrown.
`property_value`	The value to which the variable should be set.

Usage Notes

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.

- ORA-04157: if the value of the passed property_name is NULL or an empty string.
- ORA-04108: if the current container, the current user, or the currently enabled roles are different from those in effect at the time of context creation.
- ORA-04153: if the context handle is invalid.

127.3.13 GET_AVAILABLE_LANGUAGES Function

This function returns the set of available MLE languages.

Syntax

```
DBMS_MLE.GET_AVAILABLE_LANGUAGES()
  RETURN languages_t;
```

Return Values

A set of available MLE languages as a table of language identifiers as they can be used as an argument to DBMS_MLE.EVAL().

127.3.14 GET_CTX_ERROR_STACK Function

This function retrieves the JavaScript stack trace for the most recent application error in the given execution context. You need the EXECUTE DYNAMIC MLE privilege to execute this function.

Syntax

```
DBMS_MLE.GET_CTX_ERROR_STACK(
   context_handle            IN            context_handle_t)
RETURN error_frames_t;
```

Parameters

Table 127-13 GET_CTX_ERROR_STACK Function Parameters

Parameter	Description
context_handle	The handle to an MLE context in the current session.

Return Values

A collection of error stack frames, each of type error_frame_t. An empty collection is returned if there is no error stack to report.

Usage Notes

The Function may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges
- ORA-04153: if the context handle is invalid.

127.3.15 GET_ERROR_STACK Function

This function retrieves the JavaScript stack trace for the most recent application error in the given module (and optional environment) call. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.GET_ERROR_STACK(
   module_name            IN            VARCHAR2,
   env_name               IN            VARCHAR2 DEFAULT '')
RETURN error_frames_t;
```

Parameters

Table 127-14 GET_ERROR_STACK Function Parameters

Parameter	Description
`module_name`	The name of the MLE module.
`env_name`	The name of the MLE environment. Optional.

Return Values

A collection of error stack frames, each of type error_frame_t. An empty collection is returned if there is no error stack to report.

Usage Notes

The Function may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges
- `ORA-04170`: if the module name or environment name is invalid

127.3.16 IMPORT_FROM_MLE Procedure

This procedure retrieves the value of the named property from the MLE context and converts it to the requested PL/SQL type. You need the `EXECUTE DYNAMIC MLE` privilege to execute this procedure.

Syntax

```
DBMS_MLE.IMPORT_FROM_MLE (
   context_handle    IN      context_handle_t,
   property_name     IN      VARCHAR2,
   target            OUT     NOCOPY BINARY_INTEGER);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle    IN      context_handle_t,
   property_name     IN      VARCHAR2,
   target            OUT     NOCOPY BINARY_DOUBLE);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle    IN      context_handle_t,
   property_name     IN      VARCHAR2,
```

```
   target           OUT     NOCOPY BINARY_FLOAT);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY BLOB);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY BOOLEAN);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY CLOB CHARACTER SET ANY_CS);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY DATE);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY DSINTERVAL_UNCONSTRAINED);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     JSON);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY NUMBER);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY TIMESTAMP_TZ_UNCONSTRAINED);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY TIMESTAMP_UNCONSTRAINED);

DBMS_MLE.IMPORT_UROWID (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY UROWID);

DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY VARCHAR2 CHARACTER SET ANY_CS);
```

```
DBMS_MLE.IMPORT_FROM_MLE (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     NOCOPY YMINTERVAL_UNCONSTRAINED);

DBMS_MLE.IMPORT_CHAR (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     CHAR CHARACTER SET ANY_CS);

DBMS_MLE.IMPORT_RAW (
   context_handle   IN      context_handle_t,
   property_name    IN      VARCHAR2,
   target           OUT     RAW);
```

Parameters

Table 127-15 IMPORT_FROM_MLE Procedure Parameters

Parameter	Description
`context_handle`	The handle to an MLE context in the current session.
`property_name`	The name of the property to be retrieved. If the `property_name` is `NULL` or an empty string, `ORA-04157` error is thrown.
`target`	A PL/SQL variable into which the retrieved property is stored.

Usage Notes

The `IMPORT_FROM_MLE` procedure may throw the following exceptions:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04108`: if the current container, the current user, or the currently enabled roles are different from those in effect at the time of context creation.
- `ORA-04153`: if the context handle is invalid.
- `ORA-04157`: if the value of the passed property_name is `NULL` or an empty string.
- `ORA-04205`: the value cannot be converted to the target PL/SQL type.
- `ORA-06502`: the buffer of the PL/SQL variable is too small to hold the retrieved value.

127.3.17 PARSE_DEBUG_OUTPUT Function

Given a BLOB containing MLE debug output in the Java Heap Dump format, returns a textual representation of debug output.

Syntax

```
DBMS_MLE.PARSE_DEBUG_OUTPUT(
   debugoutput          IN          BLOB)
RETURN JSON;
```

Parameters

Table 127-16 PARSE_DEBUG_OUTPUT Function Parameters

Parameter	Description
`debugoutput`	MLE debug output in the Java Heap Dump format.

Return Values

The function returns a JSON representation of the debug information. The output is an array of `DebugPointData` objects.

Usage Notes

The procedure may raise the following errors:

- `ORA-04163`: if the input is not in the Java Heap Dump format.
- `ORA-04166`: if the debug output is invalid.

> **See Also:**
>
> *Oracle Database JavaScript Developer's Guide* for more information about analyzing debug output

127.3.18 SET_CTX_STDERR Procedure

This procedure redirects the `stderr` stream of the MLE context to the given `CLOB`. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_CTX_STDERR(
      context_handle  IN          context_handle_t,
      sink            IN OUT      NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 127-17 SET_CTX_STDERR Procedure Parameters

Parameter	Description
`context_handle`	Handle to an MLE context in the current session.
`sink`	The CLOB sink to redirect stderr to. Providing a `NULL` value will result in `ORA-06530` error.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to the new sink.

The procedure may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges.
- ORA-04153: if the context handle is invalid.
- ORA-06530: if the sink is NULL.

127.3.19 SET_CTX_STDERR_TO_DBMS_OUTPUT Procedure

This procedure redirects the stderr stream of the MLE context to DBMS_OUTPUT. You need the EXECUTE DYNAMIC MLE privilege to execute this function.

Syntax

```
DBMS_MLE.SET_CTX_STDERR_TO_DBMS_OUTPUT(
   context_handle      IN      context_handle_t);
```

Parameters

Table 127-18 SET_CTX_STDERR_TO_DBMS_OUTPUT Procedure Parameters

Parameter	Description
context_handle	Handle to an MLE context in the current session.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to DBMS_OUTPUT.

The procedure may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges.
- ORA-04153: if the context handle is invalid.

127.3.20 SET_CTX_STDOUT Procedure

This procedure redirects the stdout stream of the MLE context to the given CLOB. You need the EXECUTE DYNAMIC MLE privilege to execute this function.

Syntax

```
DBMS_MLE.SET_CTX_STDOUT(
   context_handle      IN context_handle_t,
   sink                IN OUT NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 127-19 SET_CTX_STDOUT Procedure Parameters

Parameter	Description
context_handle	The handle to an MLE context in the current session.
sink	The CLOB sink to redirect stdout to. Providing a NULL value will result in an ORA-06530 error.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to the new sink.

This procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04153`: if the context handle is invalid.
- `ORA-06530`: if the sink is `NULL`.

127.3.21 SET_CTX_STDOUT_TO_DBMS_OUTPUT Procedure

This procedure redirects the `stdout` stream of the MLE context to `DBMS_OUTPUT`. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_CTX_STDOUT_TO_DBMS_OUTPUT(
   context_handle      IN      context_handle_t);
```

Parameters

Table 127-20 SET_CTX_STDOUT_TO_DBMS_OUTPUT Procedure Parameters

Parameter	Description
`context_handle`	The handle to an MLE context in the current session.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to `DBMS_OUTPUT`.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04153`: if the context handle is invalid.

127.3.22 SET_ICS_STDERR Procedure

This procedure redirects the `stderr` stream of the inlined MLE call specification context to the given `CLOB` for the calling user in the current session. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_ICS_STDERR(
      name            IN          VARCHAR2,
      sink            IN OUT      NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 127-21 SET_ICS_STDERR Procedure Parameters

Parameter	Description
`name`	The name of the inlined MLE call specification.
`sink`	The `CLOB` to redirect stderr to.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to the new sink.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04043`: if the call specification does not exist.
- `ORA-06530`: if the sink is `NULL`.

127.3.23 SET_ICS_STDERR_TO_DBMS_OUTPUT Procedure

This procedure redirects the `stderr` stream of the inlined MLE call specification context to `DBMS_OUTPUT` for the calling user in the current session. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_ICS_STDERR_TO_DBMS_OUTPUT(
   name       IN       VARCHAR2);
```

Parameters

Table 127-22 SET_ICS_STDERR_TO_DBMS_OUTPUT Procedure Parameters

Parameter	Description
`name`	The name of the inlined MLE call specification.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to `DBMS_OUTPUT`.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04043`: if the call specification does not exist.

127.3.24 SET_ICS_STDOUT Procedure

This procedure redirects the `stdout` stream of an inlined MLE call specification context in the current session to the given `CLOB`.

Syntax

```
DBMS_MLE.SET_ICS_STDOUT(
    name          IN          VARCHAR2,
    sink          IN          OUT NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 127-23 SET_ICS_STDOUT Procedure Parameters

Parameter	Description
`name`	The name of the inlined MLE call specification.
`sink`	The `CLOB` sink to redirect `stdout` to. Providing a `NULL` value will result in an `ORA-06530` error.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to the new sink.

This procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges.
- `ORA-04043`: if the call specification does not exist.
- `ORA-06530`: if the sink is `NULL`.

127.3.25 SET_ICS_STDOUT_TO_DBMS_OUTPUT Procedure

This procedure redirects the `stdout` stream of the inlined call specification context to `DBMS_OUTPUT` for the calling user in the current session. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_ICS_STDOUT_TO_DBMS_OUTPUT(
   name       IN       VARCHAR2);
```

Parameters

Table 127-24 SET_ICS_STDOUT_TO_DBMS_OUTPUT Procedure Parameters

Parameter	Description
`name`	The name of the inlined call specification.

Usage Notes

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to `DBMS_OUTPUT`.

The procedure may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges.
- ORA-04043: if the call specification does not exist.

127.3.26 SET_STDERR Procedure

This procedure redirects the stderr stream of MLE contexts to the given CLOB. You need the EXECUTE DYNAMIC MLE privilege to execute this function.

Syntax

```
DBMS_MLE.SET_STDERR(
      sink             IN OUT       NOCOPY CLOB CHARACTER SET ANY_CS);

DBMS_MLE.SET_STDERR(
      module_name      IN           VARCHAR2,
      sink             IN OUT       NOCOPY CLOB CHARACTER SET ANY_CS);

DBMS_MLE.SET_STDERR(
      module_name      IN           VARCHAR2,
      env_name         IN           VARCHAR2,
      sink             IN OUT       NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 127-25 SET_STDERR Procedure Parameters

Parameter	Description
module_name	The name of the MLE module.
env_name	The name of the MLE environment.
sink	The CLOB to redirect stdout to.

Usage Notes

When called without parameters, this procedure applies to all existing contexts and contexts created in the future. Otherwise, only the context associated with the given module (or module and environment combination) is affected.

If no environment is specified, the context defined by the given module and the built-in environment is used.

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to the new sink.

The procedure may raise the following errors:

- ORA-01031: if the caller does not have sufficient privileges.
- ORA-04103: if the module does not exist.
- ORA-04105: if the environment does not exist.
- ORA-06530: if the sink is NULL.

127.3.27 SET_STDERR_TO_DBMS_OUTPUT Procedure

This procedure redirects the `stderr` stream of MLE contexts to `DBMS_OUTPUT`. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_STDERR_TO_DBMS_OUTPUT();

DBMS_MLE.SET_STDERR_TO_DBMS_OUTPUT(
   module_name          IN          VARCHAR2);

DBMS_MLE.SET_STDERR_TO_DBMS_OUTPUT(
   module_name          IN          VARCHAR2,
   env_name             IN          VARCHAR2);
```

Parameters

Table 127-26 SET_STDERR_TO_DBMS_OUTPUT Function Parameters

Parameter	Description
`module_name`	The name of the MLE module.
`env_name`	The name of the MLE environment.

Usage Notes

When called without parameters, this procedure applies to all existing contexts and contexts created in the future. Otherwise, only the context associated with the given module (or module and environment combination) is affected.

If no environment is specified, the context defined by the given module and the built-in environment is used.

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to `DBMS_OUTPUT`.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges
- `ORA-04103`: if the module does not exist
- `ORA-04105`: if the environment does not exist

127.3.28 SET_STDOUT Procedure

This procedure redirects the `stdout` stream of MLE contexts to the given `CLOB`. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_STDOUT(
   sink           IN OUT   NOCOPY CLOB CHARACTER SET ANY_CS);
```

```
DBMS_MLE.SET_STDOUT(
   module_name    IN       VARCHAR2,
   sink           IN OUT   NOCOPY CLOB CHARACTER SET ANY_CS);

DBMS_MLE.SET_STDOUT(
   module_name    IN       VARCHAR2,
   env_name       IN       VARCHAR2,
   sink           IN OUT   NOCOPY CLOB CHARACTER SET ANY_CS);
```

Parameters

Table 127-27 SET_STDOUT Procedure Parameters

Parameter	Description
`module_name`	The name of the MLE module.
`env_name`	The name of the MLE environment.
`sink`	The CLOB to redirect stdout to.

Usage Notes

When called without parameters, this procedure applies to all existing contexts and contexts created in the future. Otherwise, only the context associated with the given module (or module and environment combination) is affected.

If no environment is specified, the context defined by the given module and the built-in environment is used.

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to the new sink.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges
- `ORA-04103`: if the module does not exist
- `ORA-04105`: if the environment does not exist
- `ORA-06530`: if the sink is `NULL`.

127.3.29 SET_STDOUT_TO_DBMS_OUTPUT Procedure

This procedure redirects the `stdout` stream of MLE contexts to `DBMS_OUTPUT`. You need the `EXECUTE DYNAMIC MLE` privilege to execute this function.

Syntax

```
DBMS_MLE.SET_STDOUT_TO_DBMS_OUTPUT();

DBMS_MLE.SET_STDOUT_TO_DBMS_OUTPUT(
      module_name      IN      VARCHAR2);

DBMS_MLE.SET_STDOUT_TO_DBMS_OUTPUT(
      module_name      IN      VARCHAR2,
      env_name         IN      VARCHAR2);
```

Parameters

Table 127-28 SET_STDOUT_TO_DBMS_OUTPUT Function Parameters

Parameter	Description
`module_name`	The name of the MLE module.
`env_name`	The name of the MLE environment.

Usage Notes

When called without parameters, this procedure applies to all existing contexts and contexts created in the future. Otherwise, only the context associated with the given module (or module and environment combination) is affected.

If no environment is specified, the context defined by the given module and the built-in environment is used.

Any output that was buffered so far gets flushed to the pre-existing sink before redirecting to `DBMS_OUTPUT`.

The procedure may raise the following errors:

- `ORA-01031`: if the caller does not have sufficient privileges
- `ORA-04103`: if the module does not exist
- `ORA-04105`: if the environment does not exist

128
DBMS_MONITOR

The `DBMS_MONITOR` package enables you to use PL/SQL for controlling additional tracing and statistics gathering.

The chapter contains the following topics:

- Summary of DBMS_MONITOR Subprograms

128.1 Summary of DBMS_MONITOR Subprograms

This table lists the `DBMS_MONITOR` subprograms and briefly describes them.

Table 128-1 DBMS_MONITOR Package Subprograms

Subprogram	Description
CLIENT_ID_STAT_DISABLE Procedure	Disables statistic gathering previously enabled for a given Client Identifier
CLIENT_ID_STAT_ENABLE Procedure	Enables statistic gathering for a given Client Identifier
CLIENT_ID_TRACE_DISABLE Procedure	Disables the trace previously enabled for a given Client Identifier globally for the database
CLIENT_ID_TRACE_ENABLE Procedure	Enables the trace for a given Client Identifier globally for the database
DATABASE_TRACE_DISABLE Procedure	Disables SQL trace for the whole database or a specific instance
DATABASE_TRACE_ENABLE Procedure	Enables SQL trace for the whole database or a specific instance
SERV_MOD_ACT_STAT_DISABLE Procedure	Disables statistic gathering enabled for a given combination of Service Name, `MODULE` and `ACTION`
SERV_MOD_ACT_STAT_ENABLE Procedure	Enables statistic gathering for a given combination of Service Name, `MODULE` and `ACTION`
SERV_MOD_ACT_TRACE_DISABLE Procedure	Disables the trace for `ALL` enabled instances for a or a given combination of Service Name, `MODULE` and `ACTION` name globally
SERV_MOD_ACT_TRACE_ENABLE Procedure	Enables SQL tracing for a given combination of Service Name, `MODULE` and `ACTION` globally unless an `instance_name` is specified
SESSION_TRACE_DISABLE Procedure	Disables the previously enabled trace for a given database session identifier (SID) on the local instance
SESSION_TRACE_ENABLE Procedure	Enables the trace for a given database session identifier (SID) on the local instance

128.1.1 CLIENT_ID_STAT_DISABLE Procedure

This procedure will disable statistics accumulation for all instances and remove the accumulated results from `V$CLIENT_STATS` view enabled by the `CLIENT_ID_STAT_ENABLE` Procedure.

Syntax

```
DBMS_MONITOR.CLIENT_ID_STAT_DISABLE(
   client_id            IN   VARCHAR2);
```

Parameters

Table 128-2 CLIENT_ID_STAT_DISABLE Procedure Parameters

Parameter	Description
client_id	Client Identifier for which statistic aggregation is disabled

Examples

To disable accumulation:

```
EXECUTE DBMS_MONITOR.CLIENT_ID_STAT_DISABLE('janedoe');
```

128.1.2 CLIENT_ID_STAT_ENABLE Procedure

This procedure enables statistic gathering for a given Client Identifier.

Statistics gathering is global for the database and persistent across instance starts and restarts. That is, statistics are enabled for all instances of the same database, including restarts. Statistics are viewable through `V$CLIENT_STATS` views.

Syntax

```
DBMS_MONITOR.CLIENT_ID_STAT_ENABLE(
   client_id            IN   VARCHAR2);
```

Parameters

Table 128-3 *CLIENT_ID_STAT_ENABLE Procedure Parameters*

Parameter	Description
client_id	Client Identifier for which statistic aggregation is enabled

Examples

To enable statistic accumulation for a client with a given client ID:

```
EXECUTE DBMS_MONITOR.CLIENT_ID_STAT_ENABLE('janedoe');
```

128.1.3 CLIENT_ID_TRACE_DISABLE Procedure

This procedure will disable tracing enabled by the `CLIENT_ID_TRACE_ENABLE` Procedure.

Syntax

```
DBMS_MONITOR.CLIENT_ID_TRACE_DISABLE(
 client_id     IN  VARCHAR2);
```

Parameters

Table 128-4 CLIENT_ID_TRACE_DISABLE Procedure Parameters

Parameter	Description
`client_id`	Client Identifier for which SQL tracing is disabled

Examples

```
EXECUTE DBMS_MONITOR.CLIENT_ID_TRACE_DISABLE ('janedoe');
```

128.1.4 CLIENT_ID_TRACE_ENABLE Procedure

This procedure will enable the trace for a given client identifier globally for the database.

Syntax

```
DBMS_MONITOR.CLIENT_ID_TRACE_ENABLE(
 client_id     IN  VARCHAR2,
 waits         IN  BOOLEAN DEFAULT TRUE,
 binds         IN  BOOLEAN DEFAULT FALSE,
 plan_stat     IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 128-5 CLIENT_ID_TRACE_ENABLE Procedure Parameters

Parameter	Description
`client_id`	Database Session Identifier for which SQL tracing is enabled
`waits`	If `TRUE`, wait information is present in the trace
`binds`	If `TRUE`, bind information is present in the trace
`plan_stat`	Frequency at which we dump row source statistics. Value should be 'NEVER', 'FIRST_EXECUTION' (equivalent to `NULL`) or 'ALL_EXECUTIONS'.

Usage Notes

- The trace will be written to multiple trace files because more than one Oracle shadow process can work on behalf of a given client identifier.
- The tracing is enabled for all instances and persistent across restarts.

Examples

```
EXECUTE DBMS_MONITOR.CLIENT_ID_TRACE_ENABLE('janedoe', TRUE,
FALSE);
```

128.1.5 DATABASE_TRACE_DISABLE Procedure

This procedure disables SQL trace for the whole database or a specific instance.

Syntax

```
DBMS_MONITOR.DATABASE_TRACE_DISABLE(
   instance_name  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 128-6 DATABASE_TRACE_DISABLE Procedure Parameters

Parameter	Description
instance_name	Disables tracing for the named instance

128.1.6 DATABASE_TRACE_ENABLE Procedure

This procedure enables SQL trace for the whole database or a specific instance.

Syntax

```
DBMS_MONITOR.DATABASE_TRACE_ENABLE(
   waits          IN BOOLEAN DEFAULT TRUE,
   binds          IN BOOLEAN DEFAULT FALSE,
   instance_name  IN VARCHAR2 DEFAULT NULL,
   plan_stat      IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 128-7 DATABASE_TRACE_ENABLE Procedure Parameters

Parameter	Description
waits	If TRUE, wait information will be present in the trace
binds	If TRUE, bind information will be present in the trace
instance_name	If set, restricts tracing to the named instance
plan_stat	Frequency at which we dump row source statistics. Value should be 'NEVER', 'FIRST_EXECUTION' (equivalent to NULL) or 'ALL_EXECUTIONS'.

128.1.7 SERV_MOD_ACT_STAT_DISABLE Procedure

This procedure will disable statistics accumulation and remove the accumulated results from `V$SERV_MOD_ACT_STATS` view.

Statistics disabling is persistent for the database. That is, service statistics are disabled for instances of the same database (plus `dblinks` that have been activated as a result of the enable).

Syntax

```
DBMS_MONITOR.SERV_MOD_ACT_STAT_DISABLE(
   service_name     IN VARCHAR2,
   module_name      IN VARCHAR2,
   action_name      IN VARCHAR2 DEFAULT ALL_ACTIONS);
```

Parameters

Table 128-8 SERV_MOD_ACT_STAT_DISABLE Procedure Parameters

Parameter	Description
`service_name`	Name of the service for which statistic aggregation is disabled
`module_name`	Name of the `MODULE`. An additional qualifier for the service. It is a required parameter.
`action_name`	Name of the `ACTION`. An additional qualifier for the Service and `MODULE` name. Omitting the parameter (or supplying `ALL_ACTIONS` constant) means enabling aggregation for all Actions for a given Service/ MODULE combination. In this case, statistics are aggregated on the module level.

Usage Notes

Regarding statistics gathering, when you change the module or action, the change takes effect when the next user call is executed in the session. For example, if a module is set to 'module 1' in a session, and the module is reset to 'module 2' in a user call in the session, then the module remains 'module 1' during this user call. The module is changed to 'module 2' in the next user call in the session.

128.1.8 SERV_MOD_ACT_STAT_ENABLE Procedure

This procedure enables statistic gathering for a given combination of Service Name, `MODULE` and `ACTION`.

Calling this procedure enables statistic gathering for a hierarchical combination of Service name, `MODULE` name, and `ACTION` name on all instances for the same database. Statistics are accessible by means of the `V$SERV_MOD_ACT_STATS` view.

Syntax

```
DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE(
   service_name     IN VARCHAR2,
   module_name      IN VARCHAR2,
   action_name      IN VARCHAR2 DEFAULT ALL_ACTIONS);
```

Parameters

Table 128-9 SERV_MOD_ACT_STAT_ENABLE Procedure Parameters

Parameter	Description
service_name	Name of the service for which statistic aggregation is enabled
module_name	Name of the MODULE. An additional qualifier for the service. It is a required parameter.
action_name	Name of the ACTION. An additional qualifier for the Service and MODULE name. Omitting the parameter (or supplying ALL_ACTIONS constant) means enabling aggregation for all Actions for a given Service/ MODULE combination. In this case, statistics are aggregated on the module level.

Usage Notes

Enabling statistic aggregation for the given combination of Service/Module/Action names is slightly complicated by the fact that the Module/Action values can be empty strings which are indistinguishable from NULLs. For this reason, we adopt the following conventions:

A special constant (unlikely to be a real action names) is defined:

```
ALL_ACTIONS constant VARCHAR2 := '###ALL_ACTIONS';
```

Using ALL_ACTIONS for an action specification means that aggregation is enabled for all actions with a given module name, while using NULL (or empty string) means that aggregation is enabled for an action whose name is an empty string.

Regarding statistics gathering, when you change the module or action, the change takes effect when the next user call is executed in the session. For example, if a module is set to 'module 1' in a session, and the module is reset to 'module 2' in a user call in the session, then the module remains 'module 1' during this user call. The module is changed to 'module 2' in the next user call in the session.

Examples

To enable statistic accumulation for a given combination of Service name and MODULE:

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE( 'APPS1','PAYROLL');
```

To enable statistic accumulation for a given combination of Service name, MODULE and ACTION:

```
EXECUTE
DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE('APPS1','GLEDGER','DEBIT_ENTRY');
```

If both of the preceding commands are issued, statistics are accumulated as follows:

- For the APPS1 service, because accumulation for each Service Name is the default.
- For all actions in the PAYROLL Module.
- For the DEBIT_ENTRY Action within the GLEDGER Module.

128.1.9 SERV_MOD_ACT_TRACE_DISABLE Procedure

This procedure will disable the trace at `ALL` enabled instances for a given combination of Service Name, `MODULE`, and `ACTION` name globally.

Syntax

```
DBMS_MONITOR.SERV_MOD_ACT_TRACE_DISABLE(
   service_name     IN  VARCHAR2,
   module_name      IN  VARCHAR2,
   action_name      IN  VARCHAR2 DEFAULT ALL_ACTIONS,
   instance_name    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 128-10 SERV_MOD_ACT_TRACE_DISABLE Procedure Parameters

Parameter	Description
`service_name`	Name of the service for which tracing is disabled.
`module_name`	Name of the `MODULE`. An additional qualifier for the service
`action_name`	Name of the `ACTION`. An additional qualifier for the Service and `MODULE` name.
`instance_name`	If set, this restricts tracing to the named `instance_name`

Usage Notes

Specifying `NULL` for the `module_name` parameter means that statistics will no longer be accumulated for the sessions which do not set the `MODULE` attribute.

Examples

To enable tracing for a Service named APPS1:

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE('APPS1',
            DBMS_MONITOR.ALL_MODULES,  DBMS_MONITOR.ALL_ACTIONS,TRUE,
FALSE,NULL);
```

To disable tracing specified in the previous step:

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_DISABLE('APPS1');
```

To enable tracing for a given combination of `Service` and `MODULE` (all `ACTION`s):

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE('APPS1','PAYROLL',
            DBMS_MONITOR.ALL_ACTIONS,TRUE,FALSE,NULL);
```

To disable tracing specified in the previous step:

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_DISABLE('APPS1','PAYROLL');
```

128.1.10 SERV_MOD_ACT_TRACE_ENABLE Procedure

This procedure will enable SQL tracing for a given combination of Service Name, `MODULE` and `ACTION` globally unless an `instance_name` is specified.

Syntax

```
DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(
   service_name    IN VARCHAR2,
   module_name     IN VARCHAR2 DEFAULT ANY_MODULE,
   action_name     IN VARCHAR2 DEFAULT ANY_ACTION,
   waits           IN BOOLEAN DEFAULT TRUE,
   binds           IN BOOLEAN DEFAULT FALSE,
   instance_name   IN VARCHAR2 DEFAULT NULL,
   plan_stat       IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 128-11 SERV_MOD_ACT_TRACE_ENABLE Procedure Parameters

Parameter	Description
`service_name`	Name of the service for which SQL trace is enabled
`module_name`	Name of the `MODULE` for which SQL trace is enabled. An optional additional qualifier for the service. If omitted, SQL trace is enabled or all modules and actions in a given service.
`action_name`	Name of the `ACTION` for which SQL trace is enabled. An optional additional qualifier for the Service and `MODULE` name. If omitted, SQL trace is enabled for all actions in a given module.
`waits`	If `TRUE`, wait information is present in the trace
`binds`	If `TRUE`, bind information is present in the trace
`instance_name`	If set, this restricts tracing to the named `instance_name`
`plan_stat`	Frequency at which we dump row source statistics. Value should be 'NEVER', 'FIRST_EXECUTION' (equivalent to `NULL`) or 'ALL_EXECUTIONS'.

Usage Notes

- The procedure enables a trace for a given combination of Service, `MODULE` and `ACTION` name. The specification is strictly hierarchical: Service Name or Service Name/`MODULE`, or Service Name, MODULE, and `ACTION` name must be specified. Omitting a qualifier behaves like a wild-card, so that not specifying an `ACTION` means all `ACTION`s. Using the `ALL_ACTIONS` constant achieves the same purpose.
- This tracing is useful when an application `MODULE` and optionally known `ACTION` is experiencing poor service levels.
- By default, tracing is enabled globally for the database. The `instance_name` parameter is provided to restrict tracing to named instances that are known, for example, to exhibit poor service levels.
- Tracing information is present in multiple trace files and you must use the `trcsess` tool to collect it into a single file.

- Specifying NULL for the module_name parameter means that statistics will be accumulated for the sessions which do not set the MODULE attribute.

Examples

To enable tracing for a Service named APPS1:

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE('APPS1',
            DBMS_MONITOR.ALL_MODULES,  DBMS_MONITOR.ALL_ACTIONS,TRUE,
FALSE,NULL);
```

To enable tracing for a given combination of Service and MODULE (all ACTIONs):

```
EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE('APPS1','PAYROLL',
            DBMS_MONITOR.ALL_ACTIONS,TRUE,FALSE,NULL);
```

128.1.11 SESSION_TRACE_DISABLE Procedure

This procedure will disable the trace for a given database session at the local instance.

Syntax

```
DBMS_MONITOR.SESSION_TRACE_DISABLE(
   session_id      IN     BINARY_INTEGER DEFAULT NULL,
   serial_num      IN     BINARY_INTEGER DEFAULT NULL);
```

Parameters

Table 128-12 SESSION_TRACE_DISABLE Procedure Parameters

Parameter	Description
session_id	Database Session Identifier for which SQL trace is disabled
serial_num	Serial number for this session

Usage Notes

If serial_num is NULL but session_id is specified, a session with a given session_id is no longer traced irrespective of its serial number. If both session_id and serial_num are NULL, the current user session is no longer traced. It is illegal to specify NULL session_id and non-NULL serial_num. In addition, the NULL values are default and can be omitted.

Examples

To enable tracing for a client with a given client session ID:

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(7,4634, TRUE, FALSE);
```

To disable tracing specified in the previous step:

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_DISABLE(7,4634);;
```

128.1.12 SESSION_TRACE_ENABLE Procedure

This procedure enables a SQL trace for the given Session ID on the local instance

Syntax

```
DBMS_MONITOR.SESSION_TRACE_ENABLE(
    session_id   IN  BINARY_INTEGER DEFAULT NULL,
    serial_num   IN  BINARY_INTEGER DEFAULT NULL,
    waits        IN  BOOLEAN DEFAULT TRUE,
    binds        IN  BOOLEAN DEFAULT FALSE,
    plan_stat    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 128-13 SESSION_TRACE_ENABLE Procedure Parameters

Parameter	Description
session_id	Client Identifier for which SQL trace is enabled. If omitted (or NULL), the user's own session is assumed.
serial_num	Serial number for this session. If omitted (or NULL), only the session ID is used to determine a session.
waits	If TRUE, wait information is present in the trace
binds	If TRUE, bind information is present in the trace
plan_stat	Frequency at which we dump row source statistics. Value should be 'NEVER', 'FIRST_EXECUTION' (equivalent to NULL) or 'ALL_EXECUTIONS'.

Usage Notes

The procedure enables a trace for a given database session, and is still useful for client/server applications. The trace is enabled only on the instance to which the caller is connected, since database sessions do not span instances. This tracing is strictly local to an instance.

If serial_num is NULL but session_id is specified, a session with a given session_id is traced irrespective of its serial number. If both session_id and serial_num are NULL, the current user session is traced. It is illegal to specify NULL session_id and non-NULL serial_num. In addition, the NULL values are default and can be omitted.

Examples

To enable tracing for a client with a given client session ID:

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(7,4634, TRUE, FALSE);
```

To disable tracing specified in the previous step:

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_DISABLE(7,4634);
```

Either

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(5);
```

or

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(5, NULL);
```

traces the session with session ID of 5, while either

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE();
```

or

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(NULL, NULL);
```

traces the current user session. Also,

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(NULL, NULL, TRUE, TRUE);
```

traces the current user session including waits and binds. The same can be also expressed using keyword syntax:

```
EXECUTE DBMS_MONITOR.SESSION_TRACE_ENABLE(binds=>TRUE);
```

129
DBMS_MVIEW

DBMS_MVIEW enables you to understand capabilities for materialized views and potential materialized views, including their rewrite availability. It also enables you to refresh materialized views that are not part of the same refresh group and purge logs.

> **Note:**
>
> DBMS_MVIEW is a synonym for DBMS_SNAPSHOT.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information about using materialized views in a data warehousing environment

This chapter contains the following topics:

- Operational Notes
- Security Model
- Rules and Limits
- Summary of DBMS_MVIEW Subprograms

129.1 DBMS_MVIEW Operational Notes

If a query is less than 256 characters long, you can invoke EXPLAIN_REWRITE using the EXECUTE command from SQL*Plus. Otherwise, the recommended method is to use a PL/SQL BEGIN..END block, as shown in the examples in /rdbms/demo/smxrw.sql.

129.2 DBMS_MVIEW Security Model

The DBMS_MVIEW package consists of a number of materialized view-related subprograms, each of which has different functionality and privilege requirements.

The privilege model is generally based on the invoker's right. Each package subprogram is executed by first checking the privileges against the invoker. If all the required privileges are met, the subprogram will be executed. Otherwise, an insufficient privileges error will be thrown.

129.3 DBMS_MVIEW Rules and Limits

The DBMS_MVIEW.EXPLAIN_REWRITE procedure cannot accept queries longer than 32627 characters. These restrictions also apply when passing the defining query of a materialized view to the DBMS_MVIEW.EXPLAIN_MVIEW procedure.

129.4 Summary of DBMS_MVIEW Subprograms

This table lists the `DBMS_MVIEW` subprograms and briefly describes them.

Table 129-1 DBMS_MVIEW Package Subprograms

Subprogram	Description
BEGIN_TABLE_REORGANIZATION Procedure	Performs a process to preserve materialized view data needed for refresh
END_TABLE_REORGANIZATION Procedure	Ensures that the materialized view data for the master table is valid and that the master table is in the proper state
ESTIMATE_MVIEW_SIZE Procedure	Estimates the size of a materialized view that you might create, in bytes and rows
EXPLAIN_MVIEW Procedure	Explains what is possible with a materialized view or potential materialized view
EXPLAIN_REWRITE Procedure	Explains why a query failed to rewrite or why the optimizer chose to rewrite a query with a particular materialized view or materialized views
I_AM_A_REFRESH Function	Returns the value of the `I_AM_REFRESH` package state
PMARKER Function	Returns a partition marker from a rowid, and is used for Partition Change Tracking (PCT)
PURGE_DIRECT_LOAD_LOG Procedure	Purges rows from the direct loader log after they are no longer needed by any materialized views (used with data warehousing)
PURGE_LOG Procedure	Purges rows from the materialized view log
PURGE_MVIEW_FROM_LOG Procedure	Purges rows from the materialized view log
REFRESH Procedures	Refreshes one or more materialized views that are not members of the same refresh group
REFRESH_ALL_MVIEWS Procedure	Refreshes all materialized views that do not reflect changes to their master table or master materialized view
REFRESH_DEPENDENT Procedures	Refreshes all table-based materialized views that depend on a specified master table or master materialized view, or list of master tables or master materialized views
REGISTER_MVIEW Procedure	Enables the administration of individual materialized views
UNREGISTER_MVIEW Procedure	Enables the administration of individual materialized views once invoked at a master site or master materialized view site to unregister a materialized view

129.4.1 BEGIN_TABLE_REORGANIZATION Procedure

This procedure performs a process to preserve materialized view data needed for refresh. It must be called before a master table is reorganized.

Syntax

```
DBMS_MVIEW.BEGIN_TABLE_REORGANIZATION (
   tabowner     IN   VARCHAR2,
   tabname      IN   VARCHAR2);
```

Parameters

Table 129-2 BEGIN_TABLE_REORGANIZATION Procedure Parameters

Parameter	Description
tabowner	Owner of the table being reorganized
tabname	Name of the table being reorganized

129.4.2 END_TABLE_REORGANIZATION Procedure

This procedure ensures that the materialized view data for the master table is valid and that the master table is in the proper state. It must be called after a master table is reorganized.

Syntax

```
DBMS_MVIEW.END_TABLE_REORGANIZATION (
   tabowner    IN   VARCHAR2,
   tabname     IN   VARCHAR2);
```

Parameters

Table 129-3 END_TABLE_REORGANIZATION Procedure Parameters

Parameter	Description
tabowner	Owner of the table being reorganized
tabname	Name of the table being reorganized

129.4.3 ESTIMATE_MVIEW_SIZE Procedure

This procedure estimates the size of a materialized view that you might create, in bytes and number of rows.

Syntax

```
DBMS_MVIEW.ESTIMATE_MVIEW_SIZE (
   stmt_id       IN  VARCHAR2,
   select_clause IN  VARCHAR2,
   num_rows      OUT NUMBER,
   num_bytes     OUT NUMBER);
```

Parameters

Table 129-4 ESTIMATE_MVIEW_SIZE Procedure Parameters

Parameter	Description
stmt_id	Arbitrary string used to identify the statement in an EXPLAIN PLAN
select_clause	The SELECT statement to be analyzed
num_rows	Estimated cardinality
num_bytes	Estimated number of bytes

129.4.4 EXPLAIN_MVIEW Procedure

This procedure enables you to learn what is possible with a materialized view or potential materialized view. For example, you can determine if a materialized view is fast refreshable and what types of query rewrite you can perform with a particular materialized view.

Using this procedure is straightforward. You simply call `DBMS_MVIEW.EXPLAIN_MVIEW`, passing in as parameters the schema and materialized view name for an existing materialized view. Alternatively, you can specify the `SELECT` string or `CREATE MATERIALIZED VIEW` statement for a potential materialized view. The materialized view or potential materialized view is then analyzed and the results are written into either a table called `MV_CAPABILITIES_TABLE`, which is the default, or to an array called `MSG_ARRAY`.

The procedure is overloaded:

- The first version is for explaining an existing or potential materialized view with output to `MV_CAPABILITIES_TABLE`.
- The second version is for explaining an existing or potential materialized view with output to a `VARRAY`.

Syntax

```
DBMS_MVIEW.EXPLAIN_MVIEW (
   mv            IN VARCHAR2,
   statement_id  IN VARCHAR2:= NULL);

DBMS_MVIEW.EXPLAIN_MVIEW (
   mv          IN VARCHAR2,
   msg_array   OUT SYS.ExplainMVArrayType);
```

Parameters

Table 129-5 EXPLAIN_MVIEW Procedure Parameters

Parameter	Description
`mv`	The name of an existing materialized view (optionally qualified with the owner name separated by a ".") or a `SELECT` statement or a `CREATE MATERIALIZED VIEW` statement for a potential materialized view.
`statement_id`	A client-supplied unique identifier to associate output rows with specific invocations of `EXPLAIN_MVIEW`
`msg_array`	The PL/SQL `VARRAY` that receives the output. Use this parameter to direct `EXPLAIN_MVIEW`'s output to a PL/SQL VARRAY rather than `MV_CAPABILITIES_TABLE`.

Usage Notes

- You must run the `utlxmv.sql` script to create `MV_CAPABILITIES_TABLE` in the current schema prior to calling `EXPLAIN_MVIEW` except when you direct output to a `VARRAY`. The script is found in the `ADMIN` directory.

- In Oracle database version 19.18, the following capability is added and is displayed while explaining a materialized view using `dbms_mview.explain_mview`. It indicates that the materialized view is eligible for LPCT fast refresh:

  ```
  REFRESH_FAST_LPT
  ```

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information about refresh methods for materialized view .

129.4.5 EXPLAIN_REWRITE Procedure

This procedure enables you to learn why a query failed to rewrite, or, if it rewrites, which materialized views will be used.

Using the results from the procedure, you can take the appropriate action needed to make a query rewrite if at all possible. The query specified in the `EXPLAIN_REWRITE` statement is never actually executed.

A demo file, `xrwutl.sql`, is available to help format the output from `EXPLAIN_REWRITE`.

Syntax

You can obtain the output from `DBMS_MVIEW.EXPLAIN_REWRITE` in two ways. The first is to use a table, while the second is to create a `VARRAY`. The following shows the basic syntax for using an output table:

```
DBMS_MVIEW.EXPLAIN_REWRITE (
    query           VARCHAR2,
    mv              VARCHAR2(30),
    statement_id    VARCHAR2(30));
```

You can create an output table called `REWRITE_TABLE` by executing the `utlxrw.sql` script.

The `query` parameter is a text string representing the SQL query. The parameter, `mv`, is a fully qualified materialized view name in the form of `schema.mv`. This is an optional parameter. When it is not specified, `EXPLAIN_REWRITE` returns any relevant messages regarding all the materialized views considered for rewriting the given query. When `schema` is omitted and only `mv` is specified, `EXPLAIN_REWRITE` looks for the materialized view in the current schema.

If you want to direct the output of `EXPLAIN_REWRITE` to a `VARRAY` instead of a table, you should call the procedure as follows:

```
DBMS_MVIEW.EXPLAIN_REWRITE (
    query           [VARCHAR2 | CLOB],
    mv               VARCHAR2(30),
    output_array     SYS.RewriteArrayType);
```

Note that if the query is less than 256 characters long, `EXPLAIN_REWRITE` can be easily invoked with the `EXECUTE` command from SQL*Plus. Otherwise, the recommended method is to use a PL/SQL `BEGIN... END` block, as shown in the examples in `/rdbms/demo/smxrw*`.

You can also use `EXPLAIN_REWRITE` with multiple materialized views, in which case the syntax will be the same as with a single materialized view, except that the materialized views are

specified by a comma-delimited string. For example, to find out whether a given set of materialized views `mv1`, `mv2`, and `mv3` could be used to rewrite the query, `query_txt`, and, if not, why not, use `EXPLAIN_REWRITE` as follows:

```
DBMS_MVIEW.EXPLAIN_REWRITE(query_txt, 'mv1, mv2, mv3')
```

See *Oracle Database Data Warehousing Guide* for more information on using the `EXPLAIN_REWRITE` procedure.

Parameters

Table 129-6 EXPLAIN_REWRITE Procedure Parameters

Parameter	Description
`query`	SQL `SELECT` statement to be explained
`mv`	The fully qualified name of an existing materialized view in the form of `SCHEMA.MV`. For multiple materialized views, you can provide a comma-delimited list of names.
`statement_id`	A client-supplied unique identifier to distinguish output messages
`msg_array`	The PL/SQL `VARRAY` that receives the output. Use this parameter to direct `EXPLAIN_REWRITE`'s output to a PL/SQL `VARRAY`.

Usage Notes

To obtain the output into a table, you must run the `utlxrw.sql` script before calling `EXPLAIN_REWRITE`. This script creates a table named `REWRITE_TABLE` in the current schema.

129.4.6 I_AM_A_REFRESH Function

This function returns the value of the `I_AM_REFRESH` package state.

Syntax

```
DBMS_MVIEW.I_AM_A_REFRESH
   RETURN BOOLEAN;
```

Return Values

A return value of `true` indicates that all local replication triggers for materialized views are effectively disabled in this session because each replication trigger first checks this state. A return value of `false` indicates that these triggers are enabled.

129.4.7 PMARKER Function

This function returns a partition marker from a rowid. It is used for Partition Change Tracking (PCT).

Syntax

```
DBMS_MVIEW.PMARKER(
   rid IN ROWID)
 RETURN NUMBER;
```

Parameters

Table 129-7 PMARKER Function Parameters

Parameter	Description
rid	The rowid of a row entry in a master table

129.4.8 PURGE_DIRECT_LOAD_LOG Procedure

This procedure removes entries from the direct loader log after they are no longer needed for any known materialized view. This procedure usually is used in environments using Oracle's data warehousing technology.

Syntax

```
DBMS_MVIEW.PURGE_DIRECT_LOAD_LOG();
```

129.4.9 PURGE_LOG Procedure

This procedure purges rows from the materialized view log.

Syntax

```
DBMS_MVIEW.PURGE_LOG (
   master        IN   VARCHAR2,
   num           IN   BINARY_INTEGER := 1,
   flag          IN   VARCHAR2       := 'NOP');
```

Parameters

Table 129-8 PURGE_LOG Procedure Parameters

Parameter	Description
master	Name of the master table or master materialized view.

num	Number of least recently refreshed materialized views whose rows you want to remove from materialized view log. For example, the following statement deletes rows needed to refresh the two least recently refreshed materialized views:
	`DBMS_MVIEW.PURGE_LOG('master_table', 2);`
	To delete all rows in the materialized view log, indicate a high number of materialized views to disregard, as in this example:
	`DBMS_MVIEW.PURGE_LOG('master_table',9999);`
	This statement completely purges the materialized view log that corresponds to `master_table` if fewer than 9999 materialized views are based on `master_table`. A simple materialized view whose rows have been purged from the materialized view log must be completely refreshed the next time it is refreshed.

flag	Specify `delete` to guarantee that rows are deleted from the materialized view log for at least one materialized view. This parameter can override the setting for the parameter `num`. For example, the following statement deletes rows from the materialized view log that has dependency rows in the least recently refreshed materialized view: `DBMS_MVIEW.PURGE_LOG('master_table',1,'delete');`

129.4.10 PURGE_MVIEW_FROM_LOG Procedure

This procedure is called on the master site or master materialized view site to delete the rows in materialized view refresh related data dictionary tables maintained at the master for the specified materialized view identified by `mview_id` or the combination of `mviewowner`, `mviewname`, and `mviewsite`.

If the materialized view specified is the oldest materialized view to have refreshed from any of the master tables or master materialized views, then the materialized view log is also purged. This procedure does not unregister the materialized view.

Syntax

```
DBMS_MVIEW.PURGE_MVIEW_FROM_LOG (
   mview_id       IN   BINARY_INTEGER);

DBMS_MVIEW.PURGE_MVIEW_FROM_LOG (
   mviewowner     IN   VARCHAR2,
   mviewname      IN   VARCHAR2,
   mviewsite      IN   VARCHAR2);
```

> **Note:**
>
> This procedure is overloaded. The parameter `mview_id` is mutually exclusive with the three remaining parameters: `mviewowner`, `mviewname`, and `mviewsite`.

Parameters

Table 129-9 PURGE_MVIEW_FROM_LOG Procedure Parameters

Parameter	Description
`mview_id`	If you want to execute this procedure based on the identification of the target materialized view, specify the materialized view identification using the `mview_id` parameter. Query the `DBA_BASE_TABLE_MVIEWS` view at the materialized view log site for a listing of materialized view IDs. Executing this procedure based on the materialized view identification is useful if the target materialized view is not listed in the list of registered materialized views (`DBA_REGISTERED_MVIEWS`).
`mviewowner`	If you do not specify an `mview_id`, enter the owner of the target materialized view using the `mviewowner` parameter. Query the `DBA_REGISTERED_MVIEWS` view at the materialized view log site to view the materialized view owners.
`mviewname`	If you do not specify an `mview_id`, enter the name of the target materialized view using the `mviewname` parameter. Query the `DBA_REGISTERED_MVIEWS` view at the materialized view log site to view the materialized view names.
`mviewsite`	If you do not specify an `mview_id`, enter the site of the target materialized view using the `mviewsite` parameter. Query the `DBA_REGISTERED_MVIEWS` view at the materialized view log site to view the materialized view sites.

Usage Notes

If there is an error while purging one of the materialized view logs, the successful purge operations of the previous materialized view logs are not rolled back. This is to minimize the

size of the materialized view logs. In case of an error, this procedure can be invoked again until all the materialized view logs are purged.

129.4.11 REFRESH Procedures

This procedure refreshes a list of materialized views.

Syntax

```
DBMS_MVIEW.REFRESH (
   { list                 IN     VARCHAR2,
   | tab                  IN     DBMS_UTILITY.UNCL_ARRAY,}
   method                 IN     VARCHAR2       := NULL,
   rollback_seg           IN     VARCHAR2       := NULL,
   push_deferred_rpc      IN     BOOLEAN        := true,
   refresh_after_errors   IN     BOOLEAN        := false,
   purge_option           IN     BINARY_INTEGER := 1,
   parallelism            IN     BINARY_INTEGER := 0,
   heap_size              IN     BINARY_INTEGER := 0,
   atomic_refresh         IN     BOOLEAN        := true,
   nested                 IN     BOOLEAN        := false,
   out_of_place           IN     BOOLEAN        := false,
   skip_ext_data             IN     BOOLEAN        := false);
```

> **Note:**
>
> This procedure is overloaded. The `list` and `tab` parameters are mutually exclusive.

Parameters

Table 129-10 REFRESH Procedure Parameters

Parameter	Description
`list | tab`	Comma-delimited list of materialized views that you want to refresh. (Synonyms are not supported.) These materialized views can be located in different schemas and have different master tables or master materialized views. However, all of the listed materialized views must be in your local database. Alternatively, you may pass in a PL/SQL index-by table of type `DBMS_UTILITY.UNCL_ARRAY`, where each element is the name of a materialized view.

Table 129-10 (Cont.) REFRESH Procedure Parameters

Parameter	Description
method	A string of refresh methods indicating how to refresh the listed materialized views. An `f` indicates fast refresh, `?` indicates force refresh, `C` or `c` indicates complete refresh, and `A` or `a` indicates always refresh. `A` and `C` are equivalent. `P` or `p` refreshes by recomputing the rows in the materialized view affected by changed partitions in the detail tables. `L` or `l` indicates LPCT refresh, that is, recomputing the rows in the materialized view as per the logical partition of the base tables. **Note:** "L" and "?" refreshes attempt a combined PCT and LCPT refresh. (physical and logical partitions). If a materialized view does not have a corresponding refresh method (that is, if more materialized views are specified than refresh methods), then that materialized view is refreshed according to its default refresh method. For example, consider the following `EXECUTE` statement within SQL*Plus: `DBMS_MVIEW.REFRESH ('countries_mv,regions_mv,hr.employees_mv','cf');` This statement performs a complete refresh of the `countries_mv` materialized view, a fast refresh of the `regions_mv` materialized view, and a default refresh of the `hr.employees` materialized view.
rollback_seg	Name of the materialized view site rollback segment to use while refreshing materialized views
push_deferred_rpc	Used by updatable materialized views only. Set this parameter to `true` if you want to push changes from the materialized view to its associated master tables or master materialized views before refreshing the materialized view. Otherwise, these changes may appear to be temporarily lost.
refresh_after_errors	If this parameter is `true`, an updatable materialized view continues to refresh even if there are outstanding conflicts logged in the `DEFERROR` view for the materialized view's master table or master materialized view. If this parameter is `true` and `atomic_refresh` is `false`, this procedure continues to refresh other materialized views if it fails while refreshing a materialized view.
purge_option	If you are using the parallel propagation mechanism (in other words, parallelism is set to 1 or greater), 0 means do not purge, 1 means lazy purge, and 2 means aggressive purge. In most cases, lazy purge is the optimal setting. Set purge to aggressive to trim the queue if multiple master replication groups are pushed to different target sites, and updates to one or more replication groups are infrequent and infrequently pushed. If all replication groups are infrequently updated and pushed, then set this parameter to `0` and occasionally execute `PUSH` with this parameter set to `2` to reduce the queue.
parallelism	0 specifies serial propagation. *n* > 1 specifies parallel propagation with *n* parallel processes. 1 specifies parallel propagation using only one parallel process.

Table 129-10 (Cont.) REFRESH Procedure Parameters

Parameter	Description
heap_size	Maximum number of transactions to be examined simultaneously for parallel propagation scheduling. Oracle automatically calculates the default setting for optimal performance. Note: Do not set this parameter unless directed to do so by Oracle Support Services.
atomic_refresh	If this parameter is set to true, then the list of materialized views is refreshed in a single transaction. All of the refreshed materialized views are updated to a single point in time. If the refresh fails for any of the materialized views, none of the materialized views are updated. If this parameter is set to false, then each of the materialized views is refreshed non-atomically in separate transactions. As part of complete refresh, if truncate is used (non-atomic refresh), unique index rebuild is executed. INDEX REBUILD automatically computes statistics. Thus, statistics are updated for truncated tables.
nested	If true, then perform nested refresh operations for the specified set of materialized views. Nested refresh operations refresh all the depending materialized views and the specified set of materialized views based on a dependency order to ensure the nested materialized views are truly fresh with respect to the underlying base tables.
out_of_place	If true, then it performs an out-of-place refresh. The default is false. This parameter uses the four methods of refresh (F, P, C, ?). So, for example, if you specify F and out_of_place = true, then an out-of-place fast refresh will be attempted. Similarly, if you specify P and out_of_place = true, then out-of-place PCT refresh will be attempted.
skip_ext_data	Provides you an option to skip the MV data refresh corresponding to the external partitions.

Usage Notes

- The preference of LPCT refresh on the materialized view is:
  1. If the materialized view is both PCT and LPCT enabled, a combined LPCT+PCT refresh is performed.
  2. If the materialized view is only LPCT enabled, LPCT refresh is performed.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information about refresh methods for materialized view .

129.4.12 REFRESH_ALL_MVIEWS Procedure

This procedure refreshes all materialized views that have certain properties

All materialized views with the following properties are refreshed:

- The materialized view has not been refreshed since the most recent change to a master table or master materialized view on which it depends.

- The materialized view and all of the master tables or master materialized views on which it depends are local.
- The materialized view is in the view `DBA_MVIEWS`.

This procedure is intended for use with data warehouses.

Syntax

```
DBMS_MVIEW.REFRESH_ALL_MVIEWS (
   number_of_failures     OUT   BINARY_INTEGER,
   method                 IN    VARCHAR2         := NULL,
   rollback_seg           IN    VARCHAR2         := NULL,
   refresh_after_errors   IN    BOOLEAN          := false,
   atomic_refresh         IN    BOOLEAN          := true,
   out_of_place           IN    BOOLEAN          := false);
```

Parameters

Table 129-11 REFRESH_ALL_MVIEWS Procedure Parameters

Parameter	Description
`number_of_failures`	Returns the number of failures that occurred during processing
`method`	A single refresh method indicating the type of refresh to perform for each materialized view that is refreshed. `F` or `f` indicates fast refresh, `?` indicates force refresh, `C` or c indicates complete refresh, and `A` or a indicates always refresh. `A` and `C` are equivalent. If no method is specified, a materialized view is refreshed according to its default refresh method. `P` or `p` refreshes by recomputing the rows in the materialized view affected by changed partitions in the detail tables. `L` or `l` indicates LPCT refresh, that is, recomputing the rows in the materialized view as per the logical partition of the base tables. **Note:** "L" and "?" refreshes attempt a combined PCT and LCPT refresh. (physical and logical partitions).
`rollback_seg`	Name of the materialized view site rollback segment to use while refreshing materialized views
`refresh_after_errors`	If this parameter is `true`, an updatable materialized view continues to refresh even if there are outstanding conflicts logged in the `DEFERROR` view for the materialized view's master table or master materialized view. If this parameter is `true` and `atomic_refresh` is `false`, this procedure continues to refresh other materialized views if it fails while refreshing a materialized view.
`atomic_refresh`	If this parameter is set to `true`, then the refreshed materialized views are refreshed in a single transaction. All of the refreshed materialized views are updated to a single point in time. If the refresh fails for any of the materialized views, none of the materialized views are updated. If this parameter is set to `false`, then each of the materialized views is refreshed non-atomically in separate transactions.

Table 129-11 (Cont.) REFRESH_ALL_MVIEWS Procedure Parameters

Parameter	Description
out_of_place	If true, then it performs an out-of-place refresh. The default is false. This parameter uses the four methods of refresh (F, P, C, ?). So, for example, if you specify F and out_of_place = true, then an out-of-place fast refresh will be attempted. Similarly, if you specify P and out_of_place = true, then out-of-place PCT refresh will be attempted.

129.4.13 REFRESH_DEPENDENT Procedures

This procedure refreshes all materialized views that have certain properties.

Materialized views with the following properties are refreshed:

- The materialized view depends on a master table or master materialized view in the list of specified masters.
- The materialized view has not been refreshed since the most recent change to a master table or master materialized view on which it depends.
- The materialized view and all of the master tables or master materialized views on which it depends are local.
- The materialized view is in the view DBA_MVIEWS.

This procedure is intended for use with data warehouses.

Syntax

```
DBMS_MVIEW.REFRESH_DEPENDENT (
   number_of_failures     OUT    BINARY_INTEGER,
   { list                 IN     VARCHAR2,
   | tab                  IN     DBMS_UTILITY.UNCL_ARRAY,}
   method                 IN     VARCHAR2    := NULL,
   rollback_seg           IN     VARCHAR2    := NULL,
   refresh_after_errors   IN     BOOLEAN     := false,
   atomic_refresh         IN     BOOLEAN     := true,
   nested                 IN     BOOLEAN     := false,
   out_of_place           IN     BOOLEAN     := false);
```

> **Note:**
>
> This procedure is overloaded. The list and tab parameters are mutually exclusive.

Parameters

Table 129-12 REFRESH_DEPENDENT Procedure Parameters

Parameter	Description
number_of_failures	Returns the number of failures that occurred during processing

Table 129-12 (Cont.) REFRESH_DEPENDENT Procedure Parameters

Parameter	Description
`list | tab`	Comma-delimited list of master tables or master materialized views on which materialized views can depend. (Synonyms are not supported.) These tables and the materialized views that depend on them can be located in different schemas. However, all of the tables and materialized views must be in your local database. Alternatively, you may pass in a PL/SQL index-by table of type `DBMS_UTILITY.UNCL_ARRAY`, where each element is the name of a table.
`method`	A string of refresh methods indicating how to refresh the dependent materialized views. All of the materialized views that depend on a particular table are refreshed according to the refresh method associated with that table. `F` or `f` indicates fast refresh, ? indicates force refresh, `C` or `c` indicates complete refresh, and `A` or `a` indicates always refresh. `A` and `C` are equivalent. `P` or `p` refreshes by recomputing the rows in the materialized view affected by changed partitions in the detail tables. `L` or `l` indicates LPCT refresh, that is, recomputing the rows in the materialized view as per the logical partition of the base tables. **Note:** "L" and "?" refreshes attempt a combined PCT and LCPT refresh. (physical and logical partitions). If a table does not have a corresponding refresh method (that is, if more tables are specified than refresh methods), then any materialized view that depends on that table is refreshed according to its default refresh method. For example, the following `EXECUTE` statement within SQL*Plus: `DBMS_MVIEW.REFRESH_DEPENDENT ('employees,deptartments,hr.regions','cf');` performs a complete refresh of the materialized views that depend on the `employees` table, a fast refresh of the materialized views that depend on the `departments` table, and a default refresh of the materialized views that depend on the `hr.regions` table.
`rollback_seg`	Name of the materialized view site rollback segment to use while refreshing materialized views
`refresh_after_errors`	If this parameter is `true`, an updatable materialized view continues to refresh even if there are outstanding conflicts logged in the `DEFERROR` view for the materialized view's master table or master materialized view. If this parameter is `true` and `atomic_refresh` is `false`, this procedure continues to refresh other materialized views if it fails while refreshing a materialized view.
`atomic_refresh`	If this parameter is set to `true`, then the refreshed materialized views are refreshed in a single transaction. All of the refreshed materialized views are updated to a single point in time. If the refresh fails for any of the materialized views, none of the materialized views are updated. If this parameter is set to `false`, then each of the materialized views is refreshed non-atomically in separate transactions.
`nested`	If `true`, then perform nested refresh operations for the specified set of tables. Nested refresh operations refresh all the depending materialized views of the specified set of tables based on a dependency order to ensure the nested materialized views are truly fresh with respect to the underlying base tables.

Table 129-12 (Cont.) REFRESH_DEPENDENT Procedure Parameters

Parameter	Description
out_of_place	If true, then it performs an out-of-place refresh. The default is false. This parameter uses the four methods of refresh (F, P, C, ?). So, for example, if you specify F and out_of_place = true, then an out-of-place fast refresh will be attempted. Similarly, if you specify P and out_of_place = true, then out-of-place PCT refresh will be attempted.

129.4.14 REGISTER_MVIEW Procedure

This procedure enables the administration of individual materialized views. It is invoked at a master site or master materialized view site to register a materialized view.

Note that, typically, a materialized view is registered automatically during materialized view creation. You should only run this procedure to manually register a materialized view if the automatic registration failed or if the registration information was deleted.

Syntax

```
DBMS_MVIEW.REGISTER_MVIEW (
   mviewowner  IN   VARCHAR2,
   mviewname   IN   VARCHAR2,
   mviewsite   IN   VARCHAR2,
   mview_id    IN   DATE | BINARY_INTEGER,
   flag        IN   BINARY_INTEGER,
   qry_txt     IN   VARCHAR2,
   rep_type    IN   BINARY_INTEGER := DBMS_MVIEW.REG_UNKNOWN);
```

Parameters

Table 129-13 REGISTER_MVIEW Procedure Parameters

Parameter	Description
mviewowner	Owner of the materialized view.
mviewname	Name of the materialized view.
mviewsite	Name of the materialized view site for a materialized view registering at an Oracle database version 8.x and higher master site or master materialized view site. This name should not contain any double quotes.
mview_id	The identification number of the materialized view. Specify an Oracle database version 8.x and higher materialized view as a BINARY_INTEGER. Specify an Oracle database version 7 materialized view registering at an Oracle database version 8.x and higher master sites or master materialized view sites as a DATE.

Table 129-13 (Cont.) REGISTER_MVIEW Procedure Parameters

Parameter	Description
flag	A constant that describes the properties of the materialized view being registered. Valid constants that can be assigned include the following: DBMS_MVIEW.REG_ROWID_MVIEW for a rowid materialized view DBMS_MVIEW.REG_PRIMARY_KEY_MVIEW for a primary key materialized view DBMS_MVIEW.REG_OBJECT_ID_MVIEW for an object id materialized view DBMS_MVIEW.REG_FAST_REFRESHABLE_MVIEW for a materialized view that can be fast refreshed DBMS_MVIEW.REG_UPDATABLE_MVIEW for a materialized view that is updatable A materialized view can have more than one of these properties. In this case, use the plus sign (+) to specify more than one property. For example, if a primary key materialized view can be fast refreshed, you can enter the following for this parameter: DBMS_MVIEW.REG_PRIMARY_KEY_MVIEW + DBMS_MVIEW.REG_FAST_REFRESHABLE_MVIEW You can determine the properties of a materialized view by querying the ALL_MVIEWS data dictionary view.
qry_txt	The first 32,000 bytes of the materialized view definition query.
rep_type	Version of the materialized view. Valid constants that can be assigned include the following: DBMS_MVIEW.REG_V7_SNAPSHOT if the materialized view is at an Oracle database version 7 site • DBMS_MVIEW.REG_V8_SNAPSHOT reg_repapi_snapshot if the materialized view is at an Oracle database version 8.x or higher site DBMS_MVIEW.REG_UNKNOWN (the default) if you do not know whether the materialized view is at an Oracle database version 7 site or an Oracle database version 8.x (or higher) site

Usage Notes

This procedure is invoked at the master site or master materialized view site by a remote materialized view site using a remote procedure call. If REGISTER_MVIEW is called multiple times with the same mviewowner, mviewname, and mviewsite, then the most recent values for mview_id, flag, and qry_txt are stored. If a query exceeds the maximum VARCHAR2 size, then qry_txt contains the first 32000 characters of the query and the remainder is truncated. When invoked manually, the value of mview_id must be looked up in the materialized view data dictionary views by the person who calls the procedure.

129.4.15 UNREGISTER_MVIEW Procedure

This procedure enables the administration of individual materialized views. It is invoked at a master site or master materialized view site to unregister a materialized view.

Syntax

```
DBMS_MVIEW.UNREGISTER_MVIEW (
   mviewowner      IN   VARCHAR2,
   mviewname       IN   VARCHAR2,
   mviewsite       IN   VARCHAR2);
```

Parameters

Table 129-14 UNREGISTER_MVIEW Procedure Parameters

Parameters	Description
`mviewowner`	Owner of the materialized view
`mviewname`	Name of the materialized view
`mviewsite`	Name of the materialized view site

130
DBMS_MVIEW_STATS

`DBMS_MVIEW_STATS` package provides an interface to manage the collection and retention of statistics for materialized view refresh operations.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for information about managing and using materialized view refresh statistics

This chapter contains the following topics:

- DBMS_MVIEW_STATS Overview
- DBMS_MVIEW_STATS Security Model
- Summary of DBMS_MVIEW_STATS Subprograms

130.1 DBMS_MVIEW_STATS Overview

You can use the procedures contained in the `DBMS_MVIEW_STATS` package to manage the collection and retention of statistics for materialized view refresh operations. This includes the level and granularity at which these statistics are collected and the duration for which they are retained in the database.

You can also set database level system defaults for the parameters that control statistics collection.

130.2 DBMS_MVIEW_STATS Security Model

Refer to the Usage Notes section in each subprogram for information about the privileges required to use the subprogram.

130.3 Summary of DBMS_MVIEW_STATS Subprograms

This table lists the `DBMS_MVIEW_STATS` subprograms and briefly describes them.

Table 130-1 DBMS_MVIEW_STATS Package Subprograms

Subprogram	Description
PURGE_REFRESH_STATS Procedure	Purges the statistics of materialized view refresh operations that are older than the specified retention period.

Table 130-1 (Cont.) DBMS_MVIEW_STATS Package Subprograms

Subprogram	Description
SET_MVREF_STATS_PARAMS Procedure	Sets the values of parameters that define the collection level and retention period for materialized view refresh statistics. You can set the values either at the database level or for individual materialized views.
SET_SYSTEM_DEFAULT Procedure	Sets the system default value of a refresh statistics parameter. The two refresh statistics parameters are collection level and the retention period.

130.3.1 PURGE_REFRESH_STATS Procedure

This procedure purges refresh statistics that are older than the specified retention period for the specified materialized views.

This procedure forces a purge of refresh statistics without altering the retention period defined for the specified materialized views.

Syntax

```
DBMS_MVIEW_STATS.PURGE_REFRESH_STATS (
        mv_list                 IN      VARACHAR2,
        retention_period        IN      NUMBER);
```

Parameters

Table 130-2 PURGE_REFRESH_STATS Procedure Parameters

Parameter	Description
`mv_list`	The fully-qualified name of an existing materialized view in the form of `schema_name.mv_name`. Use a comma-separated list to specify multiple materialized views. Specify `NULL` to purge materialized view refresh statistics for all materialized views in the database.
`retention_period`	The number of days for which refresh statistics must be preserved in the data dictionary. Statistics for materialized view refresh operations that are older than the retention period are purged from the data dictionary. The retention period specified in this procedure overrides the retention period that may have been set previously either at the database level or for specified materialized views. Specify `NULL` to use the purging policy defined by the automatic statistics purge. Specify –1 to purge all refresh statistics.

Usage Notes

To invoke this procedure, you need either the `SYSDBA` privilege or privileges on every materialized view that is specified in `mv_list`.

130.3.2 SET_MVREF_STATS_PARAMS Procedure

This procedure sets the collection level and retention period for materialized view refresh statistics. You can set these properties either for individual materialized views or for all materialized views in the database.

Syntax

```
DBMS_MVIEW_STATS.SET_MVREF_STATS_PARAMS (
        mv_list            IN    VARACHAR2,
        collection_level   IN    VARCHAR2 DEFAULT NULL,
        retention_period   IN    NUMBER DEFAULT NULL);
```

Parameters

Table 130-3 SET_MVREF_STATS_PARAMS Procedure Parameters

Parameter	Description
mv_list	The fully-qualified name of an existing materialized view in the form of schema_name.mv_name. Use a comma-separated list to specify multiple materialized views. Specify NULL to set properties for all existing materialized views in the database.
collection_level	Specifies the level of detail used when collecting refresh statistics for the materialized views specified in mv_list. Set one of the following values for collection_level: • NONE: No materialized view refresh statistics are collected. • TYPICAL: Only basic refresh statistics are collected and stored for the materialized views specified in mv_list. • ADVANCED: Detailed refresh statistics are collected and stored for materialized view specified in mv_list. If this parameter is set to NULL, then the system default value for collection_level (set using SET_SYSTEM_DEFAULT) is used.
retention_period	Specifies the retention period, in days, for the refresh statistics of the materialized views specified in mv_list. Statistics that are older than the retention period are automatically purged from the data dictionary. Valid values are between 1 and 1365000. If this parameter is set to NULL, then the system default value for retention_period (set using SET_SYSTEM_DEAFULT) is used. Set retention_period to -1 to specify that refresh statistics for the materialized views in mv_list must never be purged.

Usage Notes

To set the collection level or retention period of one or more materialized views, you must have privileges on those materialized views. To set the collection level or retention period for all materialized views in the database, you must have either the SYSDBA privilege or privileges on every materialized view in the database.

To set the system-level default values for statistics collection level and retention period, use the SET_SYSTEM_DEAFULT procedure.

Use the DBA_MVREF_STATS_PARAMS view to determine the currently-set retention period and collection level for materialized view statistics collection.

To disable refresh statistics collection for all materialized views in the database, use the following:

```
DBMS_MVIEW_STATS.SET_MVREF_STATS_PARAMS (NULL, 'NONE', NULL);
```

Note that the parameters set using SET_MVREF_STATS_PARAMS only affect materialized views that exist in the database at the time the procedure is run. Any new materialized views created after this procedure is run will use the system default values for collection_level and retention_period.

130.3.3 SET_SYSTEM_DEFAULT Procedure

This procedure sets system-wide defaults that manage the collection and retention of materialized view refresh statistics. All newly-created materialized views use these defaults until the parameters are reset explicitly using the SET_MVREF_STATS_PARAMS procedure.

Syntax

```
DBMS_MVIEW_STATS.SET_SYSTEM_DEFAULT (
        parameter_name         IN   VARCHAR2,
        value              IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 130-4 SET_SYSTEM_DEFAULT Procedure Parameters

Parameter	Description
parameter_name	The name of the materialized view refresh statistics parameter whose system default value is being set. The parameters that can be set are: • COLLECTION_LEVEL: Specifies the level of detail for collecting materialized view refresh statistics. • RETENTION_PERIOD: Specifies the duration, in days, for which refresh statistics are retained in the data dictionary
value	The value of the materialized view refresh statistics parameter. The valid values for COLLECTION_LEVEL are: • NONE: No refresh statistics are collected for the refresh operation. This is the default setting. • TYPICAL: Only basic refresh statistics are collected for the refresh operation. • ADVANCED: Detailed refresh statistics are collected for the refresh operation. The valid values for RETENTION_PERIOD are: • -1 • Numbers between 1 and 1365000 The default value for retention_period is 31. If you specify NULL for any of the parameters, then the system default setting for that parameter is used.

Usage Notes

You must have `SYSDBA` privilege to invoke this procedure.

Use the `DBA_MVREF_STATS_SYS_DEFAULTS` view to display the current default settings for materialized view refresh statistics collection.

131
DBMS_NETWORK_ACL_ADMIN

The `DBMS_NETWORK_ACL_ADMIN` package provides the interface to administer the network Access Control List (ACL).

The chapter contains the following topics:

- Overview
- Deprecated Subprograms
- Security Model
- Constants
- Exceptions
- Examples
- Summary of DBMS_NETWORK_ACL_ADMIN Subprograms

> **See Also:**
>
> For more information, see "Managing Fine-grained Access to External Network Services" in *Oracle Database Security Guide*

131.1 DBMS_NETWORK_ACL_ADMIN Overview

The `DBMS_NETWORK_ACL_ADMIN` package provides the interface to administer the network access control lists (ACL).

ACLs are used to control access by users to external network services and resources from the database through PL/SQL network utility packages including UTL_TCP , UTL_HTTP , UTL_SMTP and UTL_INADDR .

131.2 DBMS_NETWORK_ACL_ADMIN Deprecated Subprograms

Oracle recommends that you do not use deprecated subprograms in new applications. Support for deprecated features is for backward compatibility only

The following subprograms are deprecated with release Oracle Database 12*c*:

- ADD_PRIVILEGE Procedure
- ASSIGN_ACL Procedure
- ASSIGN_WALLET_ACL Procedure
- CHECK_PRIVILEGE Function
- CHECK_PRIVILEGE_ACLID Function

- CREATE_ACL Procedure
- DELETE_PRIVILEGE Procedure
- DROP_ACL Procedure
- UNASSIGN_ACL Procedure
- UNASSIGN_WALLET_ACL Procedure

131.3 DBMS_NETWORK_ACL_ADMIN Security Model

The `EXECUTE` privilege on the `DBMS_NETWORK_ACL_ADMIN` package is granted to the `DBA` role and to the `EXECUTE_CATALOG_ROLE` by default.

131.4 DBMS_NETWORK_ACL_ADMIN Constants

The `DBMS_NETWORK_ACL_ADMIN` package defines constants to use specifying parameter values.

These are shown in the following table.

Table 131-1 DBMS_NETWORK_ACL_ADMIN Constants

Constant	Type	Value	Description
IP_ADDR_MASK	VARCHAR2(80)	'([[:digit:]]+\.){3}[[:digit:]]+'	IP address mask: xxx.xxx.xxx.xxx
IP_SUBNET_MASK	VARCHAR2(80)	'([[:digit:]]+\.){0,3}*'	IP subnet mask: xxx.xxx...*
HOSTNAME_MASK	VARCHAR2(80)	'[^\.\:\/*]+(\.[^\.\:\/*]+)*'	Hostname mask: ???.???.???...???
DOMAIN_MASK	VARCHAR2(80)	'*(\.[^\.\:\/*]+)*'	Domain mask: *.???.???...???

131.5 DBMS_NETWORK_ACL_ADMIN Exceptions

The following table lists the exceptions raised by the `DBMS_NETWORK_ACL_ADMIN` package.

Table 131-2 DBMS_NETWORK_ACL_ADMIN Exceptions

Exception	Error Code	Description
ACE_ALREADY_EXISTS	24243	ACE already exists
EMPTY_ACL	24246	Empty ACL
ACL_NOT_FOUND	46114	ACL not found
ACL_ALREADY_EXISTS	46212	ACL already exists
INVALID_ACL_PATH	46059	Invalid ACL path
INVALID_HOST	24244	Invalid host
INVALID_PRIVILEGE	24245	Invalid privilege
INVALID_WALLET_PATH	29248	Invalid wallet path
BAD_ARGUMENT	29261	Bad argument

Table 131-2 (Cont.) DBMS_NETWORK_ACL_ADMIN Exceptions

Exception	Error Code	Description
UNRESOLVED_PRINCIPAL	46238	Unresolved principal
PRIVILEGE_NOT_GRANTED	01927	Privilege not granted

131.6 DBMS_NETWORK_ACL_ADMIN Examples

Grant the `connect` and `resolve` privileges for host `www.us.example.com` to `SCOTT`.

Example1

```
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE(
  host => 'www.us.example.com',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('connect', 'resolve'),
                       principal_name => 'scott',
                       principal_type => xs_acl.ptype_db));
```

Example 2

Revoke the `resolve` privilege for host `www.us.example.com` from `SCOTT`.

```
dbms_network_acl_admin.remove_host_ace(
  host => 'www.us.example.com',
  ace  =>  xs$ace_type(privilege_list => xs$name_list('resolve'),
                       principal_name => 'scott',
                       principal_type => xs_acl.ptype_db));
```

Example 3

Grant the `use_client_certificates` and `use_passwords` privileges for wallet `file:/example/wallets/hr_wallet` to `SCOTT`.

```
dbms_network_acl_admin.append_wallet_ace(
  wallet_path => 'file:/example/wallets/hr_wallet',
  ace         =>  xs$ace_type(privilege_list => xs$name_list('use_client_certificates',
'use_passwords'),
                              principal_name => 'scott',
                              principal_type => xs_acl.ptype_db));
```

Example 4

Revoke the `use_passwords` privilege for wallet `file:/example/wallets/hr_wallet` from `SCOTT`.

```
dbms_network_acl_admin.remove_wallet_ace(
  wallet_path => 'file:/example/wallets/hr_wallet',
  ace         =>  xs$ace_type(privilege_list => xs$name_list('use_passwords'),
                              principal_name => 'scott',
                              principal_type => xs_acl.ptype_db));
```

Example 5

The `CONTAINS_HOST` in the `DBMS_NETWORK_ACL_UTLILITY` package determines if a host is contained in a domain. It can be used in conjunction with the `DBA_HOST_ACE` view to determine the users and their privilege assignments to access a network host.For example, for access to `www.us.example.com`:

```
SELECT HOST, LOWER_PORT, UPPER_PORT,
       ACE_ORDER, PRINCIPAL, PRINCIPAL_TYPE,
       GRANT_TYPE, INVERTED_PRINCIPAL, PRIVILEGE,
       START_DATE, END_DATE
  FROM (SELECT ACES.*,
DBMS_NETWORK_ACL_UTILITY.CONTAINS_HOST('www.us.example.com',
                                                    HOST) PRECEDENCE
          FROM DBA_HOST_ACES ACES)
 WHERE PRECEDENCE IS NOT NULL
 ORDER BY PRECEDENCE DESC,
          LOWER_PORT NULLS LAST,
          UPPER_PORT NULLS LAST,
          ACE_ORDER;

HOST               LOWER_PORT UPPER_PORT ACE_ORDER PRINCIPAL PRINCIPAL_TYPE   GRANT_TYPE
INVERTED_PRINCIPAL PRIVILEGE START_DATE END_DATE
------------------ ---------- ---------- --------- --------- ---------------- ----------
------------------ ---------- ---------- --------
www.us.example.com         80         80         1 SCOTT     DATABASE USER    GRANT
NO                 HTTP
www.us.example.com         80         80         2 ADAMS     DATABASE USER    GRANT
NO                 HTTP
*                                                  1 HQ_DBA    DATABASE USER    GRANT
NO                 CONNECT
*                                                  1 HQ_DBA    DATABASE USER    GRANT
NO                 RESOLVE
```

Example 6

For example, for `HQ_DBA`'s own permission to access to `www.us.example.com`:

```
SELECT HOST, LOWER_PORT, UPPER_PORT, PRIVILEGE, STATUS
  FROM (SELECT ACES.*,
DBMS_NETWORK_ACL_UTILITY.CONTAINS_HOST('www.us.example.com',
                                                    HOST) PRECEDENCE
          FROM USER_HOST_ACES ACES)
 WHERE PRECEDENCE IS NOT NULL
 ORDER BY PRECEDENCE DESC,
          LOWER_PORT NULLS LAST,
          UPPER_PORT NULLS LAST;


HOST               LOWER_PORT UPPER_PORT PRIVILEGE STATUS
------------------ ---------- ---------- --------- -------
*                                        CONNECT   GRANTED
*                                        RESOLVE   GRANTED
```

131.7 Summary of DBMS_NETWORK_ACL_ADMIN Subprograms

This table lists and briefly describes the `DBMS_NETWORK_ACL_ADMIN` package subprograms.

Table 131-3 DBMS_NETWORK_ACL_ADMIN Package Subprograms

Subprogram	Description
ADD_PRIVILEGE Procedure	[DEPRECATED] Adds a privilege to grant or deny the network access to the user in an access control list (ACL)

Table 131-3 (Cont.) DBMS_NETWORK_ACL_ADMIN Package Subprograms

Subprogram	Description
APPEND_HOST_ACE Procedure	Appends an access control entry (ACE) to the access control list (ACL) of a network host.
APPEND_HOST_ACL Procedure	Appends access control entries (ACE) of an access control list (ACL) to the ACL of a network host
APPEND_WALLET_ACE Procedure	Appends an access control entry (ACE) to the access control list (ACL) of a wallet
APPEND_WALLET_ACL Procedure	Appends access control entries (ACE) of an access control list (ACL) to the ACL of a wallet
ASSIGN_ACL Procedure	[DEPRECATED] Assigns an access control list (ACL) to a host computer, domain, or IP subnet, and if specified, the TCP port range.
ASSIGN_WALLET_ACL Procedure	[DEPRECATED] Assigns an access control list (ACL) to a wallet
CHECK_PRIVILEGE Function	[DEPRECATED] Checks if a privilege is granted or denied the user in an access control list (ACL)
CHECK_PRIVILEGE_ACLID Function	[DEPRECATED] Checks if a privilege is granted to or denied from the user in an ACL by specifying the object ID of the access control list
CREATE_ACL Procedure	[DEPRECATED] Creates an access control list (ACL) with an initial privilege setting
DELETE_PRIVILEGE Procedure	[DEPRECATED] Deletes a privilege in an access control list (ACL)
DROP_ACL Procedure	[DEPRECATED] Drops an access control list (ACL)
REMOVE_HOST_ACE Procedure	Removes privileges from access control entries (ACE) in the access control list (ACL) of a network host matching the given ACE
REMOVE_WALLET_ACE Procedure	Removes privileges from access control entries (ACE) in the access control list (ACL) of a wallet matching the given ACE
SET_HOST_ACL Procedure	Sets the access control list (ACL) of a network host which controls access to the host from the database
SET_WALLET_ACL Procedure	Sets the access control list (ACL) of a wallet which controls access to the wallet from the database
UNASSIGN_ACL Procedure	[DEPRECATED] Unassigns the access control list (ACL) currently assigned to a network host
UNASSIGN_WALLET_ACL Procedure	[DEPRECATED] Unassigns the access control list (ACL) currently assigned to a wallet

131.7.1 ADD_PRIVILEGE Procedure

This procedure adds a privilege to grant or deny the network access to the user. The access control entry (ACE) is created if it does not exist.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the APPEND_HOST_ACE Procedure and the APPEND_WALLET_ACE Procedure.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.ADD_PRIVILEGE (
   acl             IN VARCHAR2,
   principal       IN VARCHAR2,
   is_grant        IN BOOLEAN,
   privilege       IN VARCHAR2,
   position        IN PLS_INTEGER DEFAULT NULL,
   start_date      IN TIMESTAMP WITH TIMESTAMP DEFAULT NULL,
   end_date        IN TIMESTAMP WITH TIMESTAMP DEFAULT NULL );
```

Parameters

Table 131-4 ADD_PRIVILEGE Function Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "/sys/acls"
`principal`	Principal (database user or role) to whom the privilege is granted or denied. Case sensitive.
`is_grant`	Privilege is granted or denied.
`privilege`	Network privilege to be granted or denied
`position`	Position (1-based) of the ACE. If a non-`NULL` value is given, the privilege will be added in a new ACE at the given position and there should not be another ACE for the principal with the same `is_grant` (grant or deny). If a `NULL` value is given, the privilege will be added to the ACE matching the principal and the `is_grant` if one exists, or to the end of the ACL if the matching ACE does not exist.
`start_date`	Start date of the access control entry (ACE). When specified, the ACE will be valid only on and after the specified date. The `start_date` will be ignored if the privilege is added to an existing ACE.
`end_date`	End date of the access control entry (ACE). When specified, the ACE expires after the specified date. The `end_date` must be greater than or equal to the `start_date`. The `end_date` will be ignored if the privilege is added to an existing ACE.

Usage Notes

To remove the permission, use the DELETE_PRIVILEGE Procedure.

Examples

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.ADD_PRIVILEGE(
        acl         => 'us-example-com-permissions.xml',
        principal   => 'ST_USERS',
        is_grant    =>  TRUE,
        privilege   => 'connect')
END;
```

131.7.2 APPEND_HOST_ACE Procedure

This procedure appends an access control entry (ACE) to the access control list (ACL) of a network host. The ACL controls access to the given host from the database and the ACE specifies the privileges granted to or denied from the specified principal.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACE (
   host         IN VARCHAR2,
   lower_port   IN PLS_INTEGER DEFAULT NULL,
   upper_port   IN PLS_INTEGER DEFAULT NULL,
   ace          IN XS$ACE_TYPE);
```

Parameters

Table 131-5 APPEND_HOST_ACE Function Parameters

Parameter	Description
`host`	The host, which can be the name or the IP address of the host. You can use a wildcard to specify a domain or a IP subnet. The host or domain name is case-insensitive.
`lower_port`	Lower bound of an optional TCP port range
`upper_port`	Upper bound of an optional TCP port range. If `NULL`, `lower_port` is assumed.
`ace`	The ACE

Usage Notes

- Duplicate privileges in the matching ACE in the host ACL will be skipped.
- To remove the ACE, use the REMOVE_HOST_ACE Procedure.
- A host's ACL takes precedence over its domains' ACLs. For a given host, say `www.us.example.com`, the following domains are listed in decreasing precedence:
  - `www.us.example.com`
  - `*.us.example.com`
  - `*.example.com`
  - `*.com`
  - `*`
- An IP address' ACL takes precedence over its subnets' ACLs. For a given IP address, say `192.168.0.100`, the following subnets are listed in decreasing precedence:

  - 192.168.0.100
  - 192.168.0.*
  - 192.168.*
  - 192.*
  - *

- An ACE with a "resolve" privilege can be appended only to a host's ACL without a port range.
- When ACEs with "connect" privileges are appended to a host's ACLs with and without a port range, the one appended to the host with a port range takes precedence.
- When specifying a TCP port range of a host, it cannot overlap with other existing port ranges of the host.
- If the ACL is shared with another host or wallet, a copy of the ACL will be made before the ACL is modified.

> **See Also:**
>
> *Oracle Database Real Application Security Administrator's and Developer's Guide* for more information about the `XS$ACE_TYPE` object type

131.7.3 APPEND_HOST_ACL Procedure

This procedure appends access control entries (ACE) of an access control list (ACL) to the ACL of a network host.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.APPEND_HOST_ACL (
   host         IN VARCHAR2,
   lower_port   IN PLS_INTEGER DEFAULT NULL,
   upper_port   IN PLS_INTEGER DEFAULT NULL,
   acl          IN VARCHAR2);
```

Parameters

Table 131-6 APPEND_HOST_ACL Function Parameters

Parameter	Description
`host`	The host, which can be the name or the IP address of the host. You can use a wildcard to specify a domain or a IP subnet. The host or domain name is case-insensitive.
`lower_port`	Lower bound of an optional TCP port range
`upper_port`	Upper bound of an optional TCP port range. If `NULL`, `lower_port` is assumed.
`acl`	The ACL from which to append

Usage Notes

- Duplicate privileges in the matching ACE in the host ACL will be skipped.

- To remove the ACE, use the REMOVE_HOST_ACE Procedure.
- A host's ACL takes precedence over its domains' ACLs. For a given host, say `www.us.example.com`, the following domains are listed in decreasing precedence:
  - `www.us.example.com`
  - `*.us.example.com`
  - `*.example.com`
  - `*.com`
  - `*`
- An IP address' ACL takes precedence over its subnets' ACLs. For a given IP address, say `192.168.0.100`, the following subnets are listed in decreasing precedence:
  - `192.168.0.100`
  - `192.168.0.*`
  - `192.168.*`
  - `192.*`
  - `*`
- An ACE with a "resolve" privilege can be appended only to a host's ACL without a port range.
- When ACEs with "connect" privileges are appended to a host's ACLs with and without a port range, the one appended to the host with a port range takes precedence.
- When specifying a TCP port range of a host, it cannot overlap with other existing port ranges of the host.- If the ACL is shared with another host or wallet, a copy of the ACL will be made before the ACL is modified.

131.7.4 APPEND_WALLET_ACE Procedure

This procedure appends an access control entry (ACE) to the access control list (ACL) of a wallet. The ACL controls access to the given wallet from the database and the ACE specifies the privileges granted to or denied from the specified principal.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.APPEND_WALLET_ACE (
   wallet_path     IN VARCHAR2,
   ace             IN XS$ACE_TYPE);
```

Parameters

Table 131-7 APPEND_WALLET_ACE Function Parameters

Parameter	Description
`wallet_path`	Directory path of the wallet. The path is case-sensitive of the format `file:`*`directory-path`*.
`ace`	The ACE

Usage Notes

- Duplicate privileges in the matching ACE in the host ACL will be skipped.

- To remove the ACE, use the REMOVE_WALLET_ACE Procedure.
- If the ACL is shared with another host or wallet, a copy of the ACL is made before the ACL is modified.

> **See Also:**
>
> *Oracle Database Real Application Security Administrator's and Developer's Guide* for more information about the `XS$ACE_TYPE` object type

131.7.5 APPEND_WALLET_ACL Procedure

This procedure appends access control entries (ACE) of an access control list (ACL) to the ACL of a wallet.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.APPEND_WALLET_ACL (
   wallet_path    IN VARCHAR2,
   acl            IN VARCHAR2);
```

Parameters

Table 131-8 APPEND_WALLET_ACL Function Parameters

Parameter	Description
`wallet_path`	Directory path of the wallet. The path is case-sensitive of the format `file:`*`directory-path`*.
`ace`	The ACL from which to append

Usage Notes

- Duplicate privileges in the matching ACE in the host ACL will be skipped.
- To remove the ACE, use REMOVE_WALLET_ACE.
- If the ACL is shared with another host or wallet, a copy of the ACL is made before the ACL is modified.

131.7.6 ASSIGN_ACL Procedure

This procedure assigns an access control list (ACL) to a host computer, domain, or IP subnet, and if specified, the TCP port range.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the APPEND_HOST_ACE Procedure and the APPEND_WALLET_ACE Procedure.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.ASSIGN_ACL (
   acl         IN VARCHAR2,
   host        IN VARCHAR2,
   lower_port  IN PLS_INTEGER DEFAULT NULL,
   upper_port  IN PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 131-9 ASSIGN_ACL Function Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "`/sys/acls`".
`host`	Host to which the ACL is to be assigned. The host can be the name or the IP address of the host. A wildcard can be used to specify a domain or a IP subnet. The host or domain name is case-insensitive.
`lower_port`	Lower bound of a TCP port range if not `NULL`
`upper_port`	Upper bound of a TCP port range. If `NULL`, `lower_port` is assumed.

Usage Notes

- Only one ACL can be assigned to any host computer, domain, or IP subnet, and if specified, the TCP port range. When you assign a new access control list to a network target, Oracle Database unassigns the previous access control list that was assigned to the same target. However, Oracle Database does not drop the access control list. You can drop the access control list by using the DROP_ACL Procedure. To remove an access control list assignment, use the UNASSIGN_ACL Procedure.
- The ACL assigned to a domain takes a lower precedence than the other ACLs assigned sub-domains, which take a lower precedence than the ACLs assigned to the individual hosts. So for a given host, for example, "www.us.example.com", the following domains are listed in decreasing precedences:

  - www.us.example.com

  - *.us.example.com

  - *.example.com

  - *.com

  - *

  In the same way, the ACL assigned to an subnet takes a lower precedence than the other ACLs assigned smaller subnets, which take a lower precedence than the ACLs assigned to the individual IP addresses. So for a given IP address, for example, "192.168.0.100", the following subnets are listed in decreasing precedences:

  - 192.168.0.100

  - 192.168.0.*

  - 192.168.*

  - 192.*

  - *

- The port range is applicable only to the "connect" privilege assignments in the ACL. The "resolve" privilege assignments in an ACL have effects only when the ACL is assigned to a host without a port range.

  For the "connect" privilege assignments, an ACL assigned to the host without a port range takes a lower precedence than other ACLs assigned to the same host with a port range.
- When specifying a TCP port range, both `lower_port` and `upper_port` must not be `NULL` and `upper_port` must be greater than or equal to `lower_port`. The port range must not overlap with any other port ranges for the same host assigned already.
- To remove the assignment, use UNASSIGN_ACL Procedure.

Examples

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.ASSIGN_ACL(
     acl         => 'us-example-com-permissions.xml',
     host        => '*.us.example.com',
     lower_port  => 80);
END;
```

131.7.7 ASSIGN_WALLET_ACL Procedure

This procedure assigns an access control list (ACL) to a wallet.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the APPEND_HOST_ACE Procedure and the APPEND_WALLET_ACE Procedure.

Syntax

```
UTL_HTTP.ASSIGN_WALLET_ACL (
   acl          IN  VARCHAR2,
   wallet_path  IN  VARCHAR2);
```

Parameters

Table 131-10 ASSIGN_WALLET_ACL Procedure Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "`/sys/acls`"
`wallet_path`	Directory path of the wallet to which the ACL is to be assigned. The path is case-sensitive and of the format `file:`*`directory-path`*.

Usage Notes

To remove the assignment, use the UNASSIGN_WALLET_ACL Procedure.

Examples

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.CREATE_ACL(
    acl         => 'wallet-acl.xml',
    description => 'Wallet ACL',
    principal   => 'SCOTT',
    is_grant    => TRUE,
    privilege   => 'use-client-certificates');

  DBMS_NETWORK_ACL_ADMIN.ADD_PRIVILEGE(
    acl         => 'wallet-acl.xml',
    principal   => 'SCOTT',
    is_grant    => TRUE,
    privilege   => 'use-passwords');

  DBMS_NETWORK_ACL_ADMIN.ASSIGN_WALLET_ACL(
    acl         => 'wallet-acl.xml',
    wallet_path => 'file:/example/wallets/test_wallet');
END;
```

131.7.8 CHECK_PRIVILEGE Function

This function checks if a privilege is granted or denied the user in an ACL.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. The procedure remains available in the package only for reasons of backward compatibility.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.CHECK_PRIVILEGE (
   acl             IN VARCHAR2,
   user            IN VARCHAR2,
   privilege       IN VARCHAR2)
  RETURN NUMBER;
```

Parameters

Table 131-11 CHECK_PRIVILEGE Function Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "/sys/acls".
`user`	User to check against. If the user is `NULL`, the invoker is assumed. The username is case-sensitive as in the `USERNAME` column of the `ALL_USERS` view.
`privilege`	Network privilege to check

Return Values

Returns 1 when the privilege is granted; 0 when the privilege is denied; `NULL` when the privilege is neither granted or denied.

Examples

```
SELECT DECODE(
  DBMS_NETWORK_ACL_ADMIN.CHECK_PRIVILEGE(
       'us-example-com-permissions.xml', 'SCOTT', 'resolve'),
  1, 'GRANTED', 0, 'DENIED', NULL) PRIVILEGE
FROM DUAL;
```

131.7.9 CHECK_PRIVILEGE_ACLID Function

This function checks if a privilege is granted to or denied from the user in an ACL by specifying the object ID of the access control list.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. The procedure remains available in the package only for reasons of backward compatibility.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.CHECK_PRIVILEGE_ACLID (
   aclid           IN RAW,
   user            IN VARCHAR2 DEFAULT NULL)
   privilege       IN VARCHAR2,
 RETURN NUMBER;
```

Parameters

Table 131-12 CHECK_PRIVILEGE_ACLID Function Parameters

Parameter	Description
`aclid`	Object ID of the ACL
`user`	User to check against. If the user is `NULL`, the invoker is assumed. The username is case-sensitive as in the `USERNAME` column of the `ALL_USERS` view.
`privilege`	Network privilege to check

Return Values

Returns 1 when the privilege is granted; 0 when the privilege is denied; `NULL` when the privilege is neither granted or denied.

131.7.10 CREATE_ACL Procedure

This deprecated procedure creates an access control list (ACL) with an initial privilege setting. An ACL must have at least one privilege setting. The ACL has no access control effect unless it is assigned to the network target.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the APPEND_HOST_ACE Procedure and the APPEND_WALLET_ACE Procedure.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.CREATE_ACL (
   acl             IN VARCHAR2,
   description     IN VARCHAR2,
   principal       IN VARCHAR2,
   is_grant        IN BOOLEAN,
   privilege       IN VARCHAR2,
   start_date      IN TIMESTAMP WITH TIMEZONE DEFAULT NULL,
   end_date        IN TIMESTAMP WITH TIMEZONE DEFAULT NULL );
```

Parameters

Table 131-13 CREATE_ACL Procedure Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "/sys/acls".
`description`	Description attribute in the ACL
`principal`	Principal (database user or role) to whom the privilege is granted or denied. Case sensitive.
`is_grant`	Privilege is granted or not (denied)
`privilege`	Network privilege to be granted or denied - '`connect` | `resolve`' (case sensitive). A database user needs the `connect` privilege to an external network host computer if he or she is connecting using the `UTL_TCP`, `UTL_HTTP`, `UTL_SMTP`, and `UTL_MAIL` utility packages. To resolve a host name that was given a host IP address, or the IP address that was given a host name, with the `UTL_INADDR` package, grant the database user the `resolve` privilege.
`start_date`	Start date of the access control entry (ACE). When specified, the ACE is valid only on and after the specified date.
`end_date`	End date of the access control entry (ACE). When specified, the ACE expires after the specified date. The `end_date` must be greater than or equal to the `start_date`.

Usage Notes

To drop the access control list, use the DROP_ACL Procedure.

Examples

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.CREATE_ACL(
     acl          => 'us-example-com-permissions.xml',
     description  => 'Network permissions for *.us.example.com',
     principal    => 'SCOTT',
     is_grant     => TRUE,
     privilege    => 'connect');
END;
```

131.7.11 DELETE_PRIVILEGE Procedure

This deprecated procedure deletes a privilege in an access control list.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the REMOVE_HOST_ACE Procedure and the REMOVE_WALLET_ACE Procedure.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.DELETE_PRIVILEGE (
   acl           IN VARCHAR2,
   principal     IN VARCHAR2,
   is_grant      IN BOOLEAN DEFAULT NULL,
   privilege     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 131-14 DELETE_PRIVILEGE Function Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "/sys/acls".
`principal`	Principal (database user or role) for whom all the ACE will be deleted
`is_grant`	Privilege is granted or not (denied). If a `NULL` value is given, the deletion is applicable to both granted or denied privileges.
`privilege`	Network privilege to be deleted. If a `NULL` value is given, the deletion is applicable to all privileges.

Examples

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.DELETE_PRIVILEGE(
        acl         => 'us-example-com-permissions.xml',
        principal   => 'ST_USERS')
END;
```

131.7.12 DROP_ACL Procedure

This **deprecated procedure** drops an access control list (ACL).

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. The procedure remains available in the package only for reasons of backward compatibility.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.DROP_ACL (
   acl            IN VARCHAR2);
```

Parameters

Table 131-15 DROP_ACL Procedure Parameters

Parameter	Description
acl	Name of the ACL. Relative path will be relative to "/sys/acls".

Examples

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.DROP_ACL(
      acl => 'us-example-com-permissions.xml');
END;
```

131.7.13 REMOVE_HOST_ACE Procedure

This procedure removes privileges from access control entries (ACE) in the access control list (ACL) of a network host matching the given ACE.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.REMOVE_HOST_ACE (
   host               IN VARCHAR2,
   lower_port         IN PLS_INTEGER DEFAULT NULL,
   upper_port         IN PLS_INTEGER DEFAULT NULL,
   ace                IN XS$ACE_TYPE,
   remove_empty_acl   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 131-16 REMOVE_HOST_ACE Function Parameters

Parameter	Description
host	The host, which can be the name or the IP address of the host. You can use a wildcard to specify a domain or a IP subnet. The host or domain name is case-insensitive.
lower_port	Lower bound of an optional TCP port range

Table 131-16 (Cont.) REMOVE_HOST_ACE Function Parameters

Parameter	Description
upper_port	Upper bound of an optional TCP port range. If NULL, lower_port is assumed.
ace	The ACE
remove_empty_acl	Whether to remove the ACL when it becomes empty when the ACE is removed

Usage Notes

If the ACL is shared with another host or wallet, a copy of the ACL is made before the ACL is modified.

131.7.14 REMOVE_WALLET_ACE Procedure

This procedure removes privileges from access control entries (ACE) in the access control list (ACL) of a wallet matching the given ACE.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.REMOVE_WALLET_ACE (
   wallet_path        IN VARCHAR2,
   ace                IN XS$ACE_TYPE,
   remove_empty_acl   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 131-17 REMOVE_WALLET_ACE Function Parameters

Parameter	Description
wallet_path	Directory path of the wallet. The path is case-sensitive of the format file:*directory-path*.
ace	The ACE
remove_empty_acl	Whether to remove the ACL when it becomes empty when the ACE is removed

Usage Notes

If the ACL is shared with another host or wallet, a copy of the ACL is made before the ACL is modified.

131.7.15 SET_HOST_ACL Procedure

This procedure sets the access control list (ACL) of a network host which controls access to the host from the database.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.SET_HOST_ACL (
   host         IN VARCHAR2,
   lower_port   IN PLS_INTEGER DEFAULT NULL,
```

```
   upper_port   IN PLS_INTEGER DEFAULT NULL,
   acl          IN VARCHAR2);
```

Parameters

Table 131-18 SET_HOST_ACL Function Parameters

Parameter	Description
`host`	The host, which can be the name or the IP address of the host. You can use a wildcard to specify a domain or a IP subnet. The host or domain name is case-insensitive.
`lower_port`	Lower bound of an optional TCP port range
`upper_port`	Upper bound of an optional TCP port range. If `NULL`, `lower_port` is assumed.
`acl`	The ACL. `NULL` to unset the host's ACL.

Usage Notes

A host's ACL is created and set on-demand when an access control entry (ACE) is appended to the host's ACL. Users are discouraged from setting a host's ACL manually.

131.7.16 SET_WALLET_ACL Procedure

This procedure sets the access control list (ACL) of a wallet which controls access to the wallet from the database.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.SET_WALLET_ACL (
   wallet_path    IN VARCHAR2,
   acl            IN VARCHAR2);
```

Parameters

Table 131-19 SET_WALLET_ACL Function Parameters

Parameter	Description
`wallet_path`	Directory path of the wallet. The path is case-sensitive of the format `file:`*`directory-path`*.
`acl`	The ACL. `NULL` to unset the host's ACL.

Usage Notes

A wallet's ACL is created and set on-demand when an access control entry (ACE) is appended to the wallet's ACL. Users are discouraged from setting a wallet's ACL manually.

131.7.17 UNASSIGN_ACL Procedure

This deprecated procedure unassigns the access control list (ACL) currently assigned to a network host.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the REMOVE_HOST_ACE Procedure and the REMOVE_WALLET_ACE Procedure.

Syntax

```
DBMS_NETWORK_ACL_ADMIN.UNASSIGN_ACL (
   acl          IN VARCHAR2 DEFAULT NULL,
   host         IN VARCHAR2 DEFAULT NULL,
   lower_port   IN PLS_INTEGER DEFAULT NULL,
   upper_port   IN PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 131-20 UNASSIGN_ACL Function Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "/sys/acls". If ACL is `NULL`, any ACL assigned to the host is unassigned.
`host`	Host from which the ACL is to be removed. The host can be the name or the IP address of the host. A wildcard can be used to specify a domain or a IP subnet. The host or domain name is case-insensitive. If host is `NULL`, the ACL will be unassigned from any host. If both `host` and `acl` are `NULL`, all ACLs assigned to any hosts are unassigned.
`lower_port`	Lower bound of a TCP port range if not `NULL`
`upper_port`	Upper bound of a TCP port range. If `NULL`, `lower_port` is assumed.

Examples

```
BEGIN
   DBMS_NETWORK_ACL_ADMIN.UNASSIGN_ACL(
     host        => '*.us.example.com',
     lower_port  => 80);
END;
```

131.7.18 UNASSIGN_WALLET_ACL Procedure

This deprecated procedure unassigns the access control list (ACL) currently assigned to a wallet.

> **Note:**
>
> This procedure is deprecated in Oracle Database 12*c*. While the procedure remains available in the package for reasons of backward compatibility, Oracle recommends using the REMOVE_HOST_ACE Procedure and the REMOVE_WALLET_ACE Procedure.

Syntax

```
UTL_HTTP.UNASSIGN_WALLET_ACL (
   acl          IN  VARCHAR2 DEFAULT NULL,
   wallet_path  IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 131-21 UNASSIGN_WALLET_ACL Procedure Parameters

Parameter	Description
`acl`	Name of the ACL. Relative path will be relative to "`/sys/acls`". If `acl` is `NULL`, any ACL assigned to the wallet is unassigned
`wallet_path`	Directory path of the wallet to which the ACL is assigned. The path is case-sensitive and of the format `file:`*`directory-path`*. If both `acl` and `wallet_path` are `NULL`, all ACLs assigned to any wallets are unassigned.

Examples

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.UNASSIGN_WALLET_ACL(
    acl         => 'wallet-acl.xml',
    wallet_path => 'file:/example/wallets/test_wallet');
END;
```

132
DBMS_NETWORK_ACL_UTILITY

The `DBMS_NETWORK_ACL_UTILITY` package provides the utility functions to facilitate the evaluation of access control list (ACL) assignments governing TCP connections to network hosts.

> **See Also:**
>
> For more information, see ""Managing Fine-grained Access to External Network Services"" in *Oracle Database Security Guide*

The chapter contains the following topics:

- Security Model
- Examples
- Summary of DBMS_NETWORK_ACL_UTILITY Subprograms

132.1 DBMS_NETWORK_ACL_UTILITY Security Model

`EXECUTE` on the `DBMS_NETWORK_ACL_UTILITY` package is granted to `PUBLIC`.

132.2 DBMS_NETWORK_ACL_UTILITY Examples

The CONTAINS_HOST Function in this package indicates if a domain or subnet contains a given host or IP address.

It can be used in conjunction with the CHECK_PRIVILEGE_ACLID Function in the DBMS_NETWORK_ACL_ADMIN package to determine the privilege assignments affecting a user's permission to access a network host. The return value of the `CONTAINS_HOST` Function in can also be used to order the ACL assignments by their precedence.

Example 1

For example, for SCOTT's permission to connect to www.hr.example.com:

```
SELECT host, lower_port, upper_port, acl,
   DECODE(
       DBMS_NETWORK_ACL_ADMIN.CHECK_PRIVILEGE_ACLID(aclid, 'SCOTT', 'connect'),
          1, 'GRANTED', 0, 'DENIED', NULL) privilege
   FROM (SELECT host, acl, aclid, lower_port, upper_port,
             DBMS_NETWORK_ACL_UTILITY.CONTAINS_HOST('www.hr.example.com', host)
               precedence
         FROM dba_network_acls)
WHERE precedence > 0
ORDER BY precedence DESC, lower_port nulls LAST;


  HOST                LOWER_PORT UPPER_PORT  ACL                 PRIVILEGE
```

```
-------------------- ---------- ---------- -------------------- ---------
www.hr.example.com           80         80  /sys/acls/www.xml    GRANTED
www.hr.example.com         3000       3999  /sys/acls/www.xml    GRANTED
www.hr.example.com                         /sys/acls/www.xml    GRANTED
*.hr.example.com                           /sys/acls/all.xml
*.example.com                              /sys/acls/all.xml
```

Example 2

For example, for `SCOTT`'s permission to do domain name resolution for www.hr.example.com:

```
SELECT host, acl,
  DECODE(
    DBMS_NETWORK_ACL_ADMIN.CHECK_PRIVILEGE_ACLID(aclid, 'SCOTT', 'resolve'),
      1, 'GRANTED', 0, 'DENIED', null) privilege
  FROM (SELECT host, acl, aclid,
               DBMS_NETWORK_ACL_UTILITY.CONTAINS_HOST('www.hr.example.com', host)
                 precedence
          FROM dba_network_acls
         WHERE lower_port IS NULL AND upper_port IS NULL)
 WHERE precedence > 0
 ORDER BY precedence DESC;

HOST                 ACL                          PRIVILEGE
-------------------- ---------------------------- ---------
www.hr.example.com   /sys/acls/hr-www.xml         GRANTED
*.hr.example.com     /sys/acls/hr-domain.xml
*.example.com        /sys/acls/corp-domain.xml
```

Note that the "resolve" privilege takes effect only in ACLs assigned without any port range (when `lower_port` and `upper_port` are `NULL`). For this reason, the example does not include `lower_port` and `upper_port` columns in the query.

Related Topics

- CONTAINS_HOST Function
  This function determines if the given host is equal to or contained in the given host, domain, or subnet. It handles different representation of the same IP address or subnet. For example, an IPv4-mapped IPv6 address is considered equal to the IPv4-native address it represents. It does not perform domain name resolution when evaluating the host or domain.

132.3 Summary of DBMS_NETWORK_ACL_UTILITY Subprograms

This table lists and briefly describes the `DBMS_NETWORK_ACL_UTILITY` package subprograms.

Table 132-1 DBMS_NETWORK_ACL_UTILITY Package Subprograms

Subprogram	Description
CONTAINS_HOST Function	Determines if the given host is equal to or contained in the given host, domain, or subnet
DOMAIN_LEVEL Function	Returns the domain level of the given host name, domain, or subnet

Table 132-1 (Cont.) DBMS_NETWORK_ACL_UTILITY Package Subprograms

Subprogram	Description
DOMAINS Function	For a given host, this function returns the domains whose ACL assigned is used to determine if a user has the privilege to access the given host or not.
EQUALS_HOST Function	Determines if the two given hosts, domains, or subnets are equal

132.3.1 CONTAINS_HOST Function

This function determines if the given host is equal to or contained in the given host, domain, or subnet. It handles different representation of the same IP address or subnet. For example, an IPv4-mapped IPv6 address is considered equal to the IPv4-native address it represents. It does not perform domain name resolution when evaluating the host or domain.

Syntax

```
DBMS_NETWORK_ACL_UTILITY.CONTAINS_HOST (
   host      IN    VARCHAR2,
   domain    IN    VARCHAR2)
 RETURN NUMBER;
```

Parameters

Table 132-2 CONTAINS_HOST Function Parameters

Parameter	Description
`host`	Network host
`domain`	Network host, domain, or subnet

Return Values

Returns a non-`NULL` value if the given host is equal to or contained in the related host, domain, or subnet:

- If `domain` is a hostname, returns the level of its domain + 1
- If `domain` is a domain name, returns the domain level
- If `domain` is an IP address or subnet, return the number of significant address bits of the IP address or subnet
- If domain is the wildcard "*", returns 0

The non-`NULL` value returned indicates the precedence of the domain or subnet for ACL assignment. The higher the value, the higher is the precedence. `NULL` will be returned if the host is not equal to or contained in the given host, domain or subnet.

Examples

```
SELECT host, acl, precedence
  FROM (select host, acl,
               DBMS_NETWORK_ACL_UTILITY.CONTAINS_HOST('192.0.2.3', host)
                 precedence
          FROM dba_network_acls)
```

```
 WHERE precedence > 0
 ORDER BY precedence DESC;

HOST                 ACL                        PRECEDENCE
-------------------- -------------------------- ----------
192.0.2.3            /sys/acls/hr-www.xml               32
::ffff:192.0.2.0/120 /sys/acls/hr-domain.xml            24
::ffff:192.0.0.0/104 /sys/acls/corp-domain.xml           8
```

132.3.2 DOMAIN_LEVEL Function

This function returns the domain level of the given host name, domain, or subnet.

Syntax

```
DBMS_NETWORK_ACL_UTILITY.DOMAIN_LEVEL (
    host  IN VARCHAR2)
  RETURN NUMBER;
```

Parameters

Table 132-3 DOMAIN_LEVEL Function Parameters

Parameter	Description
host	Network host, domain, or subnet

Return Values

The domain level of the given host, domain, or subnet.

Usage Notes

Note that this function cannot handle IPv6 addresses and subnets, and subnets in CIDR notation.

Examples

```
SELECT host, acl, domain_level
  FROM (select host, acl,
               DBMS_NETWORK_ACL_UTILITY.DOMAIN_LEVEL(host) domain_level
          FROM dba_network_acls)
 order by domain_level desc;

HOST                 ACL                          DOMAIN_LEVEL
-------------------- ---------------------------- ------------
www.hr.example.com   /sys/acls/hr-www.xml                    4
*.hr.example.com     /sys/acls/hr-domain.xml                 3
*.example.com        /sys/acls/corp-domain.xml               2
```

132.3.3 DOMAINS Function

For a given host, this function returns the domains whose ACL assigned determines if a user has the privilege to access the given host or not. When the IP address of the host is given, return the subnets instead.

Syntax

```
DBMS_NETWORK_ACL_UTILITY.DOMAINS (
    host  IN VARCHAR2)
  RETURN DOMAIN_TABLE PIPELINED;
```

Parameters

Table 132-4 DOMAINS Function Parameters

Parameter	Description
host	Network host

Return Values

The domains or subnets for the given host.

Usage Notes

Note that this function cannot handle IPv6 addresses. Nor can it generate subnets of arbitrary number of prefix bits for an IPv4 address.

Examples

```
select * from table(dbms_network_acl_utility.domains('www.hr.example.com'));

DOMAINS
--------------------------
www.hr.example.com
*.hr.example.com
*.example.com
*.com
*
```

132.3.4 EQUALS_HOST Function

This function determines if the two given hosts, domains, or subnets are equal. It handles different representation of the same IP address or subnet. For example, an IPv4-mapped IPv6 address is considered equal to the IPv4- native address it represents. It does not perform domain name resolution when comparing the two hosts or domains.

Syntax

```
DBMS_NETWORK_ACL_UTILITY.EQUALS_HOST (
   host1    IN    VARCHAR2,
   host2    IN    VARCHAR2)
 RETURN NUMBER;
```

Parameters

Table 132-5 EQUALS_HOST Function Parameters

Parameter	Description
host1	Network host, domain, or subnet to compare
host2	Network host, domain, or subnet to compare

Return Values

1 if the two hosts, domains, or subnets are equal. 0 otherwise.

Examples

```
SELECT host, acl
  FROM dba_network_acls
 WHERE DBMS_NETWORK_ACL_UTILITY.EQUALS_HOST('192.0.2.*', host) = 1;

HOST                 ACL
-------------------- ----------------------------
::ffff:192.0.2.0/120   /sys/acls/hr-domain.xml
```

133
DBMS_ODCI

DBMS_ODCI package contains a single user function related to the use of Data Cartridges.

This chapter contains the following topic:

- Summary of DBMS_ODCI Subprograms

> **See Also:**
>
> - *Oracle Database Data Cartridge Developer's Guide*

133.1 Summary of DBMS_ODCI Subprograms

The DBMS_ODCI package has one subprogram, the ESTIMATE_CPU_UNITS function.

Table 133-1 ***DBMS_ODCI Package Subprograms***

Subprogram	Description
ESTIMATE_CPU_UNITS Function	Returns the approximate number of CPU instructions (in thousands) corresponding to a specified time interval (in seconds)

133.1.1 ESTIMATE_CPU_UNITS Function

This function returns the approximate number of CPU instructions (in thousands) corresponding to a specified time interval (in seconds). This information can be used to associate the CPU cost with a user-defined function for the extensible optimizer.

The function takes as input the elapsed time of the user function, measures CPU units by multiplying the elapsed time by the processor speed of the machine, and returns the approximate number of CPU instructions that should be associated with the user function. For a multiprocessor machine, ESTIMATE_CPU_UNITS considers the speed of a single processor.

Syntax

```
DBMS_ODCI.ESTIMATE_CPU_UNITS(
   elapsed_time     NUMBER)
  RETURN NUMBER;
```

Parameters

Parameter	Description
elapsed_time	The elapsed time in seconds that it takes to execute a function.

Usage Notes

When associating CPU cost with a user-defined function, use the full number of CPU units rather than the number of *thousands* of CPU units returned by `ESTIMATE_CPU_UNITS`; multiply the number returned by `ESTIMATE_CPU_UNITS` by 1,000.

134
DBMS_OPTIM_BUNDLE

The DBMS_OPTIM_BUNDLE package is used to manage the installed but disabled module bug fixes which cause an execution plan change.

This chapter contains the following topics:

- DBMS_OPTIM_BUNDLE Security Model
- Summary of DBMS_OPTIM_BUNDLE Subprograms

134.1 DBMS_OPTIM_BUNDLE Security Model

The DBMS_OPTIM_BUNDLE package runs with the privileges of the user starting the package. DBA role is required to run this package.

134.2 Summary of DBMS_OPTIM_BUNDLE Subprograms

This table lists the DBMS_OPTIM_BUNDLE subprograms and briefly describes them.

Table 134-1 DBMS_OPTIM_BUNDLE Package Subprograms

Subprogram	Description
ENABLE_OPTIM_FIXES Procedure	This procedure enables or disables fixes with plan changes up to the latest installed release update.
GETBUGSFORBUNDLE Procedure	This procedure displays execution plan bug fixes applied as part of release updates.
LISTBUNDLESWITHFCFIXES Procedure	This procedure lists the release update names and release update IDs of release updates with fix control fixes.
SET_FIX_CONTROLS Procedure	This procedure enables or disables a list of fixes with _fix_controls.

134.2.1 ENABLE_OPTIM_FIXES Procedure

The DBMS_OPTIM_BUNDLE subprogram, ENABLE_OPTIM_FIXES procedure enables or disables fixes with plan changes up to the latest installed release update.

Syntax

```
DBMS_OPTIM_BUNDLE.ENABLE_OPTIM_FIXES (
   action                          IN  VARCHAR2  DEFAULT 'OFF',
   scope                           IN  VARCHAR2  DEFAULT 'MEMORY',
   current_setting_precedence      IN  VARCHAR2  DEFAULT 'YES');
```

Parameters

Table 134-2 ENABLE_OPTIM_FIXES Procedure Parameters

Parameter	Description
action	Enables or disables all of the installed but disabled execution plan bug fixes up to and including the current release update. The possible values are: • ON-enables the bundle fixes. • OFF-disables bundle fixes by setting the _fix_controls to the value 0. It will not remove _fix_control entries from SPFILE. The default value is OFF.
scope	The scope of enabling or disabling the installed but disabled execution plan bug fixes. The possible values are: • MEMORY • SPFILE • BOTH • INITORA MEMORY/SPFILE/BOTH: These three input values enables or disables the fixes in a given scope. INITORA: This input value will just display manual command syntax that the user needs to manually enter into the database's init.ora file in order to enable or disable execution plan bug fixes. When using SCOPE=INITORA in this way, the current_setting_precedence field has no significance.
current_setting_precedence	Sets the precedence of environment settings or release update settings, where these settings are in conflict. The possible values are: • YES-the current environment settings take precedence in case of conflict • NO-release update settings take precedence in case of conflict The default value is YES.

Examples

To enable all of the installed but disabled execution plan bug fixes up to and including those from the current release update:

```
SQL> execute dbms_optim_bundle.enable_optim_fixes('ON','BOTH', 'YES');
```

This instruction would enable all fixes, in both MEMORY and in the persistant SPFILE, with a precedence of current setting.

Exceptions

The following exceptions are raised by the ENABLE_OPTIM_FIXES Procedure:

- ORA-20001: Bad input value
- ORA-20002: Internal or other errors

134.2.2 GETBUGSFORBUNDLE Procedure

The DBMS_OPTIM_BUNDLE subprogram, GETBUGSFORBUNDLE procedure displays execution plan bug fixes applied as part of release updates.

Syntax

```
DBMS_OPTIM_BUNDLE.GETBUGSFORBUNDLE (
   bundleid            IN  NUMBER  DEFAULT NULL);
```

Parameters

Table 134-3 GETBUGSFORBUNDLE Procedure Parameters

Parameter	Description
bundleid	The release update ID. If NULL, the execution plan bug fixes from the latest bundle is displayed. If a bundle ID is specified, then the execution plan bug fixes up to the specified bundle is displayed. The default value is NULL.

Examples

To view a listing of the installed but disabled execution plan bug fixes from the most recent release update applied:

```
SQL> set serveroutput on;
SQL> execute dbms_optim_bundle.getbugsforbundle;

19.8.0.0.200714DBRU:
   Bug: 29304314, fix_controls: 29304314
   Bug: 29930457, fix_controls: 29930457
```

The above example lists the installed but disabled execution plan bug fixes from the 19.8.0.0.200714DBRU release.

To view a listing of the installed but disabled execution plan bug fixes from release updates up to and including release update 171017:

```
SQL> execute dbms_optim_bundle.getbugsforbundle(171017);
```

Exceptions

The following exceptions are raised by the GETBUGSFORBUNDLE Procedure:

- ORA-20001: Bad input value
- ORA-20002: Internal or other errors

134.2.3 LISTBUNDLESWITHFCFIXES Procedure

The DBMS_OPTIM_BUNDLE subprogram, LISTBUNDLESWITHFCFIXES procedure lists the release update names and release update IDs of release updates with fix control fixes.

Syntax

```
DBMS_OPTIM_BUNDLE.LISTBUNDLESWITHFCFIXES ( );
```

Examples

To view the release update names and release update IDs:

```
SQL> set serveroutput on
SQL> exec dbms_optim_bundle.listBundlesWithFCFixes;
   bundleId: 190719, bundleName: 19.4.0.0.190719DBRU
   bundleId: 191015, bundleName: 19.5.0.0.191015DBRU
   bundleId: 200414, bundleName: 19.7.0.0.200414DBRU
   bundleId: 200714, bundleName: 19.8.0.0.200714DBRU
   bundleId: 201020, bundleName: 19.9.0.0.201020DBRU
   bundleId: 210119, bundleName: 19.10.0.0.210119DBRU
   bundleId: 210420, bundleName: 19.11.0.0.210420DBRU
   bundleId: 210720, bundleName: 19.12.0.0.210720DBRU

   PL/SQL procedure successfully completed.
```

Exceptions

The following exception is raised by the `LISTBUNDLESWITHFCFIXES Procedure`:

`ORA-20002`: Internal or other errors

134.2.4 SET_FIX_CONTROLS Procedure

The `DBMS_OPTIM_BUNDLE` subprogram, `SET_FIX_CONTROLS` procedure enables or disables a list of fixes with `_fix_controls`. The fixes can be present in a base version, in a release update, or in a one-off release. This procedure appends the new fix control settings to the existing ones.

Syntax

```
DBMS_OPTIM_BUNDLE.SET_FIX_CONTROLS (
   fix_control_string              IN  VARCHAR2,
   sid                             IN  VARCHAR2  DEFAULT '*',
   scope                           IN  VARCHAR2  DEFAULT 'MEMORY',
   current_setting_precedence      IN  VARCHAR2  DEFAULT 'YES');
```

Parameters

Table 134-4 SET_FIX_CONTROLS Procedure Parameters

Parameter	Description
`fix_control_string`	Comma separated list of `fix_control:value` pair. For example: `'13329748:0,20355502:4, 27060221:1'`
`sid`	The name of the instance on which `fix_control` setting need to be made. Acceptable values are `*` or valid `instance_name` from `sys.gv$instance`. The default value is `*`.
`scope`	The scope of enabling or disabling the installed but disabled execution plan bug fixes. The possible values are: • `MEMORY` • `SPFILE` • `BOTH` The default value is `MEMORY`.

Table 134-4 (Cont.) SET_FIX_CONTROLS Procedure Parameters

Parameter	Description
current_setting_precedence	Precedence of current setting over user supplied setting. The possible values are: • YES–the current environment settings take precedence in case of conflict • NO–input settings take precedence in case of conflict The default value is YES.

Exceptions

The following exceptions are raised by the SET_FIX_CONTROLS Procedure:

- ORA-20001: Bad input value
- ORA-20002: Internal or other errors

135
DBMS_OUTLN

The `DBMS_OUTLN` package, synonymous with `OUTLN_PKG`, contains the functional interface for subprograms associated with the management of stored outlines.

> **See Also:**
>
> For more information about using the `DBMS_OUTLN` package, see *Oracle Database SQL Tuning Guide*.

> **Note:**
>
> Stored outlines will be desupported in a future release in favor of SQL plan management. In Oracle Database 11g Release 1 (11.1), stored outlines continue to function as in past releases. However, Oracle strongly recommends that you use SQL plan management for new applications. SQL plan management creates SQL plan baselines, which offer superior SQL performance and stability compared with stored outlines.If you have existing stored outlines, please consider migrating them to SQL plan baselines by using the LOAD_PLANS_FROM_CURSOR_CACHE Functions or LOAD_PLANS_FROM_SQLSET Function of the `DBMS_SPM` package. When the migration is complete, you should disable or remove the stored outlines.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_OUTLN Subprograms

135.1 DBMS_OUTLN Overview

A stored outline is the stored data that pertains to an execution plan for a given SQL statement. It enables the optimizer to repeatedly re-create execution plans that are equivalent to the plan originally generated along with the outline.The data stored in an outline consists, in part, of a set of hints that are used to achieve plan stability.

Stored outlines will be de-supported in a future release in favor of SQL plan management. As of 11g R1, stored outlines continue to function as in past releases, but Oracle strongly recommends that you use SQL plan management for new applications. SQL plan management creates SQL plan baselines, which offer superior SQL performance and stability compared with stored outlines.If you have existing stored outlines, please consider migrating them to SQL plan baselines by using the LOAD_PLANS_FROM_CURSOR_CACHE Functions or the LOAD_PLANS_FROM_SQLSET Function of the DBMS_SPM package. When the migration is complete, you should disable or remove the stored outlines.

135.2 DBMS_OUTLN Security Model

DBMS_OUTLN contains management procedures that should be available to appropriate users only. EXECUTE privilege is not extended to the general user community unless the DBA explicitly does so.

PL/SQL functions that are available for outline management purposes can be executed only by users with EXECUTE privilege on the procedure (or package).

135.3 Summary of DBMS_OUTLN Subprograms

This table lists the DBMS_OUTLN subprograms and briefly describes them.

Table 135-1 *DBMS_OUTLN Package Subprograms*

Subprogram	Description
CLEAR_USED Procedure	Clears the outline 'used' flag
CREATE_OUTLINE Procedure	Generates outlines from the shared cursor identified by hash value and child number
DROP_BY_CAT Procedure	Drops outlines that belong to a specified category
DROP_UNUSED Procedure	Drops outlines that have never been applied in the compilation of a SQL statement
EXACT_TEXT_SIGNATURES Procedure	Updates outline signatures to those that compute based on exact text matching
UPDATE_BY_CAT Procedure	Changes the category of outlines in one category to a new category
UPDATE_SIGNATURES Procedure	Updates outline signatures to the current version's signature

135.3.1 CLEAR_USED Procedure

This procedure clears the outline 'used' flag.

Syntax

```
DBMS_OUTLN.CLEAR_USED (
   name     IN     VARCHAR2);
```

Parameters

Table 135-2 CLEAR_USED Procedure Parameters

Parameter	Description
name	Name of the outline.

135.3.2 CREATE_OUTLINE Procedure

This procedure generates an outline by reparsing the SQL statement from the shared cursor identified by hash value and child number.

Syntax

```
DBMS_OUTLN.CREATE_OUTLINE (
   hash_value    IN NUMBER,
   child_number  IN NUMBER,
   category      IN VARCHAR2 DEFAULT 'DEFAULT');
```

Parameters

Table 135-3 CREATE_OUTLINE Procedure Parameters

Parameter	Description
hash_value	Hash value identifying the target shared cursor.
child_number	Child number of the target shared cursor.
category	Category in which to create outline (optional).

135.3.3 DROP_BY_CAT Procedure

This procedure drops outlines that belong to a particular category. While outlines are put into the DEFAULT category unless otherwise specified, users have the option of grouping their outlines into groups called categories.

Syntax

```
DBMS_OUTLN.DROP_BY_CAT (
   cat VARCHAR2);
```

Parameters

Table 135-4 DROP_BY_CAT Procedure Parameters

Parameter	Description
cat	Category of outlines to drop.

Usage Notes

This procedure purges a category of outlines in a single call.

Examples

This example drops all outlines in the DEFAULT category:

```
DBMS_OUTLN.DROP_BY_CAT('DEFAULT');
```

135.3.4 DROP_UNUSED Procedure

This procedure drops outlines that have never been applied in the compilation of a SQL statement.

Syntax

```
DBMS_OUTLN.DROP_UNUSED;
```

Usage Notes

You can use `DROP_UNUSED` for outlines generated by an application for one-time use SQL statements created as a result of dynamic SQL. These outlines are never used and take up valuable disk space.

135.3.5 EXACT_TEXT_SIGNATURES Procedure

This procedure updates outline signatures to those that compute based on exact text matching.

Syntax

```
DBMS_OUTLN.EXACT_TEXT_SIGNATURES;
```

Usage Notes

This procedure is relevant only for downgrading an outline to 8.1.6 or earlier.

135.3.6 UPDATE_BY_CAT Procedure

This procedure changes the category of all outlines in one category to a new category.

Syntax

```
DBMS_OUTLN.UPDATE_BY_CAT (
   oldcat    VARCHAR2 default 'DEFAULT',
   newcat    VARCHAR2 default 'DEFAULT');
```

Parameters

Table 135-5 UPDATE_BY_CAT Procedure Parameters

Parameter	Description
`oldcat`	The current category of outlines.
`newcat`	The new category of outlines.

135.3.7 UPDATE_SIGNATURES Procedure

This procedure updates outline signatures to the current version's signature.

Syntax

```
DBMS_OUTLN.UPDATE_SIGNATURES;
```

Usage Notes

You should execute this procedure if you have imported outlines generated in an earlier release to ensure that the signatures are compatible with the current release's computation algorithm.

136
DBMS_OUTPUT

The `DBMS_OUTPUT` package enables you to send messages from stored procedures, packages, and triggers. The package is especially useful for displaying PL/SQL debugging information.

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Exceptions
- Rules and Limits
- Examples
- Data Structures
- Summary of DBMS_OUTPUT Subprograms

136.1 DBMS_OUTPUT Overview

The package is typically used for debugging, or for displaying messages and reports to SQL*DBA or SQL*Plus (such as are produced by applying the SQL command `DESCRIBE` to procedures).

The PUT Procedure and PUT_LINE Procedure in this package enable you to place information in a buffer that can be read by another trigger, procedure, or package. In a separate PL/SQL procedure or anonymous block, you can display the buffered information by calling the GET_LINE Procedure and GET_LINES Procedure.

If the package is disabled, all calls to subprograms are ignored. In this way, you can design your application so that subprograms are available only when a client is able to process the information.

136.2 DBMS_OUTPUT Security Model

The `dbmsotpt.sql` script must be run as user `SYS`. This creates the public synonym `DBMS_OUTPUT`, and `EXECUTE` permission on this package is granted to `public`.

136.3 DBMS_OUTPUT Operational Notes

The following operational notes apply to DBMS_OUTPUT.

- If you do not call `GET_LINE`, or if you do not display the messages on your screen in SQL*Plus, the buffered messages are ignored.
- SQL*Plus calls `GET_LINES` after issuing a SQL statement or anonymous PL/SQL calls.
- Typing `SET SERVEROUTPUT ON` in SQL*Plus has the effect of invoking

```
DBMS_OUTPUT.ENABLE (buffer_size => NULL);
```

with no limit on the output.

- You should generally avoid having application code invoke either the DISABLE Procedure or ENABLE Procedure because this could subvert the attempt of an external tool like SQL*Plus to control whether or not to display output.

> **Note:**
>
> Messages sent using `DBMS_OUTPUT` are not actually sent until the sending subprogram or trigger completes. There is no mechanism to flush output during the execution of a procedure.

136.4 DBMS_OUTPUT Exceptions

`DBMS_OUTPUT` subprograms raise the application error `ORA-20000` and return errors.

The output procedures can return the following errors:

Table 136-1 DBMS_OUTPUT Errors

Error	Description
`ORU-10027:`	Buffer overflow
`ORU-10028:`	Line length overflow

136.5 DBMS_OUTPUT Rules and Limits

The following are limits on DBMS_OUTPUT line and buffer size.

- The maximum line size is 32767 bytes.
- The default buffer size is 20000 bytes. The minimum size is 2000 bytes and the maximum is unlimited.

136.6 DBMS_OUTPUT Examples

This topic contains three examples of using DBMS_OUTPUT.

Example 1: Using a Trigger to Produce Output

You can use a trigger to print out some output from the debugging process. For example, you could code the trigger to invoke:

```
DBMS_OUTPUT.PUT_LINE('I got here:'||:new.col||' is the new value');
```

If you have enabled the `DBMS_OUTPUT` package, then the text produced by this `PUT_LINE` would be buffered, and you could, after executing the statement (presumably some `INSERT`, `DELETE`, or `UPDATE` that caused the trigger to fire), retrieve the line of information. For example:

```
BEGIN
   DBMS_OUTPUT.GET_LINE(:buffer, :status);
END;
```

You could then optionally display the buffer on the screen. You repeat calls to `GET_LINE` until `status` comes back as nonzero. For better performance, you should use calls to GET_LINES Procedure which can return an array of lines.

Example 2: Debugging Stored Procedures and Triggers

The `DBMS_OUTPUT` package is commonly used to debug stored procedures and triggers. This package can also be used to enable you to retrieve information about an object and format this output, as shown in "Example 3: Retrieving Information About an Object".

This function queries the employee table and returns the total salary for a specified department. The function includes several calls to the `PUT_LINE` procedure:

```
CREATE FUNCTION dept_salary (dnum NUMBER) RETURN NUMBER IS
   CURSOR emp_cursor IS
      SELECT sal, comm FROM emp WHERE deptno = dnum;
   total_wages    NUMBER(11, 2) := 0;
   counter        NUMBER(10) := 1;
BEGIN

   FOR emp_record IN emp_cursor LOOP
      emp_record.comm := NVL(emp_record.comm, 0);
      total_wages := total_wages + emp_record.sal
         + emp_record.comm;
      DBMS_OUTPUT.PUT_LINE('Loop number = ' || counter ||
         '; Wages = '|| TO_CHAR(total_wages));    /* Debug line */
      counter := counter + 1; /* Increment debug counter */
   END LOOP;
   /* Debug line */
   DBMS_OUTPUT.PUT_LINE('Total wages = ' ||
     TO_CHAR(total_wages));
   RETURN total_wages;

END dept_salary;
```

Assume the `EMP` table contains the following rows:

```
EMPNO          SAL     COMM     DEPT
-----        ------- -------- -------
1002           1500      500      20
1203           1000               30
1289           1000               10
1347           1000      250      20
```

Assume the user executes the following statements in SQL*Plus:

```
SET SERVEROUTPUT ON
VARIABLE salary NUMBER;
EXECUTE :salary := dept_salary(20);
```

The user would then see the following information displayed in the output pane:

```
Loop number = 1; Wages = 2000
Loop number = 2; Wages = 3250
Total wages = 3250

PL/SQL procedure successfully executed.
```

Example 3: Retrieving Information About an Object

In this example, the user has used the `EXPLAIN PLAN` command to retrieve information about the execution plan for a statement and has stored it in `PLAN_TABLE`. The user has also assigned a statement ID to this statement. The example `EXPLAIN_OUT` procedure retrieves the information from this table and formats the output in a nested manner that more closely depicts the order of steps undergone in processing the SQL statement.

```
 /*****************************************************************/
/* Create EXPLAIN_OUT procedure. User must pass STATEMENT_ID to */
/* to procedure, to uniquely identify statement.                 */
/*****************************************************************/
CREATE OR REPLACE PROCEDURE explain_out
   (statement_id IN VARCHAR2) AS

   -- Retrieve information from PLAN_TABLE into cursor EXPLAIN_ROWS.

   CURSOR explain_rows IS
      SELECT level, id, position, operation, options,
         object_name
      FROM plan_table
      WHERE statement_id = explain_out.statement_id
      CONNECT BY PRIOR id = parent_id
         AND statement_id = explain_out.statement_id
      START WITH id = 0
       ORDER BY id;

BEGIN

   -- Loop through information retrieved from PLAN_TABLE:

   FOR line IN explain_rows LOOP

      -- At start of output, include heading with estimated cost.

      IF line.id = 0 THEN
         DBMS_OUTPUT.PUT_LINE ('Plan for statement '
            || statement_id
            || ', estimated cost = ' || line.position);
      END IF;

      -- Output formatted information. LEVEL determines indention level.

      DBMS_OUTPUT.PUT_LINE (lpad(' ',2*(line.level-1)) ||
         line.operation || ' ' || line.options || ' ' ||
         line.object_name);
   END LOOP;

END;
```

> **See Also:**
>
> UTL_FILE

136.7 DBMS_OUTPUT Data Structures

The `DBMS_OUTPUT` package declares 2 collection types for use with the GET_LINES Procedure.

TABLE Types

CHARARR Table Type

OBJECT Types

DBMSOUTPUT_LINESARRAY Object Type

Related Topics

- GET_LINES Procedure
  This procedure retrieves an array of lines from the buffer.

136.7.1 CHARARR Table Type

This package type is to be used with the GET_LINES Procedure to obtain text submitted through the PUT Procedure and PUT_LINE Procedure.

Syntax

```
TYPE CHARARR IS TABLE OF VARCHAR2(32767) INDEX BY BINARY_INTEGER;
```

Related Topics

- GET_LINES Procedure
  This procedure retrieves an array of lines from the buffer.
- PUT Procedure
  This procedure places a partial line in the buffer.
- PUT_LINE Procedure
  This procedure places a line in the buffer.

136.7.2 DBMS_OUTPUT DBMSOUTPUT_LINESARRAY Object Type

This type, defined outside the package, is to be used with the GET_LINES Procedure to obtain text submitted through the PUT Procedure and PUT_LINE Procedure.

Syntax

```
TYPE DBMSOUTPUT_LINESARRAY IS
     VARRAY(2147483647) OF VARCHAR2(32767);
```

Related Topics

- GET_LINES Procedure
  This procedure retrieves an array of lines from the buffer.
- PUT Procedure
  This procedure places a partial line in the buffer.
- PUT_LINE Procedure
  This procedure places a line in the buffer.

136.8 Summary of DBMS_OUTPUT Subprograms

This table lists the DBMS_OUTPUT subprograms and briefly describes them.

Table 136-2 DBMS_OUTPUT Package Subprograms

Subprogram	Description
DISABLE Procedure	Disables message output
ENABLE Procedure	Enables message output
GET_LINE Procedure	Retrieves one line from buffer
GET_LINES Procedure	Retrieves an array of lines from buffer
NEW_LINE Procedure	Terminates a line created with PUT
PUT Procedure	Places a partial line in the buffer
PUT_LINE Procedure	Places line in buffer

> **Note:**
>
> The PUT Procedure that take a number are obsolete and, while currently supported, are included in this release for legacy reasons only.

136.8.1 DISABLE Procedure

This procedure disables calls to PUT, PUT_LINE, NEW_LINE, GET_LINE, and GET_LINES, and purges the buffer of any remaining information.

As with the ENABLE Procedure, you do not need to call this procedure if you are using the SERVEROUTPUT option of SQL*Plus.

Syntax

```
DBMS_OUTPUT.DISABLE;
```

Pragmas

```
pragma restrict_references(disable,WNDS,RNDS);
```

136.8.2 ENABLE Procedure

This procedure enables calls to PUT, PUT_LINE, NEW_LINE, GET_LINE, and GET_LINES.

Calls to these procedures are ignored if the DBMS_OUTPUT package is not activated.

Syntax

```
DBMS_OUTPUT.ENABLE (
   buffer_size IN INTEGER DEFAULT 20000);
```

Pragmas

```
pragma restrict_references(enable,WNDS,RNDS);
```

Parameters

Table 136-3 ENABLE Procedure Parameters

Parameter	Description
buffer_size	Upper limit, in bytes, the amount of buffered information. Setting buffer_size to NULL specifies that there should be no limit.

Usage Notes

- It is not necessary to call this procedure when you use the SET SERVEROUTPUT option of SQL*Plus.
- If there are multiple calls to ENABLE, then buffer_size is the last of the values specified. The maximum size is 1,000,000, and the minimum is 2,000 when the user specifies buffer_size (NOT NULL).
- NULL is expected to be the usual choice. The default is 20,000 for backwards compatibility with earlier database versions that did not support unlimited buffering.

136.8.3 GET_LINE Procedure

This procedure retrieves a single line of buffered information.

Syntax

```
DBMS_OUTPUT.GET_LINE (
   line    OUT VARCHAR2,
   status  OUT INTEGER);
```

Parameters

Table 136-4 GET_LINE Procedure Parameters

Parameter	Description
line	Returns a single line of buffered information, excluding a final newline character. You should declare the actual for this parameter as VARCHAR2 (32767) to avoid the risk of "ORA-06502: PL/SQL: numeric or value error: character string buffer too small".
status	If the call completes successfully, then the status returns as 0. If there are no more lines in the buffer, then the status is 1.

Usage Notes

- You can choose to retrieve from the buffer a single line or an array of lines. Call the GET_LINE procedure to retrieve a single line of buffered information. To reduce the number of calls to the server, call the GET_LINES procedure to retrieve an array of lines from the buffer.
- You can choose to automatically display this information if you are using SQL*Plus by using the special SET SERVEROUTPUT ON command.
- After calling GET_LINE or GET_LINES, any lines not retrieved before the next call to PUT, PUT_LINE, or NEW_LINE are discarded to avoid confusing them with the next message.

136.8.4 GET_LINES Procedure

This procedure retrieves an array of lines from the buffer.

Syntax

```
DBMS_OUTPUT.GET_LINES (
   lines       OUT     CHARARR,
   numlines    IN OUT  INTEGER);

DBMS_OUTPUT.GET_LINES (
   lines       OUT     DBMSOUTPUT_LINESARRAY,
   numlines    IN OUT INTEGER);
```

Parameters

Table 136-5 GET_LINES Procedure Parameters

Parameter	Description
`lines`	Returns an array of lines of buffered information. The maximum length of each line in the array is 32767 bytes. It is recommended that you use the `VARRAY` overload version in a 3GL host program to execute the procedure from a PL/SQL anonymous block.
`numlines`	Number of lines you want to retrieve from the buffer. After retrieving the specified number of lines, the procedure returns the number of lines actually retrieved. If this number is less than the number of lines requested, then there are no more lines in the buffer.

Usage Notes

- You can choose to retrieve from the buffer a single line or an array of lines. Call the `GET_LINE` procedure to retrieve a single line of buffered information. To reduce the number of calls to the server, call the `GET_LINES` procedure to retrieve an array of lines from the buffer.
- You can choose to automatically display this information if you are using SQL*Plus by using the special `SET SERVEROUTPUT ON` command.
- After calling `GET_LINE` or `GET_LINES`, any lines not retrieved before the next call to `PUT`, `PUT_LINE`, or `NEW_LINE` are discarded to avoid confusing them with the next message.

136.8.5 NEW_LINE Procedure

This procedure puts an end-of-line marker.

The GET_LINE Procedure and the GET_LINES Procedure return "lines" as delimited by "newlines". Every call to the PUT_LINE Procedure or NEW_LINE Procedure generates a line that is returned by `GET_LINE(S)`.

Syntax

```
DBMS_OUTPUT.NEW_LINE;
```

136.8.6 PUT Procedure

This procedure places a partial line in the buffer.

> **Note:**
>
> The `PUT` procedure that takes a `NUMBER` is obsolete and, while currently supported, is included in this release for legacy reasons only.

Syntax

```
DBMS_OUTPUT.PUT (
    item IN VARCHAR2);
```

Parameters

Table 136-6 PUT Procedure Parameters

Parameter	Description
`item`	Item to buffer.

Exceptions

Table 136-7 PUT Procedure Exceptions

Error	Description
`ORA-20000, ORU-10027:`	Buffer overflow, limit of <`buf_limit`> bytes.
`ORA-20000, ORU-10028:`	Line length overflow, limit of 32767 bytes for each line.

Usage Notes

- You can build a line of information piece by piece by making multiple calls to `PUT`, or place an entire line of information into the buffer by calling `PUT_LINE`.
- When you call `PUT_LINE` the item you specify is automatically followed by an end-of-line marker. If you make calls to `PUT` to build a line, then you must add your own end-of-line marker by calling `NEW_LINE`. `GET_LINE` and `GET_LINES` do not return lines that have not been terminated with a newline character.
- If your lines exceed the line limit, you receive an error message.
- Output that you create using `PUT` or `PUT_LINE` is buffered. The output cannot be retrieved until the PL/SQL program unit from which it was buffered returns to its caller.

  For example, SQL*Plus does not display `DBMS_OUTPUT` messages until the PL/SQL program completes. There is no mechanism for flushing the `DBMS_OUTPUT` buffers within the PL/SQL program.

  ```
  SQL> SET SERVEROUTPUT ON
  SQL> BEGIN
  ```

```
2  DBMS_OUTPUT.PUT_LINE ('hello');
3  DBMS_LOCK.SLEEP (10);
4  END;
```

136.8.7 PUT_LINE Procedure

This procedure places a line in the buffer.

> **Note:**
>
> The `PUT_LINE` procedure that takes a `NUMBER` is obsolete and, while currently supported, is included in this release for legacy reasons only.

Syntax

```
DBMS_OUTPUT.PUT_LINE (
   item IN VARCHAR2);
```

Parameters

Table 136-8 PUT_LINE Procedure Parameters

Parameter	Description
`item`	Item to buffer.

Exceptions

Table 136-9 PUT_LINE Procedure Exceptions

Error	Description
`ORA-20000, ORU-10027:`	Buffer overflow, limit of <`buf_limit`> bytes.
`ORA-20000, ORU-10028:`	Line length overflow, limit of 32767 bytes for each line.

Usage Notes

- You can build a line of information piece by piece by making multiple calls to `PUT`, or place an entire line of information into the buffer by calling `PUT_LINE`.
- When you call `PUT_LINE` the item you specify is automatically followed by an end-of-line marker. If you make calls to `PUT` to build a line, then you must add your own end-of-line marker by calling `NEW_LINE`. `GET_LINE` and `GET_LINES` do not return lines that have not been terminated with a newline character.
- If your lines exceeds the line limit, you receive an error message.
- Output that you create using `PUT` or `PUT_LINE` is buffered. The output cannot be retrieved until the PL/SQL program unit from which it was buffered returns to its caller.

  For example, SQL*Plus does not display `DBMS_OUTPUT` messages until the PL/SQL program completes. There is no mechanism for flushing the `DBMS_OUTPUT` buffers within the PL/SQL program. For example:

```
SQL> SET SERVEROUTPUT ON
SQL> BEGIN
   2 DBMS_OUTPUT.PUT_LINE ('hello');
   3 DBMS_LOCK.SLEEP (10);
   4 END;
```

137
DBMS_PARALLEL_EXECUTE

The `DBMS_PARALLEL_EXECUTE` package enables incremental update of table data in parallel.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Views
- Exceptions
- Examples
- Summary of DBMS_PARALLEL_EXECUTE Subprograms

> **See Also:**
>
> - *Oracle Database Development Guide*
> - *Oracle Database Reference*

137.1 DBMS_PARALLEL_EXECUTE Overview

This package lets you incrementally update table data in parallel, in two high-level steps.

1. Group sets of rows in the table into smaller-sized chunks.
2. Run a user-specified statement on these chunks in parallel, and commit when finished processing each chunk.

This package introduces the notion of *parallel execution task*. This task groups the various steps associated with the parallel execution of a PL/SQL block, which is typically updating table data.

All of the package subroutines (except the GENERATE_TASK_NAME Function and the TASK_STATUS Procedure) perform a commit.

137.2 DBMS_PARALLEL_EXECUTE Security Model

`DBMS_PARALLEL_EXECUTE` is a `SYS`-owned package which is granted to `PUBLIC`.

Users who have the `ADM_PARALLEL_EXECUTE_TASK` role can perform administrative routines (qualified by the prefix `ADM_`) and access the DBA view.

Apart from the administrative routines, all the subprograms refer to tasks owned by the current user.

To execute chunks in parallel, you must have `CREATE JOB` system privilege.

The CHUNK_BY_SQL, RUN_TASK, and RESUME_TASK subprograms require a query, and are executed using DBMS_SQL. Invokers of the DBMS_SQL interface must ensure that no query contains SQL injection.

137.3 DBMS_PARALLEL_EXECUTE Constants

The DBMS_PARALLEL_EXECUTE package uses the constants described in these two tables.

Table 137-1 DBMS_PARALLEL_EXECUTE Constants - Chunk Status Value

Constant	Type	Value	Description
ASSIGNED	NUMBER	1	Chunk has been assigned for processing
PROCESSED	NUMBER	2	Chunk has been processed successfully
PROCESSED_WITH_ERROR	NUMBER	3	Chunk has been processed, but an error occurred during processing
UNASSIGNED	NUMBER	0	Chunk is unassigned

Table 137-2 DBMS_PARALLEL_EXECUTE Constants - Task Status Value

Constant	Type	Value	Description
CHUNKED	NUMBER	5	Table associated with the task has been chunked, but none of the chunk has been assigned for processing
CHUNKING	NUMBER	2	Table associated with the task is being chunked
CHUNKING_FAILED	NUMBER	3	Chunking failed
CRASHED	NUMBER	9	Only applicable if parallel execution is used, this occurs if a job secondary process crashes or if the database crashes during EXECUTE, leaving a chunk in ASSIGNED or UNASSIGNED state.
CREATED	NUMBER	1	The task has been created by the CREATE_TASK Procedure
FINISHED	NUMBER	7	All chunks processed without error
FINISHED_WITH_ERROR	NUMBER	8	All chunks processed, but with errors in some cases
NO_CHUNKS	NUMBER	4	Table associated with the task has no chunks created
PROCESSING	NUMBER	6	Part of the chunk assigned for processing, or which has been processed

Note:

Use constants instead of absolute values, because absolute values might change in future.

137.4 DBMS_PARALLEL_EXECUTE Views

The `DBMS_PARALLEL_EXECUTE` package uses the following views.

- DBA_PARALLEL_EXECUTE_CHUNKS
- DBA_PARALLEL_EXECUTE_TASKS
- USER_PARALLEL_EXECUTE_CHUNKS
- USER_PARALLEL_EXECUTE_TASKS

137.5 DBMS_PARALLEL_EXECUTE Exceptions

The following table lists the exceptions raised by `DBMS_PARALLEL_EXECUTE`.

Table 137-3 Exceptions Raised by DBMS_PARALLEL_EXECUTE

Exception	Error Code	Description
`CHUNK_NOT_FOUND`	29499	Specified chunk does not exist
`DUPLICATE_TASK_NAME`	29497	Same task name has been used by an existing task
`INVALID_STATE_FOR_CHUNK`	29492	Attempts to chunk a table that is not in `CREATED` or `CHUNKING_FAILED` state
`INVALID_STATE_FOR_REDSUME`	29495	Attempts to resume execution, but the task is not in `FINISHED_WITH_ERROR` or `CRASHED` state
`INVALID_STATE_FOR_RUN`	29494	Attempts to execute the task that is not in `CHUNKED` state
`INVALID_STATUS`	29493	Attempts to set an invalid value to the chunk status
`INVALID_TABLE`	29491	Attempts to chunk a table by rowid in cases in which the table is not a physical table, or the table is an IOT
`MISSING_ROLE`	29490	User does not have the necessary `ADM_PARALLEL_EXECUTE` role
`TASK_NOT_FOUND`	29498	Specified `task_name` does not exist

137.6 DBMS_PARALLEL_EXECUTE Examples

The following examples run on the Human Resources (HR) schema of the Oracle Database Sample Schemas. They requires that the HR schema be created with the `JOB SYSTEM` privilege.

Chunk by ROWID

This example shows the most common usage of this package. After calling the RUN_TASK Procedure, it checks for errors and reruns in the case of error.

```
DECLARE
  l_sql_stmt VARCHAR2(1000);
  l_try NUMBER;
  l_status NUMBER;
BEGIN

  -- Create the TASK
  DBMS_PARALLEL_EXECUTE.CREATE_TASK ('mytask');
```

```
  -- Chunk the table by ROWID
  DBMS_PARALLEL_EXECUTE.CREATE_CHUNKS_BY_ROWID('mytask', 'HR', 'EMPLOYEES', true, 100);

  -- Execute the DML in parallel
  l_sql_stmt := 'update EMPLOYEES e
      SET e.salary = e.salary + 10
      WHERE rowid BETWEEN :start_id AND :end_id';
  DBMS_PARALLEL_EXECUTE.RUN_TASK('mytask', l_sql_stmt, DBMS_SQL.NATIVE,
                                 parallel_level => 10);

  -- If there is an error, RESUME it for at most 2 times.
  L_try := 0;
  L_status := DBMS_PARALLEL_EXECUTE.TASK_STATUS('mytask');
  WHILE(l_try < 2 and L_status != DBMS_PARALLEL_EXECUTE.FINISHED)
  LOOP
    L_try := l_try + 1;
    DBMS_PARALLEL_EXECUTE.RESUME_TASK('mytask');
    L_status := DBMS_PARALLEL_EXECUTE.TASK_STATUS('mytask');
  END LOOP;

  -- Done with processing; drop the task
  DBMS_PARALLEL_EXECUTE.DROP_TASK('mytask');

END;
/
```

Chunk by User-Provided SQL

A user can specify a chunk algorithm by using the CREATE_CHUNKS_BY_SQL Procedure. This example shows that rows with the same `manager_id` are grouped together and processed in one chunk.

```
DECLARE
  l_chunk_sql VARCHAR2(1000);
  l_sql_stmt VARCHAR2(1000);
  l_try NUMBER;
  l_status NUMBER;
BEGIN

  -- Create the TASK
  DBMS_PARALLEL_EXECUTE.CREATE_TASK ('mytask');

  -- Chunk the table by MANAGER_ID
  l_chunk_sql := 'SELECT distinct manager_id, manager_id FROM employees';
  DBMS_PARALLEL_EXECUTE.CREATE_CHUNKS_BY_SQL('mytask', l_chunk_sql, false);

  -- Execute the DML in parallel
  --   the WHERE clause contain a condition on manager_id, which is the chunk
  --   column. In this case, grouping rows is by manager_id.
  l_sql_stmt := 'update EMPLOYEES e
      SET e.salary = e.salary + 10
      WHERE manager_id between :start_id and :end_id';
  DBMS_PARALLEL_EXECUTE.RUN_TASK('mytask', l_sql_stmt, DBMS_SQL.NATIVE,
                                  parallel_level => 10);

  -- If there is error, RESUME it for at most 2 times.
  L_try := 0;
  L_status := DBMS_PARALLEL_EXECUTE.TASK_STATUS('mytask');
  WHILE(l_try < 2 and L_status != DBMS_PARALLEL_EXECUTE.FINISHED)
  Loop
    L_try := l_try + 1;
    DBMS_PARALLEL_EXECUTE.RESUME_TASK('mytask');
```

```
    L_status := DBMS_PARALLEL_EXECUTE.TASK_STATUS('mytask');
  END LOOP;

  -- Done with processing; drop the task
  DBMS_PARALLEL_EXECUTE.DROP_TASK('mytask');

end;
/
```

Executing Chunks in an User-defined Framework

You can execute chunks in a self-defined framework without using the RUN_TASK Procedure. This example shows how to use GET_ROWID_CHUNK Procedure, `EXECUTE IMMEDIATE`, SET_CHUNK_STATUS Procedure to execute the chunks.

```
DECLARE
  l_sql_stmt varchar2(1000);
  l_try number;
  l_status number;
  l_chunk_id number;
  l_start_rowid rowid;
  l_end_rowid rowid;
  l_any_rows boolean;
  CURSOR c1 IS SELECT chunk_id
               FROM user_parallel_execute_chunks
               WHERE task_name = 'mytask'
                 AND STATUS IN (DBMS_PARALLEL_EXECUTE.PROCESSED_WITH_ERROR,
                                DBMS_PARALLEL_EXECUTE.ASSIGNED);
BEGIN

  -- Create the Objects, task, and chunk by ROWID
  DBMS_PARALLEL_EXECUTE.CREATE_TASK ('mytask');
  DBMS_PARALLEL_EXECUTE.CREATE_CHUNKS_BY_ROWID('mytask', 'HR', 'EMPLOYEES', true, 100);

  l_sql_stmt := 'update EMPLOYEES e
      SET e.salary = e.salary + 10
      WHERE rowid BETWEEN :start_id AND :end_id';

  -- Execute the DML in his own framework
  --
  -- Process each chunk and commit.
  -- After processing one chunk, repeat this process until
  -- all the chunks are processed.
  --
  <<main_processing>>
  LOOP
      --
      -- Get a chunk to process; if there is nothing to process, then exit the
      -- loop;
      --
      DBMS_PARALLEL_EXECUTE.GET_ROWID_CHUNK('mytask',
                                                          l_chunk_id,
                                                          l_start_rowid,
                                                          l_end_rowid,
                                                          l_any_rows);
      IF (l_any_rows = false) THEN EXIT; END IF;

      --
      -- The chunk is specified by start_id and end_id.
      -- Bind the start_id and end_id and then execute it
      --
      -- If no error occured, set the chunk status to PROCESSED.
```

```
    --
    -- Catch any exception. If an exception occured, store the error num/msg
    -- into the chunk table and then continue to process the next chunk.
    --
    BEGIN
      EXECUTE IMMEDIATE l_sql_stmt using l_start_rowid, l_end_rowid;
      DBMS_PARALLEL_EXECUTE.SET_CHUNK_STATUS('mytask',l_chunk_id,
        DBMS_PARALLEL_EXECUTE.PROCESSED);
    EXCEPTION WHEN OTHERS THEN
      DBMS_PARALLEL_EXECUTE.SET_CHUNK_STATUS('mytask', l_chunk_id,
        DBMS_PARALLEL_EXECUTE.PROCESSED_WITH_ERROR, SQLCODE, SQLERRM);
    END;

    --
    -- Finished processing one chunk; Commit here
    --
    COMMIT;
END LOOP;
```

137.7 Summary of DBMS_PARALLEL_EXECUTE Subprograms

This table lists the `DBMS_PARALLEL_EXECUTE` subprograms and briefly describes them.

Table 137-4 DBMS_PARALLEL_EXECUTE Package Subprograms

Subprogram	Description
ADM_DROP_CHUNKS Procedure	Drops all chunks of the specified task owned by the specified owner
ADM_DROP_TASK Procedure	Drops the task of the given user and all related chunks
ADM_TASK_STATUS Function	Returns the task status
ADM_STOP_TASK Procedure	Stops the task of the given owner and related job secondary processes
CREATE_TASK Procedure	Creates a task for the current user
CREATE_CHUNKS_BY_NUMBER_COL Procedure	Chunks the table associated with the given task by the specified column
CREATE_CHUNKS_BY_ROWID Procedure	Chunks the table associated with the given task by `ROWID`
CREATE_CHUNKS_BY_SQL Procedure	Chunks the table associated with the given task by means of a user-provided `SELECT` statement
DROP_TASK Procedure	Drops the task and all related chunks
DROP_CHUNKS Procedure	Drops the task's chunks
GENERATE_TASK_NAME Function	Returns a unique name for a task
GET_NUMBER_COL_CHUNK Procedure	Picks an unassigned `NUMBER` chunk and changes it to `ASSIGNED`
GET_ROWID_CHUNK Procedure	Picks an unassigned `ROWID` chunk and changes it to `ASSIGNED`
PURGE_PROCESSED_CHUNKS Procedure	Deletes all the processed chunks whose status is `PROCESSED` or `PROCESSED_WITH_ERROR`
RESUME_TASK Procedures	Retries the given the task if the RUN_TASK Procedure finished with an error, or resumes the task if a crash occurred.

Table 137-4 (Cont.) DBMS_PARALLEL_EXECUTE Package Subprograms

Subprogram	Description
RUN_TASK Procedure	Executes the specified SQL statement on the chunks in parallel
SET_CHUNK_STATUS Procedure	Sets the status of the chunk
STOP_TASK Procedure	Stops the task and related job secondary processes
TASK_STATUS Procedure	Returns the task status

137.7.1 ADM_DROP_CHUNKS Procedure

This procedure drops all chunks of the specified task owned by the specified owner.

Syntax

```
DBMS_PARALLEL_EXECUTE.ADM_DROP_CHUNKS (
   task_owner      IN  VARCHAR2,
   task_name       IN  VARCHAR2);
```

Parameters

Table 137-5 ADM_DROP_CHUNKS Procedure Parameters

Parameter	Description
task_owner	Owner of the task
task_name	Name of the task

137.7.2 ADM_DROP_TASK Procedure

This procedure drops the task of the specified user and all related chunks.

Syntax

```
DBMS_PARALLEL_EXECUTE.ADM_DROP_TASK (
   task_owner      IN  VARCHAR2,
   task_name       IN  VARCHAR2);
```

Parameters

Table 137-6 ADM_DROP_TASK Procedure Parameters

Parameter	Description
task_owner	Owner of the task
task_name	Name of the task

137.7.3 ADM_TASK_STATUS Function

This function returns the task status.

Syntax

```
DBMS_PARALLEL_EXECUTE.ADM_TASK_STATUS  (
   task_owner      IN  VARCHAR2,
   task_name       IN  VARCHAR2)
 RETURN NUMBER;
```

Parameters

Table 137-7 ADM_TASK_STATUS Function Parameters

Parameter	Description
task_owner	Owner of the task
task_name	Name of the task

137.7.4 ADM_STOP_TASK Procedure

This procedure stops the task of the specified owner and related job secondary processes.

Syntax

```
DBMS_PARALLEL_EXECUTE.ADM_STOP_TASK (
   task_owner      IN  VARCHAR2,
   task_name       IN  VARCHAR2);
```

Parameters

Table 137-8 ADM_STOP_TASK Procedure Parameters

Parameter	Description
task_owner	Owner of the task
task_name	Name of the task

137.7.5 CREATE_TASK Procedure

This procedure creates a task for the current user. The pairing of `task_name` and `current_user` must be unique.

Syntax

```
DBMS_PARALLEL_EXECUTE.CREATE_TASK (
   task_name        IN   VARCHAR2,
   comment          IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 137-9 CREATE_TASK Procedure Parameters

Parameter	Description
task_name	Name of the task. The task_name can be any string in which related length must be less than or equal to 128 bytes.
comment	Comment field. The comment must be less than 4000 bytes.

137.7.6 CREATE_CHUNKS_BY_NUMBER_COL Procedure

This procedure chunks the table (associated with the specified task) by the specified column. The specified column must be a NUMBER column. This procedure takes the MIN and MAX value of the column, and then divides the range evenly according to chunk_size.

The chunks are:

```
START_ID                          END_ID
--------------------------        --------------------------
min_id_val                        min_id_val+1*chunk_size-1
min_id_val+1*chunk_size           min_id_val+2*chunk_size-1
…                                 …
min_id_val+i*chunk_size           max_id_val
```

Syntax

```
DBMS_PARALLEL_EXECUTE.CREATE_CHUNKS_BY_NUMBER_COL (
   task_name       IN  VARCHAR2,
   table_owner     IN  VARCHAR2,
   table_name      IN  VARCHAR2,
   table_column    IN  VARCHAR2,
   chunk_size      IN  NUMBER);
```

Parameters

Table 137-10 CREATE_CHUNKS_BY_NUMBER_COL Procedure Parameters

Parameter	Description
task_name	Name of the task
table_owner	Owner of the table
table_name	Name of the table
table_column	Name of the NUMBER column
chunk_size	Range of each chunk

137.7.7 CREATE_CHUNKS_BY_ROWID Procedure

This procedure chunks the table (associated with the specified task) by ROWID.

num_row and num_block are approximate guidance for the size of each chunk. The table to be chunked must be a physical table with physical ROWID having views and table functions. Index-organized tables are not allowed.

Syntax

```
DBMS_PARALLEL_EXECUTE.CREATE_CHUNKS_BY_ROWID (
   task_name       IN  VARCHAR2,
   table_owner     IN  VARCHAR2,
   table_name      IN  VARCHAR2,
   by_row          IN  BOOLEAN,
   chunk_size      IN  NUMBER);
```

Parameters

Table 137-11 CREATE_CHUNKS_BY_ROWID Procedure Parameters

Parameter	Description
task_name	Name of the task
table_owner	Owner of the table
table_name	Name of the table
by_row	TRUE if chunk_size refers to the number of rows, otherwise, chunk_size refers to the number of blocks
chunk_size	Approximate number of rows/blocks to process for each commit cycle

137.7.8 CREATE_CHUNKS_BY_SQL Procedure

This procedure chunks the table (associated with the specified task) by means of a user-provided SELECT statement.

The SELECT statement that returns the range of each chunk must have two columns: start_id and end_id. If the task is to chunk by ROWID, then the two columns must be of ROWID type. If the task is to chunk the table by NUMBER column, then the two columns must be of NUMBER type. The procedure provides the flexibility to users who want to deploy user-defined chunk algorithms.

Syntax

```
DBMS_PARALLEL_EXECUTE.CREATE_CHUNKS_BY_SQL (
   task_name       IN  VARCHAR2,
   sql_stmt        IN  CLOB,
   by_rowid        IN  BOOLEAN);
```

Parameters

Table 137-12 CREATE_CHUNKS_BY_SQL Procedure Parameters

Parameter	Description
task_name	Name of the task
sql_stmt	SQL that returns the chunk ranges
by_rowid	TRUE if the table is chunked by rowids

137.7.9 DROP_TASK Procedure

This procedure drops the task and all related chunks.

Syntax

```
DBMS_PARALLEL_EXECUTE.DROP_TASK (
   task_name       IN VARCHAR2);
```

Parameters

Table 137-13 DROP_TASK Procedure Parameters

Parameter	Description
task_name	Name of the task

137.7.10 DROP_CHUNKS Procedure

This procedure drops the task's chunks.

Syntax

```
DBMS_PARALLEL_EXECUTE.DROP_CHUNKS (
   task_name       IN VARCHAR2);
```

Parameters

Table 137-14 DROP_CHUNKS Procedure Parameters

Parameter	Description
task_name	Name of the task

137.7.11 GENERATE_TASK_NAME Function

This function returns a unique name for a task.

The name is of the form *prefix*N where N is a number from a sequence. If no prefix is specified, the generated name is, by default, TASK$_1, TASK$_2, TASK$_3, and so on. If 'SCOTT' is specified as the prefix, the name is SCOTT1, SCOTT2, and so on.

Syntax

```
DBMS_PARALLEL_EXECUTE.GENERATE_TASK_NAME (
   prefix      IN      VARCHAR2 DEFAULT 'TASK$_')
 RETURN VARCHAR2;
```

Parameters

Table 137-15 GENERATE_TASK_NAME Function Parameters

Parameter	Description
prefix	The prefix to use when generating the task name

137.7.12 GET_NUMBER_COL_CHUNK Procedure

This procedure picks an unassigned `NUMBER` chunk and changes it to `ASSIGNED`. If there are no more chunks to assign, `any_rows` is set to `FALSE`. Otherwise, the `chunk_id`, `start`, and `end_id` of the chunk are returned as `OUT` parameters.

The chunk information in `DBMS_PARALLEL_EXECUTE_CHUNKS$` is updated as follows: `STATUS` becomes `ASSIGNED`; `START_TIMESTAMP` records the current time; `END_TIMESTAMP` is cleared.

> **See Also:**
>
> Views

Syntax

```
DBMS_PARALLEL_EXECUTE.GET_NUMBER_COL_CHUNK (
   task_name       IN VARCHAR2,
   chunk_id        OUT NUMBER,
   start_id        OUT NUMBER,
   end_id          OUT NUMBER,
   any_rows        OUT BOOLEAN);
```

Parameters

Table 137-16 GET_NUMBER_COL_CHUNK Procedure Parameters

Parameter	Description
`task_name`	Name of the task
`chunk_id`	Chunk ID of the chunk
`start_id`	ID of the start row in the returned range
`end_id`	ID of the end row in the returned range
`any_rows`	Indicates if there could be any rows to process in the range

Usage Notes

If the task is chunked by `ROWID`, then use `get_rowid_range`. If the task is chunked by `NUMBER` column, then use `get_number_col_range`. If you make the wrong function call, the returning `chunk_id` and `any_rows` have valid values but `start_id` and `end_id` are `NULL`.

137.7.13 GET_ROWID_CHUNK Procedure

This procedure picks an unassigned `ROWID` chunk and changes it to `ASSIGNED`.

If there are no more chunks to assign, `any_rows` is set to `FALSE`. Otherwise, the `chunk_id`, `start`, and `end_rowid` of the chunk are returned as `OUT` parameters. The chunk info in `DBMS_PARALLEL_EXECUTE_CHUNKS$` is updated as follows: `STATUS` becomes `ASSIGNED`; `START_TIMESTAMP` records the current time; `END_TIMESTAMP` is cleared.

> **See Also:**
>
> Views

Syntax

```
DBMS_PARALLEL_EXECUTE.GET_ROWID_CHUNK (
   task_name       IN VARCHAR2,
   chunk_id        OUT NUMBER,
   start_rowid     OUT ROWID,
   end_rowid       OUT ROWID,
   any_rows        OUT BOOLEAN);
```

Parameters

Table 137-17 GET_ROWID_CHUNK Procedure Parameters

Parameter	Description
`task_name`	Name of the task
`chunk_id`	Chunk ID of the chunk
`start_rowid`	Start rowid in the returned range
`end_rowid`	End rowid in the returned range
`any_rows`	Indicates that the range could include rows to process

Usage Notes

If the task is chunked by `ROWID`, then use `get_rowid_range`. If the task is chunked by `NUMBER` column, then use `get_number_col_range`. If you make the wrong function call, the returning `chunk_id` and `any_rows` will still have valid values but `start_id` and `end_id` are `NULL`.

137.7.14 PURGE_PROCESSED_CHUNKS Procedure

This procedure deletes all the processed chunks whose status is `PROCESSED` or `PROCESSED_WITH_ERROR`.

Syntax

```
DBMS_PARALLEL_EXECUTE.PURGE_PROCESSED_CHUNKS (
   task_name       IN VARCHAR2);
```

Parameters

Table 137-18 PURGE_PROCESSED_CHUNKS Procedure Parameters

Parameter	Description
`task_name`	Name of the task

137.7.15 RESUME_TASK Procedures

This procedure retries the specified the task if the RUN_TASK Procedure finished with an error, or resumes the task if a crash occurred.

You can only invoke this procedure if the task is in a `CRASHED` or `FINISHED_WITH_ERROR` state.

For a crashed serial execution, the state remains in `PROCESSING`. The `FORCE` option allows you to resume any task in `PROCESSING` state. However, it is your responsibility to determine that a crash has occurred.

The procedure resumes processing the chunks which have not been processed. Also, chunks which are in `PROCESSED_WITH_ERROR` or `ASSIGNED` (due to crash) state are processed because those chunks did not commit.

This procedure takes the same argument as the RUN_TASK Procedure. The overload which takes `task_name` as the only input argument re-uses the arguments provided in the previous invoking of the RUN_TASK Procedure or RESUME_TASK Procedures.

> **See Also:**
>
> Table 137-2

Syntax

```
DBMS_PARALLEL_EXECUTE.RESUME_TASK (
   task_name                      IN  VARCHAR2,
   sql_stmt                       IN  CLOB,
   language_flag                  IN  NUMBER,
   edition                        IN  VARCHAR2  DEFAULT NULL,
   apply_crossedition_trigger     IN  VARCHAR2  DEFAULT NULL,
   fire_apply_trigger             IN  BOOLEAN   DEFAULT TRUE,
   parallel_level                 IN  NUMBER    DEFAULT 0,
   job_class                      IN  VARCHAR2  DEFAULT 'DEFAULT_JOB_CLASS',
   force                          IN  BOOLEAN   DEFAULT FALSE);

DBMS_PARALLEL_EXECUTE.RESUME_TASK (
   task_name                      IN  VARCHAR2,
   force                          IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 137-19 RESUME_TASK Procedure Parameters

Parameter	Description
`task_name`	Name of the task
`sql_stmt`	SQL statement; must have `:start_id` and `:end_id` placeholders
`language_flag`	Determines how Oracle handles the SQL statement. The following options are recognized: • `V6` (or 0) specifies version 6 behavior • `NATIVE` (or 1) specifies normal behavior for the database to which the program is connected • `V7` (or 2) specifies Oracle database version 7 behavior

Table 137-19 (Cont.) RESUME_TASK Procedure Parameters

Parameter	Description
edition	Specifies the edition in which to run the statement. Default is the current edition.
apply_crossedition_trigger	Specifies the unqualified name of a forward crossedition trigger that is to be applied to the specified SQL. The name is resolved using the edition and current_schema setting in which the statement is to be executed. The trigger must be owned by the user who executes the statement.
fire_apply_trigger	Indicates whether the specified apply_crossedition_trigger is itself to be executed, or only to used as be a guide in selecting other triggers
parallel_level	Number of parallel jobs; zero if run in serial; NULL uses the default parallelism
job_class	If running in parallel, the jobs all belong to the specified job class
force	If TRUE, do not raise an error if the status is PROCESSING.

Examples

Suppose the chunk table contains the following chunk ranges:

```
START_ID                       END_ID
---------------------------    ---------------------------
1                              10
11                             20
21                             30
```

And the specified SQL statement is:

```
UPDATE employees
      SET salary = salary + 10
      WHERE e.employee_id  BETWEEN :start_id AND :end_id
```

This procedure executes the following statements in parallel:

```
UPDATE employees
      SET salary =.salary + 10  WHERE employee_id BETWEEN 1 and 10;
      COMMIT;

UPDATE employees
      SET salary =.salary + 10  WHERE employee_id between 11 and 20;
      COMMIT;

UPDATE employees
      SET salary =.salary + 10  WHERE employee_id between 21 and 30;
      COMMIT;
```

Related Topics

- RUN_TASK Procedure
  This procedure executes the specified statement (sql_stmt) on the chunks in parallel.

137.7.16 RUN_TASK Procedure

This procedure executes the specified statement (`sql_stmt`) on the chunks in parallel.

It commits after processing each chunk.

The specified statement must have two placeholders called `start_id` and `end_id`, respectively, which represent the range of the chunk to be processed. The type of each placeholder must be `ROWID` where `ROWID`-based chunking was used, or `NUMBER` where `NUMBER`-based chunking was used.

Syntax

```
DBMS_PARALLEL_EXECUTE.RUN_TASK (
   task_name                      IN  VARCHAR2,
   sql_stmt                       IN  CLOB,
   language_flag                  IN  NUMBER,
   edition                        IN  VARCHAR2  DEFAULT NULL,
   apply_crossedition_trigger     IN  VARCHAR2  DEFAULT NULL,
   fire_apply_trigger             IN  BOOLEAN   DEFAULT TRUE,
   parallel_level                 IN  NUMBER    DEFAULT 0,
   job_class                      IN  VARCHAR2  DEFAULT 'DEFAULT_JOB_CLASS');
```

Parameters

Table 137-20 RUN_TASK Procedure Parameters

Parameter	Description
`task_name`	Name of the task
`sql_stmt`	SQL statement; must have `:start_id` and `:end_id` placeholders
`language_flag`	Determines how Oracle handles the SQL statement. The following options are recognized: • `V6` (or 0) specifies version 6 behavior • `NATIVE` (or 1) specifies normal behavior for the database to which the program is connected • `V7` (or 2) specifies Oracle database version 7 behavior
`edition`	Specifies the edition in which to run the statement. Default is the current edition.
`apply_crossedition_trigger`	Specifies the unqualified name of a forward crossedition trigger that is to be applied to the specified SQL. The name is resolved using the edition and `current_schema` setting in which the statement is to be executed. The trigger must be owned by the user executes the statement.
`fire_apply_trigger`	Indicates whether the specified `apply_crossedition_trigger` is itself to be executed, or only a guide to be used in selecting other triggers.
`parallel_level`	Number of parallel jobs; zero if run in serial; `NULL` uses the default parallelism.
`job_class`	If running in parallel, the jobs belong to the specified job class

Usage Notes

- The SQL statement is executed as the current user.

- Since this subprogram is subject to reexecution on error, you need to take great care in submitting a statement to `RUN_TASK` that is not idempotent.
- Chunks can be executed in parallel by `DBMS_SCHEDULER` job secondary processes. Therefore, parallel execution requires the `CREATE JOB` system privilege. The job secondary processes are created under the current user. The default number of job secondary processes is computed as the product of the Oracle parameters `cpu_count` and `parallel_threads_per_cpu`. On a Real Application Clusters installation, the number of job secondary processes is the sum of individual settings on each node in the cluster. This procedure returns only when all the chunks are processed. In parallel cases, this procedure returns only when all the secondary processes are finished.

Examples

Suppose the chunk table contains the following chunk ranges:

```
START_ID                          END_ID
--------------------------        --------------------------
1                                 10
11                                20
21                                30
```

And the specified SQL statement is:

```
UPDATE employees
      SET salary = salary + 10
      WHERE e.employee_id  BETWEEN :start_id AND :end_id
```

This procedure executes the following statements in parallel:

```
UPDATE employees
      SET salary =.salary + 10  WHERE employee_id BETWEEN 1 and 10;
      COMMIT;

UPDATE employees
      SET salary =.salary + 10  WHERE employee_id between 11 and 20;
      COMMIT;

UPDATE employees
      SET salary =.salary + 10  WHERE employee_id between 21 and 30;
      COMMIT;
```

137.7.17 SET_CHUNK_STATUS Procedure

This procedure sets the status of the chunk.

The `START_TIMESTAMP` and `END_TIMESTAMP` of the chunk is updated according to the new status:

```
Value of the new Status              Side Effect
--------------------------           --------------------------
UNASSIGNED                           START_TIMESTAMP and END_TIMESTAMP
                                     will be cleared
ASSIGNED                             START_TIMESTAMP will be the current time
                                     and END_TIMESTAMP will be cleared.
PROCESSED or PROCESSED_WITH_ERROR    The current time will be recorded
                                     in END_TIMESTAMP
```

> **See Also:**
>
> Views

Syntax

```
DBMS_PARALLEL_EXECUTE.SET_CHUNK_STATUS (
   task_name       IN VARCHAR2,
   chunk_id        OUT NUMBER,
   status          IN  NUMBER,
   err_num         IN  NUMBER   DEFAULT NULL,
   err_msg         IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 137-21 SET_CHUNK_STATUS Procedure Parameters

Parameter	Description
task_name	Name of the task
chunk_id	Chunk_id of the chunk
status	Status of the chunk: UNASSIGNED, ASSIGNED, PROCESSED PROCESSED_WITH_ERROR
err_num	Error code returned during the processing of the chunk
err_msg	Error message returned during the processing of the chunk

137.7.18 STOP_TASK Procedure

This procedure stops the task and related secondary processes.

Syntax

```
DBMS_PARALLEL_EXECUTE.STOP_TASK (
   task_name       IN VARCHAR2);
```

Parameters

Table 137-22 STOP_TASK Procedure Parameters

Parameter	Description
task_name	Name of the task

137.7.19 TASK_STATUS Procedure

This procedure returns the task status.

Syntax

```
DBMS_PARALLEL_EXECUTE.TASK_STATUS (
   task_name       IN VARCHAR2);
```

Parameters

Table 137-23 TASK_STATUS Procedure Parameters

Parameter	Description
`task_name`	Name of the task

138
DBMS_PART

The `DBMS_PART` package provides an interface for maintenance and management operations on partitioned objects.

> **See Also:**
>
> - *Oracle Database Reference* for related views

This chapter contains the following topics:

- Security Model
- Operational Notes
- Summary of DBMS_PART Subprograms

138.1 DBMS_PART Security Model

`DBMS_PART` is an invoker's rights package, running with the privileges of the user.

138.2 DBMS_PART Operational Notes

The following operational notes apply to DBMS_PART.

- `DBMS_PART` ignores all the errors that it runs into during the cleanup process.
- To display the message `PL/SQL procedure executed successfully` requires at least one cleanup operation to be successful.

138.3 Summary of DBMS_PART Subprograms

This table briefly describes the subprograms of `DBMS_PART` package.

Table 138-1 DBMS_PART Package Subprograms

Subprogram	Description
CLEANUP_GIDX Procedure	Gathers the list of global indexes where optimized asynchronous index maintenance has taken place to clean up entries pointing to data segments that no longer exist
CLEANUP_GIDX_JOB Procedure	
CLEANUP_ONLINE_OP Procedure	Cleans up failed online move operations

138.3.1 CLEANUP_GIDX Procedure

As a consequence of prior partition maintenance operations with asynchronous global index maintenance, global indexes can contain entries pointing to data segments that no longer exist. These stale index rows will not cause any correctness issues or corruptions during any operation on the table or index, whether these are queries, DMLs, DDLs or analyze. This procedure will identify and cleanup these global indexes to ensure efficiency in terms of storage and performance.

Syntax

```
DBMS_PART.CLEANUP_GIDX (
   schema_name_in    IN   VARCHAR2 DEFAULT NULL,
   table_name_in     IN   VARCHAR2 DEFAULT NULL,
   parallel          IN   VARCHAR2 DEFAULT NULL,
   options           IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 138-2 CLEANUP_GIDX Function Parameters

Parameter	Description
schema_name_in	Non-NULL processes only indexes on tables in the given schema
table_name_in	Non-NULL processes only indexes on the given table in the given schema (schema_name_in must be non-NULL if table_name_in is non-NULL)
parallel	The parallel degree to use for the ALTER INDEX DDLs.
options	The following options are supported: • CLEANUP_ORPHANS: implies that 'cleanup only' mechanism is used. • COALESCE: implies that 'coalesce cleanup' mechanism is used.

138.3.2 CLEANUP_GIDX_JOB Procedure

This procedure will identify and cleanup these global indexes to ensure efficiency in terms of storage and performance.

Syntax

```
DBMS_PART.CLEANUP_GIDX_JOB (
   parallel          IN   VARCHAR2 DEFAULT NULL,
   options           IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 138-3 CLEANUP_GIDX_JOB Function Parameters

Parameter	Description
parallel	The parallel degree to use for the ALTER INDEX DDLs.
options	The following options are supported: • CLEANUP_ORPHANS: implies that 'cleanup only' mechanism is used. • COALESCE: implies that coalesce cleanup mechanism is used.

138.3.3 CLEANUP_ONLINE_OP Procedure

There are many possible points of failure when performing `ALTER TABLE ... MOVE PARTITION ... ONLINE` operations. This procedure pro-actively cleans up such failed online move operations instead of waiting for the background process (`SMON`) to do so.

Syntax

```
DBMS_PART.CLEANUP_ONLINE_OP (
   schema_name       IN   VARCHAR2 DEFAULT NULL,
   table_name        IN   VARCHAR2 DEFAULT NULL,
   partition_name    IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 138-4 CLEANUP_ONLINE_OP Function Parameters

Parameter	Description
`schema_name`	Name of schema
`table_name`	Name of schema
`partition_name`	Name of partition

Usage Notes

- If `schema_name`, `table_name` and `partition_name` are specified, this cleans up the failed online move operation for the specified partition.
- If `schema_name` and `table_name` are specified, this cleans up all failed online move operations for all the partitions of the specified table.
- If only `schema_name` is specified, this cleans up all failed online move operations in the schema.
- If no arguments are provided, we cleans up all the failed online move operations in the system.
- All other cases raise `ORA-20000` to inform the user of invalid inputs as arguments.

139
DBMS_PCLXUTIL

The DBMS_PCLXUTIL package provides intra-partition parallelism for creating partition-wise local indexes. DBMS_PCLXUTIL circumvents the limitation that, for local index creation, the degree of parallelism is restricted to the number of partitions as only one parallel execution server process for each partition is used.

> **See Also:**
>
> There are several rules concerning partitions and indexes. For more information, see *Oracle Database Concepts* and*Oracle Database Administrator's Guide*.

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Rules and Limits
- Summary of DBMS_PCLXUTIL Subprograms

139.1 DBMS_PCLXUTIL Overview

DBMS_PCLXUTIL uses the DBMS_JOB package to provide a greater degree of parallelism for creating a local index for a partitioned table. This is achieved by asynchronous inter-partition parallelism using the background processes (with DBMS_SCHEDULER), in combination with intra-partition parallelism using the parallel execution server.

DBMS_PCLXUTIL works with both range and range-hash composite partitioning.

The DBMS_PCLXUTIL package can be used during the following DBA tasks:

1. Local index creation

   The procedure BUILD_PART_INDEX assumes that the dictionary information for the local index already exists. This can be done by issuing the create index SQL command with the UNUSABLE option.

   ```
   CREATE INDEX <idx_name> on <tab_name>(...) local(...) unusable;
   ```

   This causes the dictionary entries to be created without "building" the index itself, the time consuming part of creating an index. Now, invoking the procedure BUILD_PART_INDEX causes a concurrent build of local indexes with the specified degree of parallelism.

   ```
   EXECUTE dbms_pclxutil.build_part_index(4,4,<tab_name>,<idx_name>,FALSE);
   ```

   For composite partitions, the procedure automatically builds local indexes for all subpartitions of the composite table.

2. Local index maintenance

   By marking desired partitions usable or unusable, the `BUILD_PART_INDEX` procedure also enables selective rebuilding of local indexes. The `force_opt` parameter provides a way to override this and build local indexes for all partitions.

   ```
   ALTER INDEX <idx_name> local(...) unusable;
   ```

   Rebuild only the desired (sub)partitions (that are marked unusable):

   ```
   EXECUTE dbms_pclxutil.build_part_index(4,4,<tab_name>,<idx_name>,FALSE);
   ```

   Rebuild all (sub)partitions using `force_opt` = `TRUE`:

   ```
   EXECUTE dbms_pclxutil.build_part_index(4,4,<tab_name>,<idx_name>,TRUE);
   ```

   A progress report is produced, and the output appears on screen when the program is ended (because the `DBMS_OUTPUT` package writes messages to a buffer first, and flushes the buffer to the screen only upon termination of the program).

139.2 DBMS_PCLXUTIL Security Model

This utility can be run only as table owner, and not as any other user.

139.3 DBMS_PCLXUTIL Operational Notes

`DBMS_PCLXUTIL` submits a job for each partition. It is the responsibility of the user/dba to control the number of concurrent jobs by setting the `INIT.ORA` parameter `JOB_QUEUE_PROCESSES` correctly. There is minimal error checking for correct syntax. Any errors are reported in the job queue process trace files.

139.4 DBMS_PCLXUTIL Rules and Limits

Because `DBMS_PCLXUTIL` uses the `DBMS_JOB` package, you must be aware of the following limitations pertaining to `DBMS_JOB`:

- You must decide appropriate values for the `job_queue_processes` initialization parameter. Clearly, if the job processes are not started before calling `BUILD_PART_INDEX`(), then the package will not function properly. The background processes are specified by the following `init.ora` parameters:

  ```
  job_queue_processes=n   #the number of background processes = n
  ```

- Failure conditions are reported only in the trace files (a `DBMS_JOB` limitation), making it impossible to give interactive feedback to the user. This package prints a failure message, removes unfinished jobs from the queue, and requests the user to take a look at the `j*.trc` trace files.

> **Note:**
>
> For range partitioning, the minimum compatibility mode is 8.0; for range-hash composite partitioning, the minimum compatibility mode is 8*i*.

139.5 Summary of DBMS_PCLXUTIL Subprograms

The DBMS_PCLXUTIL package has one subprogram, the BUILD_PART_INDEX procedure.

Table 139-1 DBMS_PCLXUTIL Package Subprograms

Subprogram	Description
BUILD_PART_INDEX Procedure	Provides intra-partition parallelism for creating partition-wise local indexes

139.5.1 BUILD_PART_INDEX Procedure

This procedure provides intra-partition parallelism for creating partition-wise local indexes.

Syntax

```
DBMS_PCLXUTIL.BUILD_PART_INDEX (
   jobs_per_batch  IN NUMBER   DEFAULT 1,
   procs_per_job   IN NUMBER   DEFAULT 1,
   tab_name        IN VARCHAR2 DEFAULT NULL,
   idx_name        IN VARCHAR2 DEFAULT NULL,
   force_opt       IN BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 139-2 BUILD_PART_INDEX Procedure Parameters

Parameter	Description
jobs_per_batch	The number of concurrent partition-wise "local index builds".
procs_per_job	The number of parallel execution servers to be utilized for each local index build (1 <= procs_per_job <= max_slaves).
tab_name	The name of the partitioned table (an exception is raised if the table does not exist or not partitioned).
idx_name	The name given to the local index (an exception is raised if a local index is not created on the table tab_name).
force_opt	If TRUE, then force rebuild of all partitioned indexes; otherwise, rebuild only the partitions marked 'UNUSABLE'.

Usage Notes

This utility can be run only as table owner, and not as any other user.

Examples

Suppose a table PROJECT is created with two partitions PROJ001 and PROJ002, along with a local index IDX.

A call to the procedure BUILD_PART_INDEX(2,4,'PROJECT','IDX',TRUE) produces the following output:

```
SQLPLUS> EXECUTE dbms_pclxutil.build_part_index(2,4,'PROJECT','IDX',TRUE);
Statement processed.
```

```
INFO: Job #21 created for partition PROJ002 with 4 slaves
INFO: Job #22 created for partition PROJ001 with 4 slaves
```

140
DBMS_PDB

The `DBMS_PDB` package provides an interface to examine and manipulate data about pluggable databases (PDBs) in a multitenant container database (CDB). It also contains an interface that specifies which database objects are application common objects. You can migrate a PDB from one CDB to another CDB. After the migration is complete, all backups of the PDB before the migration are available and usable in the destination CDB.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

> **See Also:**
>
> - *Oracle Database Administrator's Guide* for information about creating and managing PDBs and CDBs
> - *Oracle XML DB Developer's Guide* for information about configuring protocol ports and DNS mappings
> - *Oracle Database SQL Language Reference* for information about creating PDBs
> - *Oracle Database Security Guide* regarding how to create audit policies in a multitenant environment

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_PDB Subprograms

140.1 DBMS_PDB Overview

A multitenant container database (CDB) is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs). The `DBMS_PDB` package provides an interface to examine and manipulate data about pluggable databases (PDBs). The subprograms in this package can also set a database object to one of the following types of application common

objects in an application container: data-linked object, extended data-linked object, or metadata-linked object.

A data-linked application common object stores data in the application root that can be accessed and modified by all of the application PDBs in the application container. For an extended data-linked object, each application PDB can create its own specific data while sharing the common data in the application root. Therefore, with an extended data-linked object, only the data stored in the application root is common for all application PDBs. A metadata-linked application common object stores the metadata for the specific object, such as a table, in the application root, so that the containers in the application container have the same structure for the object but different data. This package also contains a procedure that specifies that a database object is not an application common object.

Typically, the subprograms in this package are used in the following cases:

- An application that is installed in a PDB is migrated to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.
- An application is installed in an application root using an installation program that does not allow setting application common objects.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

140.2 DBMS_PDB Security Model

Users must have the `EXECUTE` privilege to run the procedures of `DBMS_PDB` package.

140.3 Summary of DBMS_PDB Subprograms

This table lists and describes `DBMS_PDB` package subprograms.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

Table 140-1 DBMS_PDB Package Subprograms

Subprogram	Description
CHECK_PLUG_COMPATIBILITY Function	Uses an XML file describing a pluggable database (PDB) to determine whether it may be plugged into a given multitenant container database (CDB).
DESCRIBE Procedure	Generates an XML file describing the specified pluggable database (PDB).
EXPORTRMANBACKUP Procedure	Needs to be called only when a non-CDB is migrated as PDB. This procedure needs to be executed in non-CDB database. For PDb relocation, it is automatically called during unplugin.
RECOVER Procedure	Generates an XML file describing a pluggable database by using data files belonging to the pluggable database (PDB).
REMOVE_LINK Procedure	Specifies that a database object is not an application common object.
SET_DATA_LINKED Procedure	Sets a database object to a data-linked application common object
SET_EXT_DATA_LINKED Procedure	Sets a database object to an extended data-linked application common object.
SET_METADATA_LINKED Procedure	Sets a database object to a metadata-linked application common object.
SET_PROFILE_EXPLICIT Procedure	Sets a profile as an application common profile in an application container.
SET_ROLE_EXPLICIT Procedure	Sets a role as an application common role in an application container.
SET_USER_EXPLICIT Procedure	This procedure sets a local user as an application common user in an application container.

140.3.1 CHECK_PLUG_COMPATIBILITY Function

This function uses an XML file describing a pluggable database (PDB) to determine whether it may be plugged into a given multitenant container database (CDB).

Syntax

```
DBMS_PDB.CHECK_PLUG_COMPATIBILITY (
   pdb_descr_file    IN   VARCHAR2,
   pdb_name          IN   VARCHAR2 DEFAULT NULL)
 RETURN BOOLEAN;
```

Parameters

Table 140-2 *CHECK_PLUG_COMPATIBILITY Procedure Parameters*

Parameter	Description
pdb_descr_file	Path of the XML file that will contain description of a PDB
pdb_name	Name which will be given to the PDB represented by pdb_descr_file when plugged into a given CDB. If not specified, the name will be extracted from pdb_descr_file.

Return Values

`TRUE` if the PDB described by `pdb_descr_file` is compatible with the given CDB, `FALSE` otherwise. If this function returns `FALSE`, then query the `PDB_PLUG_IN_VIOLATIONS` data dictionary view to find information about the errors that are found.

> **See Also:**
>
> *Oracle Database Reference* for information about the `PDB_PLUG_IN_VIOLATIONS` view

140.3.2 DESCRIBE Procedure

This procedure generates an XML file describing the specified pluggable database (PDB). This file can then be passed to the CHECK_PLUG_COMPATIBILITY Function to determine if the PDB described by the XML file may be plugged into a given multitenant container database (CDB).

Syntax

```
DBMS_PDB.DESCRIBE (
   pdb_descr_file    IN   VARCHAR2,
   pdb_name          IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 140-3 *DESCRIBE Procedure Parameters*

Parameter	Description
`pdb_descr_file`	Path of the XML file that will contain description of a PDB
`pdb_name`	Name of a PDB to be described. A remote database is specified by including `@dblink`.

Usage Notes

- If `pdb_name` is omitted, the PDB to which the session is connected will be described.
- If `pdb_name` is omitted, and the session is connected to the Root, an error will be returned.

Related Topics

- CHECK_PLUG_COMPATIBILITY Function
  This function uses an XML file describing a pluggable database (PDB) to determine whether it may be plugged into a given multitenant container database (CDB).

140.3.3 EXPORTRMANBACKUP Procedure

This procedure exports the RMAN backup information that belongs to a pluggable database to its dictionary before unplug so that pre-plugin backups can be used. This procedure needs to

be called only for non-CDB to PDB migration. For PDB relocation, this procedure need not be called.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

Syntax

```
DBMS_PDB.EXPORTRMANBACKUP (
   pdb_name      IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 140-4 EXPORTRMANBACKUP Procedure Parameters

Parameter	Description
pdb_name	Name of a pluggable database whose backup information needs to be exported. Omit this parameter if you are connected to a pluggable database or a non-cdb.

Usage Notes

- The PDB must be opened in `read/write` mode.
- If the database is non-CDB, then `pdb_name` must be omitted.
- If the `pdb_name` is omitted, then the pluggable database to which the session is connected is exported.
- If the `pdb_name` is omitted and the session is connected to the `root`, an error is returned.

140.3.4 RECOVER Procedure

This procedure generates an XML file describing a pluggable database by using data files belonging to the pluggable database. This XML file can then be used to plug the pluggable database into a multitenant container database (CDB) using the `CREATE PLUGGABLE DATABASE` statement.

Use this procedure when an XML file describing a pluggable database is corrupted or lost.

Syntax

```
DBMS_PDB.RECOVER (
   pdb_descr_file    IN   VARCHAR2,
   pdb_name          IN   VARCHAR2,
   filenames         IN   VARCHAR2);
```

Parameters

Table 140-5 *RECOVER Procedure Parameters*

Parameter	Description
pdb_descr_file	Path of the XML file that contains description of a pluggable database
pdb_name	Name of a pluggable database
filenames	Comma-separated list of datafile paths and/or directories containing datafiles for the pluggable database

140.3.5 REMOVE_LINK Procedure

This procedure specifies that a database object is not an application common object. In an application container, application common objects are shared between multiple containers.

Syntax

```
DBMS_PDB.REMOVE_LINK (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 140-6 *REMOVE_LINK Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.
edition_name	The name of the edition for the database object.

See Also:

Oracle Database Administrator's Guide for information about migrating an application to an application container

140.3.6 SET_DATA_LINKED Procedure

This procedure sets a database object to a data-linked application common object. In an application container, data-linked application common objects store data in the application root only, and the data can be accessed by all of the application PDBs in the application container.

The data in a data-linked application common object can be modified only in the application root.

You can use this procedure to set data-linked application common objects when you migrate an application that is installed in a PDB to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.

Syntax

```
DBMS_PDB.SET_DATA_LINKED (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 140-7 *SET_DATA_LINKED Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.
edition_name	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

140.3.7 SET_EXT_DATA_LINKED Procedure

This procedure sets a database object to an extended data-linked application common object. In an application container, for an extended data-linked object, each application PDB can create its own specific data while sharing the common data in the application root. Therefore, only the data stored in the application root is common for all application PDBs.

You can use this procedure to set extended data-linked application common objects when you migrate an application that is installed in a PDB to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.

Syntax

```
DBMS_PDB.SET_EXT_DATA_LINKED (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
```

```
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 140-8 *SET_EXT_DATA_LINKED Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.
edition_name	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

140.3.8 SET_METADATA_LINKED Procedure

This procedure sets a database object to a metadata-linked application common object. In an application container, metadata-linked application common objects store the metadata for specific objects, such as tables, so that the containers that share the application common object have the same structure but different data.

You can use this procedure to set metadata-linked application common objects when you migrate an application that is installed in a PDB to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.

Syntax

```
DBMS_PDB.SET_METADATA_LINKED (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 140-9 *SET_METADATA_LINKED Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.

Table 140-9 (Cont.) *SET_METADATA_LINKED Procedure Parameters*

Parameter	Description
`namespace`	The namespace of the database object. The `NAMESPACE` column of the `DBA_OBJECTS` view shows the namespace of an object.
`edition_name`	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

140.3.9 SET_PROFILE_EXPLICIT Procedure

This procedure sets a profile as an application common profile in an application container. This procedure is intended for migrating a profile from a previous release to an application container in the current release.

This procedure must be invoked in an application install, patch, upgrade, or uninstall operation in an application root.

Syntax

```
DBMS_PDB.SET_PROFILE_EXPLICIT (
   profile_name IN VARCHAR2);
```

Parameters

Table 140-10 *SET_PROFILE_EXPLICIT Procedure Parameters*

Parameter	Description
`profile_name`	The name of the profile.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container.

140.3.10 SET_ROLE_EXPLICIT Procedure

This procedure sets a role as an application common role in an application container. This procedure is intended for migrating a role from a previous release to an application container in the current release.

This procedure must be invoked in an application install, patch, upgrade, or uninstall operation in an application root.

Syntax

```
DBMS_PDB.SET_ROLE_EXPLICIT (
   role_name IN VARCHAR2);
```

Parameters

Table 140-11 *SET_ROLE_EXPLICIT Procedure Parameters*

Parameter	Description
role_name	The name of the role.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container.

140.3.11 SET_USER_EXPLICIT Procedure

This procedure sets a local user as an application common user in an application container.

This procedure must be invoked in an application install, patch, upgrade, or uninstall operation in an application root.

Syntax

```
DBMS_PDB.SET_USER_EXPLICIT (
   user_name IN VARCHAR2);
```

Parameters

Table 140-12 *SET_USER_EXPLICIT Procedure Parameters*

Parameter	Description
user_name	The name of the user.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container.

141
DBMS_PDB_ALTER_SHARING

In an application container with a pre-installed application, the `DBMS_PDB_ALTER_SHARING` package provides an interface to set database objects as application common objects or to specify that a database object is not an application common object.

> **Note:**
>
> This package provides a subset of the subprograms in the `DBMS_PDB` package. Users who do not have `EXECUTE` privilege on the `DBMS_PDB` package can be granted `EXECUTE` privilege on this package to run these subprograms.

> **See Also:**
>
> - *Oracle Database Administrator's Guide* for information about migrating an application to an application container
> - *Oracle XML DB Developer's Guide* for information about configuring protocol ports and DNS mappings
> - *Oracle Database SQL Language Reference* for information about creating PDBs
> - *Oracle Database Security Guide* regarding how to create audit policies in a multitenant environment

This chapter contains the following topics:

- DBMS_PDB_ALTER_SHARING Overview
- DBMS_PDB_ALTER_SHARING Security Model
- Summary of DBMS_PDB_ALTER_SHARING Subprograms

141.1 DBMS_PDB_ALTER_SHARING Overview

The subprograms in this package can set a database object to one of the following types of application common objects in an application container: data-linked object, extended data-linked object, or metadata-linked object.

A data-linked application common object stores data in the application root that can be accessed and modified by all of the application PDBs in the application container. For an extended data-linked object, each application PDB can create its own specific data while sharing the common data in the application root. Therefore, with an extended data-linked object, only the data stored in the application root is common for all application PDBs. A metadata-linked application common object stores the metadata for the specific object, such as a table, in the application root, so that the containers in the application container have the same structure for the object but different data. This package also contains a procedure that specifies that a database object is not an application common object.

You can use the subprograms in this package when you

Typically, the subprograms in this package are used in the following cases:

- An application that is installed in a PDB is migrated to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.
- An application is installed in an application root using an installation program that does not allow setting application common objects.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

141.2 DBMS_PDB_ALTER_SHARING Security Model

This package provides a subset of the subprograms in the `DBMS_PDB` package. User who do not have `EXECUTE` privilege on the `DBMS_PDB` package can be granted `EXECUTE` privilege on this package to run these subprograms.

Security on this package can be controlled in either of the following ways:

- Granting `EXECUTE` on this package to selected users or roles.
- Granting `EXECUTE_CATALOG_ROLE` to selected users or roles.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

141.3 Summary of DBMS_PDB_ALTER_SHARING Subprograms

This table lists and briefly describes `DBMS_PDB_ALTER_SHARING` package subprograms.

Table 141-1 DBMS_PDB_ALTER_SHARING Subprograms

Subprogram	Description
REMOVE_LINK Procedure	Specifies that a database object is not an application common object
SET_DATA_LINKED Procedure	Sets a database object to a data-linked application common object
SET_EXT_DATA_LINKED Procedure	Sets a database object to an extended data-linked application common object
SET_METADATA_LINKED Procedure	Sets a database object to a metadata-linked application common object
SET_PROFILE_EXPLICIT Procedure	Sets a profile as an application common profile in an application container.

Table 141-1 (Cont.) DBMS_PDB_ALTER_SHARING Subprograms

Subprogram	Description
SET_ROLE_EXPLICIT Procedure	Sets a role as an application common role in an application container.
SET_USER_EXPLICIT Procedure	Sets a local user as an application common user in an application container.

141.3.1 REMOVE_LINK Procedure

This procedure specifies that a database object is not an application common object. In an application container, application common objects are shared between multiple containers.

Syntax

```
DBMS_PDB_ALTER_SHARING.REMOVE_LINK (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 141-2 *REMOVE_LINK Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.
edition_name	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

141.3.2 SET_DATA_LINKED Procedure

This procedure sets a database object to a data-linked application common object. In an application container, data-linked application common objects store data in the application root only, and the data can be accessed by all of the application PDBs in the application container. The data in a data-linked application common object can be modified only in the application root.

You can use this procedure to set data-linked application common objects when you migrate an application that is installed in a PDB to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an

application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.

Syntax

```
DBMS_PDB_ALTER_SHARING.SET_DATA_LINKED (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 141-3 *SET_DATA_LINKED Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.
edition_name	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

141.3.3 SET_EXT_DATA_LINKED Procedure

This procedure sets a database object to an extended data-linked application common object. In an application container, for an extended data-linked object, each application PDB can create its own specific data while sharing the common data in the application root. Therefore, only the data stored in the application root is common for all application PDBs.

You can use this procedure to set extended data-linked application common objects when you migrate an application that is installed in a PDB to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.

Syntax

```
DBMS_PDB_ALTER_SHARING.SET_EXT_DATA_LINKED (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 141-4 *SET_EXT_DATA_LINKED Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.
edition_name	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

141.3.4 SET_METADATA_LINKED Procedure

This procedure sets a database object to a metadata-linked application common object. In an application container, metadata-linked application common objects store the metadata for specific objects, such as tables, so that the containers that share the application common object have the same structure but different data.

You can use this procedure to set metadata-linked application common objects when you migrate an application that is installed in a PDB to an application container. The application can be migrated to the application root or to an application PDB. For example, you can migrate an application installed in a PDB plugged into an Oracle Database 12*c* Release 1 (12.1) CDB to an application container in an Oracle Database 12*c* Release 2 (12.2) CDB.

Syntax

```
DBMS_PDB_ALTER_SHARING.SET_METADATA_LINKED (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   namespace    IN NUMBER,
   edition_name IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 141-5 *SET_METADATA_LINKED Procedure Parameters*

Parameter	Description
schema_name	The name of the schema that owns the database object.
object_name	The name of the database object.
namespace	The namespace of the database object. The NAMESPACE column of the DBA_OBJECTS view shows the namespace of an object.

Table 141-5 (Cont.) *SET_METADATA_LINKED Procedure Parameters*

Parameter	Description
edition_name	The name of the edition for the database object.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container

141.3.5 SET_PROFILE_EXPLICIT Procedure

This procedure sets a profile as an application common profile in an application container. This procedure is intended for migrating a profile from a previous release to an application container in the current release.

This procedure must be invoked in an application install, patch, upgrade, or uninstall operation in an application root.

Syntax

```
DBMS_PDB_ALTER_SHARING.SET_PROFILE_EXPLICIT (
   profile_name IN VARCHAR2);
```

Parameters

Table 141-6 *SET_PROFILE_EXPLICIT Procedure Parameters*

Parameter	Description
profile_name	The name of the profile.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container.

141.3.6 SET_ROLE_EXPLICIT Procedure

This procedure sets a role as an application common role in an application container. This procedure is intended for migrating a role from a previous release to an application container in the current release.

This procedure must be invoked in an application install, patch, upgrade, or uninstall operation in an application root.

Syntax

```
DBMS_PDB_ALTER_SHARING.SET_ROLE_EXPLICIT (
   role_name IN VARCHAR2);
```

Parameters

Table 141-7 *SET_ROLE_EXPLICIT Procedure Parameters*

Parameter	Description
role_name	The name of the role.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container.

141.3.7 SET_USER_EXPLICIT Procedure

This procedure sets a local user as an application common user in an application container.

This procedure must be invoked in an application install, patch, upgrade, or uninstall operation in an application root.

Syntax

```
DBMS_PDB_ALTER_SHARING.SET_USER_EXPLICIT (
   user_name IN VARCHAR2);
```

Parameters

Table 141-8 *SET_USER_EXPLICIT Procedure Parameters*

Parameter	Description
user_name	The name of the user.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about migrating an application to an application container.

142
DBMS_PERF

The DBMS_PERF package provides and interface to generate active reports for monitoring database performance

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information about "Avoiding SQL Injection in PL/SQL"

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_PERF Subprograms

142.1 DBMS_PERF Overview

The DBMS_PERF package provides an interface for generating database performance reports. All subprograms return an active report and these reports can be generated at the system level, session level or at SQL level.

142.2 DBMS_PERF Security Model

The DBMS_PERF package requires the DBA role.

142.3 Summary of DBMS_PERF Subprograms

This table lists the `DBMS_PERF` subprograms and briefly describes them.

Table 142-1 DBMS_PERF Package Subprograms

Subprogram	Description
REPORT_PERFHUB Function	Generates a composite active performance report of the entire database system for a specified time period
REPORT_SESSION Function	Generates a performance report for a specific database session where a session is identified by `inst_id`, `sid`, and `serial_num`.
REPORT_SQL Function	Generates an active performance report for a particular SQL statement identified by its `sql_id`.

142.3.1 REPORT_PERFHUB Function

This function generates a composite active performance report of the entire database system for a specified time period.

Syntax

```
DBMS_PERF.REPORT_PERFHUB (
   is_realtime          IN NUMBER   DEFAULT NULL,
   outer_start_time     IN DATE     DEFAULT NULL,
   outer_end_time       IN DATE     DEFAULT NULL,
   selected_start_time  IN DATE     DEFAULT NULL,
   selected_end_time    IN DATE     DEFAULT NULL,
   inst_id              IN NUMBER   DEFAULT NULL,
   dbid                 IN NUMBER   DEFAULT NULL,
   monitor_list_detail  IN NUMBER   DEFAULT NULL,
   workload_sql_detail  IN NUMBER   DEFAULT NULL,
   addm_task_detail     IN NUMBER   DEFAULT NULL,
   report_reference     IN VARCHAR2 DEFAULT NULL,
   report_level         IN VARCHAR2 DEFAULT NULL,
   type                 IN VARCHAR2 DEFAULT 'ACTIVE',
   base_path            IN VARCHAR2 DEFAULT NULL);
 RETURN CLOB;
```

Parameters

Table 142-2 *REPORT_PERFHUB Function Parameters*

Parameter	Description
`is_realtime`	If 1, then real-time. If `NULL` (default) or `0`, then historical mode.
`outer_start_time`	Start time of outer period shown in the time selector. If `NULL` (default): • If `is_realtime=0` (historical), then 24 hours before `outer_end_time`. • If `is_realtime=1` (realtime mode), then 1 hour before `outer_end_time`.
`outer_end_time`	End time of outer period shown in the time selector. If `NULL` (default), then latest AWR snapshot. • If `is_realtime=0` (historical), then the latest AWR snapshot • If `is_realtime=1` (realtime mode), this is the current time (and any input is ignored)
`selected_start_time`	Start time period of selection. If `NULL` (default) • If `is_realtime=0`, then 1 hour before `selected_end_time` • If `is_realtime=1`, then 5 minutes before `selected_end_time`
`selected_end_time`	End time period of selection. If `NULL` (default) • If `is_realtime=0`, then latest AWR snapshot • `If is_realtime=1`, then current time
`inst_id`	Instance ID to for which to retrieve data • If `-1`, then current instance • If number is specified, then for that instance • If `NULL` (default), then all instances

Table 142-2 (Cont.) *REPORT_PERFHUB Function Parameters*

Parameter	Description
dbid	DBID to query. • If NULL, then current DBID. • If is_realtime=1, then DBID must be the local DBID.
monitor_list_detail	Top N in SQL monitor list for which to retrieve SQL monitor details. • If NULL (default), then retrieves top 10 • If 0, then retrieves no monitor list details
workload_sql_detail	Top N in Workload Top SQL list to retrieve monitor details, • If NULL (default), then retrieves top 10 • If 0, then retrieves no monitor list details
addm_task_detail	Maximum N latest ADDM tasks to retrieve • If NULL (default), retrieves available data but no more than N • If 0, then retrieves no ADDM task details
report_reference	Must be NULL when used from SQL*Plus.
report_level	'typical' will get all tabs in performance hub
type	Report type: • 'ACTIVE' (default) • 'xml' returns XML
base_path	URL path for HTML resources since flex HTML requires access to external files. This is only valid for type='ACTIVE' and is typically not used. Default value will retrieve the required files from OTN.

Usage Notes

- Once a time period is selected, the performance information is collected and presented based on performance subject areas.
- The time period can be real-time or historical.
- When real-time data is selected, more granular data is presented because data points are available every minute.
- When historical data is selected, more detailed data (broken down by different metrics) is presented, but the data points are averaged out to the Automatic Workload Repository (AWR) interval (usually an hour).
- Different tabs are available in the Performance Hub, depending on whether is_real-time is 1 for real time mode or 0 for historical mode.

142.3.2 REPORT_SESSION Function

This function produces a performance report for a specific database session where a session is identified by inst_id, sid, andserial_num.

If any of those parameters are missing, then the report is for the current session.

The session-level performance report contains the following tabs:

- Summary - This tab contains key identifiers and attributes of the session along with a summary of its activity data. It also contains a list of SQLs, PLSQL blocks and Database

Operations (DBOP) executed by that session that were monitored by Real-time SQL Monitoring.

- Activity - This tab shows activity broken down by wait classes for this session. The data used for this chart is fetched from Active Session History (ASH).
- Metrics - This tab shows charts for certain key metrics for the selected session over time and is only available in historical mode. Some of the metrics shown are CPU usage, PGA usage, IO Throughput and IO Requests.

Syntax

```
DBMS_PERF.REPORT_SESSION (
    inst_id              IN NUMBER   DEFAULT NULL,
    sid                  IN NUMBER   DEFAULT NULL,
    serial               IN NUMBER   DEFAULT NULL,
    is_realtime          IN NUMBER   DEFAULT NULL,
    outer_start_time     IN DATE     DEFAULT NULL,
    outer_end_time       IN DATE     DEFAULT NULL,
    selected_start_time  IN DATE     DEFAULT NULL,
    selected_end_time    IN DATE     DEFAULT NULL,
    dbid                 IN NUMBER   DEFAULT NULL,
    monitor_list_detail  IN NUMBER   DEFAULT NULL,
    report_reference     IN VARCHAR2 DEFAULT NULL,
    report_level         IN VARCHAR2 DEFAULT NULL,
    type                 IN VARCHAR2 DEFAULT 'ACTIVE',
    base_path            IN VARCHAR2 DEFAULT NULL)
  RETURN CLOB;
```

Parameters

Table 142-3 *REPORT_SESSION Function Parameters*

Parameter	Description
inst_id	Instance ID to for which to retrieve data. If NULL (default), then instance of current session.
sid	Session ID for which to retrieve performance. If NULL, uses current session.
serial	Serial# of session. If NULL, then the serial# of the specified sid is used provided the session is connected.
is_realtime	If 1, then real-time. If NULL (default) or 0, then historical mode.
outer_start_time	Start time of outer period shown in the time selector. If NULL (default): • If is_realtime=0 (historical), then 24 hours before outer_end_time. • If is_realtime=1 (realtime mode), then 1 hour before outer_end_time.
outer_end_time	End time of outer period shown in the time selector. If NULL (default), then latest AWR snapshot. • If is_realtime=0 (historical), then the latest AWR snapshot • If is_realtime=1 (realtime mode), this is the current time (and any input is ignored)
selected_start_time	Start time period of selection. If NULL (default) • If is_realtime=0, then 1 hour before selected_end_time • If is_realtime=1, then 5 minutes before selected_end_time

Table 142-3 (Cont.) *REPORT_SESSION Function Parameters*

Parameter	Description
selected_end_time	End time period of selection. If NULL (default) • If is_realtime=0, then latest AWR snapshot • If is_realtime=1, then current time
dbid	DBID to query. • If NULL, then current DBID. • If is_realtime=1, then DBID must be the local DBID.\
monitor_list_detail	Top N in SQL monitor list for which to retrieve SQL monitor details. • If NULL (default), then retrieves top 10 • If 0, then retrieves no monitor list details
report_reference	Must be NULL when used from SQL*Plus.
report_level	'typical' will get all tabs in the session hub (or session details)
type	Report type: • 'ACTIVE' (default) • 'xml' returns XML
base_path	URL path for HTML resources since flex HTML requires access to external files

142.3.3 REPORT_SQL Function

This function generates an active performance report for a particular SQL statement identified by its sql_id.

The SQL-level performance report contains the following tabs:

- Summary - This tab contains an overview of the SQL statement with key attributes like the SQL text, user name, sessions executing it, and related information. It also contains a Plans tab which shows statistics and activity for each distinct plan for this SQL statement found in memory and in the AWR.
- Activity - This tab shows activity broken down by wait classes for this SQL statement. The data used for this chart is fetched from Active Session History (ASH).
- Execution Statistics - This tab shows statistics and activity for each distinct plan for this statement along with a graphical and tabular representation of the plan.
- Monitored SQL - All executions of this SQL statement that were monitored by Real-time SQL Monitoring are listed in this tab.
- Plan Control - This tab shows information about SQL Profiles and SQL Plan Baselines if they exist for this SQL statement.
- Historical Statistics - This tab is available only in Historical mode. It contains statistics, such as number of executions, number of I/Os, rows processed, and other information produced over time for different execution plans. This information is retrieved from AWR.

Syntax

```
DBMS_PERF.REPORT_SQL (
    sql_id               IN varchar2 default null,
    is_realtime          IN number   default null,
    outer_start_time     IN date     default null,
```

```
    outer_end_time       IN date     default null,
    selected_start_time  IN date     default null,
    selected_end_time    IN date     default null,
    inst_id              IN number   default null,
    dbid                 IN number   default null,
    monitor_list_detail  IN number   default null,
    report_reference     IN varchar2 default null,
    report_level         IN varchar2 default null,    type                 IN varchar2
default 'ACTIVE',
    base_path            IN varchar2 default null);
  RETURN CLOB;
```

Parameters

Table 142-4 *REPORT_SQL Function Parameters*

Parameter	Description
sql_id	SQL_ID for which to retrieve performance. If NULL, gets SQL details for the last executed SQL statement.
is_realtime	If 1, then real-time. If NULL (default) or 0, then historical mode.
outer_start_time	Start time of outer period shown in the time selector. If NULL (default): • If is_realtime=0 (historical), then 24 hours before outer_end_time. • If is_realtime=1 (realtime mode), then 1 hour before outer_end_time.
outer_end_time	End time of outer period shown in the time selector. If NULL (default), then latest AWR snapshot. • If is_realtime=0 (historical), then the latest AWR snapshot • If is_realtime=1 (realtime mode), this is the current time (and any input is ignored)
selected_start_time	Start time period of selection. If NULL (default) • If is_realtime=0, then 1 hour before selected_end_time • If is_realtime=1, then 5 minutes before selected_end_time
selected_end_time	End time period of selection. If NULL (default) • If is_realtime=0, then latest AWR snapshot • If is_realtime=1, then current time
inst_id	Instance ID to for which to retrieve data. If NULL (default), then instance of current session.
dbid	DBID to query. • If NULL, then current DBID. • If is_realtime=1, then DBID must be the local DBID.\
monitor_list_detail	Top N in SQL monitor list for which to retrieve SQL monitor details. • If NULL (default), then retrieves top 10 • If 0, then retrieves no monitor list details
report_reference	Must be NULL when used from SQL*Plus.
report_level	'typical' will get all tabs in performance hub
type	Report type: • 'ACTIVE' (default) • 'xml' returns XML
base_path	URL path for HTML resources since flex HTML requires access to external files

143
DBMS_PIPE

The `DBMS_PIPE` package lets two or more sessions in the same instance communicate. Oracle pipes are similar in concept to the pipes used in UNIX, but Oracle pipes are not implemented using the operating system pipe mechanisms.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operational Notes
- Exceptions
- Examples
- Summary of DBMS_PIPE Subprograms

143.1 DBMS_PIPE Overview

Pipe functionality has several potential applications: external service interface, independent transactions, alerters (non-transactional), debugging, and concentrator.

- External service interface: You can communicate with user-written services that are external to the RDBMS. This can be done effectively in a shared server process, so that several instances of the service are executing simultaneously. Additionally, the services are available asynchronously. The requestor of the service does not need to block a waiting reply. The requestor can check (with or without time out) at a later time. The service can be written in any of the 3GL languages that Oracle supports.
- Independent transactions: The pipe can communicate to a separate session which can perform an operation in an independent transaction (such as logging an attempted security violation detected by a trigger).
- Alerters (non-transactional): You can post another process without requiring the waiting process to poll. If an "after-row" or "after-statement" trigger were to alert an application, then the application would treat this alert as an indication that the data probably changed. The application would then read the data to get the current value. Because this is an "after" trigger, the application would want to do a "`SELECT FOR UPDATE`" to make sure it read the correct data.
- Debugging: Triggers and stored procedures can send debugging information to a pipe. Another session can keep reading out of the pipe and display it on the screen or write it to a file.
- Concentrator: This is useful for multiplexing large numbers of users over a fewer number of network connections, or improving performance by concentrating several user-transactions into one DBMS transaction.

143.2 DBMS_PIPE Security Model

To use `DBMS_PIPE` and its subprograms, you must be granted the `EXECUTE` privilege on the package. Security can further be achieved by creating pipes using the `private` parameter in the `CREATE_PIPE` function and by writing common packages that only expose particular features or pipenames to particular users or roles.

Depending upon your security requirements, you may choose to use either public pipes or private pipes, which are described in DBMS_PIPE Operational Notes.

The `DBMS_PIPE` package uses invoker's rights, meaning the operations of sending and receiving messages run in the invoker's schema.

In order to use `DBMS_PIPE` messages with Cloud Object stores, you must have the `EXECUTE` privilege on the `DBMS_PIPE` package, the `DBMS_CLOUD` package, and on the Credential Object for accessing the object store URI.

> **See Also:**
>
> - DBMS_CLOUD for information about that package and about working with cloud service credentials

143.3 DBMS_PIPE Constants

This is the maximum time to wait attempting to send or receive a message.

```
maxwait   constant integer := 86400000; /* 1000 days */
```

143.4 DBMS_PIPE Operational Notes

Except for persistent pipes, information sent through Oracle pipes is buffered in the system global area (SGA). All information in pipes is lost when the instance is shut down.

> **WARNING:**
>
> Pipes are independent of transactions. Be careful using pipes when transaction control can be affected.

The operation of `DBMS_PIPE` is considered with regard to the following topics:

- Public Pipes
- Writing and Reading Pipes
- Private Pipes
- Singleton Pipes
- Persistent Pipes

Public Pipes

You may create a public pipe either implicitly or explicitly. For *implicit* public pipes, the pipe is automatically created when it is referenced for the first time, and it disappears when it no longer contains data. Because the pipe descriptor is stored in the SGA, there is some space usage overhead until the empty pipe is aged out of the cache.

You create an *explicit* public pipe by calling the `CREATE_PIPE` function with the `private` flag set to `FALSE`. You must deallocate explicitly-created pipes by calling the `REMOVE_PIPE` function.

The domain of a public pipe is the schema in which it was created, either explicitly or implicitly.

Reading and Writing Pipes

Each public pipe works asynchronously. Any number of schema users can write to a public pipe, as long as they have `EXECUTE` permission on the `DBMS_PIPE` package, and they know the name of the public pipe. However, once buffered information is read by one user, it is emptied from the buffer, and is not available for other readers of the same pipe.

The sending session builds a message using one or more calls to the `PACK_MESSAGE` procedure. This procedure adds the message to the session's local message buffer. The information in this buffer is sent by calling the `SEND_MESSAGE` function, designating the pipe name to be used to send the message. When `SEND_MESSAGE` is called, all messages that have been stacked in the local buffer are sent.

A process that wants to receive a message calls the `RECEIVE_MESSAGE` function, designating the pipe name from which to receive the message. The process then calls the `UNPACK_MESSAGE` procedure to access each of the items in the message.

Private Pipes

You explicitly create a private pipe by calling the `CREATE_PIPE` function. Once created, the private pipe persists in shared memory until you explicitly deallocate it by calling the `REMOVE_PIPE` function. A private pipe is also deallocated when the database instance is shut down.

You cannot create a private pipe if an implicit pipe exists in memory and has the same name as the private pipe you are trying to create. In this case, `CREATE_PIPE` returns an error.

Access to a private pipe is restricted to:

- Sessions running under the same userid as the creator of the pipe
- Stored subprograms executing in the same userid privilege domain as the pipe creator
- Users connected as `SYSDBA`

An attempt by any other user to send or receive messages on the pipe, or to remove the pipe, results in an immediate error. Any attempt by another user to create a pipe with the same name also causes an error.

As with public pipes, you must first build your message using calls to `PACK_MESSAGE` before calling `SEND_MESSAGE`. Similarly, you must call `RECEIVE_MESSAGE` to retrieve the message before accessing the items in the message by calling `UNPACK_MESSAGE`.

Singleton Pipes

Singleton pipes provide the ability to cache a single message in the shared memory of the current database instance, allowing high throughput concurrent reads of a message across database sessions. In a RAC database, pipes are not synchronized across database

instances. Each database instance has its own pipe, private to the memory of the database instance.

A single message can be cached in a singleton pipe, which can be comprised of multiple fields up to a total message size of 32,767 bytes. DBMS_PIPE supports the ability to pack multiple attributes in a message using the PACK_MESSAGE procedure.

Singleton pipes can be public or private and implicit or explicit. Using a private singleton pipe, the message can only be received by sessions with the same user as the creator of the pipe. The message in a public singleton pipe can be received by any database session with EXECUTE privilege on the DBMS_PIPE package.

Singleton pipes cache the message in the pipe until it is invalidated or purged. Explicit invalidation by purging the pipe is accomplished by using the PURGE procedure or by overwriting the message using the SEND_MESSAGE function. Automatic invalidation happens once the shelflife time has elapsed, which is specified as part of the CREATE_PIPE function and the SEND_MESSAGE function.

A user-defined cache function can be used to automatically populate the message in a singleton pipe and can be specified when reading a message using the RECEIVE_MESSAGE function. If the message is not in cache or the message shelflife time has elapsed, the singleton pipe automatically populates a new message in the pipe. The cache function simplifies the use of a pipe by avoiding condition logic on failure to receive a message from an empty pipe. It also ensures there is no cache miss when reading messages from a singleton pipe.

> **Note:**
>
> When using a cache function with RECEIVE_MESSAGE, the time spent executing the cache function is not included towards the overall timeout specified for receiving messages.

Singleton pipes do not get evicted from Oracle Database memory, ensuring maximum cache hits. An explicit singleton pipe exists in database memory until it is removed using the REMOVE_PIPE procedure or the database restarts. An implicit singleton pipe exists in database memory until there is one cached message in the pipe.

Persistent Pipes

You can optionally use Cloud Object stores for persistent message storage. The process for sending and receiving messages from a pipe is largely the same whether you use local or cloud storage, the difference being with cloud storage, a credential object and associated cloud URI must be provided. DBMS_PIPE uses the DBMS_CLOUD package to integrate with object storage, meaning DBMS_PIPE supports all object stores and credential types that are allowed with DBMS_CLOUD.

Storing pipe messages in object storage allows two or more databases in the same or different regions to communicate using DBMS_PIPE messages and removes the max number of messages that can be stored in a pipe. This means that a database instance open on one Real Application Cluster (RAC) instance can reliably use DBMS_PIPE for communicating between processes running on different cluster instances.

A locking mechanism is implemented to allow only one sending or receiving operation to be working in the cloud object store at a time. A lock file is created when a process begins using the cloud store and then is deleted once the operation is complete. Only when there is no lock

present is a process able to begin working in the cloud store. Processes will retry until the given max wait time is reached. If a lock has been left behind by an unknown process that is no longer operating, the lock will be deleted and remade by the current process trying to access the cloud store after 15 minutes.

> **Note:**
>
> When accessing `DBMS_PIPE` across different databases using Cloud Object stores, Oracle recommends using the `CREATE_PIPE` function before sending or receiving a message to ensure that the pipe is created with the required access permissions of a private or public pipe.

> **See Also:**
>
> - *Using Oracle Autonomous Database Serverless* for supported URI formats
> - *Using Oracle Autonomous Database Serverless* for `DBMS_CLOUD` subprograms and REST APIs
> - *Using Oracle Autonomous Database Serverless* for information about supported credential types

143.5 DBMS_PIPE Exceptions

`DBMS_PIPE` package subprograms can return the errors listed in the following table.

Table 143-1 DBMS_PIPE Errors

Error	Description
`ORA-23321:`	Pipename may not be null. This can be returned by the `CREATE_PIPE` function, or any subprogram that takes a pipe name as a parameter.
`ORA-23322:`	Insufficient privilege to access pipe. This can be returned by any subprogram that references a private pipe in its parameter list.

143.6 DBMS_PIPE Examples

These examples show use of DBMS_PIPE in debugging PL/SQL, debugging Pro*C, executing system commands, and an external service interface.

Example 1: Debugging - PL/SQL

This example shows the procedure that a PL/SQL program can call to place debugging information in a pipe.

```
CREATE OR REPLACE PROCEDURE debug (msg VARCHAR2) AS
    status  NUMBER;
BEGIN
  DBMS_PIPE.PACK_MESSAGE(LENGTH(msg));
  DBMS_PIPE.PACK_MESSAGE(msg);
  status := DBMS_PIPE.SEND_MESSAGE('plsql_debug');
```

```
  IF status != 0 THEN
    raise_application_error(-20099, 'Debug error');
  END IF;
END debug;
```

Example 2: Debugging - Pro*C

The following Pro*C code receives messages from the `PLSQL_DEBUG` pipe in the previous example, and displays the messages. If the Pro*C session is run in a separate window, then it can be used to display any messages that are sent to the debug procedure from a PL/SQL program executing in a separate session.

```
#include <stdio.h>
#include <string.h>

EXEC SQL BEGIN DECLARE SECTION;
   VARCHAR username[20];
   int     status;
   int     msg_length;
   char    retval[2000];
EXEC SQL END DECLARE SECTION;

EXEC SQL INCLUDE SQLCA;

void sql_error();

main()
{

-- Prepare username:
   strcpy(username.arr, "HR/<password>");
   username.len = strlen(username.arr);

   EXEC SQL WHENEVER SQLERROR DO sql_error();
   EXEC SQL CONNECT :username;

   printf("connected\n");

-- Start an endless loop to look for and print messages on the pipe:
   FOR (;;)
   {
      EXEC SQL EXECUTE
         DECLARE
            len INTEGER;
            typ INTEGER;
            sta INTEGER;
            chr VARCHAR2(2000);
         BEGIN
            chr := '';
            sta := dbms_pipe.receive_message('plsql_debug');
            IF sta = 0 THEN
               DBMS_PIPE.UNPACK_MESSAGE(len);
               DBMS_PIPE.UNPACK_MESSAGE(chr);
            END IF;
            :status := sta;
            :retval := chr;
            IF len IS NOT NULL THEN
               :msg_length := len;
            ELSE
               :msg_length := 2000;
            END IF;
         END;
```

```
      END-EXEC;
      IF (status == 0)
         printf("\n%.*s\n", msg_length, retval);
      ELSE
         printf("abnormal status, value is %d\n", status);
   }
}

void sql_error()
{
   char msg[1024];
   int rlen, len;
   len = sizeof(msg);
   sqlglm(msg, &len, &rlen);
   printf("ORACLE ERROR\n");
   printf("%.*s\n", rlen, msg);
   exit(1);
}
```

Example 3: Execute System Commands

This example shows PL/SQL and Pro*C code let a PL/SQL stored procedure (or anonymous block) call PL/SQL procedures to send commands over a pipe to a Pro*C program that is listening for them.

The Pro*C program sleeps and waits for a message to arrive on the named pipe. When a message arrives, the Pro*C program processes it, carrying out the required action, such as executing a UNIX command through the *system*() call or executing a SQL command using embedded SQL.

`DAEMON.SQL` is the source code for the PL/SQL package. This package contains procedures that use the `DBMS_PIPE` package to send and receive message to and from the Pro*C daemon. Note that full handshaking is used. The daemon always sends a message back to the package (except in the case of the `STOP` command). This is valuable, because it allows the PL/SQL procedures to be sure that the Pro*C daemon is running.

You can call the DAEMON packaged procedures from an anonymous PL/SQL block using SQL*Plus or Enterprise Manager. For example:

```
SQLPLUS> variable rv number
SQLPLUS> execute :rv := DAEMON.EXECUTE_SYSTEM('ls -la');
```

On a UNIX system, this causes the Pro*C daemon to execute the command *system*("*ls -la*").

Remember that the daemon needs to be running first. You might want to run it in the background, or in another window beside the SQL*Plus or Enterprise Manager session from which you call it.

The `DAEMON.SQL` also uses the `DBMS_OUTPUT` package to display the results. For this example to work, you must have execute privileges on this package.

DAEMON.SQL Example. This is the code for the PL/SQL `DAEMON` package:

```
CREATE OR REPLACE PACKAGE daemon AS
  FUNCTION execute_sql(command VARCHAR2,
                       timeout NUMBER DEFAULT 10)
    RETURN NUMBER;

  FUNCTION execute_system(command VARCHAR2,
                          timeout NUMBER DEFAULT 10)
    RETURN NUMBER;
```

```
  PROCEDURE stop(timeout NUMBER DEFAULT 10);
END daemon;
/
CREATE OR REPLACE PACKAGE BODY daemon AS

  FUNCTION execute_system(command VARCHAR2,
                          timeout NUMBER DEFAULT 10)
  RETURN NUMBER IS

    status       NUMBER;
    result       VARCHAR2(20);
    command_code NUMBER;
    pipe_name    VARCHAR2(30);
  BEGIN
    pipe_name := DBMS_PIPE.UNIQUE_SESSION_NAME;

    DBMS_PIPE.PACK_MESSAGE('SYSTEM');
    DBMS_PIPE.PACK_MESSAGE(pipe_name);
    DBMS_PIPE.PACK_MESSAGE(command);
    status := DBMS_PIPE.SEND_MESSAGE('daemon', timeout);
    IF status <> 0 THEN
      RAISE_APPLICATION_ERROR(-20010,
        'Execute_system: Error while sending.  Status = ' ||
         status);
    END IF;

    status := DBMS_PIPE.RECEIVE_MESSAGE(pipe_name, timeout);
    IF status <> 0 THEN
      RAISE_APPLICATION_ERROR(-20011,
        'Execute_system: Error while receiving.
         Status = ' || status);
    END IF;

    DBMS_PIPE.UNPACK_MESSAGE(result);
    IF result <> 'done' THEN
      RAISE_APPLICATION_ERROR(-20012,
        'Execute_system: Done not received.');
    END IF;

    DBMS_PIPE.UNPACK_MESSAGE(command_code);
    DBMS_OUTPUT.PUT_LINE('System command executed.  result = ' ||
                         command_code);
    RETURN command_code;
  END execute_system;

  FUNCTION execute_sql(command VARCHAR2,
                       timeout NUMBER DEFAULT 10)
  RETURN NUMBER IS

    status       NUMBER;
    result       VARCHAR2(20);
    command_code NUMBER;
    pipe_name    VARCHAR2(30);

  BEGIN
    pipe_name := DBMS_PIPE.UNIQUE_SESSION_NAME;

    DBMS_PIPE.PACK_MESSAGE('SQL');
    DBMS_PIPE.PACK_MESSAGE(pipe_name);
    DBMS_PIPE.PACK_MESSAGE(command);
    status := DBMS_PIPE.SEND_MESSAGE('daemon', timeout);
    IF status <> 0 THEN
```

```
      RAISE_APPLICATION_ERROR(-20020,
        'Execute_sql: Error while sending.  Status = ' || status);
    END IF;

    status := DBMS_PIPE.RECEIVE_MESSAGE(pipe_name, timeout);

    IF status <> 0 THEN
      RAISE_APPLICATION_ERROR(-20021,
        'execute_sql: Error while receiving.
         Status = ' || status);
    END IF;

    DBMS_PIPE.UNPACK_MESSAGE(result);
    IF result <> 'done' THEN
      RAISE_APPLICATION_ERROR(-20022,
        'execute_sql: done not received.');
    END IF;

    DBMS_PIPE.UNPACK_MESSAGE(command_code);
    DBMS_OUTPUT.PUT_LINE
        ('SQL command executed.  sqlcode = ' || command_code);
    RETURN command_code;
  END execute_sql;

  PROCEDURE stop(timeout NUMBER DEFAULT 10) IS
    status NUMBER;
  BEGIN
    DBMS_PIPE.PACK_MESSAGE('STOP');
    status := DBMS_PIPE.SEND_MESSAGE('daemon', timeout);
    IF status <> 0 THEN
      RAISE_APPLICATION_ERROR(-20030,
        'stop: error while sending.  status = ' || status);
    END IF;
  END stop;
END daemon;
```

daemon.pc Example. This is the code for the Pro*C daemon. You must precompile this using the Pro*C Precompiler, Version 1.5.x or later. You must also specify the `USERID` and `SQLCHECK` options, as the example contains embedded PL/SQL code.

> **Note:**
>
> To use a `VARCHAR` output host variable in a PL/SQL block, you must initialize the length component before entering the block.

```
proc iname=daemon userid=HR/<password> sqlcheck=semantics
```

Then C-compile and link in the normal way.

```
#include <stdio.h>
#include <string.h>

EXEC SQL INCLUDE SQLCA;

EXEC SQL BEGIN DECLARE SECTION;
  char *uid = "HR/<password>";
  int status;
  VARCHAR command[20];
```

```
  VARCHAR value[2000];
  VARCHAR return_name[30];
EXEC SQL END DECLARE SECTION;

void
connect_error()
{
  char msg_buffer[512];
  int msg_length;
  int buffer_size = 512;

  EXEC SQL WHENEVER SQLERROR CONTINUE;
  sqlglm(msg_buffer, &buffer_size, &msg_length);
  printf("Daemon error while connecting:\n");
  printf("%.*s\n", msg_length, msg_buffer);
  printf("Daemon quitting.\n");
  exit(1);
}

void
sql_error()
{
  char msg_buffer[512];
  int msg_length;
  int buffer_size = 512;

  EXEC SQL WHENEVER SQLERROR CONTINUE;
  sqlglm(msg_buffer, &buffer_size, &msg_length);
  printf("Daemon error while executing:\n");
  printf("%.*s\n", msg_length, msg_buffer);
  printf("Daemon continuing.\n");
}
main()
{
command.len = 20; /*initialize length components*/
value.len = 2000;
return_name.len  = 30;
  EXEC SQL WHENEVER SQLERROR DO connect_error();
  EXEC SQL CONNECT :uid;
  printf("Daemon connected.\n");

  EXEC SQL WHENEVER SQLERROR DO sql_error();
  printf("Daemon waiting...\n");
  while (1) {
    EXEC SQL EXECUTE
      BEGIN
        :status := DBMS_PIPE.RECEIVE_MESSAGE('daemon');
        IF :status = 0 THEN
          DBMS_PIPE.UNPACK_MESSAGE(:command);
        END IF;
      END;
    END-EXEC;
    IF (status == 0)
    {
      command.arr[command.len] = '\0';
      IF (!strcmp((char *) command.arr, "STOP"))
      {
        printf("Daemon exiting.\n");
        break;
      }

      ELSE IF (!strcmp((char *) command.arr, "SYSTEM"))
```

```
  {
    EXEC SQL EXECUTE
      BEGIN
        DBMS_PIPE.UNPACK_MESSAGE(:return_name);
        DBMS_PIPE.UNPACK_MESSAGE(:value);
      END;
    END-EXEC;
    value.arr[value.len] = '\0';
    printf("Will execute system command '%s'\n", value.arr);

    status = system(value.arr);
    EXEC SQL EXECUTE
      BEGIN
        DBMS_PIPE.PACK_MESSAGE('done');
        DBMS_PIPE.PACK_MESSAGE(:status);
        :status := DBMS_PIPE.SEND_MESSAGE(:return_name);
      END;
    END-EXEC;

    IF (status)
    {
      printf
       ("Daemon error while responding to system command.");
      printf("  status: %d\n", status);
    }
  }
  ELSE IF (!strcmp((char *) command.arr, "SQL")) {
    EXEC SQL EXECUTE
      BEGIN
        DBMS_PIPE.UNPACK_MESSAGE(:return_name);
        DBMS_PIPE.UNPACK_MESSAGE(:value);
      END;
    END-EXEC;
    value.arr[value.len] = '\0';
    printf("Will execute sql command '%s'\n", value.arr);

    EXEC SQL WHENEVER SQLERROR CONTINUE;
    EXEC SQL EXECUTE IMMEDIATE :value;
    status = sqlca.sqlcode;

    EXEC SQL WHENEVER SQLERROR DO sql_error();
    EXEC SQL EXECUTE
      BEGIN
        DBMS_PIPE.PACK_MESSAGE('done');
        DBMS_PIPE.PACK_MESSAGE(:status);
        :status := DBMS_PIPE.SEND_MESSAGE(:return_name);
      END;
    END-EXEC;

    IF (status)
    {
      printf("Daemon error while responding to sql command.");
      printf("  status: %d\n", status);
    }
  }
  ELSE
  {
    printf
      ("Daemon error: invalid command '%s' received.\n",
        command.arr);
  }
}
```

```
  ELSE
  {
    printf("Daemon error while waiting for signal.");
    printf("  status = %d\n", status);
  }
}
EXEC SQL COMMIT WORK RELEASE;
exit(0);
```

Example 4: External Service Interface

Put the user-written 3GL code into an OCI or Precompiler program. The program connects to the database and executes PL/SQL code to read its request from the pipe, computes the result, and then executes PL/SQL code to send the result on a pipe back to the requestor.

Below is an example of a stock service request. The recommended sequence for the arguments to pass on the pipe for all service requests is:

```
  protocol_version      VARCHAR2        - '1', 10 bytes or less
  returnpipe            VARCHAR2        - 30 bytes or less
  service               VARCHAR2        - 30 bytes or less
  arg1                  VARCHAR2/NUMBER/DATE
     ...
  argn                  VARCHAR2/NUMBER/DATE
```

The recommended format for returning the result is:

```
  success               VARCHAR2        - 'SUCCESS' if OK,
                                          otherwise error message
  arg1                  VARCHAR2/NUMBER/DATE
     ...
  argn                  VARCHAR2/NUMBER/DATE
```

The "stock price request server" would do, using OCI or PRO* (in pseudo-code):

```
<loop forever>
  BEGIN dbms_stock_server.get_request(:stocksymbol); END;
  <figure out price based on stocksymbol (probably from some radio
        signal), set error if can't find such a stock>
  BEGIN dbms_stock_server.return_price(:error, :price); END;
```

A client would do:

```
BEGIN :price := stock_request('YOURCOMPANY'); end;
```

The stored procedure, `dbms_stock_server`, which is called by the preceding "stock price request server" is:

```
CREATE OR REPLACE PACKAGE dbms_stock_server IS
  PROCEDURE get_request(symbol OUT VARCHAR2);
  PROCEDURE return_price(errormsg IN VARCHAR2, price IN VARCHAR2);
END;

CREATE OR REPLACE PACKAGE BODY dbms_stock_server IS
  returnpipe    VARCHAR2(30);

  PROCEDURE returnerror(reason VARCHAR2) IS
    s INTEGER;
  BEGIN
    dbms_pipe.pack_message(reason);
    s := dbms_pipe.send_message(returnpipe);
    IF s <> 0 THEN
      raise_application_error(-20000, 'Error:' || to_char(s) ||
```

```
        ' sending on pipe');
    END IF;
  END;

  PROCEDURE get_request(symbol OUT VARCHAR2) IS
    protocol_version VARCHAR2(10);
    s                INTEGER;
    service          VARCHAR2(30);
  BEGIN
    s := dbms_pipe.receive_message('stock_service');
    IF s <> 0 THEN
      raise_application_error(-20000, 'Error:' || to_char(s) ||
        'reading pipe');
    END IF;
    dbms_pipe.unpack_message(protocol_version);
    IF protocol_version <> '1' THEN
      raise_application_error(-20000, 'Bad protocol: ' ||
        protocol_version);
    END IF;
    dbms_pipe.unpack_message(returnpipe);
    dbms_pipe.unpack_message(service);
    IF service != 'getprice' THEN
      returnerror('Service ' || service || ' not supported');
    END IF;
    dbms_pipe.unpack_message(symbol);
  END;

  PROCEDURE return_price(errormsg in VARCHAR2, price in VARCHAR2) IS
    s INTEGER;
  BEGIN
    IF errormsg is NULL THEN
      dbms_pipe.pack_message('SUCCESS');
      dbms_pipe.pack_message(price);
    ELSE
      dbms_pipe.pack_message(errormsg);
    END IF;
    s := dbms_pipe.send_message(returnpipe);
    IF s <> 0 THEN
      raise_application_error(-20000, 'Error:'||to_char(s)||
        ' sending on pipe');
    END IF;
  END;
END;
```

The procedure called by the client is:

```
CREATE OR REPLACE FUNCTION stock_request (symbol VARCHAR2)
    RETURN VARCHAR2 IS
  s        INTEGER;
  price    VARCHAR2(20);
  errormsg VARCHAR2(512);
BEGIN
  dbms_pipe.pack_message('1');  -- protocol version
  dbms_pipe.pack_message(dbms_pipe.unique_session_name); -- return pipe
  dbms_pipe.pack_message('getprice');
  dbms_pipe.pack_message(symbol);
  s := dbms_pipe.send_message('stock_service');
  IF s <> 0 THEN
    raise_application_error(-20000, 'Error:'||to_char(s)||
      ' sending on pipe');
  END IF;
```

```
  s := dbms_pipe.receive_message(dbms_pipe.unique_session_name);
  IF s <> 0 THEN
    raise_application_error(-20000, 'Error:'||to_char(s)||
      ' receiving on pipe');
  END IF;
  dbms_pipe.unpack_message(errormsg);
  IF errormsg <> 'SUCCESS' THEN
    raise_application_error(-20000, errormsg);
  END IF;
  dbms_pipe.unpack_message(price);
  RETURN price;
END;
```

You would typically only `GRANT EXECUTE` on `DBMS_STOCK_SERVICE` to the stock service application server, and would only `GRANT EXECUTE` on `stock_request` to those users allowed to use the service.

> **See Also:**
>
> DBMS_ALERT

143.7 Summary of DBMS_PIPE Subprograms

This table lists the `DBMS_PIPE` subprograms and briefly describes them.

Table 143-2 DBMS_PIPE Package Subprograms

Subprogram	Description
CREATE_PIPE Function	Creates a pipe (necessary for private pipes)
GET_CREDENTIAL_NAME Function	Returns the globally set `credential_name` variable
GET_LOCATION_URI Function	Returns the globally set `location_uri` variable
NEXT_ITEM_TYPE Function	Returns datatype of next item in buffer
PACK_MESSAGE Procedures	Builds message in local buffer
PURGE Procedure	Purges contents of named pipe
RECEIVE_MESSAGE Function	Copies message from named pipe into local buffer
REMOVE_PIPE Function	Removes the named pipe
RESET_BUFFER Procedure	Purges contents of local buffer
SEND_MESSAGE Function	Sends message on named pipe: This implicitly creates a public pipe if the named pipe does not exist
SET_CREDENTIAL_NAME Procedure	Sets the global `credential_name` variable
SET_LOCATION_URI Procedure	Sets the global `location_uri` variable
UNIQUE_SESSION_NAME Function	Returns unique session name
UNPACK_MESSAGE Procedures	Accesses next item in buffer

143.7.1 CREATE_PIPE Function

This function explicitly creates a public or private pipe. If the `private` flag is `TRUE`, then the pipe creator is assigned as the owner of the private pipe.

Explicitly-created pipes can only be removed by calling `REMOVE_PIPE`, or by shutting down the instance.

Syntax

```
DBMS_PIPE.CREATE_PIPE (
   pipename     IN VARCHAR2,
   maxpipesize  IN INTEGER DEFAULT 65536,
   private      IN BOOLEAN DEFAULT TRUE,
   singleton    IN BOOLEAN DEFAULT FALSE,
   shelflife    IN INTEGER DEFAULT 0)
RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(create_pipe,WNDS,RNDS);
```

Parameters

Table 143-3 CREATE_PIPE Function Parameters

Parameter	Description
`pipename`	Name of the pipe you are creating. You must use this name when you call `SEND_MESSAGE` and `RECEIVE_MESSAGE`. This name must be unique across the instance. Caution: Do not use pipe names beginning with `ORA$`. These are reserved for use by procedures provided by Oracle. Pipename should not be longer than 128 bytes, and is case insensitive. At this time, the name cannot contain Globalization Support characters.
`maxpipesize`	The maximum size allowed for the pipe, in bytes. The total size of all of the messages on the pipe cannot exceed this amount. The message is blocked if it exceeds this maximum. The default `maxpipesize` is 65536 bytes. The `maxpipesize` for a pipe becomes a part of the characteristics of the pipe and persists for the life of the pipe. Callers of `SEND_MESSAGE` with larger values cause the `maxpipesize` to be increased. Callers with a smaller value use the existing, larger value.
`private`	Uses the default, `TRUE`, to create a private pipe. Public pipes can be implicitly created when you call `SEND_MESSAGE`.
`singleton`	Use the value `TRUE` to indicate that the pipe should be created as a singleton pipe. Singleton pipes cannot be persistent pipes. The default value of `singleton` is `FALSE`.

Table 143-3 (Cont.) CREATE_PIPE Function Parameters

Parameter	Description
shelflife	Only applicable to singleton pipes, this parameter is optionally used to set the expiration time in seconds of a message cached in a singleton pipe. Once the shelflife time is exceeded, the message is no longer accessible from the pipe. The shelflife can be used for implicit invalidation of the message in a singleton pipe. The default value of 0 indicates that the message will not expire. The shelflife of a message in a singleton pipe can also be specified when sending a message. See SEND_MESSAGE Function.

Return Values

Table 143-4 CREATE_PIPE Function Return Values

Return	Description
0	Successful. If the pipe already exists and the user attempting to create it is authorized to use it, then Oracle returns 0, indicating success, and any data already in the pipe remains. If a user connected as SYSDBA/SYSOPER re-creates a pipe, then Oracle returns status 0, but the ownership of the pipe remains unchanged.
6	Failed to convert existing pipe to a singleton pipe. An implicit pipe with more than one existing message cannot be converted to a singleton pipe. For an explicit pipe that is not a singleton pipe, SEND_MESSAGE cannot send a message with the singleton parameter set to TRUE.
7	A non-zero value was given for shelflife and the pipe is not a singleton pipe.
ORA-23322	Failure due to naming conflict. If a pipe with the same name exists and was created by a different user, then Oracle signals error ORA-23322, indicating the naming conflict.

Exceptions

Table 143-5 CREATE_PIPE Function Exception

Exception	Description
Null pipe name	Permission error: Pipe with the same name already exists, and you are not allowed to use it.

143.7.2 GET_CREDENTIAL_NAME Function

This function retrieves the global `credential_name` variable to be used as the default credential with the cloud. The `GET_CREDENTIAL_NAME` function is only applicable to persistent pipes with messages stored in Cloud Object Storage.

Syntax

```
DBMS_PIPE.GET_CREDENTIAL_NAME ()
RETURN VARCHAR2;
```

Return Values

This function returns the globally set `credential_name` variable. The variable is set to `NULL` by default.

143.7.3 GET_LOCATION_URI Function

This function retrieves the global `location_uri` variable to be used as the default location URI with the cloud. The `GET_LOCATION_URI` function is only applicable to persistent pipes with messages stored in Cloud Object Storage.

Syntax

```
DBMS_PIPE.GET_LOCATION_URI ()
RETURN VARCHAR2;
```

Return Values

This function returns the globally set `location_uri` variable. The variable is set to `NULL` by default.

143.7.4 NEXT_ITEM_TYPE Function

This function determines the datatype of the next item in the local message buffer.

After you have called `RECEIVE_MESSAGE` to place pipe information in a local buffer, call `NEXT_ITEM_TYPE`.

Syntax

```
DBMS_PIPE.NEXT_ITEM_TYPE
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(next_item_type,WNDS,RNDS);
```

Return Values

Table 143-6 NEXT_ITEM_TYPE Function Return Values

Return	Description
0	No more items
6	NUMBER
9	VARCHAR2
11	ROWID
12	DATE
23	RAW

143.7.5 PACK_MESSAGE Procedures

This procedure builds your message in the local message buffer.

To send a message, first make one or more calls to `PACK_MESSAGE`. Then, call `SEND_MESSAGE` to send the message in the local buffer on the named pipe.

The procedure is overloaded to accept items of type `VARCHAR2`, `NCHAR`, `NUMBER`, `DATE`., `RAW` and `ROWID` items. In addition to the data bytes, each item in the buffer requires one byte to indicate its type, and two bytes to store its length. One additional byte is needed to terminate the message.The overhead for all types other than `VARCHAR` is 4 bytes.

Syntax

```
DBMS_PIPE.PACK_MESSAGE (
   item  IN  VARCHAR2);

DBMS_PIPE.PACK_MESSAGE (
   item  IN  NCHAR);

DBMS_PIPE.PACK_MESSAGE (
   item  IN  NUMBER);

DBMS_PIPE.PACK_MESSAGE (
   item  IN  DATE);

DBMS_PIPE.PACK_MESSAGE_RAW (
   item  IN  RAW);

DBMS_PIPE.PACK_MESSAGE_ROWID (
   item  IN  ROWID);
```

Pragmas

```
pragma restrict_references(pack_message,WNDS,RNDS);
pragma restrict_references(pack_message_raw,WNDS,RNDS);
pragma restrict_references(pack_message_rowid,WNDS,RNDS);
```

Parameters

Table 143-7 PACK_MESSAGE Procedure Parameters

Parameter	Description
item	Item to pack into the local message buffer.

Usage Notes

In Oracle database version 8.x, the char-set-id (2 bytes) and the char-set-form (1 byte) are stored with each data item. Therefore, the overhead when using Oracle database version 8.x is 7 bytes.

When you call SEND_MESSAGE to send this message, you must indicate the name of the pipe on which you want to send the message. If this pipe already exists, then you must have sufficient privileges to access this pipe. If the pipe does not already exist, then it is created automatically.

Exceptions

ORA-06558 is raised if the message buffer overflows (currently 4096 bytes). Each item in the buffer takes one byte for the type, two bytes for the length, plus the actual data. There is also one byte needed to terminate the message.

143.7.6 PURGE Procedure

This procedure empties the contents of the named pipe.

An empty implicitly-created pipe is aged out of the shared global area according to the least-recently-used algorithm. Thus, calling PURGE lets you free the memory associated with an implicitly-created pipe.

Syntax

```
DBMS_PIPE.PURGE (
   pipename  IN  VARCHAR2);
```

Pragmas

```
pragma restrict_references(purge,WNDS,RNDS);
```

Parameters

Table 143-8 PURGE Procedure Parameters

Parameter	Description
pipename	Name of pipe from which to remove all messages. The local buffer may be overwritten with messages as they are discarded. Pipename should not be longer than 128 bytes, and is case-insensitive.

Usage Notes

Because `PURGE` calls `RECEIVE_MESSAGE`, the local buffer might be overwritten with messages as they are purged from the pipe. Also, you can receive an `ORA-23322` (insufficient privileges) error if you attempt to purge a pipe with which you have insufficient access rights.

Exceptions

Permission error if pipe belongs to another user.

143.7.7 RECEIVE_MESSAGE Function

This function copies the message into the local message buffer.

Syntax

```
DBMS_PIPE.RECEIVE_MESSAGE (
   pipename     IN VARCHAR2,
   timeout      IN INTEGER      DEFAULT maxwait,
   cache_func   IN VARCHAR2     DEFAULT NULL)
RETURN INTEGER;


DBMS_PIPE.RECEIVE_MESSAGE (
   pipename           IN VARCHAR2,
   timeout            IN INTEGER      DEFAULT maxwait,
   credential_name    IN VARCHAR2     DEFAULT NULL,
   location_uri       IN VARCHAR2     DEFAULT NULL)
RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(receive_message,WNDS,RNDS);
```

Parameters

Table 143-9 RECEIVE_MESSAGE Function Parameters

Parameter	Description
`pipename`	Name of the pipe on which you want to receive a message. Names beginning with `ORA$` are reserved for use by Oracle
`timeout`	Time to wait for a message, in seconds. The default value is the constant `MAXWAIT`, which is defined as 86400000 (1000 days). A timeout of 0 lets you read without blocking.

Table 143-9 (Cont.) RECEIVE_MESSAGE Function Parameters

Parameter	Description
cache_func	Only applicable to singleton pipes, cache_func is the cache function name used to automatically cache a message in a singleton pipe and can be created as either a PL/SQL function or an embedded function in a PL/SQL package. If specified, the cache function is invoked as the current session user invoking the RECEIVE_MESSAGE function so the current user must have privilege on the function. The name of the function must be fully qualified with the owner schema: • OWNER.FUNCTION_NAME • OWNER.PACKAGE.FUNCTION_NAME Note that the time spent executing the cache function is not included towards the overall timeout specified for receiving messages.
credential_name	The credential name for the cloud store used to store messages. This parameter is only applicable to persistent pipes. The default value is NULL. A passed parameter takes precedence over the package argument's value. Credentials require EXECUTE and READ/WRITE privileges.
location_uri	The location URL for the cloud store being used to store messages. This parameter is only applicable to persistent pipes. The location_uri parameter is a global variable that has a value of NULL by default. A passed parameter takes precedence over the global variable's value. If a credential name is specified, a location_uri must also be provided.

Return Values

Table 143-10 RECEIVE_MESSAGE Function Return Values

Return	Description
0	Success
1	Timed out. If the pipe was implicitly-created and is empty, then it is removed.
2	Record in the pipe is too large for the buffer. (This should not happen.)
3	An interrupt occurred.
8	The cache function was specified on a non-singleton pipe.
ORA-23322	User has insufficient privileges to read from the pipe.

Usage Notes

To receive a message from a pipe, first call RECEIVE_MESSAGE. When you receive a message, it is removed from the pipe; hence, a message can only be received once (unless using a singleton pipe). For implicitly-created pipes, the pipe is removed after the last record is removed from the pipe. An implicit singleton pipe exists in database memory until there is one cached message in the pipe.

If the pipe that you specify when you call `RECEIVE_MESSAGE` does not already exist, then Oracle implicitly creates the pipe and waits to receive the message. If the message does not arrive within a designated timeout interval, then the call returns and the pipe is removed.

After receiving the message, you must make one or more calls to `UNPACK_MESSAGE` to access the individual items in the message. The `UNPACK_MESSAGE` procedure is overloaded to unpack items of type `DATE`, `NUMBER`, `VARCHAR2`, and there are two additional procedures to unpack `RAW` and `ROWID` items. If you do not know the type of data that you are attempting to unpack, then call `NEXT_ITEM_TYPE` to determine the type of the next item in the buffer.

> **Note:**
>
> If performing cross-database messaging, the `CREATE_PIPE` function must be called before the first time you attempt to receive a message in the new database. The name of this pipe and its properties must be the same as the pipe the message was sent on in the other database.

Persistent messages are guaranteed to either be written or read by exactly one process. This prevents message content inconsistency due to concurrent writes and reads. Using a persistent messaging pipe, `DBMS_PIPE` allows only one operation, sending a message or receiving a message to be active at a given time. However, if an operation is not possible due to an ongoing operation, the process retries periodically until the `timeout` value is reached.

If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:

- If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

Cache Function Parameter

Singleton pipes support the cache function to automatically cache a message in the pipe in case of the following two scenarios:

- The singleton pipe is empty
- The message in the singleton pipe is invalid because the `shelflife` time has elapsed

The cache function simplifies the use of singleton pipes by avoiding condition logic on failure to receive a message from an empty pipe and ensures there is no cache miss.

The name of the function should be fully qualified with the owner schema:

- `OWNER.FUNCTION_NAME`
- `OWNER.PACKAGE.FUNCTION_NAME`

To use a cache function, the current session user that invokes `DBMS_PIPE.RECEIVE_MESSAGE` must have required privileges to execute the cache function.

Cache Function Syntax

```
CREATE OR REPLACE FUNCTION cache_function_name(
   pipename IN VARCHAR2
) RETURN INTEGER;
```

Cache Function Parameters

Parameter	Data Type	Description
pipename	VARCHAR2	Name of the singleton pipe

Cache Function Return Values

Return	Description
0	Success
Non-zero	Failure value returned from DBMS_PIPE.RECEIVE_MESSAGE

Define a cache function to provide encapsulation and abstraction of complexity from the reader sessions of a singleton pipe. The typical operations within a cache function would be:

- Create a singleton pipe for an explicit pipe using DBMS_PIPE.CREATE_PIPE
- Create the message to cache in the singleton pipe
- Send the message to the singleton pipe, optionally specifying a shelflife for the implicit message

Exceptions

Table 143-11 RECEIVE_MESSAGE Function Exceptions

Exception	Description
Null pipe name	Permission error. Insufficient privilege to remove the record from the pipe. The pipe is owned by someone else.

Example

```
DECLARE
   l_status INTEGER;
BEGIN
   l_status := DBMS_PIPE.RECEIVE_MESSAGE(pipename   => 'MY_PIPE1',
                                         timeout    => 1,
                                         cache_func =>
'MY_USER.MY_CACHE_FUNC');
END;
/
```

143.7.8 RESET_BUFFER Procedure

This procedure resets the PACK_MESSAGE and UNPACK_MESSAGE positioning indicators to 0.

Because all pipes share a single buffer, you may find it useful to reset the buffer before using a new pipe. This ensures that the first time you attempt to send a message to your pipe, you do not inadvertently send an expired message remaining in the buffer.

Syntax

```
DBMS_PIPE.RESET_BUFFER;
```

Pragmas

```
pragma restrict_references(reset_buffer,WNDS,RNDS);
```

143.7.9 REMOVE_PIPE Function

This function removes explicitly-created pipes.

Pipes created implicitly by SEND_MESSAGE are automatically removed when empty. However, pipes created explicitly by CREATE_PIPE are removed only by calling REMOVE_PIPE, or by shutting down the instance. All unconsumed records in the pipe are removed before the pipe is deleted.

This is similar to calling PURGE on an implicitly-created pipe.

Syntax

```
DBMS_PIPE.REMOVE_PIPE (
   pipename  IN  VARCHAR2)
RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(remove_pipe,WNDS,RNDS);
```

Parameters

Table 143-12 REMOVE_PIPE Function Parameters

Parameter	Description
pipename	Name of pipe that you want to remove.

Return Values

Table 143-13 REMOVE_PIPE Function Return Values

Return	Description
0	Success If the pipe does not exist, or if the pipe already exists and the user attempting to remove it is authorized to do so, then Oracle returns 0, indicating success, and any data remaining in the pipe is removed.

Table 143-13 (Cont.) REMOVE_PIPE Function Return Values

Return	Description
ORA-23322	Insufficient privileges. If the pipe exists, but the user is not authorized to access the pipe, then Oracle signals error ORA-23322, indicating insufficient privileges.

Exceptions

Table 143-14 REMOVE_PIPE Function Exception

Exception	Description
Null pipe name	Permission error: Insufficient privilege to remove pipe. The pipe was created and is owned by someone else.

143.7.10 SEND_MESSAGE Function

This function sends a message on the named pipe.

The message is contained in the local message buffer, which was filled with calls to PACK_MESSAGE. You can create a pipe explicitly using CREATE_PIPE, otherwise, it is created implicitly.

To create an implicit singleton pipe, set the singleton parameter to TRUE. The following arguments are applicable to singleton pipes:

- singleton: Indicates that the pipe should be created as a singleton pipe (default: FALSE).
- shelflife: Optionally specify a shelflife expiration of a cached message in the singleton pipe. It can be used for implicit invalidation of the message in a singleton pipe. This argument is applicable for implicit as well as explicit singleton pipes. A shelflife value specified in the SEND_MESSAGE Function overwrites the shelflife specified for the explicit singleton pipe in the CREATE_PIPE Function and will be the default for any new messages cached in the singleton pipe.

Syntax

```
DBMS_PIPE.SEND_MESSAGE (
    pipename     IN VARCHAR2,
    timeout      IN INTEGER DEFAULT MAXWAIT,
    maxpipesize  IN INTEGER DEFAULT 65536,
    singleton    IN BOOLEAN DEFAULT FALSE,
    shelflife    IN INTEGER DEFAULT 0)
  RETURN INTEGER;


DBMS_PIPE.SEND_MESSAGE (
    pipename         IN VARCHAR2,
    timeout          IN INTEGER DEFAULT MAXWAIT,
    credential_name  IN VARCHAR2 DEFAULT NULL,
    location_uri     IN VARCHAR2 DEFAULT NULL)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(send_message,WNDS,RNDS);
```

Parameters

Table 143-15 SEND_MESSAGE Function Parameters

Parameter	Description
pipename	Name of the pipe on which you want to place the message. If you are using an explicit pipe, then this is the name that you specified when you called CREATE_PIPE. Caution: Do not use pipe names beginning with 'ORA$'. These names are reserved for use by procedures provided by Oracle. Pipename should not be longer than 128 bytes, and is case-insensitive. At this time, the name cannot contain Globalization Support characters.
timeout	Time to wait while attempting to place a message on a pipe, in seconds. The default value is the constant MAXWAIT, which is defined as 86400000 (1000 days).
maxpipesize	Maximum size allowed for the pipe, in bytes. The total size of all the messages on the pipe cannot exceed this amount. The message is blocked if it exceeds this maximum. The default is 65536 bytes. The maxpipesize for a pipe becomes a part of the characteristics of the pipe and persists for the life of the pipe. Callers of SEND_MESSAGE with larger values cause the maxpipesize to be increased. Callers with a smaller value simply use the existing, larger value. Specifying maxpipesize as part of the SEND_MESSAGE procedure eliminates the need for a separate call to open the pipe. If you created the pipe explicitly, then you can use the optional maxpipesize parameter to override the creation pipe size specifications.
singleton	Specifying singleton as TRUE indicates that the implicit pipe should be created as a singleton pipe. This argument is not required if the pipe is explicitly created as a singleton pipe using the CREATE_PIPE function. The default value of singleton is FALSE.
shelflife	Only applicable to singleton pipes, this parameter is optionally used to set the expiration time in seconds of a message cached in an implicit or explicit singleton pipe. Once the shelflife time is exceeded, the message is no longer accessible from the pipe. The shelflife can be used for implicit invalidation of the message in a singleton pipe. The default value of 0 indicates that the message will not expire. Specifying shelflife as part of the SEND_MESSAGE procedure overwrites the shelflife value specified for an explicit singleton pipe using CREATE_PIPE and will be the default for any new messages cached in the singleton pipe.
credential_name	The credential name for the cloud store used to store messages. This parameter is only applicable to persistent pipes. The default value is NULL. A passed parameter takes precedence over the package argument's value. Credentials require the EXECUTE and READ/WRITE privileges.

Table 143-15 (Cont.) SEND_MESSAGE Function Parameters

Parameter	Description
location_uri	The location URL for the cloud store being used to store messages. This parameter is only applicable to persistent pipes. The location_uri parameter is a global variable that has a value of NULL by default. A passed parameter takes precedence over the global variable's value. If a credential name is specified, a location_uri must also be provided.

Return Values

Table 143-16 SEND_MESSAGE Function Return Values

Return	Description
0	Success. If the pipe already exists and the user attempting to create it is authorized to use it, then Oracle returns 0, indicating success, and any data already in the pipe remains. If a user connected as SYSDBS/SYSOPER re-creates a pipe, then Oracle returns status 0, but the ownership of the pipe remains unchanged.
1	Timed out. This procedure can timeout either because it cannot get a lock on the pipe, or because the pipe remains too full to be used. If the pipe was implicitly-created and is empty, then it is removed.
3	An interrupt occurred. If the pipe was implicitly created and is empty, then it is removed.
6	Failed to convert existing pipe to a singleton pipe. An implicit pipe with more than one existing message cannot be converted to a singleton pipe. For explicit pipe that is not a singleton pipe, SEND_MESSAGE cannot send a message with the singleton parameter set to TRUE.
7	A non-zero value was given for shelflife and the pipe is not a singleton pipe.
ORA-23322	Insufficient privileges. If a pipe with the same name exists and was created by a different user, then Oracle signals error ORA-23322, indicating the naming conflict.

Usage Notes

Persistent messages are guaranteed to either be written or read by exactly one process. This prevents message content inconsistency due to concurrent writes and reads. Using a persistent messaging pipe, DBMS_PIPE allows only one operation, sending a message or receiving a message to be active at a given time. However, if an operation is not possible due to an ongoing operation, the process retries periodically until the timeout value is reached.

If you use Oracle Cloud Infrastructure Object Storage to store messages, you can use Oracle Cloud Infrastructure Native URIs or Swift URIs. However, the location URI and the credential must match in type as follows:

- If you use a native URI format to access Oracle Cloud Infrastructure Object Storage, you must use Native Oracle Cloud Infrastructure Signing Keys authentication in the credential object.
- If you use Swift URI format to access Oracle Cloud Infrastructure Object Storage, you must use an auth token authentication in the credential object.

Exceptions

Table 143-17 SEND_MESSAGE Function Exception

Exception	Description
`Null pipe name`	Permission error. Insufficient privilege to write to the pipe. The pipe is private and owned by someone else.

143.7.11 SET_CREDENTIAL_NAME Procedure

This procedure sets the global `credential_name` variable to be used as the default credential with the cloud. The `SET_CREDENTIAL_NAME` procedure is only applicable to persistent pipes with messages stored in Cloud Object Storage.

Syntax

```
DBMS_PIPE.SET_CREDENTIAL_NAME (
   credential_name    IN VARCHAR2);
```

Parameters

Table 143-18 SET_CREDENTIAL_NAME Procedure Parameters

Parameter	Description
`credential_name`	The credential name for the cloud store used to store messages.

143.7.12 SET_LOCATION_URI Procedure

This procedure sets the global `set_location_uri` variable to be used as the default location URI with the cloud. The `SET_LOCATION_URI` procedure is only applicable to persistent pipes with messages stored in Cloud Object Storage.

Syntax

```
DBMS_PIPE.SET_LOCATION_URI (
   location_uri    IN VARCHAR2);
```

Parameters

Table 143-19 SET_LOCATION_URI Procedure

Parameter	Description
location_uri	The location URI for the cloud store used to store messages.

143.7.13 UNIQUE_SESSION_NAME Function

This function receives a name that is unique among all of the sessions that are currently connected to a database.

Multiple calls to this function from the same session always return the same value. You might find it useful to use this function to supply the PIPENAME parameter for your SEND_MESSAGE and RECEIVE_MESSAGE calls.

Syntax

```
DBMS_PIPE.UNIQUE_SESSION_NAME
  RETURN VARCHAR2;
```

Pragmas

```
pragma restrict_references(unique_session_name,WNDS,RNDS,WNPS);
```

Return Values

This function returns a unique name. The returned name can be up to 30 bytes.

143.7.14 UNPACK_MESSAGE Procedures

This procedure retrieves items from the buffer.

After you have called RECEIVE_MESSAGE to place pipe information in a local buffer, call UNPACK_MESSAGE.

> **Note:**
>
> The UNPACK_MESSAGE procedure is overloaded to return items of type VARCHAR2, NCHAR, NUMBER, or DATE. There are two additional procedures to unpack RAW and ROWID items.

Syntax

```
DBMS_PIPE.UNPACK_MESSAGE (
   item  OUT VARCHAR2);

DBMS_PIPE.UNPACK_MESSAGE (
   item  OUT NCHAR);

DBMS_PIPE.UNPACK_MESSAGE (
   item  OUT NUMBER);
```

```
DBMS_PIPE.UNPACK_MESSAGE (
   item  OUT DATE);

DBMS_PIPE.UNPACK_MESSAGE_RAW (
   item  OUT RAW);

DBMS_PIPE.UNPACK_MESSAGE_ROWID (
   item  OUT ROWID);
```

Pragmas

```
pragma restrict_references(unpack_message,WNDS,RNDS);
pragma restrict_references(unpack_message_raw,WNDS,RNDS);
pragma restrict_references(unpack_message_rowid,WNDS,RNDS);
```

Parameters

Table 143-20 UNPACK_MESSAGE Procedure Parameters

Parameter	Description
item	Argument to receive the next unpacked item from the local message buffer.

Exceptions

ORA-06556 or 06559 are generated if the buffer contains no more items, or if the item is not of the same type as that requested.

144
DBMS_PLSQL_CODE_COVERAGE

The `DBMS_PLSQL_CODE_COVERAGE` package provides an interface for the collection of code coverage data of PL/SQL applications at the basic block level.

This chapter contains the following topics:

- DBMS_PLSQL_CODE_COVERAGE Overview
- DBMS_PLSQL_CODE_COVERAGE Security Model
- DBMS_PLSQL_CODE_COVERAGE Data Structures
- Summary of DBMS_PLSQL_CODE_COVERAGE Subprograms

144.1 DBMS_PLSQL_CODE_COVERAGE Overview

The `DBMS_PLSQL_CODE_COVERAGE` package provides an interface for collecting code coverage information at the basic block level of PL/SQL applications. A basic block refers to a single entry single exit block of PL/SQL code. PL/SQL developers want to know how well their test infrastructure exercised their code. The coverage tables are created using the `CREATE_COVERAGE_TABLES` procedure.

A typical code coverage run in a session involves calls to :

- `START_COVERAGE`
- Run PL/SQL code
- `STOP_COVERAGE`

The `GET_BLOCK_MAP` function helps you calculate your total coverage.

See Also:

- *Oracle Database Development Guide* for more information about using PL/SQL basic block coverage to maintain quality
- *Oracle Database PL/SQL Language Reference* for the `COVERAGE PRAGMA` syntax and semantics
- *Oracle Database PL/SQL Language Reference* for more information about the PLSQL_OPTIMIZE_LEVEL compilation parameter

144.2 DBMS_PLSQL_CODE_COVERAGE Security Model

The user must have `EXECUTE` privilege on the DBMS_PLSQL_CODE_COVERAGE package.

The user must have `CREATE` privilege on the unit to collect coverage information about this unit.

PL/SQL basic block coverage data is collected when program units use `INTERPRETED` compilation (parameter set `PLSQL_CODE_TYPE = INTERPRETED`). PL/SQL basic block coverage data is not collected when program units use `NATIVE` compilation. You can disable the `NATIVE` compiler by setting the parameter `PLSQL_OPTIMIZE_LEVEL` <= `1`. Regardless of the compilation mode, coverage data for wrapped units is not collected.

144.3 DBMS_PLSQL_CODE_COVERAGE Constants

The `DBMS_PLSQL_CODE_COVERAGE` package provides constants that are used with the `namespace` parameter of the `GET_BLOCK_MAP` function.

These constants are described in the following table.

Table 144-1 DBMS_PLSQL_CODE_COVERAGE Constants

Name	Type	Value	Description
`function_namespace`	`NUMBER`	1	Specifies the function namespace
`package_spec_namespace`	`NUMBER`	1	Specifies the package specification namespace
`package_body_namespace`	`NUMBER`	2	Specifies the package definition (body) namespace
`procedure_namespace`	`NUMBER`	1	Specifies the procedure namespace
`trigger_namespace`	`NUMBER`	3	Specifies the trigger namespace
`type_spec_namespace`	`NUMBER`	1	Specifies the type specification namespace
`type_body_namespace`	`NUMBER`	2	Specifies the type definition (body) namespace

144.4 DBMS_PLSQL_CODE_COVERAGE Data Structures

Record Types

- MAP_REC Record Type

Table Types

- T_MAP_REC Table Type

144.4.1 MAP_REC Record Type

The MAP_REC record type defines the PL/SQL basic block location in the source code.

Syntax

```
TYPE map_rec IS RECORD (
    procedure_name VARCHAR2(32767),
    block_num      NUMBER,
    line           NUMBER,
    col            NUMBER,
```

```
    not_feasible    NUMBER);
```

Fields

Table 144-2 MAP_REC Fields

Field	Description
procedure_name	The name of the procedure containing the basic block
block_num	Identifies the basic block
line	Starting line of the basic block
col	Starting column of the basic block
not_feasible	Not_feasible marking of the basic block

144.4.2 T_MAP_REC Table Type

The T_MAP_REC table type specifies the collection of PL/SQL basic blocks in a unit.

Syntax

```
TYPE t_map_rec IS TABLE OF map_rec;
```

144.5 Summary of DBMS_PLSQL_CODE_COVERAGE Subprograms

This table lists the DBMS_PLSQL_CODE_COVERAGE subprograms and briefly describes them.

Table 144-3 DBMS_PLSQL_CODE_COVERAGE Package Subprograms

Subprogram	Description
CREATE_COVERAGE_TABLES Procedure	Creates coverage tables
GET_BLOCK_MAP Function	Gets the mapping of basic blocks to PL/SQL source
START_COVERAGE Function	Starts the coverage data collection in the user's session and returns the RUN_ID
STOP_COVERAGE Procedure	Ends the current coverage run

144.5.1 CREATE_COVERAGE_TABLES Procedure

This procedure creates the tables used for coverage data collection.

Syntax

```
DBMS_PLSQL_CODE_COVERAGE.CREATE_COVERAGE_TABLES (
   FORCE_IT     IN BOOLEAN DEFAULT FALSE);
```

Parameters

Parameter	Description
FORCE_IT	The default is to raise an error if the coverage tables already exists. If set to TRUE, the tables are dropped silently if the tables already exist, and new tables are created.

Exceptions

Table 144-4 CREATE_COVERAGE_TABLES Exceptions

Exception	Description
COVERAGE_ERROR	The FORCE_IT parameter is FALSE and the tables already exist.

144.5.2 GET_BLOCK_MAP Function

This function gets the mapping of basic blocks to PL/SQL source.

Syntax

```
DBMS_PLSQL_CODE_COVERAGE.GET_BLOCK_MAP(
                      unit_owner IN VARCHAR2,
                      unit_name  IN VARCHAR2,
                      namespace  IN POSITIVE)
  RETURN T_MAP_REC;
```

Table 144-5 Parameters

Parameter	Description
unit_owner	The owner of the unit. The unit owner is case insensitive. If the unit_owner is empty or NULL, then it defaults to the current schema.
unit_name	The unit whose mapping is to be gotten. The unit_name is case insensitive.
namespace	Namespace to which this unit_name gets resolved. See DBMS_PLSQL_CODE_COVERAGE Constants for a list of valid namespace values.

144.5.3 START_COVERAGE Function

This function starts the coverage data collection in the user's session and returns a unique identifier RUN_ID for the run.

Syntax

```
DBMS_PLSQL_CODE_COVERAGE.START_COVERAGE (
   run_comment   IN VARCHAR2)
   RETURN NUMBER;
```

Parameters

Parameter	Description
run_comment	Allows the user to name a run and identify the test.

144.5.4 STOP_COVERAGE Procedure

This procedure ends the current coverage run.

Syntax

```
DBMS_PLSQL_CODE_COVERAGE.STOP_COVERAGE;
```

Exceptions

Table 144-6 STOP_COVERAGE Exceptions

Exception	Description
COVERAGE_ERROR	An error is raised if the coverage tables do not exist.

145
DBMS_PREDICTIVE_ANALYTICS

Machine learning can discover useful information buried in vast amounts of data. However, both the programming interfaces and the machine learning expertise required to obtain these results are too complex for use by the wide audiences that can obtain benefits from using Oracle Machine Learning for SQL.

The `DBMS_PREDICTIVE_ANALYTICS` package addresses both of these complexities by automating the entire machine learning process from data preprocessing through model building to scoring new data. This package provides an important tool that makes machine learning possible for a broad audience of users, in particular, business analysts.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_PREDICTIVE_ANALYTICS Subprograms

145.1 DBMS_PREDICTIVE_ANALYTICS Overview

DBMS_PREDICTIVE_ANALYTICS automates parts of the machine learning process.

Machine learning, according to a commonly used process model, requires the following steps:

1. Understand the business problem.
2. Understand the data.
3. Prepare the data for mining.
4. Create models using the prepared data.
5. Evaluate the models.
6. Deploy and use the model to score new data.

`DBMS_PREDICTIVE_ANALYTICS` automates parts of step 3 — 5 of this process.

Predictive analytics procedures analyze and prepare the input data, create and test machine learning models using the input data, and then use the input data for scoring. The results of scoring are returned to the user. The models and supporting objects are not preserved after the operation completes.

145.2 DBMS_PREDICTIVE_ANALYTICS Security Model

The `DBMS_PREDICTIVE_ANALYTICS` package is owned by user `SYS` and is installed as part of database installation. Execution privilege on the package is granted to public. The routines in the package are run with invokers' rights (run with the privileges of the current user).

The `DBMS_PREDICTIVE_ANALYTICS` package exposes APIs which are leveraged by the Oracle Machine Learning for SQL option. Users who wish to invoke procedures in this package require the `CREATE MINING MODEL` system privilege (as well as the `CREATE TABLE` and `CREATE VIEW` system privilege).

145.3 Summary of DBMS_PREDICTIVE_ANALYTICS Subprograms

This table lists and briefly describes the `DBMS_PREDICTIVE_ANALYTICS` package subprograms.

Table 145-1 DBMS_PREDICTIVE_ANALYTICS Package Subprograms

Subprogram	Purpose
EXPLAIN Procedure	Ranks attributes in order of influence in explaining a target column.
PREDICT Procedure	Predicts the value of a target column based on values in the input data.
PROFILE Procedure	Generates rules that identify the records that have the same target value.

145.3.1 EXPLAIN Procedure

The `EXPLAIN` procedure identifies the attributes that are important in explaining the variation in values of a target column.

The input data must contain some records where the target value is known (not `NULL`). These records are used by the procedure to train a model that calculates the attribute importance.

> **Note:**
>
> `EXPLAIN` supports `DATE` and `TIMESTAMP` datatypes in addition to the numeric, character, and nested datatypes supported by Oracle Machine Learning for SQL models.
>
> Data requirements for Oracle Machine Learning for SQL are described in *Oracle Machine Learning for SQL User's Guide*

The `EXPLAIN` procedure creates a result table that lists the attributes in order of their explanatory power. The result table is described in the Usage Notes.

Syntax

```
DBMS_PREDICTIVE_ANALYTICS.EXPLAIN (
     data_table_name     IN VARCHAR2,
     explain_column_name IN VARCHAR2,
     result_table_name   IN VARCHAR2,
     data_schema_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 145-2 EXPLAIN Procedure Parameters

Parameter	Description
`data_table_name`	Name of input table or view
`explain_column_name`	Name of the column to be explained

Table 145-2 (Cont.) EXPLAIN Procedure Parameters

Parameter	Description
result_table_name	Name of the table where results are saved
data_schema_name	Name of the schema where the input table or view resides and where the result table is created. Default: the current schema.

Usage Notes

The EXPLAIN procedure creates a result table with the columns described in Table 145-3.

Table 145-3 EXPLAIN Procedure Result Table

Column Name	Datatype	Description
ATTRIBUTE_NAME	VARCHAR2(30)	Name of a column in the input data; all columns except the explained column are listed in the result table.
EXPLANATORY_VALUE	NUMBER	Value indicating how useful the column is for determining the value of the explained column. Higher values indicate greater explanatory power. Value can range from 0 to 1. An individual column's explanatory value is independent of other columns in the input table. The values are based on how strong each individual column correlates with the explained column. The value is affected by the number of records in the input table, and the relations of the values of the column to the values of the explain column. An explanatory power value of 0 implies there is no useful correlation between the column's values and the explain column's values. An explanatory power of 1 implies perfect correlation; such columns should be eliminated from consideration for PREDICT. In practice, an explanatory power equal to 1 is rarely returned.
RANK	NUMBER	Ranking of explanatory power. Rows with equal values for explanatory_value have the same rank. Rank values are not skipped in the event of ties.

Example

The following example performs an EXPLAIN operation on the SUPPLEMENTARY_DEMOGRAPHICS table of Sales History.

```
--Perform EXPLAIN operation
BEGIN
    DBMS_PREDICTIVE_ANALYTICS.EXPLAIN(
        data_table_name      => 'supplementary_demographics',
        explain_column_name  => 'home_theater_package',
        result_table_name    => 'demographics_explain_result');
END;
/
--Display results
SELECT * FROM demographics_explain_result;

ATTRIBUTE_NAME                          EXPLANATORY_VALUE       RANK
--------------------------------------- ----------------- ----------
Y_BOX_GAMES                                  .524311073          1
YRS_RESIDENCE                                .495987246          2
HOUSEHOLD_SIZE                               .146208506          3
AFFINITY_CARD                                  .0598227          4
```

```
EDUCATION                                .018462703          5
OCCUPATION                               .009721543          6
FLAT_PANEL_MONITOR                        .00013733          7
PRINTER_SUPPLIES                                  0          8
OS_DOC_SET_KANJI                                  0          8
BULK_PACK_DISKETTES                               0          8
BOOKKEEPING_APPLICATION                           0          8
COMMENTS                                          0          8
CUST_ID                                           0          8
```

The results show that `Y_BOX_GAMES`, `YRS_RESIDENCE`, and `HOUSEHOLD_SIZE` are the best predictors of `HOME_THEATER_PACKAGE`.

145.3.2 PREDICT Procedure

The `PREDICT` procedure predicts the values of a target column.

The input data must contain some records where the target value is known (not `NULL`). These records are used by the procedure to train and test a model that makes the predictions.

> **Note:**
>
> `PREDICT` supports `DATE` and `TIMESTAMP` datatypes in addition to the numeric, character, and nested datatypes supported by Oracle Machine Learning for SQL models.
>
> Data requirements for Oracle Machine Learning for SQL are described in *Oracle Machine Learning for SQL User's Guide*

The `PREDICT` procedure creates a result table that contains a predicted target value for every record. The result table is described in the Usage Notes.

Syntax

```
DBMS_PREDICTIVE_ANALYTICS.PREDICT (
    accuracy                OUT NUMBER,
    data_table_name         IN VARCHAR2,
    case_id_column_name     IN VARCHAR2,
    target_column_name      IN VARCHAR2,
    result_table_name       IN VARCHAR2,
    data_schema_name        IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 145-4 PREDICT Procedure Parameters

Parameter	Description
`accuracy`	Output parameter that returns the predictive confidence, a measure of the accuracy of the predicted values. The predictive confidence for a categorical target is the most common target value; the predictive confidence for a numerical target is the mean.
`data_table_name`	Name of the input table or view.
`case_id_column_name`	Name of the column that uniquely identifies each case (record) in the input data.

Table 145-4 (Cont.) PREDICT Procedure Parameters

Parameter	Description
target_column_name	Name of the column to predict.
result_table_name	Name of the table where results will be saved.
data_schema_name	Name of the schema where the input table or view resides and where the result table is created. Default: the current schema.

Usage Notes

The PREDICT procedure creates a result table with the columns described in Table 145-5.

Table 145-5 PREDICT Procedure Result Table

Column Name	Datatype	Description
Case ID column name	VARCHAR2 or NUMBER	The name of the case ID column in the input data.
PREDICTION	VARCHAR2 or NUMBER	The predicted value of the target column for the given case.
PROBABILITY	NUMBER	For classification (categorical target), the probability of the prediction. For regression problems (numerical target), this column contains NULL.

> **Note:**
>
> Make sure that the name of the case ID column is not 'PREDICTION' or 'PROBABILITY'.

Predictions are returned for all cases whether or not they contained target values in the input.

Predicted values for known cases may be interesting in some situations. For example, you could perform deviation analysis to compare predicted values and actual values.

Example

The following example performs a PREDICT operation and displays the first 10 predictions. The results show an accuracy of 79% in predicting whether each customer has an affinity card.

```
--Perform PREDICT operation
DECLARE
    v_accuracy NUMBER(10,9);
BEGIN
    DBMS_PREDICTIVE_ANALYTICS.PREDICT(
        accuracy             => v_accuracy,
        data_table_name      => 'supplementary_demographics',
        case_id_column_name  => 'cust_id',
        target_column_name   => 'affinity_card',
        result_table_name    => 'pa_demographics_predict_result');
    DBMS_OUTPUT.PUT_LINE('Accuracy = ' || v_accuracy);
END;
/

Accuracy = .788696903
```

```
--Display results
SELECT * FROM pa_demographics_predict_result WHERE rownum < 10;

   CUST_ID PREDICTION PROBABILITY
---------- ---------- -----------
    101501          1  .834069848
    101502          0  .991269965
    101503          0   .99978311
    101504          1  .971643388
    101505          1  .541754127
    101506          0  .803719133
    101507          0  .999999303
    101508          0  .999999987
    101509          0  .999953074
```

145.3.3 PROFILE Procedure

The `PROFILE` procedure generates rules that describe the cases (records) from the input data.

For example, if a target column `CHURN` has values 'Yes' and 'No', `PROFILE` generates a set of rules describing the expected outcomes. Each profile includes a rule, record count, and a score distribution.

The input data must contain some cases where the target value is known (not `NULL`). These cases are used by the procedure to build a model that calculates the rules.

> **Note:**
>
> `PROFILE` does not support nested types or dates.
>
> Data requirements for Oracle Machine Learning for SQL are described in *Oracle Machine Learning for SQL User's Guide*

The `PROFILE` procedure creates a result table that specifies rules (profiles) and their corresponding target values. The result table is described in the Usage Notes.

Syntax

```
DBMS_PREDICTIVE_ANALYTICS.PROFILE (
     data_table_name           IN VARCHAR2,
     target_column_name        IN VARCHAR2,
     result_table_name         IN VARCHAR2,
     data_schema_name          IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 145-6 PROFILE Procedure Parameters

Parameter	Description
`data_table_name`	Name of the table containing the data to be analyzed.
`target_column_name`	Name of the target column.
`result_table_name`	Name of the table where the results will be saved.

Table 145-6 (Cont.) PROFILE Procedure Parameters

Parameter	Description
data_schema_name	Name of the schema where the input table or view resides and where the result table is created. Default: the current schema.

Usage Notes

The PROFILE procedure creates a result table with the columns described in Table 145-7.

Table 145-7 PROFILE Procedure Result Table

Column Name	Datatype	Description
PROFILE_ID	NUMBER	A unique identifier for this profile (rule).
RECORD_COUNT	NUMBER	The number of records described by the profile.
DESCRIPTION	SYS.XMLTYPE	The profile rule. See "XML Schema for Profile Rules".

XML Schema for Profile Rules

The DESCRIPTION column of the result table contains XML that conforms to the following XSD:

```
<xs:element name="SimpleRule">
  <xs:complexType>
    <xs:sequence>
      <xs:group ref="PREDICATE"/>
      <xs:element ref="ScoreDistribution" minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:attribute name="id" type="xs:string" use="optional"/>
    <xs:attribute name="score" type="xs:string" use="required"/>
    <xs:attribute name="recordCount" type="NUMBER" use="optional"/>
  </xs:complexType>
</xs:element>
```

Example

This example generates a rule describing customers who are likely to use an affinity card (target value is 1) and a set of rules describing customers who are not likely to use an affinity card (target value is 0). The rules are based on only two predictors: education and occupation.

```
SET serveroutput ON
SET trimspool ON
SET pages 10000
SET long 10000
SET pagesize 10000
SET linesize 150
CREATE VIEW cust_edu_occ_view AS
             SELECT cust_id, education, occupation, affinity_card
             FROM sh.supplementary_demographics;
BEGIN
    DBMS_PREDICTIVE_ANALYTICS.PROFILE(
         DATA_TABLE_NAME    => 'cust_edu_occ_view',
         TARGET_COLUMN_NAME => 'affinity_card',
         RESULT_TABLE_NAME  => 'profile_result');
END;
/
```

This example generates eight rules in the result table `profile_result`. Seven of the rules suggest a target value of 0; one rule suggests a target value of 1. The `score` attribute on a rule identifies the target value.

This `SELECT` statement returns all the rules in the result table.

```
SELECT a.profile_id, a.record_count, a.description.getstringval()
 FROM profile_result a;
```

This `SELECT` statement returns the rules for a target value of 0.

```
SELECT *
  FROM profile_result t
  WHERE extractvalue(t.description, '/SimpleRule/@score') = 0;
```

The eight rules generated by this example are displayed as follows.

```
<SimpleRule id="1" score="0" recordCount="443">
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"Armed-F" "Exec." "Prof." "Protec."
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">"< Bach." "Assoc-V" "HS-grad"
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="0" recordCount="297" />
  <ScoreDistribution value="1" recordCount="146" />
</SimpleRule>

<SimpleRule id="2" score="0" recordCount="18">
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"Armed-F" "Exec." "Prof." "Protec."
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">"10th" "11th" "12th" "1st-4th" "5th-6th" "7th-8th" "9th" "Presch."
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="0" recordCount="18" />
</SimpleRule>

<SimpleRule id="3" score="0" recordCount="458">
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"Armed-F" "Exec." "Prof." "Protec."
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">"Assoc-A" "Bach."
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="0" recordCount="248" />
  <ScoreDistribution value="1" recordCount="210" />
</SimpleRule>

<SimpleRule id="4" score="1" recordCount="276">
```

```
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"Armed-F" "Exec." "Prof." "Protec."
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">"Masters" "PhD" "Profsc"
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="1" recordCount="183" />
  <ScoreDistribution value="0" recordCount="93" />
</SimpleRule>

<SimpleRule id="5" score="0" recordCount="307">
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">"Assoc-A" "Bach." "Masters" "PhD" "Profsc"
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"Crafts" "Sales" "TechSup" "Transp."
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="0" recordCount="184" />
  <ScoreDistribution value="1" recordCount="123" />
</SimpleRule>

<SimpleRule id="6" score="0" recordCount="243">
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">"Assoc-A" "Bach." "Masters" "PhD" "Profsc"
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"?" "Cleric." "Farming" "Handler" "House-s" "Machine" "Other"
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="0" recordCount="197" />
  <ScoreDistribution value="1" recordCount="46" />
</SimpleRule>

<SimpleRule id="7" score="0" recordCount="2158">
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">
        "10th" "11th" "12th" "1st-4th" "5th-6th" "7th-8th" "9th" "< Bach." "Assoc-V" "HS-grad"
        "Presch."
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"?" "Cleric." "Crafts" "Farming" "Machine" "Sales" "TechSup" " Transp."
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
  <ScoreDistribution value="0" recordCount="1819"/>
  <ScoreDistribution value="1" recordCount="339"/>
</SimpleRule>

<SimpleRule id="8" score="0" recordCount="597">
```

```
  <CompoundPredicate booleanOperator="and">
    <SimpleSetPredicate field="EDUCATION" booleanOperator="isIn">
      <Array type="string">
        "10th" "11th" "12th" "1st-4th"  "5th-6th" "7th-8th" "9th" "< Bach." "Assoc-V" "HS-grad"
        "Presch."
      </Array>
    </SimpleSetPredicate>
    <SimpleSetPredicate field="OCCUPATION" booleanOperator="isIn">
      <Array type="string">"Handler" "House-s" "Other"
      </Array>
    </SimpleSetPredicate>
  </CompoundPredicate>
<ScoreDistribution value="0" recordCount="572"/>
<ScoreDistribution value="1" recordCount="25"/>
</SimpleRule>
```

146
DBMS_PREPROCESSOR

The DBMS_PREPROCESSOR package provides an interface to print or retrieve the source text of a PL/SQL unit in its post-processed form.

This package contains the following topics:

- Overview
- Operating Notes
- Data Structures
- Summary of DBMS_PREPROCESSOR Subprograms

146.1 DBMS_PREPROCESSOR Overview

It is necessary to first understand the three styles of subprograms, in order to understand how `DBMS_PREPROCESSOR` works.

The following are the three styles of subprograms:

1. Subprograms that take a schema name, a unit type name, and the unit name.
2. Subprograms that take a `VARCHAR2` string which contains the source text of an arbitrary PL/SQL compilation unit.
3. Subprograms that take a `VARCHAR2` index-by table which contains the segmented source text of an arbitrary PL/SQL compilation unit.

Subprograms of the first style are used to print or retrieve the post-processed source text of a stored PL/SQL unit. The user must have the privileges necessary to view the original source text of this unit. The user must also specify the schema in which the unit is defined, the type of the unit, and the name of the unit. If the schema is null, then the current user schema is used. If the status of the stored unit is `VALID` and the user has the required privilege, then the post-processed source text is guaranteed to be the same as that of the unit the last time it was compiled.

Subprograms of the second or third style are used to generate post-processed source text in the current user schema. The source text is passed in as a single `VARCHAR2` string in the second style, or as a `VARCHAR2` index-by table in the third style. The source text can represent an arbitrary PL/SQL compilation unit. A typical usage is to pass the source text of an anonymous block and generate its post-processed source text in the current user schema. The third style can be useful when the source text exceeds the `VARCHAR2` length limit.

146.2 DBMS_PREPROCESSOR Operating Notes

These notes explain how DBMS_PREPROCESSOR works with the three subprogram styles.

- For subprograms of the first style, the status of the stored PL/SQL unit does not need to be `VALID`. Likewise, the source text passed in as a `VARCHAR2` string or a `VARCHAR2` index-by table may contain compile time errors. If errors are found when generating the post-processed source, the error message text will also appear at the end of the post-processed

source text. In some cases, the preprocessing can be aborted because of errors. When this happens, the post-processed source text will appear to be incomplete and the associated error message can help to indicate that an error has occurred during preprocessing.

- For subprograms of the second or third style, the source text can represent any arbitrary PL/SQL compilation unit. However, the source text of a valid PL/SQL compilation unit cannot include commonly used prefixes such as `CREATE OR REPLACE`. In general, the input source should be syntactically prepared in a way as if it were obtained from the `ALL_SOURCE` view. The following list gives some examples of valid initial syntax for some PL/SQL compilation units.

```
anonymous block   (BEGIN | DECLARE) ...
package           PACKAGE <name> ...
package body      PACKAGE BODY <name> ...
procedure         PROCEDURE <name> ...
function          FUNCTION <name> ...
type              TYPE <name> ...
type body         TYPE BODY <name> ...
trigger           (BEGIN | DECLARE) ...
```

  If the source text represents a named PL/SQL unit that is valid, that unit will not be created after its post-processed source text is generated.

- If the text of a wrapped PL/SQL unit is obtained from the `ALL_SOURCE` view, the keyword `WRAPPED` always immediately follows the name of the unit, as in this example:

```
PROCEDURE "some proc" WRAPPED
a000000
b2
...
```

  If such source text is presented to one of the GET_POST_PROCESSED_SOURCE Functions or to one of the PRINT_POST_PROCESSED_SOURCE Procedures, the exception `DBMS_PREPROCESSOR.WRAPPED_INPUT` is raised.

146.3 DBMS_PREPROCESSOR Data Structures

The `DBMS_PREPROCESSOR` package defines a `TABLE` type.

Table Types

SOURCE_LINES_T Table Type

146.3.1 DBMS_PREPROCESSOR SOURCE_LINES_T Table Type

This table type stores lines of post-processed source text. It is used to hold PL/SQL source text both before and after it is processed. It is especially useful in cases in which the amount of text exceeds 32K.

Syntax

```
TYPE source_lines_t IS
    TABLE OF VARCHAR2(32767) INDEX BY BINARY_INTEGER;
```

146.4 Summary of DBMS_PREPROCESSOR Subprograms

This table lists the `DBMS_PREPROCESSOR` subprograms and briefly describes them.

Table 146-1 DBMS_PREPROCESSOR Package Subprograms

Subprogram	Description
GET_POST_PROCESSED_SOURCE Functions	Returns the post-processed source text
PRINT_POST_PROCESSED_SOURCE Procedures	Prints post-processed source text

146.4.1 GET_POST_PROCESSED_SOURCE Functions

This overloaded function returns the post-processed source text. The different functionality of each form of syntax is presented along with the definition.

Syntax

Returns post-processed source text of a stored PL/SQL unit:

```
DBMS_PREPROCESSOR.GET_POST_PROCESSED_SOURCE (
   object_type    IN VARCHAR2,
   schema_name    IN VARCHAR2,
   object_name    IN VARCHAR2)
  RETURN source_lines_t;
```

Returns post-processed source text of a compilation unit:

```
DBMS_PREPROCESSOR.GET_POST_PROCESSED_SOURCE (
   source        IN VARCHAR2)
  RETURN source_lines_t;
```

Returns post-processed source text of an `INDEX-BY` table containing the source text of the compilation unit:

```
DBMS_PREPROCESSOR.GET_POST_PROCESSED_SOURCE (
   source        IN source_lines_t)
  RETURN source_lines_t;
```

Parameters

Table 146-2 GET_POST_PROCESSED_SOURCE Function Parameters

Parameter	Description
`object_type`	Must be one of `PACKAGE`, `PACKAGE BODY`, `PROCEDURE`, `FUNCTION`, `TYPE`, `TYPE`, `BODY` or `TRIGGER`. Case sensitive.
`schema_name`	The schema name. Case insensitive unless a quoted identifier is used. If `NULL`, use current schema.
`object_name`	The name of the object.The `object_type` is always case insensitive. Case insensitive unless a quoted identifier is used.
`source`	The source text of the compilation unit
`source_lines_t`	`INDEX-BY` table containing the source text of the compilation unit. The source text is a concatenation of all the non-`NULL` `INDEX-BY` table elements in ascending index order.

Return Values

The function returns an INDEX-BY table containing the lines of the post-processed source text starting from index 1.

Usage Notes

- Newline characters are not removed.
- Each line in the post-processed source text is mapped to a row in the INDEX-BY table.
- In the post-processed source, unselected text will have blank lines.

Exceptions

Table 146-3 GET_POST_PROCESSED_SOURCE Function Exceptions

Exception	Description
ORA-24234	Insufficient privileges or object does not exist
ORA-24235	Bad value for object type. Should be one of PACKAGE, PACKAGE BODY, PROCEDURE, FUNCTION, TYPE, TYPE, BODY or TRIGGER.
ORA-24236	The source text is empty
ORA-00931	Missing identifier. The object_name should not be NULL.
ORA-06502	Numeric or value error: • Character string buffer too small • A line is too long (> 32767 bytes)

146.4.2 PRINT_POST_PROCESSED_SOURCE Procedures

This overloaded procedure calls DBMS_OUTPUT.PUT_LINE to let you view post-processed source text. The different functionality of each form of syntax is presented along with the definition.

Syntax

Prints post-processed source text of a stored PL/SQL unit:

```
DBMS_PREPROCESSOR.PRINT_POST_PROCESSED_SOURCE (
   object_type    IN VARCHAR2,
   schema_name    IN VARCHAR2,
   object_name    IN VARCHAR2);
```

Prints post-processed source text of a compilation unit:

```
DBMS_PREPROCESSOR.PRINT_POST_PROCESSED_SOURCE (
   source        IN VARCHAR2);
```

Prints post-processed source text of an INDEX-BY table containing the source text of the compilation unit:

```
DBMS_PREPROCESSOR.PRINT_POST_PROCESSED_SOURCE (
   source        IN source_lines_t);
```

Parameters

Table 146-4 PRINT_POST_PROCESSED_SOURCE Procedure Parameters

Parameter	Description
object_type	Must be one of PACKAGE, PACKAGE BODY, PROCEDURE, FUNCTION, TYPE, TYPE, BODY or TRIGGER. Case sensitive.
schema_name	The schema name. Case insensitive unless a quoted identifier is used. If NULL, use current schema.
object_name	The name of the object.The object_type is always case insensitive. Case insensitive unless a quoted identifier is used.
source	The source text of the compilation unit
source_lines_t	INDEX-BY table containing the source text of the compilation unit. The source text is a concatenation of all the non-NULL INDEX-BY table elements in ascending index order.

Exceptions

Table 146-5 PRINT_POST_PROCESSED_SOURCE Procedure Exceptions

Exception	Description
ORA-24234	Insufficient privileges or object does not exist
ORA-24235	Bad value for object type. Should be one of PACKAGE, PACKAGE BODY, PROCEDURE, FUNCTION, TYPE, TYPE, BODY or TRIGGER.
ORA-24236	The source text is empty
ORA-00931	Missing identifier. The object_name should not be NULL.
ORA-06502	Numeric or value error: • Character string buffer too small • A line is too long (> 32767 bytes)

Usage Notes

The index-by table may contain holes. NULL elements are ignored when doing the concatenation.

147
DBMS_PRIVILEGE_CAPTURE

The `DBMS_PRIVILEGE_CAPTURE` package provides an interface to database privilege analysis.

> **See Also:**
>
> Oracle® Database Security Guide regarding on how to analyze the use of privilege grants

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Examples
- Summary of DBMS_PRIVILEGE_CAPTURE Subprograms

147.1 DBMS_PRIVILEGE_CAPTURE Overview

Database privilege analysis enables you to create a policy that records the usage of system and object privileges that have been granted to users. You then can determine the privileges that your users are using and not using. From there, you can revoke any unused privileges, thereby reducing the number of excess privilege grants for users.

By analyzing the privileges that users must have to perform specific tasks, privilege analysis policies help you to achieve a least privilege model for your users.

147.2 DBMS_PRIVILEGE_CAPTURE Security Model

The privilege analysis administrator role, `CAPTURE_ADMIN`, is granted `EXECUTE` permission on the `DBMS_PRIVILEGE_CAPTURE` package by default.

The `CAPTURE_ADMIN` role is granted to the `DBA` role during database installation.

147.3 DBMS_PRIVILEGE_CAPTURE Constants

The `DBMS_PRIVILEGE_CAPTURE` package defines several enumerated constants for specifying parameter values.

Table 147-1 DBMS_PRIVILEGE_CAPTURE Constants

Constant	Value	Type	Description
G_DATABASE	1	NUMBER	Analyzes all privilege use, except privileges used by the SYS user.
G_ROLE	2	NUMBER	Analyzes privilege use for the specified roles.
G_CONTEXT	3	NUMBER	Analyzes privilege use when the condition parameter evaluates to true.
G_ROLE_AND_CONTEXT	4	NUMBER	Analyzes privilege use for the specified roles when the condition parameter evaluates to true.

147.4 DBMS_PRIVILEGE_CAPTURE Examples

These examples illustrate using the DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE procedure to create various types of privilege analysis, like database analysis, role analysis, and context-specific analysis. The examples also illustrate combining different conditions in context-specific analysis.

```
--Create a database privilege analysis policy
BEGIN
DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE(
        name         => 'all_priv_analysis_pol',
        description  => 'database-wide policy to analyze all privileges',
        type         => DBMS_PRIVILEGE_CAPTURE.G_DATABASE);
END;

--Create a privilege analysis policy to analyze privileges from the role PUBLIC
BEGIN
DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE(
       name         => 'pub_analysis_pol',
       description  => 'Policy to record privilege use by PUBLIC',
       type         => DBMS_PRIVILEGE_CAPTURE.G_ROLE,
       roles        => role_name_list('PUBLIC'));
END;

-- Create a policy to analyze privileges from the application module, "Account
-- Payable"
BEGIN
DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE(
  name         => 'acc_pay_analysis_pol',
  type         => DBMS_PRIVILEGE_CAPTURE.G_CONTEXT,
  condition    => 'SYS_CONTEXT(''USERENV'', ''MODULE'') = ''Account Payable''');
END;

-- Create a policy that records privileges for session user APPS when running the
-- application module "Account Payable"
BEGIN
DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE(
  name         => 'acc_pay_analysis_pol',
  type         => DBMS_PRIVILEGE_CAPTURE.G_CONTEXT,
  condition    => 'SYS_CONTEXT(''USERENV'', ''MODULE'') = ''Account Payable'' AND
                   SYS_CONTEXT(''USERENV'', ''SESSION_USER'') = ''APPS''');
END;
```

147.5 Summary of DBMS_PRIVILEGE_CAPTURE Subprograms

This table lists and briefly describes the `DBMS_PRIVILEGE_CAPTURE` package subprograms.

Table 147-2 DBMS_PRIVILEGE_CAPTURE Package Subprograms

Subprogram	Description
CAPTURE_DEPENDENCY_PRIVS Procedure	Captures the privileges that are used by definer's rights and invoker's rights PL/SQL program units for compilation.
CREATE_CAPTURE Procedure	Creates a policy that specifies the conditions for analyzing privilege use.
DELETE_RUN Procedure	Deletes a privilege analysis capture run
DISABLE_CAPTURE Procedure	Stops the recording of privilege use for a specified privilege analysis policy
DROP_CAPTURE Procedure	Removes a privilege analysis policy together with the data recorded
ENABLE_CAPTURE Procedure	Starts the recording of privilege analysis for a specified privilege analysis policy
GENERATE_RESULT Procedure	Populates the privilege analysis data dictionary views with data

147.5.1 CAPTURE_DEPENDENCY_PRIVS Procedure

This procedure captures the privileges that are used by definer's rights and invoker's rights PL/SQL program units for compilation.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.CAPTURE_DEPENDENCY_PRIVS ();
```

Parameters

This procedure has no parameters.

Usage Notes

Every rerun of the `DBMS_PRIVILEGE_CAPTURE.CAPTURE_DEPENDENCY_PRIVS` procedure deletes any existing records from the privilege analysis data dictionary views. It then recaptures records based on the existing PL/SQL program units.

147.5.2 CREATE_CAPTURE Procedure

This procedure creates a privilege analysis policy that specifies the conditions for analyzing privilege use. It also optionally specifies the roles for which privilege use is to be analyzed, and the conditions under which privilege use is to be analyzed.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE (
   name            IN  VARCHAR2,
   description     IN  VARCHAR2 DEFAULT NULL,
   type            IN  NUMBER DEFAULT G_DATABASE,
```

```
   roles           IN  ROLE_NAME_LIST DEFAULT ROLE_NAME_LIST(),
   condition       IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 147-3 CREATE_CAPTURE Procedure Parameters

Parameter	Description
name	Name of the privilege analysis policy. A string of size up to 30 characters.
description	Description of the policy (up to 1024 characters)
type	Type of the privilege analysis policy. Possible values are: • G_DATABASE: Captures all privilege use in the database, except privileges used by the SYS user. • G_ROLE: Captures the use of a privilege if the privilege is part of a specified role or list of roles. • G_CONTEXT: Captures the use of a privilege if the context specified by the condition parameter evaluates to true. • G_ROLE_AND_CONTEXT: Captures the use of a privilege if the privilege is part of the specified list of roles and when the condition specified by the condition parameter is true.
roles	The roles whose privileges are to be analyzed. Required if the type is G_ROLE or G_ROLE_AND_CONTEXT.
condition	PL/SQL boolean expression containing up to 4000 characters. Required if type is G_CONTEXT or G_ROLE_AND_CONTEXT. Note that the boolean expression can only contain SYS_CONTEXT, but not other functions.

Usage Notes

- When using role-based analysis for the CREATE_CAPTURE procedure, privilege use is analyzed even if the privilege is indirectly granted to the specified role.

  For example, say role R2 contains role R1, and R1 contains privilege P1. If the privilege policy includes only role R2, any use of the P1 privilege is still analyzed, as privilege P1 is an indirect part of role R2.

- When using the condition parameter, use the following syntax for the PL/SQL expression:

  ```
  condition::= predicate | (predicate1) AND (predicate2)
  | (predicate1) OR (predicate2)
  ```

  Where,

  ```
  predicate::= sys_context(namespace, attribute) relop constant_value |
  sys_context(namespace, attribute) between constant_value and
   constant_value | sys_context(namespace, attribute) in {constant_value
  (,constant_value)* }
  ```

  Where,

  ```
  relop::= = | < | <= | > | >= | <>
  ```

- A privilege analysis policy cannot analyze the use of SYS user privileges.

147.5.3 DELETE_RUN Procedure

This procedure deletes a privilege analysis capture run.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.DELETE_RUN (
   name      IN VARCHAR2,
   run_name  IN VARCHAR2);
```

Parameters

Table 147-4 *DELETE_RUN Procedure Parameters*

Parameter	Description
`name`	Name of the privilege analysis policy with which the capture run is associated
`run_name`	Name of the capture run

Usage Notes

- You can find the names of existing privilege capture policies by querying the `DBA_PRIV_CAPTURES` data dictionary view.
- Another way to delete a capture run is to drop the policy with which the capture run is associated. Dropping the policy automatically drops its associated capture runs.
- When you drop a capture run it is no longer accessible through the privilege capture data dictionary views.

147.5.4 DISABLE_CAPTURE Procedure

This procedure stops the recording of privilege use for a specified privilege analysis policy. When a policy is disabled, privilege use meeting the policy condition is no longer recorded.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.DISABLE_CAPTURE (
   name        IN VARCHAR2);
```

Parameters

Table 147-5 DISABLE_CAPTURE Procedure Parameters

Parameter	Description
`name`	Name of the privilege analysis policy to be disabled

Usage Notes

When a privilege analysis policy is first created, it is disabled by default.

147.5.5 DROP_CAPTURE Procedure

This procedure removes a privilege analysis policy together with the data recorded. When a policy is removed, all previously recorded privilege use data associated with the policy is deleted.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.DROP_CAPTURE (
   name      IN VARCHAR2);
```

Parameters

Table 147-6 *DROP_CAPTURE Procedure Parameters*

Parameter	Description
name	Name of the privilege analysis policy to be removed

Usage Notes

- You must disable a privilege analysis policy before removing it. An enabled policy cannot be removed.
- If there are capture runs associated with this policy, then they are automatically dropped when you drop the policy.

147.5.6 ENABLE_CAPTURE Procedure

This procedure starts the recording of privilege analysis for a specified privilege analysis policy and optionally provides a capture run for this policy. After a policy is enabled, all privilege use under the policy condition is recorded.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.ENABLE_CAPTURE (
   name      IN VARCHAR2,
   run_name  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 147-7 *ENABLE_CAPTURE Procedure Parameters*

Parameter	Description
name	Name of the privilege analysis policy to be enabled
run_name	Name of the capture run to associate with this policy, less than 128 characters. Enclose exotic characters in double quotation marks.

Usage Notes

The following usage notes apply:

- When a privilege analysis policy is first created, it is disabled by default. You must run ENABLE_CAPTURE to enable the privilege analysis policy.

- You can enable only one privilege analysis policy at a time. However, a database-wide privilege analysis of the `G_DATABASE` type can be enabled together with another non `G_DATABASE` privilege analysis.
- You cannot enable the same run multiple times. For example, `run_01` cannot be used again if you want to re-enable the capture for `run_01`. Instead, create a new run.

147.5.7 GENERATE_RESULT Procedure

This procedure populates the privilege analysis data dictionary views with data.

> **See Also:**
>
> Oracle® Database Security Guide for more information about privilege analysis views.

Syntax

```
DBMS_PRIVILEGE_CAPTURE.GENERATE_RESULT (
   name        IN VARCHAR2,
   run_name    IN VARCHAR2 DEFAULT NULL,
   DEPENDENCY  IN BOOLEAN DEFAULT NULL);
```

Parameters

Table 147-8 *GENERATE_RESULT Procedure Parameters*

Parameter	Description
`name`	Name of the privilege analysis policy for which views are populated
`run_name`	Name of the capture run that is associated with the privilege analysis policy. If you omit this parameter, then the records of all created runs will be analyzed. When you specify the `run_name` parameter, only the records of that run are analyzed and all other runs are unaffected.
`dependency`	Enter `Y` (yes) or `N` (no) to indicate if PL/SQL compilation privileges, set by the `DBMS_PRIVILEGE_CAPTURE.CAPTURE_DEPENDENCY_PRIVS` procedure, should be included in the report.

Usage Notes

You must disable a privilege analysis policy before populating the privilege analysis views for the policy. You cannot invoke this subprogram on an enabled privilege analysis policy.

148
DBMS_PROCESS

The `DBMS_PROCESS` package provides an interface to manage the prespawned servers.

This chapter contains the following topics:

- DBMS_PROCESS Overview
- DBMS_PROCESS Security Model
- Summary of DBMS_PROCESS Subprograms

148.1 DBMS_PROCESS Overview

By default, Oracle Database can prespawn foreground processes to improve the performance of client connections.

A prespawned process refers to a process that has been spawned but does not have a session yet. When a user connects to the database or a service process is needed, the process performs further initialization as needed.

To manage foreground processes, use the `DBMS_PROCESS` package. The procedures in this package configure the number of foreground processes for a connection pool, start a connection pool, and stop a connection pool.

See Also:

- *Oracle Database Administrator's Guide* for a detailed description of managing a pre-created processes
- *Oracle Database Reference* for information about the `USE_DEDICATED_BROKER` parameter, which determines how dedicated servers are spawned. `TRUE` enables the listener to send the connection to a dedicated connection broker that spawns the dedicated server. `FALSE` sets the listener to spawn a dedicated server directly.

148.2 DBMS_PROCESS Security Model

You must be granted the `SYSDBA` administrative privilege to use the `DBMS_PROCESS` package.

148.3 Summary of DBMS_PROCESS Subprograms

This table lists the `DBMS_PROCESS` subprograms and briefly describes them.

Table 148-1 DBMS_PROCESS Package Subprograms

Subprogram	Description
CONFIGURE_POOL Procedure	Configures the minimum number of prespawned processes, the number of prespawned processes in a batch, and the initial number of pre-spawned processes for a foreground connection pool
START_POOL Procedure	Starts a prespawned foreground process pool.
STOP_POOL Procediure	Stops a prespawned foreground process pool.

148.3.1 CONFIGURE_POOL Procedure

This procedure configures the minimum number of prespawned processes, the number of prespawned processes in a batch, and the initial number of pre-spawned processes for a foreground connection pool.

Syntax

```
DBMS_PROCESS.CONFIGURE_POOL(
   pool_name       IN  VARCHAR2 DEFAULT "SYS_DEFAULT_FOREGROUND_POOL",
   min_count       IN  NUMBER DEFAULT 10,
   batch_count     IN  NUMBER DEFAULT 20,
   init_count      IN  NUMBER DEFAULT 0);
```

Parameters

Table 148-2 CONFIGURE_POOL Procedure Parameters

Parameter	Description
`pool_name`	Name of the foreground pool
`min_count`	Count after which a batch of prespawned processes will be spawned in the pool (default 10; maximum number of processes 64,000)
`batch_count`	Batch of prespawned servers to be spawned in the pool (default 20). The total value of the `min_count` and `batch_count` values cannot exceed the `PROCESSES` parameter value.
`init_count`	Set of prespawned servers to be spawned in startup, or before a storm of connections, spawned in batches (default 0). Do not exceed the `PROCESSES` parameter value.

Exceptions

Table 148-3 CONFIGURE_POOL Procedure Exceptions

Error	Description
ORA-456	Prespawn not enabled
ORA-457	Prespawn pool not found
ORA-458	Prespawn pool already started
ORA-459	Prespawn pool not started
ORA-460	Process pool invalid configuration value(s)

Usage Notes

- To find information about foreground processes, query the V$PROCESS_POOL dynamic view.

Example

```
BEGIN
  DBMS_PROCESS.CONFIGURE_POOL (
   pool_name   => 'hrdb_pool',
   min_count   => 40,
   batch_count => 20,
   init_count  => 10);
END;
/
```

148.3.2 START_POOL Procedure

This procedure starts a prespawned foreground process pool.

Syntax

```
DBMS_PROCESS.START_POOL(
   pool_name      IN  VARCHAR2 DEFAULT "SYS_DEFAULT_FOREGROUND_POOL");
```

Parameters

Table 148-4 START_POOL Procedure Parameters

Parameter	Description
pool_name	Name of prespawned foreground process pool. To find existing pools, query the V$PROCESS_POOL dynamic view.

Exceptions

Table 148-5 START_POOL Procedure Exceptions

Error	Description
ORA-457	Prespawn pool not found
ORA-458	Prespawn pool already started

Table 148-5 (Cont.) START_POOL Procedure Exceptions

Error	Description
ORA-459	Prespawn pool not started

Example

```
BEGIN
  DBMS_PROCESS.START_POOL (
   pool_name   => 'hrdb_pool',
END;
```

148.3.3 STOP_POOL Procedure

This procedure stops a prespawned foreground process pool.

Syntax

```
DBMS_PROCESS.STOP_POOL(
   pool_name      IN  VARCHAR2 DEFAULT "SYS_DEFAULT_FOREGROUND_POOL");
```

Parameters

Table 148-6 STOP_POOL Procedure Parameters

Parameter	Description
pool_name	Name of the prespawned foreground process pool. To find existing pools, query the V$PROCESS_POOL dynamic view.

Exceptions

Table 148-7 STOP_POOL Procedure Exceptions

Error	Description
ORA-457	Prespawn pool not found
ORA-458	Prespawn pool already started
ORA-459	Prespawn pool not started

Example

```
BEGIN
  DBMS_PROCESS.STOP_POOL (
   pool_name   => 'hrdb_pool',
END;
```

149
DBMS_PROFILER

The package provides an interface to profile existing PL/SQL applications and identify performance bottlenecks. You can then collect and persistently store the PL/SQL profiler data.

This chapter contains the following topics:

- Overview
- Security Model
- Operational Notes
- Exceptions
- Summary of DBMS_PROFILER Subprograms

149.1 DBMS_PROFILER Overview

This package enables the collection of profiler (performance) data for performance improvement or for determining code coverage for PL/SQL applications. Application developers can use code coverage data to focus their incremental testing efforts.

With this interface, you can generate profiling information for all named library units that are executed in a session. The profiler gathers information at the PL/SQL virtual machine level. This information includes the total number of times each line has been executed, the total amount of time that has been spent executing that line, and the minimum and maximum times that have been spent on a particular execution of that line.

> **Note:**
>
> It is possible to infer the code coverage figures for PL/SQL units for which data has been collected.

The profiling information is stored in database tables. This enables querying on the data: you can build customizable reports (summary reports, hottest lines, code coverage data, and so on. And you can analyze the data.

The `PROFTAB.SQL` script creates tables with the columns, datatypes, and definitions as shown in the following tables.

Table 149-1 Columns in Table PLSQL_PROFILER_RUNS

Column	Datatype	Definition
`runid`	`NUMBER PRIMARY KEY`	Unique run identifier from plsql_profiler_runnumber
`related_run`	`NUMBER`	Runid of related run (for client/server correlation)
`run_owner`	`VARCHAR2(128)`	User who started run
`run_date`	`DATE`	Start time of run

Table 149-1 (Cont.) Columns in Table PLSQL_PROFILER_RUNS

Column	Datatype	Definition
run_comment	VARCHAR2(2047)	User provided comment for this run
run_total_time	NUMBER	Elapsed time for this run in nanoseconds
run_system_info	VARCHAR2(2047)	Currently unused
run_comment1	VARCHAR2(2047)	Additional comment
spare1	VARCHAR2(256)	Unused

Table 149-2 Columns in Table PLSQL_PROFILER_UNITS

Column	Datatype	Definition
runid	NUMBER	Primary key, references plsql_profiler_runs,
unit_number	NUMBER	Primary key, internally generated library unit #
unit_type	VARCHAR2(128)	Library unit type
unit_owner	VARCHAR2(128)	Library unit owner name
unit_name	VARCHAR2(128)	Library unit name timestamp on library unit
unit_timestamp	DATE	In the future will be used to detect changes to unit between runs
total_time	NUMBER	Total time spent in this unit in nanoseconds. The profiler does not set this field, but it is provided for the convenience of analysis tools.
spare1	NUMBER	Unused
spare2	NUMBER	Unused

Table 149-3 Columns in Table PLSQL_PROFILER_DATA

Column	Datatype	Definition
runid	NUMBER	Primary key, unique (generated) run identifier
unit_number	NUMBER	Primary key, internally generated library unit number
line#	NUMBER	Primary key, not null, line number in unit
total_occur	NUMBER	Number of times line was executed
total_time	NUMBER	Total time spent executing line in nanoseconds
min_time	NUMBER	Minimum execution time for this line in nanoseconds
max_time	NUMBER	Maximum execution time for this line in nanoseconds
spare1	NUMBER	Unused
spare2	NUMBER	Unused
spare3	NUMBER	Unused
spare4	NUMBER	Unused

With Oracle database version 8.x, a sample textual report writer(profrep.sql) is provided with the PL/SQL demo scripts.

Note that prior to Oracle Database 10*g*, the DBMS_PROFILER package was not automatically loaded when the database was created, and the Oracle-supplied PROFLOAD.SQL script was used to create it. In 10*g* and beyond, the DBMS_PROFILER package is loaded automatically when the database is created, and PROFLOAD.SQL is no longer needed.

149.2 DBMS_PROFILER Security Model

The profiler only gathers data for units for which a user has CREATE privilege; you cannot use the package to profile units for which EXECUTE ONLY access has been granted. In general, if a user can debug a unit, the same user can profile it. However, a unit can be profiled whether or not it has been compiled DEBUG. Oracle advises that modules that are being profiled should be compiled DEBUG, since this provides additional information about the unit in the database.

> **Note:**
>
> DBMS_PROFILER treats any program unit that is compiled in NATIVE mode as if you do not have CREATE privilege, that is, you will not get any output.

149.3 DBMS_PROFILER Operational Notes

These notes describe a typical run, how to interpret output, and two methods of exception generation.

Typical Run

Improving application performance is an iterative process. Each iteration involves the following steps:

1. Running the application with one or more benchmark tests with profiler data collection enabled.
2. Analyzing the profiler data and identifying performance problems.
3. Fixing the problems.

The PL/SQL profiler supports this process using the concept of a *"run".* A run involves running the application through benchmark tests with profiler data collection enabled. You can control the beginning and the ending of a run by calling the START_PROFILER and STOP_PROFILER functions.

The user must first create database tables in the profiler user's schema to collect the data. The PROFTAB.SQL script creates the tables and other data structures required for persistently storing the profiler data.

Note that running PROFTAB.SQL drops the current tables. The PROFTAB.SQL script is in the RDBMS/ADMIN directory. Some PL/SQL operations, such as the first execution of a PL/SQL unit, may involve I/O to catalog tables to load the byte code for the PL/SQL unit being executed. Also, it may take some time executing package initialization code the first time a package procedure or function is called.

To avoid timing this overhead, "*warm up"* the database before collecting profile data. To do this, run the application once without gathering profiler data.

You can allow profiling across all users of a system, for example, to profile all users of a package, independent of who is using it. In such cases, the `SYSADMIN` should use a modified `PROFTAB.SQL` script which:

- Creates the profiler tables and sequence
- Grants `SELECT`/`INSERT`/`UPDATE` on those tables and sequence to all users
- Defines public synonyms for the tables and sequence

> **Note:**
>
> Do not alter the actual fields of the tables.

A typical run then involves:

- Starting profiler data collection in the run.
- Executing PL/SQL code for which profiler and code coverage data is required.
- Stopping profiler data collection, which writes the collected data for the run into database tables

> **Note:**
>
> The collected profiler data is not automatically stored when the user disconnects. You must issue an explicit call to the `FLUSH_DATA` or the `STOP_PROFILER` function to store the data at the end of the session. Stopping data collection stores the collected data.

As the application executes, profiler data is collected in memory data structures that last for the duration of the run. You can call the `FLUSH_DATA` function at intermediate points during the run to get incremental data and to free memory for allocated profiler data structures. Flushing the collected data involves storing collected data in the database tables created earlier.

> **See Also:**
>
> "FLUSH_DATA Function and Procedure".

Interpreting Output

The table `plsql_profiler_data` contains one row for each line of the source unit for which code was generated. The line# value specifies which source line. If the row exists, and the total_occur value in that row is > 0, some code associated with that line was executed. If the row exists, and total_occur value is 0, no code associated with that line was executed. If the row doesn't exist in the table, no code was generated for that line, and therefore it should not be mentioned in reports

If the source of a single statement is on a single line, any code generated for that statement will be attributed to that line number. (In some cases, such as a simple declaration, or because of optimization, no code will be needed). To get coverage information, units should be compiled with `PLSQL_OPTIMIZE_LEVEL=1`.

If a statement spans multiple lines, any code generated for that statement will be attributed to some line in the range, but it is not guaranteed that every line in the range will have code attributed to it. In such a case there will be gaps in the set of line# values. In particular, multi-line SQL-related statements may appear to be on a single line (usually the first). This is because PL/SQL passes the processed text of the cursor to the SQL engine; therefore, as far as PL/SQL is concerned, the entire SQL statement is a single indivisible operation.

When multiple statements are on the same line, the profiler will combine the occurrences for each statement. This may be confusing if a line has embedded control flow. For example, if `'then ...'` and `'else ...'` are on the same line, it will not be possible to determine whether the `'then'` or the `'else'` was taken.

In general, profiler and coverage reports are most easily interpreted if each statement is on its own line.

Two Methods of Exception Generation

Each routine in this package has two versions that allow you to determine how errors are reported.

- A function that returns success/failure as a status value and will never raise an exception
- A procedure that returns normally if it succeeds and raises an exception if it fails

In each case, the parameters of the function and procedure are identical. Only the method by which errors are reported differs. If there is an error, there is a correspondence between the error codes that the functions return, and the exceptions that the procedures raise.

To avoid redundancy, the following section only provides details about the functional form.

149.4 DBMS_PROFILER Exceptions

DBMS_PROFILER throws the exceptions described in this topic.

Table 149-4 DBMS_PROFILER Exceptions

Exception	Description
`version_mismatch`	Corresponds to error_version.
`profiler_error`	Corresponds to either "error_param" or "error_io".

A `0` return value from any function denotes successful completion; a nonzero return value denotes an error condition. The possible errors are as follows:

- 'A subprogram was called with an incorrect parameter.'

```
error_param constant binary_integer := 1;
```

- 'Data flush operation failed. Check whether the profiler tables have been created, are accessible, and that there is adequate space.'

```
error_io    constant binary_integer := 2;
```

- There is a mismatch between package and database implementation. Oracle returns this error if an incorrect version of the `DBMS_PROFILER` package is installed, and if the version of the profiler package cannot work with this database version. The only recovery is to install the correct version of the package.

```
error_version constant binary_integer := -1;
```

149.5 Summary of DBMS_PROFILER Subprograms

This table lists the DBMS_PROFILER subprograms and briefly describes them.

Table 149-5 DBMS_PROFILER Package Subprograms

Subprogram	Description
FLUSH_DATA Function and Procedure	Flushes profiler data collected in the user's session
GET_VERSION Procedure	Gets the version of this API
INTERNAL_VERSION_CHECK Function	Verifies that this version of the DBMS_PROFILER package can work with the implementation in the database
PAUSE_PROFILER Function and Procedure	Pauses profiler data collection
RESUME_PROFILER Function and Procedure	Resumes profiler data collection
START_PROFILER Functions and Procedures	Starts profiler data collection in the user's session
STOP_PROFILER Function and Procedure	Stops profiler data collection in the user's session

149.5.1 FLUSH_DATA Function and Procedure

This function flushes profiler data collected in the user's session. The data is flushed to database tables, which are expected to preexist.

> **Note:**
>
> Use the PROFTAB.SQL script to create the tables and other data structures required for persistently storing the profiler data.

Syntax

```
DBMS_PROFILER.FLUSH_DATA
  RETURN BINARY_INTEGER;

DBMS_PROFILER.FLUSH_DATA;
```

149.5.2 GET_VERSION Procedure

This procedure gets the version of this API.

Syntax

```
DBMS_PROFILER.GET_VERSION (
   major  OUT BINARY_INTEGER,
   minor  OUT BINARY_INTEGER);
```

Parameters

Table 149-6 GET_VERSION Procedure Parameters

Parameter	Description
major	Major version of DBMS_PROFILER.
minor	Minor version of DBMS_PROFILER.

149.5.3 INTERNAL_VERSION_CHECK Function

This function verifies that this version of the DBMS_PROFILER package can work with the implementation in the database.

Syntax

```
DBMS_PROFILER.INTERNAL_VERSION_CHECK
  RETURN BINARY_INTEGER;
```

149.5.4 PAUSE_PROFILER Function and Procedure

This function pauses profiler data collection.

Syntax

```
DBMS_PROFILER.PAUSE_PROFILER
  RETURN BINARY_INTEGER;

DBMS_PROFILER.PAUSE_PROFILER;
```

149.5.5 RESUME_PROFILER Function and Procedure

This function resumes profiler data collection.

Syntax

```
DBMS_PROFILER.RESUME_PROFILER
  RETURN BINARY_INTEGER;

DBMS_PROFILER.RESUME_PROFILER;
```

149.5.6 START_PROFILER Functions and Procedures

This function starts profiler data collection in the user's session.

There are two overloaded forms of the START_PROFILER function; one returns the run number of the started run, as well as the result of the call. The other does not return the run number. The first form is intended for use with GUI-based tools controlling the profiler.

Syntax

```
DBMS_PROFILER.START_PROFILER(
   run_comment    IN VARCHAR2 := sysdate,
   run_comment1   IN VARCHAR2 :='',
   run_number     OUT BINARY_INTEGER)
 RETURN BINARY_INTEGER;
```

```
DBMS_PROFILER.START_PROFILER(
   run_comment IN VARCHAR2 := sysdate,
   run_comment1 IN VARCHAR2 :='')
RETURN BINARY_INTEGER;

DBMS_PROFILER.START_PROFILER(
   run_comment   IN VARCHAR2 := sysdate,
   run_comment1  IN VARCHAR2 :='',
   run_number    OUT BINARY_INTEGER);

DBMS_PROFILER.START_PROFILER(
   run_comment IN VARCHAR2 := sysdate,
   run_comment1 IN VARCHAR2 :='');
```

Parameters

Table 149-7 START_PROFILER Function Parameters

Parameter	Description
run_comment	Each profiler run can be associated with a comment. For example, the comment could provide the name and version of the benchmark test that was used to collect data.
run_number	Stores the number of the run so you can store and later recall the run's data.
run_comment1	Allows you to make interesting comments about the run.

149.5.7 STOP_PROFILER Function and Procedure

This function stops profiler data collection in the user's session.

This function has the side effect of flushing data collected so far in the session, and it signals the end of a run.

Syntax

```
DBMS_PROFILER.STOP_PROFILER
  RETURN BINARY_INTEGER;

DBMS_PROFILER.STOP_PROFILER;
```

150
DBMS_PROPAGATION_ADM

The DBMS_PROPAGATION_ADM package, one of a set of Oracle Replication packages, provides administrative interfaces for configuring a propagation from a source queue to a destination queue.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_PROPAGATION_ADM Subprograms

150.1 DBMS_PROPAGATION_ADM Overview

This package provides interfaces to start, stop, and configure a propagation.

150.2 DBMS_PROPAGATION_ADM Security Model

Security on this package can be controlled by granting EXECUTE on this package to selected users or roles, or by granting EXECUTE_CATALOG_ROLE to selected users or roles.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted EXECUTE privilege on the package directly. It cannot be granted through a role.

When the DBMS_PROPAGATION_ADM package is used to manage an Oracle Replication configuration, it requires that the user is granted the privileges of an Oracle Replication administrator.

150.3 Summary of DBMS_PROPAGATION_ADM Subprograms

This table lists the DBMS_PROPAGATION_ADM subprograms and briefly describes them.

Table 150-1 DBMS_PROPAGATION_ADM Package Subprograms

Subprogram	Description
ALTER_PROPAGATION Procedure	Adds, alters, or removes a rule set for a propagation
CREATE_PROPAGATION Procedure	Creates a propagation and specifies the source queue, destination queue, and rule set for the propagation
DROP_PROPAGATION Procedure	Drops a propagation
START_PROPAGATION Procedure	Starts a propagation
STOP_PROPAGATION Procedure	Stops a propagation

> **Note:**
>
> All subprograms commit unless specified otherwise.

150.3.1 ALTER_PROPAGATION Procedure

This procedure adds, alters, or removes a rule set for a propagation.

Syntax

```
DBMS_PROPAGATION_ADM.ALTER_PROPAGATION(
   propagation_name          IN  VARCHAR2,
   rule_set_name             IN  VARCHAR2  DEFAULT NULL,
   remove_rule_set           IN  BOOLEAN   DEFAULT FALSE,
   negative_rule_set_name    IN  VARCHAR2  DEFAULT NULL,
   remove_negative_rule_set  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 150-2 ALTER_PROPAGATION Procedure Parameters

Parameter	Description
`propagation_name`	The name of the propagation you are altering. You must specify an existing propagation name. Do not specify an owner.
`rule_set_name`	The name of the positive rule set for the propagation. The positive rule set contains the rules that instruct the propagation to propagate messages. If you want to use a positive rule set for the propagation, then you must specify an existing rule set in the form `[schema_name.]rule_set_name`. For example, to specify a positive rule set in the `hr` schema named `prop_rules`, enter `hr.prop_rules`. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the `DBMS_RULE_ADM` package. If you specify `NULL` and the `remove_rule_set` parameter is set to `FALSE`, then the procedure retains any existing positive rule set. If you specify `NULL` and the `remove_rule_set` parameter is set to `TRUE`, then the procedure removes any existing positive rule set.
`remove_rule_set`	If `TRUE`, then the procedure removes the positive rule set for the specified propagation. If you remove a positive rule set for a propagation, and the propagation does not have a negative rule set, then the propagation propagates all messages. If you remove a positive rule set for a propagation, and a negative rule set exists for the propagation, then the propagation propagates all messages in its queue that are not discarded by the negative rule set. If `FALSE`, then the procedure retains the positive rule set for the specified propagation. If the `rule_set_name` parameter is non-`NULL`, then this parameter should be set to `FALSE`.

Table 150-2 (Cont.) ALTER_PROPAGATION Procedure Parameters

Parameter	Description
negative_rule_set_name	The name of the negative rule set for the propagation. The negative rule set contains the rules that instruct the propagation to discard messages. If you want to use a negative rule set for the propagation, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a negative rule set in the hr schema named neg_rules, enter hr.neg_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify NULL and the remove_negative_rule_set parameter is set to FALSE, then the procedure retains any existing negative rule set. If you specify NULL and the remove_negative_rule_set parameter is set to TRUE, then the procedure removes any existing negative rule set. If you specify both a positive and a negative rule set for a propagation, then the negative rule set is always evaluated first.
remove_negative_rule_set	If TRUE, then the procedure removes the negative rule set for the specified propagation. If you remove a negative rule set for a propagation, and the propagation does not have a positive rule set, then the propagation propagates all messages. If you remove a negative rule set for a propagation, and a positive rule set exists for the propagation, then the propagation propagates all messages in its queue that are not discarded by the positive rule set. If FALSE, then the procedure retains the negative rule set for the specified propagation. If the negative_rule_set_name parameter is non-NULL, then this parameter should be set to FALSE.

150.3.2 CREATE_PROPAGATION Procedure

This procedure creates a propagation and specifies the source queue, destination queue, and any rule set for the propagation. A propagation propagates messages in a local source queue to a destination queue. The destination queue might or might not be in the same database as the source queue.

Syntax

```
DBMS_PROPAGATION_ADM.CREATE_PROPAGATION(
   propagation_name           IN  VARCHAR2,
   source_queue               IN  VARCHAR2,
   destination_queue          IN  VARCHAR2,
   destination_dblink         IN  VARCHAR2  DEFAULT NULL,
   rule_set_name              IN  VARCHAR2  DEFAULT NULL,
   negative_rule_set_name     IN  VARCHAR2  DEFAULT NULL,
   queue_to_queue             IN  BOOLEAN   DEFAULT NULL,
   original_propagation_name  IN  VARCHAR2  DEFAULT NULL,
   auto_merge_threshold       IN  NUMBER    DEFAULT NULL);
```

Parameters

Table 150-3 CREATE_PROPAGATION Procedure Parameters

Parameter	Description
propagation_name	The name of the propagation you are creating. A NULL setting is not allowed. Do not specify an owner. **Note:** The propagation_name setting cannot be altered after the propagation is created.
source_queue	The name of the source queue, specified as [*schema_name*.]*queue_name*. The current database must contain the source queue. For example, to specify a source queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default.
destination_queue	The name of the destination queue, specified as [*schema_name*.]*queue_name*. For example, to specify a destination queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default.
destination_dblink	The name of the database link that will be used by the propagation. The database link is from the database that contains the source queue to the database that contains the destination queue. If NULL, then the source queue and destination queue must be in the same database. **Note:** Connection qualifiers are not allowed.
rule_set_name	The name of the positive rule set for the propagation. The positive rule set contains the rules that instruct the propagation to propagate messages. If you want to use a positive rule set for the propagation, then you must specify an existing rule set in the form [*schema_name*.]*rule_set_name*. For example, to specify a positive rule set in the hr schema named prop_rules, enter hr.prop_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify NULL, and no negative rule set exists for the propagation, then the propagation propagates all messages in its queue. If you specify NULL, and a negative rule set exists for the propagation, then the propagation propagates all messages in its queue that are not discarded by the negative rule set.

Table 150-3 (Cont.) CREATE_PROPAGATION Procedure Parameters

Parameter	Description
negative_rule_set_name	The name of the negative rule set for the propagation. The negative rule set contains the rules that instruct the propagation to discard messages. If you want to use a negative rule set for the propagation, then you must specify an existing rule set in the form [*schema_name.*]*rule_set_name*. For example, to specify a negative rule set in the hr schema named neg_rules, enter hr.neg_rules. If the schema is not specified, then the current user is the default. An error is returned if the specified rule set does not exist. You can create a rule set and add rules to it using the DBMS_RULE_ADM package. If you specify NULL, and no positive rule set exists for the propagation, then the propagation propagates all messages in its queue. If you specify NULL, and a positive rule set exists for the propagation, then the propagation propagates all messages in its queue that are not discarded by the positive rule set. If you specify both a positive and a negative rule set for a propagation, then the negative rule set is always evaluated first.
queue_to_queue	If TRUE or NULL, then the propagation is a queue to queue propagation. A queue-to-queue propagation always has its own propagation job and uses a service for automatic failover when the destination queue is a buffered queue in an Oracle Real Application Clusters (Oracle RAC) database. If FALSE, then the propagation is a queue-to-dblink propagation. A queue-to-dblink propagation can share a propagation job with other propagations that use the same database link and does not support automatic failover in an Oracle RAC environment.
original_propagation_name	Specify the original propagation name if the propagation being created is part of a split and merge operation. The split operation clones the original propagation under a new name. The name of the original propagation is important when the cloned propagation is copied back to the original stream. Specify NULL if the propagation being created is not part of a split and merge operation.
auto_merge_theshold	Specify a positive number if both of the following conditions are met: • The propagation being created is part of a split and merge operation. • The stream will be merged back to the original stream automatically. Specify NULL if either of the following conditions are met: • The propagation being created is not part of a split and merge operation. • The propagation being created is part of a split and merge operation, but the stream being split off will not be merged back to the original stream automatically.

Usage Notes

This procedure starts propagation and might create a propagation job. If this procedure creates a propagation job, then it establishes a default schedule for the propagation job. Each propagation job is an Oracle Scheduler job. You can adjust the schedule of a propagation job using Oracle Scheduler.

The user who owns the source queue is the user who propagates messages. This user must have the necessary privileges to propagate messages.

> **See Also:**
>
> DBMS_RULE_ADM

150.3.3 DROP_PROPAGATION Procedure

This procedure drops a propagation and deletes all messages for the destination queue in the source queue. This procedure also removes the schedule for propagation from the source queue to the destination queue.

Syntax

```
DBMS_PROPAGATION_ADM.DROP_PROPAGATION(
   propagation_name       IN  VARCHAR2,
   drop_unused_rule_sets  IN  BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 150-4 DROP_PROPAGATION Procedure Parameters

Parameter	Description
propagation_name	The name of the propagation you are dropping. You must specify an existing propagation name. Do not specify an owner.
drop_unused_rule_sets	If TRUE, then the procedure drops any rule sets, positive and negative, used by the specified propagation if these rule sets are not used by any other Oracle Replication client, which includes capture processes, propagations, apply processes, and messaging clients. If this procedure drops a rule set, then this procedure also drops any rules in the rule set that are not in another rule set. If FALSE, then the procedure does not drop the rule sets used by the specified propagation, and the rule sets retain their rules.

Usage Notes

When you use this procedure to drop a propagation, information about rules created for the propagation is removed from the data dictionary views for Oracle Replication rules. Information about such a rule is removed even if the rule is not in either rule set for the propagation.

The following are the data dictionary views for Oracle Replication rules:

- ALL_STREAMS_GLOBAL_RULES
- DBA_STREAMS_GLOBAL_RULES

- ALL_STREAMS_SCHEMA_RULES
- DBA_STREAMS_SCHEMA_RULES
- ALL_STREAMS_TABLE_RULES
- DBA_STREAMS_TABLE_RULES

> **Note:**
>
> When you drop a propagation, the propagation job used by the propagation is dropped automatically, if no other propagations are using the propagation job.

150.3.4 START_PROPAGATION Procedure

This procedure starts a propagation.

Syntax

```
DBMS_PROPAGATION_ADM.START_PROPAGATION(
   propagation_name  IN  VARCHAR2);
```

Parameter

Table 150-5 START_PROPAGATION Procedure Parameter

Parameter	Description
propagation_name	The name of the propagation you are starting. You must specify an existing propagation name. Do not specify an owner.

Usage Notes

The propagation status is persistently recorded. Hence, if the status is ENABLED, then the propagation is started upon database instance startup.

150.3.5 STOP_PROPAGATION Procedure

This procedure stops a propagation.

Syntax

```
DBMS_PROPAGATION_ADM.STOP_PROPAGATION(
   propagation_name  IN  VARCHAR2,
   force             IN  BOOLEAN  DEFAULT FALSE);
```

Parameter

Table 150-6 STOP_PROPAGATION Procedure Parameter

Parameter	Description
propagation_name	The name of the propagation you are stopping. You must specify an existing propagation name. Do not specify an owner.

Table 150-6 (Cont.) STOP_PROPAGATION Procedure Parameter

Parameter	Description
`force`	If `TRUE`, then the procedure stops the propagation and clears the statistics for the propagation. If `FALSE`, then the procedure stops the propagation without clearing the statistics for the propagation.

Usage Notes

The propagation status is persistently recorded. Hence, if the status is `DISABLED` or `ABORTED`, then the propagation is not started upon database instance startup.

151
DBMS_QOPATCH

The DBMS_QOPATCH package provides an interface to view the installed database patches.

This chapter contains the following topics:

- Overview
- Security Modell
- Operational Notes
- Error Messages
- Summary of DBMS_QOPATCH Subprograms

151.1 DBMS_QOPATCH Overview

The DBMS_QOPATCH package provides a PLSQL/SQL interface to view the installed database patches. The interface provides all the patch information available as part of the OPATCH LSINVENTORY -XML command. The package accesses the OUI patch inventory in real time to provide patch and meta-information.

151.2 DBMS_QOPATCH Security Model

The DBMS_QOPATCH package is created as part of SYS schema and SYS is the only user who can execute these subprograms.

151.3 DBMS_QOPATCH Operational Notes

The following operational notes apply to DBMS_QOPATCH.

- The package will work only if the database is OPEN.
- In an Oracle Real Application Clusters (RAC) environment, if the subprogram requires to fetch data from other RAC nodes it spawns a job in the other node(s) to get the data. In this case JOB_QUEUE_PROCESSES needs to be >0 for the package to fetch the data from other RAC nodes.
- If there is a delay in the job execution, the package returns ORA-20008 error.

151.4 DBMS_QOPATCH Exceptions

This table lists the exceptions raised by the DBMS_QOPATCH package.

Table 151-1 DBMS_QOPATCH Error Messages

Error Code	Description
ORA-20001	Latest XML inventory is not loaded into table.

Table 151-1 (Cont.) DBMS_QOPATCH Error Messages

Error Code	Description
ORA-20002	Directory creation failed.
ORA-20003	Configuration of a job on a node failed.
ORA-20004	Job configuration failed as node is inactive.
ORA-20005	Job is not configured with given node, instance name.
ORA-20006	Number of RAC active instances and opatch jobs configured are not same.
ORA-20007	Job configuration failed as node or instance is not active.
ORA-20008	Timed out - job execution time is more than 120Secs.
ORA-20009	Job execution failed.
ORA-20010	Node is inactive and job cannot be executed.
ORA-20011	Job name is NULL and inventory cannot be loaded.
ORA-20012	JOB_QUEUE_PROCESSES is set to zero and the inventory cannot be loaded.
ORA-20013	DBMS_QOPATCH ran mostly in non-install area.
ORA-20014	Database is not opened.
ORA-20015	Database opened as read-only.

151.5 Summary of DBMS_QOPATCH Subprograms

This table lists and briefly describes the DBMS_QOPATCH package subprograms.

Table 151-2 DBMS_QOPATCH Package Subprograms

Subprogram	Description
CHECK_PATCH_INSTALLED Function	Checks if a patch is installed.
GET_OPATCH_BUGS Function	Provides a bugs list for a patch in XML format if the patch number is given. If patch is not given then it lists all the bugs installed in all the patches in XML format.
GET_OPATCH_COUNT Function	Provides the total number of installed patches in XML format
GET_OPATCH_DATA Function	Provides top level patch information for the patch (such as Patch ID, patch creation time) in the XML element
GET_OPATCH_FILES Function	Provides the list of files modified in the given patch number in XML format
GET_OPATCH_INSTALL_INFO Function	Returns the XML element containing the ORACLE_HOME details such as patch and inventory location
GET_OPATCH_LIST Function	Provides list of patches installed as an XML element from the XML inventory
GET_OPATCH_LSINVENTORY	Returns whole opatch inventory as XML instance document.
GET_OPATCH_OLAYS Function	Provides overlay patches for a given patch as XML element

Table 151-2 (Cont.) DBMS_QOPATCH Package Subprograms

Subprogram	Description
GET_OPATCH_PREQS Function	Provides prerequisite patches for a given patch as XML element
GET_OPATCH_XSLT	Returns the style-sheet for the opatch XML inventory presentation
GET_PATCH_DETAILS Function	Displays the detailed patch information from `opatch lsinventory`.
GET_SQLPATCH_STATUS Procedure	Displays the SQL patch status by querying from SQL patch registry to produce complete patch level information
IS_PATCH_INSTALLED Function	Provides information (such as patchID, application date, and SQL patch information) on the installed patch as XML node by querying the XML inventory
OPATCH_COMPARE_CURRENT Function	Compares the current database with a list of bugs.
OPATCH_COMPARE_NODES Function	Compares the given RAC node(s) with the instance to the present connected node with current node. The return value indicates whether the data was refreshed or not.
SET_CURRENT_OPINST Procedure	Sets the node name and instance to get the inventory details specific to it in an Oracle Real Application Clusters (RAC) environment

151.5.1 CHECK_PATCH_INSTALLED Function

Checks if the patch is installed.

Syntax

```
DBMS_QOPATCH.CHECK_PATCH_INSTALLED (
   bugs     IN QOPATCH_LIST);
 RETURN VARCHAR2;
```

Parameters

Table 151-3 CHECK_PATCH_INSTALLED Parameters

Parameter	Description
`bugs`	List of patches to be checked.

151.5.2 GET_OPATCH_BUGS Function

This function provides a bugs list in a patch if the patch number is given. If a patch number is not given, it lists all the bugs in the specified XML format.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_BUGS (
   patchnum IN VARCHAR2 DEFAULT NULL);
 RETURN XMLTYPE;
```

Parameters

Table 151-4 GET_OPATCH_BUGS Function Parameters

Parameter	Description
patchnum	Patch number

151.5.3 GET_OPATCH_COUNT Function

This function provides the total number of installed patches in XML format.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_COUNT (
   patchnum IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-5 GET_OPATCH_COUNT Function Parameters

Parameter	Description
patchnum	Patch number

151.5.4 GET_OPATCH_DATA Function

This function provides top level patch information for the patch (such as Patch ID, patch creation time) in the XML element.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_DATA (
   patchnum IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-6 GET_OPATCH_DATA Function Parameters

Parameter	Description
patchnum	Patch number

151.5.5 GET_OPATCH_FILES Function

This function provides the list of files modified in the given patch number in XML format.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_FILES (
   patchnum IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-7 GET_OPATCH_FILES Function Parameters

Parameter	Description
patchnum	Patch number

151.5.6 GET_OPATCH_INSTALL_INFO Function

This function returns the XML element containing the ORACLE_HOME details such as patch and inventory location.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_INSTALL_INFO
 RETURNS XMLTYPE;
```

151.5.7 GET_OPATCH_LIST Function

This function provides list of patches installed as an XML element from the XML inventory.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_LIST
 RETURN XMLTYPE;
```

151.5.8 GET_OPATCH_LSINVENTORY

This function returns whole opatch inventory as XML instance document.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_LSINVENTORY
 RETURN XMLTYPE;
```

151.5.9 GET_OPATCH_OLAYS Function

This function provides overlay patches for a given patch as XML element.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_OLAYS (
   patchnum IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-8 GET_OPATCH_OLAYS Function Parameters

Parameter	Description
patchnum	Patch number

151.5.10 GET_OPATCH_PREQS Function

This function provides prerequisite patches for a given patch as XML element.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_PREQS (
   patchnum IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-9 GET_OPATCH_PREQS Function Parameters

Parameter	Description
patchnum	Patch number

151.5.11 GET_OPATCH_XSLT

This function returns the style-sheet for the opatch XML inventory presentation. You can use the return type of this subprogram to perform XMLTRANSFORM and the transformed result has the same appearance as opatch text output.

Syntax

```
DBMS_QOPATCH.GET_OPATCH_XSLT
 RETURN XMLTYPE;
```

151.5.12 GET_PATCH_DETAILS Function

Displays the detailed patch information from `opatch lsinventory`.

Syntax

```
DBMS_QOPATCH.GET_PATCH_DETAILS (
   patch    IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-10 GET_PATCH_DETAILS Function Parameters

Parameter	Description
patch	The patch number.

151.5.13 GET_SQLPATCH_STATUS Procedure

This procedure displays the SQL patch status by querying from SQL patch registry to produce complete patch level information. If the patch number is given, it displays the information specific to the given SQL patch, otherwise information for all SQL patches.

Syntax

```
DBMS_QOPATCH.GET_SQLPATCH_STATUS (
   patchnum IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 151-11 GET_SQLPATCH_STATUS Procedure Parameters

Parameter	Description
patchnum	Patch number

151.5.14 IS_PATCH_INSTALLED Function

This function provides information (such as patchID, application date, and SQL patch information) on the installed patch as XML node by querying the XML inventory.

Syntax

```
DBMS_QOPATCH.IS_PATCH_INSTALLED (
   patchnum IN VARCHAR2);
 RETURN XMLTYPE;
```

Parameters

Table 151-12 IS_PATCH_INSTALLED Function Parameters

Parameter	Description
patchnum	Patch number

151.5.15 OPATCH_COMPARE_CURRENT Function

Compares the current database with a list of bugs.

Syntax

```
DBMS_QOPATCH.OPATCH_COMPARE_CURRENT Function (
   bugs     IN QOPATCH_LIST);
 RETURN VARCHAR2;
```

Parameters

Table 151-13 OPATCH_COMPARE_CURRENT Function Parameters

Parameter	Description
bugs	List of bugs to compare with the current database.

151.5.16 OPATCH_COMPARE_NODES Function

Compares the given RAC node(s) with the instance to the present connected node with current node. The return value indicates whether the data was refreshed or not.

Syntax

```
DBMS_QOPATCH.OPATCH_COMPARE_NODES (
   node     IN VARCHAR2 DEFAULT NULL,
   inst     IN VARCHAR2 DEFAULT NULL);
 RETURN VARCHAR2;
```

Parameters

Table 151-14 OPATCH_COMPARE_NODES Function Parameters

Parameter	Description
node	Node name
inst	Instance name

151.5.17 SET_CURRENT_OPINST Procedure

This procedure sets the node name and instance to get the inventory details specific to it in an Oracle Real Application Clusters (RAC) environment.

Syntax

```
DBMS_QOPATCH.SET_CURRENT_OPINST (
   node_name    IN VARCHAR2 DEFAULT NULL,
   inst_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 151-15 SET_CURRENT_OPINST Procedure Parameters

Parameter	Description
node_name	Name of node
inst_name	Name of instance

152
DBMS_RANDOM

The DBMS_RANDOM package provides a built-in random number generator. DBMS_RANDOM is not intended for cryptography.

This chapter contains the following topics:

- Deprecated Subprograms
- Security Model
- Operational Notes
- Summary of DBMS_RANDOM Subprograms

152.1 DBMS_RANDOM Deprecated Subprograms

These subprograms are deprecated with Oracle Database 11*g*. Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

- INITIALIZE Procedure
- RANDOM Function
- TERMINATE Procedure

152.2 DBMS_RANDOM Security Model

This package should be installed as SYS. By default, the package is initialized with the current user name, current time down to the second, and the current session. Oracle recommends that users who need to execute this package should be given EXECUTE privilege explicitly and should not rely on PUBLIC EXECUTE privilege.

152.3 DBMS_RANDOM Operational Notes

These operational notes apply to DBMS_RANDOM.

- DBMS_RANDOM.RANDOM produces integers in [-2^^31, 2^^31).
- DBMS_RANDOM.VALUE produces numbers in [0,1) with 38 digits of precision.

DBMS_RANDOM can be explicitly initialized, but does not need to be initialized before calling the random number generator. It will automatically initialize with the date, user ID, and process ID if no explicit initialization is performed.

If this package is seeded twice with the same seed, then accessed in the same way, it will produce the same results in both cases.

In some cases, such as when testing, you may want the sequence of random numbers to be the same on every run. In that case, you seed the generator with a constant value by calling one of the overloads of DBMS_RANDOM.SEED. To produce different output for every run, simply to omit the call to "Seed" and the system will choose a suitable seed for you.

152.4 Summary of DBMS_RANDOM Subprograms

This table lists the `DBMS_RANDOM` subprograms and briefly describes them.

Table 152-1 *DBMS_RANDOM Package Subprograms*

Subprogram	Description
INITIALIZE Procedure	Initializes the package with a seed value
NORMAL Function	Returns random numbers in a normal distribution
RANDOM Function	Generates a random number
SEED Procedures	Resets the seed
STRING Function	Gets a random string
TERMINATE Procedure	Terminates package
VALUE Functions	Gets a random number, greater than or equal to 0 and less than 1, with 38 digits to the right of the decimal (38-digit precision), while the overloaded function gets a random Oracle number x, where x is greater than or equal to `low` and less than `high`

152.4.1 INITIALIZE Procedure

This deprecated procedure initializes the generator.

> **Note:**
>
> This procedure is deprecated with Release 11gR1 and, although currently supported, it should not be used.

Syntax

```
DBMS_RANDOM.INITIALIZE (
   val   IN   BINARY_INTEGER);
```

Pragmas

```
PRAGMA restrict_references (initialize, WNDS);
```

Parameters

Table 152-2 INITIALIZE Procedure Parameters

Parameter	Description
`val`	Seed number used to generate a random number

Usage Notes

This procedure is obsolete as it simply calls the SEED Procedures.

152.4.2 NORMAL Function

This function returns random numbers in a standard normal distribution.

Syntax

```
DBMS_RANDOM.NORMAL
  RETURN NUMBER;
```

Pragmas

```
PRAGMA restrict_references (normal, WNDS);
```

Return Values

Table 152-3 NORMAL Function Parameters

Parameter	Description
number	Returns a random number

152.4.3 RANDOM Function

This deprecated procedure generates a random number.

> **Note:**
>
> This function is deprecated with Release 11gR1 and, although currently supported, it should not be used.

Syntax

```
DBMS_RANDOM.RANDOM
   RETURN binary_integer;
```

Pragmas

```
PRAGMA restrict_references (random, WNDS);
```

Return Values

Table 152-4 RANDOM Function Parameters

Parameter	Description
binary_integer	Returns a random integer greater or equal to -power(2,31) and less than power(2,31)

152.4.4 SEED Procedures

This procedure resets the seed.

Syntax

```
DBMS_RANDOM.SEED (
   val  IN  BINARY_INTEGER);

DBMS_RANDOM.SEED (
   val  IN  VARCHAR2);
```

Pragmas

```
PRAGMA restrict_references (seed, WNDS);
```

Parameters

Table 152-5 SEED Procedure Parameters

Parameter	Description
`val`	Seed number or string used to generate a random number

Usage Notes

The seed can be a string up to length 2000.

152.4.5 STRING Function

This function gets a random string.

Syntax

```
DBMS_RANDOM.STRING
   opt  IN  CHAR,
   len  IN  NUMBER)
  RETURN VARCHAR2;
```

Pragmas

```
PRAGMA restrict_references (string, WNDS);
```

Parameters

Table 152-6 STRING Function Parameters

Parameter	Description
`opt`	Specifies what the returning string looks like: • 'u', 'U' - returning string in uppercase alpha characters • 'l', 'L' - returning string in lowercase alpha characters • 'a', 'A' - returning string in mixed case alpha characters • 'x', 'X' - returning string in uppercase alpha-numeric characters • 'p', 'P' - returning string in any printable characters. Otherwise the returning string is in uppercase alpha characters.

Table 152-6 (Cont.) STRING Function Parameters

Parameter	Description
len	Length of the returning string

Return Values

Table 152-7 STRING Function Return Values

Parameter	Description
VARCHAR2	Returns a VARCHAR2

152.4.6 TERMINATE Procedure

When you are finished with the package, call the TERMINATE procedure.

> **Note:**
>
> This procedure is deprecated with Release 11gR1 and, although currently supported, it should not be used.

Syntax

```
DBMS_RANDOM.TERMINATE;
```

152.4.7 VALUE Functions

The basic function gets a random number, greater than or equal to 0 and less than 1, with 38 digits to the right of the decimal (38-digit precision). Alternatively, you can get a random Oracle number x, where x is greater than or equal to low and less than high.

Syntax

```
DBMS_RANDOM.VALUE
  RETURN NUMBER;

DBMS_RANDOM.VALUE(
  low  IN  NUMBER,
  high IN  NUMBER)
RETURN NUMBER;
```

Parameters

Table 152-8 VALUE Function Parameters

Parameter	Description
low	Lowest number in a range from which to generate a random number. The number generated may be equal to low

Table 152-8 (Cont.) VALUE Function Parameters

Parameter	Description
high	Highest number below which to generate a random number. The number generated will be less than high

Return Values

Table 152-9 VALUE Function Return Values

Parameter	Description
NUMBER	Returns an Oracle Number

153
DBMS_REDACT

The `DBMS_REDACT` package provides an interface to Oracle Data Redaction, which enables you to redact data that is returned from queries issued by low-privileged users or an application.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operating Procedures
- Summary of DBMS_REDACT Subprograms

> **See Also:**
>
> - *Oracle Database Data Redaction Guide* regarding using Data Redaction to protect sensitive data

153.1 DBMS_REDACT Overview

Data redaction provides a way to define redaction policies for an application. Oracle Data Redaction provides functionality to redact data that is returned from user `SELECT` queries in an application. The redaction takes place in real time.

The Data Redaction policy applies to the querying user, depending on this user's `SYS_CONTEXT` or `XS_SYS_CONTEXT` values. This redaction process does not require that the queried data be static or unchanging, or for the entire data set to be redacted at one time in an off-line manner. Oracle Database redacts the data only for the rows specified by the user's query, not the data for the entire column. The redaction takes place immediately before the data is returned to the querying user or application.

153.2 DBMS_REDACT Security Model

If the querying user has the `EXEMPT REDACTION POLICY` privilege, redaction will not be performed. If the user does not have the `EXEMPT REDACTION POLICY` privilege, the policy expression will be evaluated in the current user's environment. If the policy expression evaluates to `TRUE`, then redaction will be performed, otherwise no redaction will be performed.

You must have the `EXECUTE` privilege on the `DBMS_REDACT` package, as well as the `ADMINISTER REDACTION POLICY` system or schema privilege, to execute its subprograms. Procedures in the interface are executed with privileges of the current user.

153.3 DBMS_REDACT Constants

The DBMS_REDACT package defines several constants for specifying parameter values.

Table 153-1 Values for function_type Parameter of DBMS_REDACT.ADD_POLICY

Constant	Value	Type	Description
FULL	1	BINARY_INTEGER	Redact to fixed values
PARTIAL	2	BINARY_INTEGER	Partial redaction, redact a portion of the column data
RANDOM	4	BINARY_INTEGER	Random redaction, each query results in a different random value
REGEXP	5	BINARY_INTEGER	Regular expression based redaction
REGEXP_WIDTH	7	BINARY_INTEGER	Regular expression based redaction that preserves the width of a column that uses a regular expression; designed for applications that use the OCI_ATTR_CHAR_SIZE attribute of the Oracle OLE DB Provider interface
NULLIFY	6	BINARY_INTEGER	Returns a null value as a redacted value
NONE	0	BINARY_INTEGER	No redaction

Table 153-2 Values for action Parameter of DBMS_REDACT.ALTER_POLICY

Constant	Value	Type	Description
ADD_COLUMN	1	BINARY_INTEGER	Add a column to the redaction policy
DROP_COLUMN	2	BINARY_INTEGER	Drop a column from the redaction policy
MODIFY_EXPRESSION	3	BINARY_INTEGER	Modify the expression of a redaction policy (the expression evaluates to a BOOLEAN value: if TRUE then redaction is applied, otherwise not)
MODIFY_COLUMN	4	BINARY_INTEGER	Modify a column in the redaction policy to change the redaction function_type or the function_parameters
SET_POLICY_DESCRIPTION	5	BINARY_INTEGER	Set a description for the redaction policy
SET_COLUMN_DESCRIPTION	6	BINARY_INTEGER	Set a description for the redaction performed on the column

153.4 DBMS_REDACT Operating Procedures

The following table presents the relationship between the type of redaction function and its parameters, based on the datatype of the column being redacted. Examples of the various format strings are provided, showing how to perform some commonplace redaction for a string

datatype (in this case, a Social Security Number (SSN)), a DATE datatype, and various examples of redaction for the number datatype.

Table 153-3 Data Redaction Function Types

function_type	function_parameters	Examples
DBMS_REDACT.NONE	-	-
DBMS_REDACT.FULL	-	-
DBMS_REDACT.NULLIFY	-	-
DBMS_REDACT.PARTIAL (for character types)	A comma-separated list, containing the following five fields (with no spaces after the commas delimiting the fields): • REDACT_PARTIAL_INPUT_FORMAT • REDACT_PARTIAL_OUTPUT_FORMAT • REDACT_PARTIAL_MASKCHAR • REDACT_PARTIAL_MASKFROM • REDACT_PARTIAL_MASKTO See Table 153-4.	'VVVFVVFVVVV,VVV-VV-VVVV,X,1,5' for redacting the first 5 digits of SSN strings like 123-45-6789, adding dashes back to format it, resulting in strings like XXX-XX-6789 'VVVFVVFVVVV,VVV VV VVVV,X,1,5' for redacting the first 5 digits of SSN strings like 123-45-6789, adding spaces to format it, resulting in strings like XXX XX 6789
DBMS_REDACT.PARTIAL (for character types), continued	REDACT_PARTIAL_INPUT_FORMAT - the input format: V for value to be possibly redacted, F for formatting character to be ignored	The REDACT_PARTIAL_INPUT_FORMAT field value VVVFVVFVVVV for matching SSN strings like 123-45-6789
DBMS_REDACT.PARTIAL (for character types), continued	REDACT_PARTIAL_OUTPUT_FORMAT - the output format: V for output of redaction. Any other character will be treated as a formatting character and output literally.	The REDACT_PARTIAL_OUTPUT_FORMAT field value VVV-VV-VVVV can be used to redact SSN strings into XXX-XX-6789 (X comes from REDACT_PARTIAL_MASKCHAR field).
DBMS_REDACT.PARTIAL (for character types), continued	REDACT_PARTIAL_MASKCHAR - the character used to redact the input	The value X for redacting SSN strings into XXX-XX-6789.
DBMS_REDACT.PARTIAL (for character types), continued	REDACT_PARTIAL_MASKFROM - specifies which V within the REDACT_PARTIAL_INPUT_FORMAT from which to start the redaction (see explanation following the next entry, REDACT_PARTIAL_MASKTO)	The value 1 for redacting SSN strings starting at the first V of REDACT_PARTIAL_INPUT_FORMAT of VVVFVVFVVVV into strings like XXX-XX-6789

Table 153-3 (Cont.) Data Redaction Function Types

function_type	function_parameters	Examples
DBMS_REDACT.PARTIAL (for character types), continued	REDACT_PARTIAL_MASKTO - specifies which V within the REDACT_PARTIAL_INPUT_FORMAT at which to end the redaction	The value 5 for redacting SSN strings up to and including the fifth V within REDACT_PARTIAL_INPUT_FORMAT of VVVFVVFVVVV into strings like XXX-XX-6789. Note how the format character '-' (corresponding to the first F within REDACT_PARTIAL_INPUT_FORMAT) is ignored as far as redaction is concerned, so the value here is 5 as opposed to 6.
DBMS_REDACT.PARTIAL (for character types), continued	The REDACT_PARTIAL_MASKFROM and REDACT_PARTIAL_MASKTO field values are specified as counts of the number of V characters in the REDACT_PARTIAL_INPUT_FORMAT field, up to and including the intended position, starting from the leftmost V. This way, REDACT_PARTIAL_MASKFROM and REDACT_PARTIAL_MASKTO are independent of the specific formatting of the data. For example, in the common use case of redacting an SSN to show only the last four digits, data like 123456789 (with REDACT_PARTIAL_INPUT_FORMAT of VVVVVVVVV) and data like 123-45-6789 (with REDACT_PARTIAL_INPUT_FORMAT of VVVFVVFVVVV), would both use REDACT_PARTIAL_MASKFROM of 1 and REDACT_PARTIAL_MASKTO of 5.	-
DBMS_REDACT.PARTIAL (for number types)	A comma-separated list, containing the following three fields (with no spaces after the commas delimiting the fields): • REDACT_PARTIAL_MASKCHAR • REDACT_PARTIAL_MASKFROM • REDACT_PARTIAL_MASKTO See Table 153-4.	'9,1,5' for redacting the first 5 digits of an SSN number 123456789 into 999996789; or '0,1,2' for redacting a number 1.23 to 0.03.

Table 153-3 (Cont.) Data Redaction Function Types

function_type	function_parameters	Examples
DBMS_REDACT.PARTIAL (for number types), continued	REDACT_PARTIAL_MASKCHAR - the character used to redact the input, in the range between 0 and 9 REDACT_PARTIAL_MASKFROM - the position, starting from 1, from which to start the redaction. The position does not include the decimal point if it is present. REDACT_PARTIAL_MASKTO - the position at which to end the redaction	-
DBMS_REDACT.PARTIAL (for datetime datatypes)	A list, containing the following five fields (concatenated so that there is no space between the fields): • REDACT_PARTIAL_DATE_MONTH • REDACT_PARTIAL_DATE_DAY • REDACT_PARTIAL_DATE_YEAR • REDACT_PARTIAL_DATE_HOUR • REDACT_PARTIAL_DATE_MINUTE • REDACT_PARTIAL_DATE_SECOND See Table 153-4.	'm12DYHMS', which changes 01-May-01 01:01:01 to 01-Dec-01 01:01:01.

Table 153-3 (Cont.) Data Redaction Function Types

function_type	function_parameters	Examples
DBMS_REDACT.PARTIAL (for datetime datatypes), continued	REDACT_PARTIAL_DATE_MONTH: 'M' (no redacting of month) or 'm#' (redact month to a specific month, if possible), where # (the month specified by its number) is between 1 and 12 REDACT_PARTIAL_DATE_DAY: 'D' (no redacting of date) or 'd#' (redact day to #, if possible), # between 1 and 31 REDACT_PARTIAL_DATE_YEAR: 'Y' (no redacting of year) or 'y#' (redact year to #, if possible), # between 1 and 9999 REDACT_PARTIAL_DATE_HOUR: 'H' (no redacting of hour) or 'h#' (redact hour to #, if possible), # between 0 and 23 REDACT_PARTIAL_DATE_MINUTE : 'M' (no redacting of minute) or 'm#' (redact minute to #, if possible), # between 0 and 59 REDACT_PARTIAL_DATE_SECOND : 'S' (no redacting of second) or 's#' (redact second to #, if possible), # between 0 and 59	-
DBMS_REDACT.REGEXP	-	-
DBMS_REDACT.REGEXP_WIDTH	-	-
DBMS_REDACT.RANDOM	-	-

Table 153-4 Format Descriptors with Component Field Names and Delimiters

Datatype	Format Descriptor for Partial redaction
Character	REDACT_PARTIAL_INPUT_FORMAT || ',' || REDACT_PARTIAL_OUTPUT_FORMAT || ',' || REDACT_PARTIAL_MASKCHAR || ',' || REDACT_PARTIAL_MASKFROM || ',' || REDACT_PARTIAL_MASKTO
Number	REDACT_PARTIAL_MASKCHAR || ',' || REDACT_PARTIAL_MASKFROM || ',' || REDACT_PARTIAL_MASKTO
Datetime	REDACT_PARTIAL_DATE_MONTH || REDACT_PARTIAL_DATE_DAY || REDACT_PARTIAL_DATE_YEAR || REDACT_PARTIAL_DATE_HOUR || REDACT_PARTIAL_DATE_MINUTE || REDACT_PARTIAL_DATE_SECOND

153.5 Summary of DBMS_REDACT Subprograms

This table lists and briefly describes the `DBMS_REDACT` package subprograms.

Table 153-5 DBMS_REDACT Package Subprograms

Subprogram	Description
ADD_POLICY Procedure	Defines a Data Redaction policy for a table or view
ALTER_POLICY Procedure	Alters a Data Redaction policy for a table or view
APPLY_POLICY_EXPR_TO_COL Procedure	Applies a Data Redaction named policy expression to a redacted column
CREATE_POLICY_EXPRESSION Procedure	Creates a Data Redaction named policy expression
DISABLE_POLICY Procedure	Disables a Data Redaction policy
DROP_POLICY Procedure	Drops a Data Redaction policy
DROP_POLICY_EXPRESSION Procedure	Drops a Data Redaction named policy expression
ENABLE_POLICY Procedure	Enables a Data Redaction policy
UPDATE_FULL_REDACTION_VALUES Procedure	Modifies the default displayed values for a Data Redaction policy for full redaction
UPDATE_POLICY_EXPRESSION Procedure	Updates a Data Redaction named policy expression

153.5.1 ADD_POLICY Procedure

This procedure defines a Data Redaction policy for a table or view.

Syntax

```
DBMS_REDACT.ADD_POLICY (
   object_schema               IN    VARCHAR2 := NULL,
   object_name                 IN    VARCHAR2,
   policy_name                 IN    VARCHAR2,
   column_name                 IN    VARCHAR2 := NULL,
   function_type               IN    BINARY_INTEGER := DBMS_REDACT.FULL,
   function_parameters         IN    VARCHAR2 := NULL,
   expression                  IN    VARCHAR2,
   enable                      IN    BOOLEAN := TRUE,
   regexp_pattern              IN    VARCHAR2 := NULL,
   regexp_replace_string       IN    VARCHAR2 := NULL,
   regexp_position             IN    BINARY_INTEGER := 1,
   regexp_occurrence           IN    BINARY_INTEGER := 0,
   regexp_match_parameter      IN    VARCHAR2 := NULL,
   policy_description          IN    VARCHAR2 := NULL,
   column_description          IN    VARCHAR2 := NULL);
```

Parameters

Table 153-6 *ADD_POLICY Procedure Parameters*

Parameter	Description
`object_schema`	Schema owning the table or view, current user if `NULL`
`object_name`	Name of table or view on which to add a Data Redaction policy
`policy_name`	Name of policy
`column_name`	[Optional] Name of one column to which the redaction policy applies. If you must redact more than one column, use the ALTER_POLICY Procedure to add the additional columns.
`function_type`	Type of redaction function to use. Possible values are: - `DBMS_REDACT.NONE` - `DBMS_REDACT.FULL` (default) - `DBMS_REDACT.NULLIFY` - `DBMS_REDACT.PARTIAL` - `DBMS_REDACT.RANDOM` - `DBMS_REDACT.REGEXP` - `DBMS_REDACT.REGEXP_WIDTH` See Table 153-1 for an overview of the meanings of these values, and for some examples of their use.

Table 153-6 (Cont.) *ADD_POLICY Procedure Parameters*

Parameter	Description
function_parameters	Parameters to the redaction function. The possible values depend on the value of the function_type provided. Use function_parameters parameter only for DBMS_REDACT.PARTIAL. Use regexp_* parameters only for DBMS_REDACT.REGEXP or DBMS_REDACT.REGEXP_WIDTH to define the Data Redaction policy. If the function_type is DBMS_REDACT.REGEXP or DBMS_REDACT.REGEXP_WIDTH, then you must omit the function_parameters parameter, and use the regexp_* parameters to define the Data Redaction policy. - DBMS_REDACT.NONE: Can be omitted entirely and defaults to NULL - DBMS_REDACT.FULL: Can be omitted entirely and defaults to NULL - DBMS_REDACT.NULLIFY: Can be omitted entirely and defaults to NULL - DBMS_REDACT.RANDOM: Can be omitted entirely and defaults to NULL - Redacting parameters for partial character redacting. For character datatypes, a comma-separated list containing these fields: • Input format: 'V' for value to be possibly redacted, 'F' for formatting character to be ignored • Output format: 'V' for output of redacting, any other characters will be treated as formatting characters. • Redact character: a character that will be used to replace the actual values. Examples are '*' and 'x'. • Starting digit position: specifies the starting (character) position to begin replacing actual values with the redacting character. The beginning of the string is position 1. Positions do not include formatting characters. • Ending digit position: specifies the ending (character) position to end redacting. An example is 'VVVFVVFVVVV,VVV-VV-VVVV,X,1,5' for redacting the first 5 digits of SSN string 123-45-6789, and adding dashes back to format it like an SSN, resulting in XXX-XX-6789. For number datatypes, a comma-separated list containing these fields: • Redact character: this is a character between '0' to '9' that will be used to replace the actual values. • Starting digit position: specifies the starting (digit) position to begin replacing actual values with the redacting character. The beginning of the string is position 1. Positions do not include the decimal point. • Ending digit position: this specifies the ending digit position to end redacting. An example is '9,1,5' for redacting the first 5 digits of Social Security Number 123456789, resulting in 999996789. For datetime datatypes, the format is a packed string (no spaces or commas) containing the following sequence of fields. Please note that each field can consist of one or more characters, and the field length depends on whether redacting is required. The one-character fields are used to specify that no redaction of that component of the datetime value is to take place. The longer fields indicate a specific

Table 153-6 (Cont.) *ADD_POLICY Procedure Parameters*

Parameter	Description
	time or date to use as the redacted value of that component of the datetime value. • Month: `'M'` (no redacting of month) or `'m#'` (redact month to a specific month, if possible), where `#` (the month specified by its number) is between `1` and `12` • Day: `'D'` (no redacting of day) or `'d#'` (redact day to `#`, if possible), `#` between `1` and `31` • Year: `'Y'` (no redacting of year) or `'y#'` (redact year to `#`, if possible), `#` between `1` and `9999` • Hour: `'H'` (no redacting of hour) or `'h#'` (redact hour to `#`, if possible), `#` between `0` and `23` • Minute: `'M'` (no redacting of minute) or `'m#'` (redact minute to `#`, if possible), `#` between `0` and `59` • Second: `'S'` (no redacting of second) or `'s#'` (redact second to `#`, if possible), `#` between `0` and `59` An example is `'m12d1y2001h1m1s1'`, which changes `02-May-13 12:30:23` to `01-Dec-01 01:01:01`. For partial character and number-redacting shortcuts, see *Oracle Database Advanced Security Guide*.
`expression`	Default boolean expression for the table or view. If this expression is used, then redaction takes place only if this policy expression evaluates to `TRUE`. The following functions are supported: • `SYS_CONTEXT` • `XS_SYS_CONTEXT` • `SUBSTR` functions (includes `SUBSTRB`, `SUBSTRC`, `SUBSTR2`, `SUBSTR4`) • `LENGTH` functions (includes `LENGTHB`, `LENGTHC`, `LENGTH2`, `LENGTH4`) • `V` (`APEX_UTIL.GET_SESSION_STATE`) • `NV` (`APEX_UTIL.GET_NUMERIC_SESSION_STATE`) • `OLS_LABEL_DOMINATES` • `DOMINATES` • `OLS_DOMINATES` • `OLS_DOM` • `DOM` • `OLS_STRICTLY_DOMINATES` • `STRICTLY_DOMINATES` • `S_DOM` • `SA_UTL.DOMINATES` • `SA_UTL.CHECK_READ` • `SA_UTL.NUMERIC_LABEL` • `CHAR_TO_LABEL` • `SA_SESSION.LABEL` See *Oracle Database Advanced Security Guide* for more information about these supported functions

Table 153-6 (Cont.) *ADD_POLICY Procedure Parameters*

Parameter	Description
`enable`	Boolean value that determines whether the Data Redaction policy is enabled on creation. The default value is `TRUE`, which means that the policy is automatically enabled upon creation. If the `enable` parameter is set to `FALSE`, the policy takes effect only when it is subsequently enabled by calling the `DBMS_REDACT.ENABLE_POLICY` procedure.
`regexp_pattern`	Regular expression pattern up to 512 bytes. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. Also, do not specify the `function_parameters` parameter when `function_type` is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. See *Oracle Database SQL Language Reference* for more information and examples on using regular expression patterns.
`regexp_replace_string`	Replacement string (up to 4000 characters in length) with up to 500 back-references to subexpressions in the form `\n`, where `n` is a number between 1 and 9. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP`
`regexp_position`	Integer counting from 1, specifies the position where the search must begin. The default is 1. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`.
`regexp_occurrence`	• Use 0 to replace all occurrences of the match. The default is 0. • Use positive integer *n* to replace the *n*-th occurence of the match. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`.
`regexp_match_parameter`	Changes the default matching behavior, possible values are a combination of 'i', 'c', 'n', 'm', 'x' Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. See *Oracle Database SQL Language Reference* for more information and examples on using regular expression match parameters.
`policy_description`	Description of redaction policy
`column_description`	Description of the column being redacted

Exceptions

- `ORA-28060` - A Data Redaction policy already exists on this column.
- `ORA-28061` - This object cannot have a Data Redaction policy defined on it.
- `ORA-28062` - The policy expression is too long.
- `ORA-28063` - The policy expression is empty.
- `ORA-28064` - The redaction function is not valid.
- `ORA-28066` - Invalid column *column*.
- `ORA-28067` - Missing or invalid column name.

- ORA-28069 - A Data Redaction policy already exists on this object.
- ORA-28073 - The column *column_name* has an unsupported datatype or attribute.
- ORA-28074 - The *field_name* field of the redaction parameters is not valid.

  The field can be any of the following:

  - REDACT_PARTIAL_INPUT_FORMAT
  - REDACT_PARTIAL_OUTPUT_FORMAT
  - REDACT_PARTIAL_MASKCHAR
  - REDACT_PARTIAL_MASKFROM
  - REDACT_PARTIAL_MASKTO
  - REDACT_PARTIAL_DATE_MONTH
  - REDACT_PARTIAL_DATE_DAY
  - REDACT_PARTIAL_DATE_YEAR
  - REDACT_PARTIAL_DATE_HOUR
  - REDACT_PARTIAL_DATE_MINUTE
  - REDACT_PARTIAL_DATE_SECOND

  See Table 153-3 and Table 153-4 for examples of the field contents and field ordering.
- ORA-28075 - The data redaction policy expression had an error.
- ORA-28076 - Empty role name specified as SYS_SESSION_ROLES attribute.
- ORA-28077 - The specified role name *role_name* for SYS_SESSION_ROLES exceeds the maximum length.
- ORA-28078 - A regular expression parameter is missing or invalid.
- ORA-28082 - The parameter *parameter* is invalid (where the possible values are function_parameters, column_description, and policy_description).
- ORA-28085 - The input and output lengths of the redaction do not match.
- ORA-28086 - The data redaction policy expression had an error
- ORA-28087 - The policy expression has an unsupported (use of) operator <operator>.
- ORA-28088 - The policy expression has an unsupported PL/SQL function function_name.
- ORA-28089 - Invalid SYS_CONTEXT namespace <namespace> in policy expression.
- ORA-28090 - Unsupported nesting of function function_name within function function_name in policy expression.
- ORA-28091 - Unsupported use of PL/SQL function function_name in policy expression.
- ORA-28092 - The parameter parameter_name with value value has an error.
- ORA-28096 - The parameter parameter_name is unexpected for ADD_POLICY.
- ORA-28097 - A data redaction policy cannot be applied to an object owned by SYS.
- ORA-28104 - Input value for *parameter* is not valid.

Usage Notes

If you specify `1=1` for the expression parameter, then it is always `TRUE`, which causes the redaction to be performed as long as the querying user is not exempt from the Oracle Data Redaction policy. Specifying `1=1` for the expression parameter will improve Data Redaction performance because Oracle Database will not evaluate the policy expression.

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

See Operating Procedures for more information regarding function types and function parameters with related examples.

A Data Redaction named policy expression that has been applied to a redacted column takes precedence over the expression defined in the `expression` parameter. To find redacted columns that are affected by named policy expressions, query the `REDACTION_EXPRESSIONS` data dictionary view.

Example

Partial redaction policy:

```
BEGIN
   DBMS_REDACT.ADD_POLICY(
     object_schema        => 'hr',
     object_name          => 'employees',
     column_name          => 'employee_id',
     policy_name          => 'redact_emp_id_nums',
     function_type        => DBMS_REDACT.PARTIAL,
     function_parameters  => '7,1,5',
     expression           => '1=1');
END;
```

Full redaction policy:

```
BEGIN
   DBMS_REDACT.ADD_POLICY(
     object_schema        => 'hr',
     object_name          => 'employees',
     column_name          => 'employee_id',
     policy_name          => 'redact_emp_ids',
     function_type        => DBMS_REDACT.FULL,
     expression           => 'SYS_CONTEXT(''SYS_SESSION_ROLES'',''CLERK'')
                                                                  = ''FALSE''');
END;
```

153.5.2 ALTER_POLICY Procedure

This procedure alters an existing Data Redaction policy for a table or view.

It alters the policy in one or more of the following ways:

- By changing the policy expression
- By changing the type of redaction for a specified column
- By changing the parameters to the redaction function for a specified column
- By adding a column to the redaction policy (the redaction type and any parameters must be specified).
- By removing a column from the redaction policy

- By changing the description of the policy
- By changing the description of the column

Syntax

```
DBMS_REDACT.ALTER_POLICY (
   object_schema               IN    VARCHAR2 := NULL,
   object_name                 IN    VARCHAR2,
   policy_name                 IN    VARCHAR2,
   action                      IN    BINARY_INTEGER := DBMS_REDACT.ADD_COLUMN,
   column_name                 IN    VARCHAR2 := NULL,
   function_type               IN    BINARY_INTEGER := DBMS_REDACT.FULL,
   function_parameters         IN    VARCHAR2 := NULL,
   expression                  IN    VARCHAR2,
   regexp_pattern              IN    VARCHAR2 := NULL,
   regexp_replace_string       IN    VARCHAR2 := NULL,
   regexp_position             IN    BINARY_INTEGER := 1,
   regexp_occurrence           IN    BINARY_INTEGER := 0,
   regexp_match_parameter      IN    VARCHAR2 := NULL,
   policy_description          IN    VARCHAR2 := NULL,
   column_description          IN    VARCHAR2 := NULL);
```

Parameters

Table 153-7 *ALTER_POLICY Procedure Parameters*

Parameter	Description
`object_schema`	Schema owning the table or view, current user if `NULL`
`object_name`	Name of table or view on which to alter a Data Redaction policy
`policy_name`	Name of policy limited to 30 bytes
`action`	Action to take. For more information see Table 153-2.
`column_name`	(Optional) Name of one column to which the redaction policy applies. Required for the following actions: • `add_column` • `drop_column` • `modify_column` • `set_column_description`

Table 153-7 (Cont.) *ALTER_POLICY Procedure Parameters*

Parameter	Description
function_type	Type of redaction function to use. Possible values are: - DBMS_REDACT.NONE - DBMS_REDACT.FULL (default) - DBMS_REDACT.NULLIFY - DBMS_REDACT.PARTIAL - DBMS_REDACT.RANDOM - DBMS_REDACT.REGEXP - DBMS_REDACT.REGEXP_WIDTH If the function_type is not DBMS_REDACT.PARTIAL, then you must omit the function_parameters parameter. If the function_type is not DBMS_REDACT.REGEXP or DBMS_REDACT.REGEXP_WIDTH, then you must omit the regexp_* parameters. If the function_type is DBMS_REDACT.REGEXP or DBMS_REDACT.REGEXP_WIDTH, then use the regexp_pattern, regexp_replace_string, regexp_position, regexp_occurrence, and regexp_match_parameter to define the Data Redaction policy. See Table 153-1 for an overview of the meanings of these values, and for some examples of their use.

Table 153-7 (Cont.) *ALTER_POLICY Procedure Parameters*

Parameter	Description
function_parameters	Parameters to the redaction function. Use only if function_type is DBMS_REDACT.PARTIAL. If the function_type is DBMS_REDACT.REGEXP or DBMS_REDACT.REGEXP_WIDTH, then you must omit the function_parameters parameter, and use the regexp_pattern, regexp_replace_string, regexp_position, regexp_occurrence, and regexp_match_parameter to define the Data Redaction policy. - If the function_type is DBMS_REDACT.NONE or DBMS_REDACT.NULLIFY, the function_parameters parameter must be omitted. - If the function_type is DBMS_REDACT.FULL or DBMS_REDACT.RANDOM, the function_parameters parameter must be omitted. - If the function_type is DBMS_REDACT.PARTIAL, the function_parameters parameter represents the redacting parameters for partial redacting. • Input format: 'V' for value to be possibly redacted, 'F' for formatting character to be ignored • Output format: 'V' for output of redaction, any other characters will be treated as formatting characters • Redact character: a character that will be used to replace the actual values. Examples are '*' and 'x' • Starting digit position: specifies the starting (character) position to begin replacing actual values with the redacting character. The beginning of the string is position 1. Positions do not include formatting characters. • Ending digit position: specifies the ending (character) position to end redacting. An example is 'VVVFVVFVVVV,VVV-VV-VVVV,X,1,5' for redacting the first 5 digits of SSN string 123-45-6789, and adding dashes back to format it like an SSN, resulting in XXX-XX-6789. For number datatypes, a comma-separated list containing these fields: • Redact character: this is a character between '0' to '9' that will be used to replace the actual values. • Starting digit position: specifies the starting (digit) position to begin replacing actual values with the redacting character. The beginning of the string is position 1. Positions do not include the decimal point. • Ending digit position: this specifies the ending digit position to end redacting. An example is '9,1,5' for redacting the first 5 digits of Social Security Number 123456789, resulting in 999996789. For datetime datatypes, the format is a packed string (no spaces or commas) containing the following sequence of fields. Please note that each field can consist of one or more characters, and the field length depends on whether redacting is required. The one-character fields are used to specify that no redaction of that component of the datetime value is to take place. The longer fields indicate a specific time or date to use as the redacted value of that component of the datetime value.

Table 153-7 (Cont.) *ALTER_POLICY Procedure Parameters*

Parameter	Description
	• Month: `'M'` (no redacting of month) or `'m#'` (redact month to a specific month, if possible), where `#` (the month specified by its number) is between `1` and `12`. • Day: `'D'` (no redacting of date) or `'d#'` (redact day to `#`, if possible), `#` between `1` and `31`. • Year: `'Y'` (no redacting of year) or `'y#'` (redact year to `#`, if possible), `#` between `1` and `9999`. • Hour: `'H'` (no redacting of hour) or `'h#'` (redact hour to `#`, if possible), `#` between `0` and `23`. • Minute: `'M'` (no redacting of minute) or `'m#'` (redact minute to `#`, if possible), `#` between `0` and `59`. • Second: `'S'` (no redacting of second) or `'s#'` (redact second to `#`, if possible), `#` between `0` and `59`. An example is `'m12DYHMS'`, which changes `01-May-01 01:01:01` to `01-Dec-01 01:01:01` For partial character and number-redacting shortcuts, see *Oracle Database Advanced Security Guide*.
`expression`	Default boolean expression for the table or view. Redaction takes place only if this policy expression evaluates to `TRUE`. The following functions are supported: • `SYS_CONTEXT` • `XS_SYS_CONTEXT` (See *Oracle Database Advanced Security Guide* • `NV` (`APEX_UTIL.GET_NUMERIC_SESSION_STATE`) • `V` (`APEX_UTIL.GET_SESSION_STATE`) • `OLS_LABEL_DOMINATES` • `DOMINATES` • `OLS_DOMINATES` • `OLS_DOM` • `DOM` • `OLS_STRICTLY_DOMINATES` • `STRICTLY_DOMINATES` • `S_DOM` • `SA_UTL.DOMINATES` • `SA_UTL.CHECK_READ` • `SA_UTL.NUMERIC_LABEL` • `CHAR_TO_LABEL` • `SA_SESSION.LABEL`
`regexp_pattern`	Regular expression pattern up to 512 bytes. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. Also, do not specify the `function_parameters` parameter when `function_type` is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. See *Oracle Database SQL Language Reference* for more information and examples on using regular expression patterns.

Table 153-7 (Cont.) *ALTER_POLICY Procedure Parameters*

Parameter	Description
`regexp_replace_string`	Replacement string (up to 4000 characters in length) with up to 500 back-references to subexpressions in the form `\n`, where `n` is a number between 1 and 9. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`.
`regexp_position`	Integer counting from 1, specifies the position where the search must begin. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. The default is 1 meaning that Oracle Database begins the search at the first character of the `column_name` data.
`regexp_occurrence`	• Use 0 to replace all occurrences of the match. 0 is the default. • Use positive integer *n* to replace the *n*-th occurence of the match. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`.
`regexp_match_parameter`	Changes the default matching behavior, possible values are a combination of 'i', 'c', 'ic', 'iq', 'n', 'm', 'x'. The behavior of this parameter is the same for this function as for the `REGEXP_REPLACE` SQL function. Use only if the `function_type` parameter is `DBMS_REDACT.REGEXP` or `DBMS_REDACT.REGEXP_WIDTH`. See *Oracle Database SQL Language Reference* for more information and examples on using regular expression match parameters.
`policy_description`	Description of redaction policy
`column_description`	Description of the column being redacted

Exceptions

- `ORA-28060` - A data redaction policy already exists on this column.
- `ORA-28061` - This object cannot have a data redaction policy defined on it.
- `ORA-28062` - The policy expression is too long.
- `ORA-28063` - The policy expression is empty.
- `ORA-28064` - The redaction function is not valid.
- `ORA-28066` - Invalid column *`column`*.
- `ORA-28067` - Missing or invalid column name.
- `ORA-28068` - The object *`object`* does not have a Data Redaction policy.
- `ORA-28069` - A Data Redaction policy already exists on this object.
- `ORA-28070` - The column *`column`* does not have a Data Redaction policy.
- `ORA-28071` - The action is not valid.
- `ORA-28072` - The specified policy name is incorrect.
- `ORA-28073` - The column *`column_name`* has an unsupported datatype or attribute.
- `ORA-28074` - The *`field_name`* field of the redaction parameters is not valid.

The field can be any of the following:

- REDACT_PARTIAL_INPUT_FORMAT
- REDACT_PARTIAL_OUTPUT_FORMAT
- REDACT_PARTIAL_MASKCHAR
- REDACT_PARTIAL_MASKFROM
- REDACT_PARTIAL_MASKTO
- REDACT_PARTIAL_DATE_MONTH
- REDACT_PARTIAL_DATE_DAY
- REDACT_PARTIAL_DATE_YEAR
- REDACT_PARTIAL_DATE_HOUR
- REDACT_PARTIAL_DATE_MINUTE
- REDACT_PARTIAL_DATE_SECOND

See Table 153-3 and Table 153-4 for examples of the field contents and field ordering.

- ORA-28075 - The data redaction policy expression had an error.
- ORA-28076 - Empty role name specified as SYS_SESSION_ROLES attribute.
- ORA-28077 - The specified role name *role_name* for SYS_SESSION_ROLES exceeds the maximum length.
- ORA-28078 - A regular expression parameter is missing or invalid.
- ORA-28082 - The parameter *parameter* is invalid (where the possible values are function_parameters, column_description, and policy_description).
- ORA-28085 - The input and output lengths of the redaction do not match.
- ORA-28086 - The data redaction policy expression had an error
- ORA-28087 - The policy expression has an unsupported (use of) operator <operator>.
- ORA-28088 - The policy expression has an unsupported PL/SQL function function_name.
- ORA-28089 - Invalid SYS_CONTEXT namespace <namespace> in policy expression.
- ORA-28090 - Unsupported nesting of function function_name within function function_name in policy expression.
- ORA-28091 - Unsupported use of PL/SQL function function_name in policy expression.
- ORA-28092 - The parameter parameter_name with value value has an error.
- ORA-28095 - The parameter parameter_name is unexpected for ACTION <action> for ALTER_POLICTY.
- ORA-28097 - A data redaction policy cannot be applied to an object owned by SYS.
- ORA-28104 - Input value for *parameter* is not valid.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

See Operating Procedures for more information regarding Function Types and Function Parameters with related examples.

A Data Redaction named policy expression that has been applied to a redacted column takes precedence over the expression defined in the `expression` parameter. To find redacted columns that are affected by named policy expressions, query the `REDACTION_EXPRESSIONS` data dictionary view.

Examples

```
BEGIN
   DBMS_REDACT.ALTER_POLICY(
     object_schema          => 'HR',
     object_name            => 'EMPLOYEES',
     policy_name            => 'redact_emp_id_nums',
     action                 => DBMS_REDACT.DROP_COLUMN,
     column_name            => 'EMAIL');
END;
```

153.5.3 APPLY_POLICY_EXPR_TO_COL Procedure

This procedure associates an Oracle Data Redaction named policy expression with a redacted column from a table or view.

Syntax

```
DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL (
   object_schema                IN VARCHAR2 := NULL,
   object_name                  IN VARCHAR2,
   column_name                  IN VARCHAR2,
   policy_expression_name       IN VARCHAR2 := NULL);
```

Parameters

Table 153-8 *APPLY_POLICY_EXPR_TO_COL Procedure Parameters*

Parameter	Description
`object_schema`	Name of the schema that contains the redacted column. If omitted, then the current schema is used.
`object_name`	Name of the object (table or view) that contains the redacted column
`column_name`	Name of the redacted column to which the policy expression is applied
`policy_expression_name`	If `NULL`, then the named policy expression associated with the redacted column `column_name` is removed.

Exceptions

- `ORA-28061` - This object cannot have a data redaction policy defined on it.
- `ORA-28066` - Invalid column `column_name`.
- `ORA-28067` - Missing or invalid column name.
- `ORA-28068` - The object *object* does not have a Data Redaction policy.
- `ORA-28070` - The column `column_name` does not have a Data Redaction policy.
- `ORA-28082` - The parameter *parameter* is invalid.
- `ORA-28092` - The parameter *parameter* with value *value* has an error.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

You can find existing Data Redaction policy expressions by querying the `REDACTION_EXPRESSIONS` data dictionary view. To find columns that have been redacted, query the `REDACTION_COLUMNS` data dictionary view.

Example

```
BEGIN
   DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL(
      object_schema           => 'OE',
      object_name             => 'CUSTOMERS',
      column_name             => 'INCOME_LEVEL',
      policy_expression_name  => 'oe_redact_pol');
END;
```

153.5.4 CREATE_POLICY_EXPRESSION Procedure

This procedure creates an Oracle Data Redaction named policy expression.

Syntax

```
DBMS_REDACT.CREATE_POLICY_EXPRESSION (
   policy_expression_name          IN    VARCHAR2,
   expression                      IN    VARCHAR2,
   policy_expression_description   IN    VARCHAR2 := NULL);
```

Parameters

Table 153-9 *CREATE_POLICY_EXPRESSION Procedure Parameters*

Parameter	Description
`policy_expression_name`	Name of the policy expression
`expression`	Definition of the policy expression.
`policy_expression_description`	Description of the policy expression

Exceptions

- `ORA-28082` - The parameter *parameter* is invalid.
- `ORA-28092` - The parameter *parameter* with value *value* has an error.

Usage Notes

If the expression parameter is `1=1`, then it is always `TRUE`, which causes the redaction to be performed as long as the querying user is not exempt from the Oracle Data Redaction policy. Specifying this parameter as `1=1` will improve Data Redaction performance because Oracle Database will not evaluate the policy expression.

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

See APPLY_POLICY_EXPR_TO_COL Procedure for how to apply policy expressions to table columns.

After you create a policy expression, you can associate it with a redacted table or view column by running the DBMS_REDACT.APPLY_POLICY_EXPR_TO_COL procedure. To find existing redacted columns, query the REDACTION_COLUMNS data dictionary view.

Example

```
BEGIN
   DBMS_REDACT.CREATE_POLICY_EXPRESSION(
     policy_expression_name          => 'oe_redact_pol',
     expression                      => 'SYS_CONTEXT(''USERENV'',''SESSION_USER'') =
''OE''''),
     policy_expression_description   => 'Enables policy for user OE ');
END;
```

153.5.5 DISABLE_POLICY Procedure

This procedure disables a Data Redaction policy.

Syntax

```
DBMS_REDACT.DISABLE_POLICY (
   object_schema               IN    VARCHAR2 := NULL,
   object_name                 IN    VARCHAR2,
   policy_name                 IN    VARCHAR2);
```

Parameters

Table 153-10 *DISABLE_POLICY Procedure Parameters*

Parameter	Description
object_schema	Schema owning the table or view, current user if NULL
object_name	Name of table or view for which to disable a Data Redaction policy
policy_name	Name of policy to be disabled

Exceptions

- ORA_28061 - This object cannot have a Data Redaction policy defined on it.
- ORA-28068 - The object *object* does not have a Data Redaction policy.
- ORA-28072 - The specified policy name is incorrect.
- ORA-28080 - The policy was already disabled.
- ORA_28104 - Input value for policy-name is not valid.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

Examples

```
BEGIN
   DBMS_REDACT.DISABLE_POLICY (
     object_schema    => 'hr',
     object_name      => 'employees',
     policy_name      => 'redact_emp_ids');
END;
```

153.5.6 DROP_POLICY Procedure

This procedure drops a Data Redaction policy by removing a redaction policy from the table or view.

Syntax

```
DBMS_REDACT.DROP_POLICY (
   object_schema                IN    VARCHAR2 := NULL,
   object_name                  IN    VARCHAR2,
   policy_name                  IN    VARCHAR2);
```

Parameters

Table 153-11 *DROP_POLICY Procedure Parameters*

Parameter	Description
`object_schema`	Schema owning the table or view, current user if `NULL`
`object_name`	Name of table or view from which to drop a Data Redaction policy
`policy_name`	Name of policy to be dropped

Exceptions

- `ORA_28061` - This object cannot have a Data Redaction policy defined on it.
- `ORA-28068` - The object *`object`* does not have a Data Redaction policy.
- `ORA-28072` - The specified policy name is incorrect.
- `ORA_28104` - Input value for `policy-name` is not valid.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

Examples

```
BEGIN
   DBMS_REDACT.DROP_POLICY (
      object_schema    =>  'hr',
      object_name      =>  'employees',
      policy_name      =>  'redact_emp_ids');
END;
```

153.5.7 DROP_POLICY_EXPRESSION Procedure

This procedure drops a named policy expression.

Syntax

```
DBMS_REDACT.DROP_POLICY_EXPRESSION (
   policy_expression_name       IN    VARCHAR2);
```

Parameters

Table 153-12 *DROP_POLICY_EXPRESSION Procedure Parameters*

Parameter	Description
policy_expression_name	Name of the policy expression

Exceptions

- ORA-28082 - The parameter *parameter* is invalid.
- ORA-28092 - The parameter *parameter* with value *value* has an error.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

You can find existing Data Redaction policy expressions by querying the REDACTION_EXPRESSIONS data dictionary view.

Example

```
BEGIN
   DBMS_REDACT.DROP_POLICY_EXPRESSION(
     policy_expression_name  => 'oe_redact_pol');
END;
```

153.5.8 ENABLE_POLICY Procedure

This procedure re-enables a Data Redaction policy.

Syntax

```
DBMS_REDACT.ENABLE_POLICY (
   object_schema                IN    VARCHAR2 := NULL,
   object_name                  IN    VARCHAR2,
   policy_name                  IN    VARCHAR2);
```

Parameters

Table 153-13 *ENABLE_POLICY Procedure Parameters*

Parameter	Description
object_schema	Schema owning the table or view, current user if NULL
object_name	Name of table or view on which to enable a Data Redaction policy
policy_name	Name of policy to be enabled

Exceptions

- ORA-28068 - The object *object* does not have a Data Redaction policy.
- ORA-28071 - The action is not valid.
- ORA-28072 - The specified policy name is incorrect.
- ORA-28079 - The policy was already enabled.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

Examples

```
BEGIN
   DBMS_REDACT.ENABLE_POLICY (
      object_schema    =>  'hr',
      object_name      =>  'employees',
      policy_name      =>  'redact_emp_ids');
END;
```

153.5.9 UPDATE_FULL_REDACTION_VALUES Procedure

This procedure modifies the redacted output that will be displayed when querying a column having a Data Redaction policy where the `function_type` is `DBMS_REDACT.FULL`. Restart the database instance to apply these changes.

Syntax

```
DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES (
   number_val       IN NUMBER                    := NULL,
   binfloat_val     IN BINARY_FLOAT              := NULL,
   bindouble_val    IN BINARY_DOUBLE             := NULL,
   char_val         IN CHAR                      := NULL,
   varchar_val      IN VARCHAR2                  := NULL,
   nchar_val        IN NCHAR                     := NULL,
   nvarchar_val     IN NVARCHAR2                 := NULL,
   date_val         IN DATE                      := NULL,
   ts_val           IN TIMESTAMP                 := NULL,
   tswtz_val        IN TIMESTAMP WITH TIME ZONE  := NULL,
   blob_val         IN BLOB                      := NULL,
   clob_val         IN CLOB                      := NULL,
   nclob_val        IN NCLOB                     := NULL,
   boolean_val      IN BOOLEAN                   := NULL);
```

Parameters

Table 153-14 ***UPDATE_FULL_REDACTION_VALUES Procedure Parameters***

Parameter	Description
`number_val`	Modifies the default value for columns of the `NUMBER` datatype
`binfloat_val`	Modifies the default value for columns of the `BINARY_FLOAT` datatype
`bindouble_val`	Modifies the default value for columns of the `BINARY_DOUBLE` datatype
`char_val`	Modifies the default value for columns of the `CHAR` datatype
`varchar_val`	Modifies the default value for columns of the `VARCHAR2` datatype
`nchar_val`	Modifies the default value for columns of the `NCHAR` datatype
`nvarchar_val`	Modifies the default value for columns of the `NVARCHAR2` datatype
`date`	Modifies the default value for columns of the `DATE` datatype
`ts_val`	Modifies the default value for columns of the `TIMESTAMP` datatype
`tswtz_val`	Modifies the default value for columns of the `TIMESTAMP WITH TIME ZONE` datatype

Table 153-14 (Cont.) *UPDATE_FULL_REDACTION_VALUES Procedure Parameters*

Parameter	Description
blob_val	Modifies the default value for columns of the BLOB datatype
clob_val	Modifies the default value for columns of the CLOB datatype
nclob_val	Modifies the default value for columns of the NCLOB datatype
boolean_val	Modifies the default value for columns of the BOOLEAN datatype

Example 153-1 Change the Current Value for Columns That Use the NUMBER Data Type

1. `EXEC DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES (number_val => 7);`
2. Close the PDB:

   ```
   ALTER PLUGGABLE DATABASE <pdb_name> CLOSE IMMEDIATE;
   ```

3. Re-open the PDB:

   ```
   ALTER PLUGGABLE DATABASE <pdb_name> OPEN;
   ```

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

153.5.10 UPDATE_POLICY_EXPRESSION Procedure

This procedure updates an Oracle Data Redaction named policy expression.

Syntax

```
DBMS_REDACT.UPDATE_POLICY_EXPRESSION (
   policy_expression_name          IN    VARCHAR2,
   expression                      IN    VARCHAR2,
   policy_expression_description   IN    VARCHAR2 := NULL);
```

Parameters

Table 153-15 *UPDATE_POLICY_EXPRESSION Procedure Parameters*

Parameter	Description
policy_expression_name	Name of the policy expression
expression	Definition of the policy expression
policy_expression_description	Description of the policy expression

Exceptions

- ORA-28082 - The parameter *parameter* is invalid.
- ORA-28092 - The parameter *parameter* with value *value* has an error.

Usage Notes

See DBMS_REDACT Security Model for a list of the required privileges for this procedure.

You can find existing policy expressions by querying the `REDACTION_EXPRESSIONS` data dictionary view.

Example

```
BEGIN
   DBMS_REDACT.UPDATE_POLICY_EXPRESSION(
     policy_expression_name         => 'oe_redact_pol',
     expression                     => 'SYS_CONTEXT(''USERENV'',''SESSION_USER'') !=
''OE''''),
     policy_expression_description  => 'Updates policy expression for oe_redact_pol');
END;
```

154
DBMS_REDEFINITION

The `DBMS_REDEFINITION` package provides an interface to perform an online redefinition of tables.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operational Notes
- Rules and Limits
- Examples
- Summary of DBMS_REDEFINITION Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more information about online redefinition of tables

154.1 DBMS_REDEFINITION Overview

To achieve online redefinition, incrementally maintainable local materialized views are used. These logs keep track of the changes to the master tables and are used by the materialized views during refresh synchronization.

154.2 DBMS_REDEFINITION Security Model

Subprograms in the `DBMS_REDEFINITION` package are run with invokers' rights (with the privileges of the current user).

There are two modes:

- In `USER` mode, the user who has the `CREATE TABLE` and `CREATE MVIEW` privileges may redefine a table residing in his own schema.
- In `FULL` mode, the user who has the `ANY` privilege may redefine tables in any schema.

154.3 DBMS_REDEFINITION Constants

The `DBMS_REDEFINITION` package defines several constants for specifying parameter values.

Table 154-1 DBMS_REDEFINITION Constants

Constant	Type	Value	Description
CONS_CONSTRAINT	PLS_INTEGER	3	Used to specify that dependent object type is a constraint
CONS_INDEX	PLS_INTEGER	2	Used to specify that dependent object type is a index
CONS_MVLOG	PLS_INTEGER	10	Used to (un)register a materialized view log, as a dependent object of the table, through the REGISTER_DEPENDENT_OBJECT Procedureand the UNREGISTER_DEPENDENT_OBJECT Procedure.
CONS_ORIG_PARAMS	PLS_INTEGER	1	Used to specify that indexes should be cloned with their original storage parameters
CONS_TRIGGER	PLS_INTEGER	4	Used to specify that dependent object type is a trigger
CONS_USE_PK	BINARY_INTEGER	1	Used to indicate that the redefinition should be done using primary keys or pseudo-primary keys (unique keys with all component columns having not-NULL constraints)
CONS_USE_ROWID	BINARY_INTEGER	2	Used to indicate that the redefinition should be done using rowids
CONS_VPD_AUTO	BINARY_INTEGER	2	Used to indicate to copy VPD policies automatically
CONS_VPD_MANUAL	BINARY_INTEGER	4	Used to indicate to copy VPD policies manually
CONS_VPD_NONE	BINARY_INTEGER	1	Used to indicate that there are no VPD policies on the original table

154.4 DBMS_REDEFINITION Operational Notes

The following operational notes apply to DBMS_REDEFINITION.

- CONS_USE_PK and CONS_USE_ROWID are constants used as input to the "options_flag" parameter in both the START_REDEF_TABLE Procedure and CAN_REDEF_TABLE Procedure. CONS_USE_ROWID is used to indicate that the redefinition should be done using rowids while CONS_USE_PK implies that the redefinition should be done using primary keys or pseudo-primary keys (which are unique keys with all component columns having NOT NULL constraints).
- CONS_INDEX, CONS_MVLOG,CONS_TRIGGER and CONS_CONSTRAINT are used to specify the type of the dependent object being (un)registered in REGISTER_DEPENDENT_OBJECT Procedure and UNREGISTER_DEPENDENT_OBJECT Procedure (parameter "dep_type").

  CONS_INDEX ==> dependent object is of type INDEX

  CONS_TRIGGER ==> dependent object is of type TRIGGER

  CONS_CONSTRAINT==> dependent object type is of type CONSTRAINT

  CONS_MVLOG ==> dependent object is of type MATERIALIZED VIEW LOG

- `CONS_ORIG_PARAMS` as used as input to the "`copy_indexes`" parameter in COPY_TABLE_DEPENDENTS Procedure. Using this parameter implies that the indexes on the original table be copied onto the interim table using the same storage parameters as that of the original index.

154.5 DBMS_REDEFINITION Rules and Limits

Various rules and limits apply to implementation of the `DBMS_REDEFINITION` package.

For more information about these, see the *Oracle Database Administrator's Guide*.

154.6 DBMS_REDEFINITION Examples

The following examples demonstrate `DBMS_REDEFINITION` functionality.

We create two tables `EMP` and `EMP_INT` as the original and the interim tables, respectively:

```
CREATE TABLE "EMP"
( "EMPNO" NUMBER(4,0) PRIMARY KEY,
"ENAME" VARCHAR2(10),
"JOB" VARCHAR2(9),
"MGR" NUMBER(4,0),
"HIREDATE" DATE,
"SAL" NUMBER(7,2),
"COMM" NUMBER(7,2),
"DEPTNO" NUMBER(2,0))

TABLESPACE myts;

CREATE TABLE "EMP_INT"
( "EMPNO" NUMBER(4,0) PRIMARY KEY,
"ENAME" VARCHAR2(10),
"JOB" VARCHAR2(9),
"MGR" NUMBER(4,0),
"HIREDATE" DATE,
"SAL" NUMBER(7,2),
"COMM" NUMBER(7,2),
"DEPTNO" NUMBER(2,0))

TABLESPACE compressed_ts;
```

Regular Multi-Step Redefinition

```
DBMS_REDEFINITION.START_REDEF_TABLE('SCOTT', 'EMP', 'EMP_INT',
ENABLE_ROLLBACK => TRUE);
DBMS_REDEFINITION.FINISH_REDEF_TABLE('SCOTT', 'EMP', 'EMP_INT');
```

Assume theDBA wants to evaluate the performance of the application for 2 days, after moving the table `EMP` from tablespace `myts` to `compressed_ts`. One can run sync_interim_table SYNC_INTERIM_TABLE Procedureto keep both the tables in sync (say, every hour).

```
DBMS_REDEFINITION.SYNC_INTERIM_TABLE('SCOTT', 'EMP', 'EMP_INT');
```

Case 1 — DBA is not happy with the performance, so decides to rollback.

```
DBMS_REDEFINITION.ROLLBACK('SCOTT', 'EMP', 'EMP_INT');
```

Case 2 — DBA is happy with the performance, so decides not to rollback.

```
DBMS_REDEFINITION.ABORT_ROLLBACK('SCOTT', 'EMP', 'EMP_INT');
```

This terminates the possibility of rollback.

Single-Step Redefinition

```
DBMS_REDEFINITION.REDEF_TABLE('SCOTT','EMP','ROW STORE COMPRESS ADVANCED',
enable_rollback => TRUE);
```

> **Note:**
>
> Online table redefinition rollback is not supported when the `REDEF_TABLE` procedure is used to redefine a table.

154.7 Summary of DBMS_REDEFINITION Subprograms

This table lists the `DBMS_REDEFINITION` subprograms and briefly describes them.

Table 154-2 DBMS_REDEFINITION Package Subprograms

Subprogram	Description
ABORT_REDEF_TABLE Procedure	Cleans up errors that occur during the redefinition process and removes all temporary objects created by the reorganization process
ABORT_ROLLBACK Procedure	Aborts rollback
ABORT_UPDATE Procedure	Aborts an update started with the `EXECUTE_UPDATE` procedure
CAN_REDEF_TABLE Procedure	Determines if a given table can be redefined online
COPY_TABLE_DEPENDENTS Procedure	Copies the dependent objects of the original table onto the interim table
EXECUTE_UPDATE Procedure	Optimizes the performance of bulk updates to a table
FINISH_REDEF_TABLE Procedure	Completes the redefinition process
REDEF_TABLE Procedure	Provides a single push-button interface that integrates several redefinition steps
REGISTER_DEPENDENT_OBJECT Procedure	Registers a dependent object (index, trigger, constraint or materialized view log) on the table being redefined and the corresponding dependent object on the interim table
ROLLBACK Procedure	Performs rollback
SET_PARAM Procedure	Sets a new value for a specified parameter used by the redefinition process identified by a redefinition ID

Table 154-2 (Cont.) DBMS_REDEFINITION Package Subprograms

Subprogram	Description
START_REDEF_TABLE Procedure	Initiates the redefinition process
SYNC_INTERIM_TABLE Procedure	Keeps the interim table synchronized with the original table
UNREGISTER_DEPENDENT_OBJECT Procedure	Unregisters a dependent object (index, trigger, constraint or materialized view log) on the table being redefined and the corresponding dependent object on the interim table

154.7.1 ABORT_REDEF_TABLE Procedure

This procedure cleans up errors that occur during the redefinition process.

This procedure can also be used to terminate the redefinition process any time after the START_REDEF_TABLE Procedure has been called and before the FINISH_REDEF_TABLE Procedure is called. This process will remove the temporary objects that are created by the redefinition process such as materialized view logs.

Syntax

```
DBMS_REDEFINITION.ABORT_REDEF_TABLE (
   uname                  IN  VARCHAR2,
   orig_table             IN  VARCHAR2,
   int_table              IN  VARCHAR2,
   part_name              IN  VARCHAR2 := NULL);
```

Parameters

Table 154-3 ABORT_REDEF_TABLE Procedure Parameters

Parameter	Description
uname	Schema name of the tables
orig_table	Name of the table to be redefined
int_table	Name of the interim table. Can take a comma-delimited list of interim table names.
part_name	Name of the partition being redefined. If redefining only a single partition of a table, specify the partition name in this parameter. NULL implies the entire table is being redefined. Can take a comma-delimited list of partition names to be redefined.

154.7.2 ABORT_ROLLBACK Procedure

This procedure aborts rollback for a table that was redefined.

When online redefinition of a table is started with the START_REDEF_TABLE procedure, rollback can be enabled for the changes performed by online redefinition of a table by setting the enable_rollback parameter to TRUE. If you want to retain the changes made by online redefinition, you can abort the rollback to clean up the database objects that enable rollback.

Syntax

```
DBMS_REDEFINITION.ABORT_ROLLBACK (
   uname           IN  VARCHAR2,
   orig_table      IN  VARCHAR2,
   int_table       IN  VARCHAR2 := NULL,
   part_name       IN  VARCHAR2 := NULL);
```

Parameters

Table 154-4 ABORT_ROLLBACK Procedure Parameters

Parameter	Description
uname	Schema name of the tables
orig_table	Name of the table to be redefined
int_table	Name of the interim table
part_name	Name of the partition being redefined

154.7.3 ABORT_UPDATE Procedure

This procedure can aborts an update started with the `EXECUTE_UPDATE` procedure in the `RDBMS_REDEFINITION` package.

Syntax

```
DBMS_REDEFINITION.ABORT_UPDATE (
   update_stmt  IN  CLOB);
```

Parameters

Table 154-5 ABORT_UPDATE Procedure Parameters

Parameter	Description
update_stmt	The SQL `UPDATE` statement to be aborted The SQL statement must exactly match the SQL statement in the `EXECUTE_UPDATE` procedure.

> **See Also:**
>
> *Oracle Database Administrator's Guide*

154.7.4 CAN_REDEF_TABLE Procedure

This procedure determines if a given table can be redefined online. This is the first step of the online redefinition process. If the table is not a candidate for online redefinition, an error message is raised.

Syntax

```
DBMS_REDEFINITION.CAN_REDEF_TABLE (
   uname         IN  VARCHAR2,
   tname        IN  VARCHAR2,
   options_flag  IN  PLS_INTEGER := 1,
   part_name     IN  VARCHAR2 := NULL);
```

Parameters

Table 154-6 CAN_REDEF_TABLE Procedure Parameters

Parameter	Description
uname	Schema name of the table
tname	Name of the table to be re-organized
options_flag	Indicates the type of redefinition method to use. • If dbms_redefinition.cons_use_pk, the redefinition is done using primary keys or pseudo-primary keys (unique keys with all component columns having NOT NULL constraints). The default method of redefinition is using primary keys. • If dbms_redefinition.cons_use_rowid, the redefinition is done using rowids.
part_name	Name of the partition being redefined. If redefining only a single partition of a table, specify the partition name in this parameter. NULL implies the entire table is being redefined.

Exceptions

If the table is not a candidate for online redefinition, an error message is raised.

154.7.5 COPY_TABLE_DEPENDENTS Procedure

This procedure clones the dependent objects of the table being redefined onto the interim table and registers the dependent objects. This procedure does not clone the already registered dependent objects.

This subprogram is used to clone the dependent objects like grants, triggers, constraints and privileges from the table being redefined to the interim table (which represents the post-redefinition table).

Syntax

```
DBMS_REDEFINITION.COPY_TABLE_DEPENDENTS(
   uname                  IN  VARCHAR2,
   orig_table             IN  VARCHAR2,
   int_table              IN  VARCHAR2,
   copy_indexes           IN  PLS_INTEGER := 1,
   copy_triggers          IN  BOOLEAN     := TRUE,
   copy_constraints       IN  BOOLEAN     := TRUE,
```

```
  copy_privileges          IN  BOOLEAN     := TRUE,
  ignore_errors            IN  BOOLEAN     := FALSE,
  num_errors               OUT PLS_INTEGER,
  copy_statistics          IN  BOOLEAN     := FALSE,
  copy_mvlog               IN  BOOLEAN     := FALSE);
```

Parameters

Table 154-7 COPY_TABLE_DEPENDENTS Procedure Parameters

Parameter	Description
uname	Schema name of the tables
orig_table	Name of the table being redefined
int_table	Name of the interim table
copy_indexes	Flag indicating whether to copy the indexes • 0 - do not copy any index • dbms_redefinition.cons_orig_params – copy the indexes using the physical parameters of the source indexes
copy_triggers	TRUE = clone triggers, FALSE = do nothing
copy_constraints	TRUE = clone constraints, FALSE = do nothing. If compatibility setting is 10.2 or higher, then clone CHECK and NOT NULL constraints
copy_privileges	TRUE = clone privileges, FALSE = do nothing
ignore_errors	TRUE = if an error occurs while cloning a particular dependent object, then skip that object and continue cloning other dependent objects. FALSE = that the cloning process should stop upon encountering an error.
num_errors	Number of errors that occurred while cloning dependent objects
copy_statistics	TRUE = copy statistics, FALSE = do nothing
copy_mvlog	TRUE = copy materialized view log, FALSE = do nothing

Usage Notes

- The user must check the column num_errors before proceeding to ensure that no errors occurred during the cloning of the objects.
- In case of an error, the user should fix the cause of the error and call the COPY_TABLE_DEPENDENTS Procedure again to clone the dependent object. Alternatively the user can manually clone the dependent object and then register the manually cloned dependent object using the REGISTER_DEPENDENT_OBJECT Procedure.
- All cloned referential constraints involving the interim tables will be created disabled (they will be automatically enabled after the redefinition) and all triggers on interim tables will not fire till the redefinition is completed. After the redefinition is complete, the cloned objects will be renamed to the corresponding pre-redefinition names of the objects (from which they were cloned from).
- It is the user's responsibility that the cloned dependent objects are unaffected by the redefinition. All the triggers will be cloned and it is the user's responsibility that the cloned triggers are unaffected by the redefinition.

154.7.6 EXECUTE_UPDATE Procedure

This procedure can optimize the performance of bulk updates to a table. Performance is optimized because the updates are not logged in the redo log.

The `EXECUTE_UPDATE` procedure automatically uses the components of online table redefinition, such an interim table, a materialized view, and a materialized view log, to enable optimized bulk updates to a table. The `EXECUTE_UPDATE` procedure also removes fragmentation of the affected rows and ensures that the update is atomic. If the bulk updates raise any errors, then you can use the `ABORT_UPDATE` procedure to undo the changes made by the `EXECUTE_UPDATE` procedure.

Syntax

```
DBMS_REDEFINITION.EXECUTE_UPDATE (
   update_stmt  IN  CLOB);
```

Parameters

Table 154-8 EXECUTE_UPDATE Procedure Parameters

Parameter	Description
`update_stmt`	The SQL `UPDATE` statement

> **See Also:**
>
> *Oracle Database Administrator's Guide*

154.7.7 FINISH_REDEF_TABLE Procedure

This procedure completes the redefinition process.

Before this step, you can create new indexes, triggers, grants, and constraints on the interim table. The referential constraints involving the interim table must be disabled. After completing this step, the original table is redefined with the attributes and data of the interim table. The original table is locked briefly during this procedure.

Syntax

```
DBMS_REDEFINITION.FINISH_REDEF_TABLE (
   uname                   IN   VARCHAR2,
   orig_table              IN   VARCHAR2,
   int_table               IN   VARCHAR2,
   part_name               IN   VARCHAR2 := NULL,
   dml_lock_timeout        IN   PLS_INTEGER := NULL,
   continue_after_errors   IN   BOOLEAN := FALSE,
   disable_rollback        IN   PLS_INTEGER := FALSE);
```

Parameters

Table 154-9 FINISH_REDEF_TABLE Procedure Parameters

Parameters	Description
uname	Schema name of the tables
orig_table	Name of the table to be redefined
int_table	Name of the interim table. Can take a comma-delimited list of interim table names.
part_name	Name of the partition being redefined. If redefining only a single partition of a table, specify the partition name in this parameter. NULL implies the entire table is being redefined. Can take a comma-delimited list of partition names to be redefined.
dml_lock_timeout	Specifies the number of seconds the procedure waits for its required locks before failing. The permissible range of values for timeout is 0 to 1,000,000. The default is NULL (wait mode).
continue_after_errors	When redefining multiple partitions allows operation execution to continue on the next partition (applies only to batched partition redefinition).
disable_rollback	When set to TRUE, disables the rollback option if it was enabled in the START_REDEF_TABLE procedure. Specifying TRUE cleans up the database objects that enable rollback.

Examples

Wait up to 600 seconds for required locks on SH.SALES:

```
EXECUTE DBMS_REDEFINITION.FINISH_REDEF_TABLE (
  'SH', 'SALES', 'INT_SALES', 600);
```

154.7.8 REDEF_TABLE Procedure

This procedure provides a single interface that integrates several redefinition steps including the CAN_REDEF_TABLE Procedure, the START_REDEF_TABLE Procedure, the COPY_TABLE_DEPENDENTS Procedure and the FINISH_REDEF_TABLE Procedure.

This procedure can change data storage properties including tablespaces (for table, partition, subpartition, index, LOB column), compress type (for table, partition, subpartition, index, LOB column) and STORE_AS clause for the LOB column.

Syntax

```
DBMS_REDEFINITION.REDEF_TABLE (
   uname                      IN  VARCHAR2,
   tname                      IN  VARCHAR2,
   table_compression_type     IN  VARCHAR2 := NULL,
   table_part_tablespace      IN  VARCHAR2 := NULL,
   index_key_compression_type IN  VARCHAR2 := NULL,
   index_tablespace           IN  VARCHAR2 := NULL,
   lob_compression_type       IN  VARCHAR2 := NULL,
   lob_tablespace             IN  VARCHAR2 := NULL,
   lob_store_as               IN  VARCHAR2 := NULL,
   refresh_dep_mviews         IN  VARCHAR2 := 'N',
   dml_lock_timeout           IN  PLS_INTEGER := NULL);
```

Parameters

Table 154-10 REDEF_TABLE Procedure Parameters

Parameter	Description
uname	Schema name of the table
tname	Name of the table to be redefined
table_compression_type	Text string of the table compression clause. NULL means there is no change.
table_part_tablespace	Tablespace name for the entire table or partitions. NULL means there is no change.
index_key_compression_type	Text string of the compression clause for all indexes on the table. NULL means there is no change.
index_tablespace	Tablespace name for all indexes on the table. NULL means there is no change.
lob_compression_type	Text string of the compression clause for all LOBs in the entire table. NULL means there is no change.
lob_tablespace	Tablespace name for all LOBs in the table. NULL means there is no change.
lob_store_as	Specifies LOB store as 'SECUREFILE' or 'BASICFILE'. NULL means there is no change.
refresh_dep_mviews	When set to 'Y', fast refresh of dependent materialized views is performed once at the end of the redefinition operation.
dml_lock_timeout	Specifies the number of seconds the procedure waits for its required locks before failing. The permissible range of values for timeout is 0 to 1,000,000. The default is NULL (wait mode).

Examples

```
BEGIN
   DBMS_REDEFINITION.REDEF_TABLE(
     uname                        => 'TABOWNER2',
     tname                        => 'EMP2',
     table_compression_type       => 'ROW STORE COMPRESS ADVANCED',
     table_part_tablespace        => 'NEWTBS',
     index_key_compression_type   => 'COMPRESS 1',
     index_tablespace             => 'NEWIDXTBS',
     lob_compression_type         => 'COMPRESS HIGH',
     lob_tablespace               => 'SLOBTBS',
     lob_store_as                 => 'SECUREFILE');
END;
```

Related Topics

- CAN_REDEF_TABLE Procedure
  This procedure determines if a given table can be redefined online. This is the first step of the online redefinition process. If the table is not a candidate for online redefinition, an error message is raised.
- START_REDEF_TABLE Procedure
  This procedure starts a table redefinition.

- COPY_TABLE_DEPENDENTS Procedure
  This procedure clones the dependent objects of the table being redefined onto the interim table and registers the dependent objects. This procedure does not clone the already registered dependent objects.
- FINISH_REDEF_TABLE Procedure
  This procedure completes the redefinition process.

> **See Also:**
>
> *Oracle Database Administrator's Guide* regarding "Performing Online Redefinition with the `REDEF_TABLE` Procedure"

154.7.9 REGISTER_DEPENDENT_OBJECT Procedure

This procedure registers a dependent object (index, trigger, constraint or materialized view log) on the table being redefined and the corresponding dependent object on the interim table.

This can be used to have the same object on each table but with different attributes. For example: for an index, the storage and tablespace attributes could be different but the columns indexed remain the same

Syntax

```
DBMS_REDEFINITION.REGISTER_DEPENDENT_OBJECT(
   uname                  IN    VARCHAR2,
   orig_table             IN    VARCHAR2,
   int_table              IN    VARCHAR2,
   dep_type               IN    PLS_INTEGER,
   dep_owner              IN    VARCHAR2,
   dep_orig_name          IN    VARCHAR2,
   dep_int_name           IN    VARCHAR2);
```

Parameters

Table 154-11 REGISTER_DEPENDENT_OBJECT Procedure Parameters

Parameters	Description
`uname`	Schema name of the tables
`orig_table`	Name of the table to be redefined
`int_table`	Name of the interim table
`dep_type`	Type of the dependent object (see Constants and Operational Notes)
`dep_owner`	Owner of the dependent object
`dep_orig_name`	Name of the original dependent object
`dep_int_name`	Name of the interim dependent object

Usage Notes

- Attempting to register an already registered object will raise an error.
- Registering a dependent object will automatically remove that object from `DBA_REDEFINITION_ERRORS` if an entry exists for that object.

154.7.10 ROLLBACK Procedure

This procedure rolls back changes to a table after online table redefinition to return the table to its original definition and preserve DML changes made to the table.

Syntax

```
DBMS_REDEFINITION.ROLLBACK (
   uname                  IN  VARCHAR2,
   orig_table             IN  VARCHAR2,
   int_table              IN  VARCHAR2 := NULL,
   part_name              IN  VARCHAR2 := NULL,
   dml_lock_timeout       IN  PLS_INTEGER := NULL,
   continue_after_errors  IN  BOOLEAN := FALSE);
```

Parameters

Table 154-12 ROLLBACK Procedure Parameters

Parameter	Description
uname	Schema name of the table to be redefined
orig_table	Name of the table to be redefined
int_table	Name of the interim table.
part_name	Name of the partition being redefined.
dml_lock_timeout	Specifies the number of seconds the procedure waits for its required locks before failing. The permissible range of values for timeout is 0 to 1,000,000. The default is NULL (wait mode).
continue_after_errors	When rolling back redefinition changes on multiple partitions, allows operation execution to continue on the next partition (applies only to batched partition redefinition).

154.7.11 SET_PARAM Procedure

This procedure sets a new value for a specified parameter used by the redefinition process identified by a redefinition ID.

> **Note:**
>
> Currently, the only value that can be changed by this procedure is the value for the of the refresh_dep_mviews parameter that is specified in the REDEF_TABLE procedure or the START_REDEF_TABLE procedure. You can determine the redefinition ID and check the value of the refresh_dep_mviews parameter for an online table redefinition operation by querying the DBA_REDEFINITION_STATUS view.

Syntax

```
DBMS_REDEFINITION.SET_PARAM (
   redefinition_id  IN  VARCHAR2,
   param_name       IN  VARCHAR2,
   param_value      IN  VARCHAR2);
```

Parameters

Table 154-13 SET_PARAM Procedure Parameters

Parameter	Description
redefinition_id	The redefinition ID that identifies the redefinition process
param_name	The parameter name
param_value	The new parameter value

> **See Also:**
>
> *Oracle Database Administrator's Guide*

154.7.12 START_REDEF_TABLE Procedure

This procedure starts a table redefinition.

Prior to calling this procedure, you must manually create an empty interim table (in the same schema as the table to be redefined) with the desired attributes of the post-redefinition table, and then call this procedure to initiate the redefinition.

Syntax

```
DBMS_REDEFINITION.START_REDEF_TABLE (
   uname                   IN  VARCHAR2,
   orig_table              IN  VARCHAR2,
   int_table               IN  VARCHAR2,
   col_mapping             IN  VARCHAR2 := NULL,
   options_flag            IN  BINARY_INTEGER := 1,
   orderby_cols            IN  VARCHAR2 := NULL,
   part_name               IN  VARCHAR2 := NULL,
   continue_after_errors   IN  BOOLEAN := FALSE,
   copy_vpd_opt            IN  BINARY_INTEGER := CONS_VPD_NONE,
   refresh_dep_mviews      IN  VARCHAR2 := 'N',
   enable_rollback         IN  BOOLEAN := FALSE);
```

Parameters

Table 154-14 START_REDEF_TABLE Procedure Parameters

Parameter	Description
uname	Schema name of the tables
orig_table	Name of the table to be redefined
int_table	Name of the interim table. Can take a comma-delimited list of interim table names.
col_mapping	Mapping information from the columns in the original table to the columns in the interim table. (This is similar to the column list on the SELECT clause of a query.) If NULL, all the columns in the original table are selected and have the same name after redefinition.

Table 154-14 (Cont.) START_REDEF_TABLE Procedure Parameters

Parameter	Description
options_flag	Indicates the type of redefinition method to use: • If dbms_redefinition.cons_use_pk, the redefinition is done using primary keys or pseudo-primary keys (unique keys with all component columns having NOT NULL constraints). The default method of redefinition is using primary keys. • If dbms_redefinition.cons_use_rowid, the redefinition is done using rowids.
orderby_cols	This optional parameter accepts the list of columns (along with the optional keyword(s) ascending/descending) with which to order by the rows during the initial instantiation of the interim table (the order by is only done for the initial instantiation and not for subsequent synchronizations)
part_name	Name of the partition being redefined. If redefining only a single partition of a table, specify the partition name in this parameter. NULL implies the entire table is being redefined. Can take a comma-delimited list of partition names to be redefined.
continue_after_errors	When redefining multiple partitions allows operation execution to continue on the next partition (applies only to batched partition redefinition)
copy_vpd_opt	Specifies how VPD policies are handled in online redefinition
refresh_dep_mviews	When set to 'Y', fast refresh of dependent materialized views is performed when the START_REDEF_TABLE procedure is run, each time the SYNC_INTERIM_TABLE procedure is run, and when the FINISH_REDEF_TABLE procedure is run.
enable_rollback	When set to TRUE, enables the rollback option. When this parameter is set to true, Oracle Database maintains the interim table created during redefinition after redefinition is complete. You can run the SYNC_INTERIM_TABLE procedure to synchronize the interim table periodically to apply DML changes made to the redefined table to the interim table. An internal materialized view and materialized view log enables maintenance of the interim table. If you decide to roll back the online table redefinition with the ROLLBACK procedure, then the interim table is synchronized, and Oracle Database switches back to it so that the table has its original definition.

Examples

Start redefinition of three partitions (sal03q1,sal03q2,sal03q3) in table 'STEVE.salestable' using three interim tables of int_salestable1, int_salestable2 and int_salestable3, respectively. The operation will continue on sal03q3 even if it fails on sal03q1.

```
DBMS_REDEFINITION.START_REDEF_TABLE(
   uname                 => 'STEVE',
   orig_table            => 'salestable',
   int_table             => 'int_salestable1, int_salestable2, int_salestable3',
   col_mapping           => NULL,
   options_flag          => DBMS_REDEFINITION.CONS_USE_ROWID,
   part_name             => 'sal03q1,sal03q2,sal03q3',
   continue_after_errors => TRUE);
```

Specify to copy VPD policies automatically:

```
EXECUTE DBMS_REDEFINITION.START_REDEF_TABLE (
  uname                 => 'SCOTT',
  orig_table            => 'T',
  int_table             => 'INT_T',
  copy_vpd_opt          => DBMS_REDEFINITION.CONS_VPD_AUTO);
```

154.7.13 SYNC_INTERIM_TABLE Procedure

This procedure keeps the interim table synchronized with the original table.

Syntax

```
DBMS_REDEFINITION.SYNC_INTERIM_TABLE (
   uname                  IN  VARCHAR2,
   orig_table             IN  VARCHAR2,
   int_table              IN  VARCHAR2,
   part_name              IN  VARCHAR2 := NULL,
   continue_after_errors  IN  BOOLEAN := FALSE);
```

Parameters

Table 154-15 SYNC_INTERIM_TABLE Procedure Parameters

Parameter	Description
uname	Schema name of the table
orig_table	Name of the table to be redefined
int_table	Name of the interim table. Can take a comma-delimited list of interim table names.
part_name	Name of the partition being redefined. If redefining only a single partition of a table, specify the partition name in this parameter. NULL implies the entire table is being redefined. Can take a comma-delimited list of partition names to be redefined.
continue_after_errors	When redefining multiple partitions allows operation execution to continue on the next partition (applies only to batched partition redefinition)

Usage Notes

- This step is useful in minimizing the amount of synchronization needed to be done by the FINISH_REDEF_TABLE Procedure before completing the online redefinition.
- This procedure can be called between long running operations (such as CREATE INDEX) on the interim table to sync it up with the data in the original table and speed up subsequent operations.

154.7.14 UNREGISTER_DEPENDENT_OBJECT Procedure

This procedure unregisters a dependent object (index, trigger, constraint or materialized view log) on the table being redefined and the corresponding dependent object on the interim table.

Syntax

```
DBMS_REDEFINITION.UNREGISTER_DEPENDENT_OBJECT(
   uname             IN VARCHAR2,
   orig_table        IN VARCHAR2,
   int_table         IN VARCHAR2,
```

```
   dep_type          IN PLS_INTEGER,
   dep_owner         IN VARCHAR2,
   dep_orig_name     IN VARCHAR2,
   dep_int_name      IN  VARCHAR2);
```

Parameters

Table 154-16 UNREGISTER_DEPENDENT_OBJECT Procedure Parameters

Parameters	Description
uname	Schema name of the tables
orig_table	Name of the table to be redefined
int_table	Name of the interim table
dep_type	Type of the dependent object
dep_owner	Owner of the dependent object
dep_orig_name	Name of the original dependent object
dep_int_name	Name of the interim dependent object

155
DBMS_REFRESH

The `DBMS_REFRESH` package enables you to create groups of materialized views that can be refreshed together to a transactionally consistent point in time. These groups are called refresh groups.

This chapter contains the following topics:

- DBMS_REFRESH Overview
- DBMS_REFRESH Security Model

155.1 DBMS_REFRESH Overview

When it is important for materialized views to be transactionally consistent with each other, you can organize them into refresh groups.

By refreshing the refresh group, you can ensure that the data in all of the materialized views in the refresh group correspond to the same transactionally consistent point in time. A materialized view in a refresh group still can be refreshed individually, but doing so nullifies the benefits of the refresh group because refreshing the materialized view individually does not refresh the other materialized views in the refresh group

155.2 DBMS_REFRESH Security Model

Users must have the `EXECUTE` privilege to run the procedures of `DBMS_REFRESH` package.

155.3 Summary of DBMS_REFRESH Subprograms

This table lists the `DBMS_REFRESH` subprograms and briefly describes them.

Table 155-1 DBMS_REFRESH Package Subprograms

Subprogram	Description
ADD Procedure	Adds materialized views to a refresh group
CHANGE Procedure	Changes the refresh interval for a refresh group
DESTROY Procedure	Removes all of the materialized views from a refresh group and deletes the refresh group
MAKE Procedure	Specifies the members of a refresh group and the time interval used to determine when to refresh the members of this group
REFRESH Procedure	Manually refreshes a refresh group
SUBTRACT Procedure	Removes materialized views from a refresh group

155.3.1 ADD Procedure

This procedure adds materialized views to a refresh group.

Syntax

```
DBMS_REFRESH.ADD (
   name     IN VARCHAR2,
   { list   IN VARCHAR2,
   | tab    IN DBMS_UTILITY.UNCL_ARRAY, }
   lax      IN BOOLEAN := FALSE);
```

> **Note:**
>
> This procedure is overloaded. The `list` and `tab` parameters are mutually exclusive.

Parameters

Table 155-2 ADD Procedure Parameters

Parameter	Description
`name`	Name of the refresh group to which you want to add members, specified as `[`*schema_name*`.]`*refresh_group_name*. If the schema is not specified, then the current user is the default.
`list`	Comma-delimited list of materialized views that you want to add to the refresh group. Synonyms are not supported. Each materialized view is specified as `[`*schema_name*`.]`*materialized_view_name*. If the schema is not specified, then the refresh group owner is the default.
`tab`	Instead of a comma-delimited list, you can supply a PL/SQL associative array of type `DBMS_UTILITY.UNCL_ARRAY`, where each element is the name of a materialized view. The first materialized view should be in position 1. The last position must be `NULL`. Each materialized view is specified as `[`*schema_name*`.]`*materialized_view_name*. If the schema is not specified, then the refresh group owner is the default.
`lax`	A materialized view can belong to only one refresh group at a time. If you are moving a materialized view from one group to another, then you must set the `lax` flag to `TRUE` to succeed. Oracle then automatically removes the materialized view from the other refresh group and updates its refresh interval to be that of its new group. Otherwise, the call to `ADD` generates an error message.

155.3.2 CHANGE Procedure

This procedure changes the refresh interval for a refresh group.

Syntax

```
DBMS_REFRESH.CHANGE (
   name                 IN VARCHAR2,
   next_date            IN DATE           := NULL,
```

```
   interval              IN VARCHAR2       := NULL,
   implicit_destroy      IN BOOLEAN        := NULL,
   rollback_seg          IN VARCHAR2       := NULL,
   push_deferred_rpc     IN BOOLEAN        := NULL,
   refresh_after_errors  IN BOOLEAN        := NULL,
   purge_option          IN BINARY_INTEGER := NULL,
   parallelism           IN BINARY_INTEGER := NULL,
   heap_size             IN BINARY_INTEGER := NULL);
```

Parameters

Table 155-3 CHANGE Procedure Parameters

Parameter	Description
`name`	Name of the refresh group for which you want to alter the refresh interval.
`next_date`	Next date that you want a refresh to occur. By default, this date remains unchanged.
`interval`	Function used to calculate the next time to refresh the materialized views in the refresh group. This interval is evaluated immediately before the refresh. Thus, select an interval that is greater than the time it takes to perform a refresh. By default, the interval remains unchanged.
`implicit_destroy`	Allows you to reset the value of the `implicit_destroy` flag. If this flag is set, then Oracle automatically deletes the group if it no longer contains any members. By default, this flag remains unchanged.
`rollback_seg`	Allows you to change the rollback segment used. By default, the rollback segment remains unchanged. To reset this parameter to use the default rollback segment, specify `NULL`, including the quotes. Specifying `NULL` without quotes indicates that you do not want to change the rollback segment currently being used.
`push_deferred_rpc`	Starting with Oracle Database 12*c* Release 2 (12.2), this parameter is ignored.
`refresh_after_errors`	Starting with Oracle Database 12*c* Release 2 (12.2), this parameter is ignored.
`purge_option`	Starting with Oracle Database 12*c* Release 2 (12.2), this parameter is ignored.
`parallelism`	• 0 specifies serial propagation. • *n* > 1 specifies parallel propagation with n parallel processes. • 1 specifies parallel propagation using only one parallel process.
`heap_size`	Maximum number of transactions to be examined simultaneously for parallel propagation scheduling. Oracle automatically calculates the default setting for optimal performance.

> **Note:**
>
> Do not set this parameter unless directed to do so by Oracle Support Services.

155.3.3 DESTROY Procedure

This procedure removes all of the materialized views from a refresh group and delete the refresh group.

Syntax

```
DBMS_REFRESH.DESTROY (
   name    IN    VARCHAR2);
```

Parameters

Table 155-4 DESTROY Procedure Parameters

Parameter	Description
name	Name of the refresh group that you want to destroy.

155.3.4 MAKE Procedure

This procedure specifies the members of a refresh group and the time interval used to determine when to refresh the members of this group.

Syntax

```
DBMS_REFRESH.MAKE (
   name                  IN    VARCHAR2
   { list                IN    VARCHAR2,
   | tab                 IN    DBMS_UTILITY.UNCL_ARRAY,}
   next_date             IN    DATE,
   interval              IN    VARCHAR2,
   implicit_destroy      IN    BOOLEAN         := FALSE,
   lax                   IN    BOOLEAN         := FALSE,
   job                   IN    BINARY_INTEGER  := 0,
   rollback_seg          IN    VARCHAR2        := NULL,
   push_deferred_rpc     IN    BOOLEAN         := TRUE,
   refresh_after_errors  IN    BOOLEAN         := FALSE
   purge_option          IN    BINARY_INTEGER  := NULL,
   parallelism           IN    BINARY_INTEGER  := NULL,
   heap_size             IN    BINARY_INTEGER  := NULL
   job_name              IN    VARCHAR2        := NULL,
   auto_commit           IN    BOOLEAN         := NULL);
```

> **Note:**
>
> This procedure is overloaded. The `list` and `tab` parameters are mutually exclusive.

Parameters

Table 155-5 MAKE Procedure Parameters

Parameter	Description
name	Unique name used to identify the refresh group, specified as [*schema_name.*]*refresh_group_name*. If the schema is not specified, then the current user is the default. Refresh groups must follow the same naming conventions as tables.
list	Comma-delimited list of materialized views that you want to refresh. Synonyms are not supported. These materialized views can be located in different schemas and have different master tables or master materialized views. However, all of the listed materialized views must be in your current database. Each materialized view is specified as [*schema_name.*]*materialized_view_name*. If the schema is not specified, then the refresh group owner is the default.
tab	Instead of a comma-delimited list, you can supply a PL/SQL associative array of names of materialized views that you want to refresh using the data type DBMS_UTILITY.UNCL_ARRAY. If the table contains the names of *n* materialized views, then the first materialized view should be in position 1 and the *n* + 1 position should be set to NULL. Each materialized view is specified as [*schema_name.*]*materialized_view_name*. If the schema is not specified, then the refresh group owner is the default.
next_date	Next date that you want a refresh to occur.
interval	Function used to calculate the next time to refresh the materialized views in the group. This field is used with the next_date value. For example, if you specify NEXT_DAY(SYSDATE+1, "MONDAY") as your interval, and if your next_date evaluates to Monday, then Oracle refreshes the materialized views every Monday. This interval is evaluated immediately before the refresh. Thus, select an interval that is greater than the time it takes to perform a refresh.
implicit_destroy	Set this to TRUE to delete the refresh group automatically when it no longer contains any members. Oracle checks this flag only when you call the SUBTRACT procedure. That is, setting this flag still enables you to create an empty refresh group.
lax	A materialized view can belong to only one refresh group at a time. If you are moving a materialized view from an existing group to a new refresh group, then you must set this to TRUE to succeed. Oracle then automatically removes the materialized view from the other refresh group and updates its refresh interval to be that of its new group. Otherwise, the call to MAKE generates an error message.
job	Needed by the Import utility. Use the default value, 0.
rollback_seg	Name of the rollback segment to use while refreshing materialized views. The default, NULL, uses the default rollback segment.
push_deferred_rpc	Starting with Oracle Database 12*c* Release 2 (12.2), this parameter is ignored.
refresh_after_errors	Starting with Oracle Database 12*c* Release 2 (12.2), this parameter is ignored.
purge_option	Starting with Oracle Database 12*c* Release 2 (12.2), this parameter is ignored.

Table 155-5 (Cont.) MAKE Procedure Parameters

Parameter	Description
parallelism	• 0 specifies serial propagation. • *n* > 1 specifies parallel propagation with n parallel processes. • 1 specifies parallel propagation using only one parallel process.
heap_size	Maximum number of transactions to be examined simultaneously for parallel propagation scheduling. Oracle automatically calculates the default setting for optimal performance. **Note:** Do not set this parameter unless directed to do so by Oracle Support Services.
job_name	This parameter is needed by the import utility. User should use the default value, NULL.
auto_commit	Supported values are NULL, TRUE, and FALSE. • NULL—allows the user to continue using DBMS_JOB. • TRUE—commit statement will be automatically issued after the job of the refresh group are created by DBMS_REFRESH.MAKE. • FALSE—user must issue a commit statement to finish the transaction after calling DBMS_REFRESH.MAKE. The default value is NULL.

Usage Notes

Import utility and export utility need CREATE JOB privilege if DBMS_SCHEDULER jobs are used.

155.3.5 REFRESH Procedure

This procedure manually refreshes a refresh group.

Syntax

```
DBMS_REFRESH.REFRESH (
   name    IN    VARCHAR2);
```

Parameters

Table 155-6 REFRESH Procedure Parameters

Parameter	Description
name	Name of the refresh group that you want to refresh manually.

155.3.6 SUBTRACT Procedure

This procedure removes materialized views from a refresh group.

Syntax

```
DBMS_REFRESH.SUBTRACT (
   name      IN    VARCHAR2,
   { list    IN    VARCHAR2,
   | tab     IN    DBMS_UTILITY.UNCL_ARRAY, }
   lax       IN    BOOLEAN := FALSE);
```

> **Note:**
>
> This procedure is overloaded. The `list` and `tab` parameters are mutually exclusive.

Parameters

Table 155-7 SUBTRACT Procedure Parameters

Parameter	Description
`name`	Name of the refresh group from which you want to remove members, specified as `[`*`schema_name`*`.]`*`refresh_group_name`*. If the schema is not specified, then the current user is the default.
`list`	Comma-delimited list of materialized views that you want to remove from the refresh group. (Synonyms are not supported.) These materialized views can be located in different schemas and have different master tables or master materialized views. However, all of the listed materialized views must be in your current database. Each materialized view is specified as `[`*`schema_name`*`.]`*`materialized_view_name`*. If the schema is not specified, then the refresh group owner is the default.
`tab`	Instead of a comma-delimited list, you can supply a PL/SQL associative array of type `DBMS_UTILITY.UNCL_ARRAY`, where each element is the name of a materialized view. The first materialized view should be in position 1. The last position must be `NULL`. Each materialized view is specified as `[`*`schema_name`*`.]`*`materialized_view_name`*. If the schema is not specified, then the refresh group owner is the default.
`lax`	Set this to `FALSE` if you want Oracle to generate an error message if the materialized view you are attempting to remove is not a member of the refresh group.

156
DBMS_REPAIR

The `DBMS_REPAIR` package contains data corruption repair procedures that enable you to detect and repair corrupt blocks in tables and indexes. You can address corruptions where possible and continue to use objects while you attempt to rebuild or repair them.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operating Notes
- Exceptions
- Examples
- Summary of DBMS_REPAIR Subprograms

> **See Also:**
>
> For detailed information about using the `DBMS_REPAIR` package, see *Oracle Database Administrator's Guide*

156.1 DBMS_REPAIR Overview

The `DBMS_REPAIR` package is intended for use by database administrators only. It is not intended for use by application developers.

156.2 DBMS_REPAIR Security Model

The package is owned by `SYS`. Execution privilege is not granted to other users.

156.3 DBMS_REPAIR Constants

The `DBMS_REPAIR` package defines several enumerated constants that should be used for specifying parameter values. Enumerated constants must be prefixed with the package name. For example, `DBMS_REPAIR.TABLE_OBJECT`.

The following table lists the parameters and the enumerated constants.

Table 156-1 DBMS_REPAIR Parameters with Enumerated Constants

Parameter	Option	Type	Description
object_type	• TABLE_OBJECT • INDEX_OBJECT • CLUSTER_OBJECT	BINARY_INTEGER	-
action	• CREATE_ACTION • DROP_ACTION • PURGE_ACTION	BINARY_INTEGER	-
table_type	• REPAIR_TABLE • ORPHAN_TABLE	BINARY_INTEGER	-
flags	• SKIP_FLAG • NOSKIP_FLAG	BINARY_INTEGER	-
object_id	• ALL_INDEX_ID := 0	BINARY_INTEGER	Clean up all objects that qualify
wait_for_lock	• LOCK_WAIT := 1 • LOCK_NOWAIT := 0	BINARY_INTEGER	Specifies whether to try getting DML locks on underlying table [[sub]partition] object

> **Note:**
>
> The default table_name will be REPAIR_TABLE when table_type is REPAIR_TABLE, and will be ORPHAN_KEY_TABLE when table_type is ORPHAN_TABLE.

156.4 DBMS_REPAIR Operating Notes

The procedure to create the ORPHAN_KEYS_TABLE is similar to the one used to create the REPAIR_TABLE.

```
CONNECT / AS SYSDBA;
EXEC DBMS_REPAIR.ADMIN_TABLES('ORPHAN_KEYS_TABLE', DBMS_REPAIR.ORPHAN_TABLE,
                              DBMS_REPAIR.CREATE_ACTION);
EXEC DBMS_REPAIR.ADMIN_TABLES('REPAIR_TABLE', DBMS_REPAIR.REPAIR_TABLE,
                              DBMS_REPAIR.CREATE_ACTION);
DESCRIBE ORPHAN_KEYS_TABLE;
DESCRIBE REPAIR_TABLE;
SELECT * FROM ORPHAN_KEYS_TABLE;
SELECT * FROM REPAIR_TABLE;
```

The DBA would create the repair and orphan keys tables once. Subsequent executions of the CHECK_OBJECT Procedure would add rows into the appropriate table indicating the types of errors found.

The name of the repair and orphan keys tables can be chosen by the user, with the following restriction: the name of the repair table must begin with the 'REPAIR_' prefix, and the name of the orphan keys table must begin with the 'ORPHAN_' prefix. The following code is also legal:

```
CONNECT / AS SYSDBA;
EXEC DBMS_REPAIR.ADMIN_TABLES('ORPHAN_FOOBAR', DBMS_REPAIR.ORPHAN_TABLE,
                              DBMS_REPAIR.CREATE_ACTION);
EXEC DBMS_REPAIR.ADMIN_TABLES('REPAIR_ABCD', DBMS_REPAIR.REPAIR_TABLE,
                              DBMS_REPAIR.CREATE_ACTION);
DESCRIBE ORPHAN_FOOBAR;
```

```
DESCRIBE REPAIR_ABCD;
SELECT * FROM ORPHAN_FOOBAR;
SELECT * FROM REPAIR_ABCD;
```

When invoking the CHECK_OBJECT Procedure the name of the repair and orphan keys tables that were created should be specified correctly, especially if the default values were not used in the ADMIN_TABLES Procedure or CREATE_ACTION.

Other actions in the ADMIN_TABLES Procedure can be used to purge/delete the REPAIR_TABLE and the ORPHAN_KEYS_TABLE.

156.5 DBMS_REPAIR Exceptions

The table in this topic describes the exceptions raised by the DDBMS_REPAIR subprograms.

Table 156-2 DBMS_REPAIR Exceptions

Exception	Description	Action
942	Reported by DBMS_REPAIR.ADMIN_TABLES during a DROP_ACTION when the specified table doesn't exist.	-
955	Reported by DBMS_REPAIR. CREATE_ACTION when the specified table already exists.	-
24120	An invalid parameter was passed to the specified DBMS_REPAIR procedure.	Specify a valid parameter value or use the parameter's default.
24122	An incorrect block range was specified.	Specify correct values for the BLOCK_START and BLOCK_END parameters.
24123	An attempt was made to use the specified feature, but the feature is not yet implemented.	Do not attempt to use the feature.
24124	An invalid ACTION parameter was specified.	Specify CREATE_ACTION, PURGE_ACTION or DROP_ACTION for the ACTION parameter.
24125	An attempt was made to fix corrupt blocks on an object that has been dropped or truncated since DBMS_REPAIR.CHECK_OBJECT was run.	Use DBMS_REPAIR.ADMIN_TABLES to purge the repair table and run DBMS_REPAIR.CHECK_OBJECT to determine whether there are any corrupt blocks to be fixed.
24127	TABLESPACE parameter specified with an ACTION other than CREATE_ACTION.	Do not specify TABLESPACE when performing actions other than CREATE_ACTION.
24128	A partition name was specified for an object that is not partitioned.	Specify a partition name only if the object is partitioned.
24129	An attempt was made to pass a table name parameter without the specified prefix.	Pass a valid table name parameter.
24130	An attempt was made to specify a repair or orphan table that does not exist.	Specify a valid table name parameter.
24131	An attempt was made to specify a repair or orphan table that does not have a correct definition.	Specify a table name that refers to a properly created table.

Table 156-2 (Cont.) DBMS_REPAIR Exceptions

Exception	Description	Action
24132	An attempt was made to specify a table name is greater than 30 characters long.	Specify a valid table name parameter.

156.6 DBMS_REPAIR Examples

This topic shows examples of DBMS_REPAIR usage.

```
/* Fix the bitmap status for all the blocks in table mytab in schema sys */

EXECUTE DBMS_REPAIR.SEGMENT_FIX_STATUS('SYS', 'MYTAB');

/* Mark block number 45, filenumber 1 for table mytab in sys schema as FULL.*/

EXECUTE DBMS_REPAIR.SEGMENT_FIX_STATUS('SYS', 'MYTAB', TABLE_OBJECT,1, 45, 1);
```

156.7 Summary of DBMS_REPAIR Subprograms

This table lists the DBMS_REPAIR subprograms and briefly describes them.

Table 156-3 DBMS_REPAIR Package Subprograms

Subprogram	Description
ADMIN_TABLES Procedure	Provides administrative functions for the DBMS_REPAIR package repair and orphan key tables, including create, purge, and drop functions
CHECK_OBJECT Procedure	Detects and reports corruptions in a table or index
DUMP_ORPHAN_KEYS Procedure	Reports on index entries that point to rows in corrupt data blocks
FIX_CORRUPT_BLOCKS Procedure	Marks blocks software corrupt that have been previously detected as corrupt by CHECK_OBJECT
ONLINE_INDEX_CLEAN Function	Performs a manual cleanup of failed or interrupted online index builds or rebuilds
REBUILD_FREELISTS Procedure	Rebuilds an object's freelists
SEGMENT_FIX_STATUS Procedure	Fixes the corrupted state of a bitmap entry
SKIP_CORRUPT_BLOCKS Procedure	Sets whether to ignore blocks marked corrupt during table and index scans or to report ORA-1578 when blocks marked corrupt are encountered

156.7.1 ADMIN_TABLES Procedure

This procedure provides administrative functions for the DBMS_REPAIR package repair and orphan key tables.

Syntax

```
DBMS_REPAIR.ADMIN_TABLES (
   table_name  IN   VARCHAR2,
   table_type  IN   BINARY_INTEGER,
```

```
   action      IN   BINARY_INTEGER,
   tablespace  IN   VARCHAR2  DEFAULT NULL);
```

Parameters

Table 156-4 ADMIN_TABLES Procedure Parameters

Parameter	Description
table_name	Name of the table to be processed. Defaults to ORPHAN_KEY_TABLE or REPAIR_TABLE based on the specified table_type. When specified, the table name must have the appropriate prefix: ORPHAN_ or REPAIR_.
table_type	Type of table; must be either ORPHAN_TABLE or REPAIR_TABLE. See "Constants".
action	Indicates what administrative action to perform. Must be either CREATE_ACTION, PURGE_ACTION, or DROP_ACTION. If the table already exists, and if CREATE_ACTION is specified, then an error is returned. PURGE_ACTION indicates to delete all rows in the table that are associated with non-existent objects. If the table does not exist, and if DROP_ACTION is specified, then an error is returned. When CREATE_ACTION and DROP_ACTION are specified, an associated view named DBA_<table_name> is created and dropped respectively. The view is defined so that rows associated with non-existent objects are eliminated. Created in the SYS schema. See "Constants".
tablespace	Indicates the tablespace to use when creating a table. By default, the SYS default tablespace is used. An error is returned if the tablespace is specified and if the action is not CREATE_ACTION.

156.7.2 CHECK_OBJECT Procedure

This procedure checks the specified objects and populates the repair table with information about corruptions and repair directives.

Validation consists of block checking all blocks in the object.

Syntax

```
DBMS_REPAIR.CHECK_OBJECT (
   schema_name       IN  VARCHAR2,
   object_name       IN  VARCHAR2,
   partition_name    IN  VARCHAR2       DEFAULT NULL,
   object_type       IN  BINARY_INTEGER DEFAULT TABLE_OBJECT,
   repair_table_name IN  VARCHAR2       DEFAULT 'REPAIR_TABLE',
   flags             IN  BINARY_INTEGER DEFAULT NULL,
   relative_fno      IN  BINARY_INTEGER DEFAULT NULL,
   block_start       IN  BINARY_INTEGER DEFAULT NULL,
   block_end         IN  BINARY_INTEGER DEFAULT NULL,
   corrupt_count     OUT BINARY_INTEGER);
```

Parameters

Table 156-5 CHECK_OBJECT Procedure Parameters

Parameter	Description
schema_name	Schema name of the object to be checked.
object_name	Name of the table or index to be checked.
partition_name	Partition or subpartition name to be checked. If this is a partitioned object, and if partition_name is not specified, then all partitions and subpartitions are checked. If this is a partitioned object, and if the specified partition contains subpartitions, then all subpartitions are checked.
object_type	Type of the object to be processed. This must be either TABLE_OBJECT (default) or INDEX_OBJECT. See "Constants".
repair_table_name	Name of the repair table to be populated. The table must exist in the SYS schema. Use the ADMIN_TABLES Procedure to create a repair table. The default name is REPAIR_TABLE.
flags	Reserved for future use.
relative_fno	Relative file number: Used when specifying a block range.
block_start	First block to process if specifying a block range. May be specified only if the object is a single table, partition, or subpartition.
block_end	Last block to process if specifying a block range. May be specified only if the object is a single table, partition, or subpartition. If only one of block_start or block_end is specified, then the other defaults to the first or last block in the file respectively.
corrupt_count	Number of corruptions reported.

Usage Notes

You may optionally specify a DBA range, partition name, or subpartition name when you want to check a portion of an object.

156.7.3 DUMP_ORPHAN_KEYS Procedure

This procedure reports on index entries that point to rows in corrupt data blocks. For each such index entry encountered, a row is inserted into the specified orphan table.

If the repair table is specified, then any corrupt blocks associated with the base table are handled in addition to all data blocks that are marked software corrupt. Otherwise, only blocks that are marked corrupt are handled.

This information may be useful for rebuilding lost rows in the table and for diagnostic purposes.

Syntax

```
DBMS_REPAIR.DUMP_ORPHAN_KEYS (
   schema_name       IN  VARCHAR2,
   object_name       IN  VARCHAR2,
   partition_name    IN  VARCHAR2       DEFAULT NULL,
   object_type       IN  BINARY_INTEGER DEFAULT INDEX_OBJECT,
   repair_table_name IN  VARCHAR2       DEFAULT 'REPAIR_TABLE',
```

```
   orphan_table_name IN  VARCHAR2       DEFAULT 'ORPHAN_KEYS_TABLE',
   flags             IN  BINARY_INTEGER DEFAULT NULL,
   key_count         OUT BINARY_INTEGER);
```

Parameters

Table 156-6 DUMP_ORPHAN_KEYS Procedure Parameters

Parameter	Description
schema_name	Schema name.
object_name	Object name.
partition_name	Partition or subpartition name to be processed. If this is a partitioned object, and if partition_name is not specified, then all partitions and subpartitions are processed. If this is a partitioned object, and if the specified partition contains subpartitions, then all subpartitions are processed.
object_type	Type of the object to be processed. The default is INDEX_OBJECT See "Constants".
repair_table_name	Name of the repair table that has information regarding corrupt blocks in the base table. The specified table must exist in the SYS schema. The ADMIN_TABLES Procedure is used to create the table.
orphan_table_name	Name of the orphan key table to populate with information regarding each index entry that refers to a row in a corrupt data block. The specified table must exist in the SYS schema. The ADMIN_TABLES Procedure is used to create the table.
flags	Reserved for future use.
key_count	Number of index entries processed.

156.7.4 FIX_CORRUPT_BLOCKS Procedure

This procedure fixes the corrupt blocks in specified objects based on information in the repair table that was previously generated by the CHECK_OBJECT Procedure.

Prior to effecting any change to a block, the block is checked to ensure the block is still corrupt. Corrupt blocks are repaired by marking the block software corrupt. When a repair is effected, the associated row in the repair table is updated with a fix timestamp.

Syntax

```
DBMS_REPAIR.FIX_CORRUPT_BLOCKS (
   schema_name       IN  VARCHAR2,
   object_name       IN  VARCHAR2,
   partition_name    IN  VARCHAR2       DEFAULT NULL,
   object_type       IN  BINARY_INTEGER DEFAULT TABLE_OBJECT,
   repair_table_name IN  VARCHAR2       DEFAULT 'REPAIR_TABLE',
   flags             IN  BINARY_INTEGER DEFAULT NULL,
   fix_count         OUT BINARY_INTEGER);
```

Parameters

Table 156-7 FIX_CORRUPT_BLOCKS Procedure Parameters

Parameter	Description
schema_name	Schema name.
object_name	Name of the object with corrupt blocks to be fixed.
partition_name	Partition or subpartition name to be processed. If this is a partitioned object, and if partition_name is not specified, then all partitions and subpartitions are processed. If this is a partitioned object, and if the specified partition contains subpartitions, then all subpartitions are processed.
object_type	Type of the object to be processed. This must be either TABLE_OBJECT (default) or INDEX_OBJECT. See "Constants".
repair_table_name	Name of the repair table with the repair directives. Must exist in the SYS schema.
flags	Reserved for future use.
fix_count	Number of blocks fixed.

Related Topics

- CHECK_OBJECT Procedure
  This procedure checks the specified objects and populates the repair table with information about corruptions and repair directives.

156.7.5 ONLINE_INDEX_CLEAN Function

This function performs a manual cleanup of failed or interrupted online index builds or rebuilds.

This action is also performed periodically by SMON, regardless of user-initiated cleanup.

This function returns TRUE if all indexes specified were cleaned up and FALSE if one or more indexes could not be cleaned up.

Syntax

```
DBMS_REPAIR.ONLINE_INDEX_CLEAN (
   object_id      IN BINARY_INTEGER DEFAULT ALL_INDEX_ID,
   wait_for_lock  IN BINARY_INTEGER DEFAULT LOCK_WAIT)
 RETURN BOOLEAN;
```

Parameters

Table 156-8 ONLINE_INDEX_CLEAN Function Parameters

Parameter	Description
object_id	Object id of index to be cleaned up. The default cleans up all object ids that qualify.

Table 156-8 (Cont.) ONLINE_INDEX_CLEAN Function Parameters

Parameter	Description
wait_for_lock	This parameter specifies whether to try getting DML locks on underlying table [[sub]partition] object. The default retries up to an internal retry limit, after which the lock get will give up. If LOCK_NOWAIT is specified, then the lock get does not retry.

156.7.6 REBUILD_FREELISTS Procedure

This procedure rebuilds the freelists for the specified object.

All free blocks are placed on the master freelist. All other freelists are zeroed.

If the object has multiple freelist groups, then the free blocks are distributed among all freelists, allocating to the different groups in round-robin fashion.

Syntax

```
DBMS_REPAIR.REBUILD_FREELISTS (
   schema_name    IN VARCHAR2,
   object_name    IN  VARCHAR2,
   partition_name IN VARCHAR2 DEFAULT NULL,
   object_type    IN BINARY_INTEGER DEFAULT TABLE_OBJECT);
```

Parameters

Table 156-9 REBUILD_FREELISTS Procedure Parameters

Parameter	Description
schema_name	Schema name.
object_name	Name of the object whose freelists are to be rebuilt.
partition_name	Partition or subpartition name whose freelists are to be rebuilt. If this is a partitioned object, and partition_name is not specified, then all partitions and subpartitions are processed. If this is a partitioned object, and the specified partition contains subpartitions, then all subpartitions are processed.
object_type	Type of the object to be processed. This must be either TABLE_OBJECT (default) or INDEX_OBJECT. See"Constants".

156.7.7 SEGMENT_FIX_STATUS Procedure

With this procedure you can fix the corrupted state of a bitmap entry. The procedure either recalculates the state based on the current contents of the corresponding block or sets the state to a specific value.

Syntax

```
DBMS_REPAIR.SEGMENT_FIX_STATUS (
   segment_owner   IN VARCHAR2,
   segment_name    IN VARCHAR2,
   segment_type    IN BINARY_INTEGER DEFAULT TABLE_OBJECT,
```

```
file_number     IN BINARY_INTEGER DEFAULT NULL,
block_number    IN BINARY_INTEGER DEFAULT NULL,
status_value    IN BINARY_INTEGER DEFAULT NULL,
partition_name  IN VARCHAR2 DEFAULT NULL,);
```

Parameters

Table 156-10 SEGMENT_FIX_STATUS Procedure Parameters

Parameter	Description
schema_owner	Schema name of the segment.
segment_name	Segment name.
partition_name	Optional. Name of an individual partition. NULL for nonpartitioned objects. Default is NULL.
segment_type	Optional Type of the segment (for example, TABLE_OBJECT or INDEX_OBJECT). Default is NULL.
file_number	(optional) The tablespace-relative file number of the data block whose status has to be fixed. If omitted, all the blocks in the segment will be checked for state correctness and fixed.
block_number	(optional) The file-relative block number of the data block whose status has to be fixed. If omitted, all the blocks in the segment will be checked for state correctness and fixed.
status_value	(optional) The value to which the block status described by the file_number and block_number will be set. If omitted, the status will be set based on the current state of the block. This is almost always the case, but if there is a bug in the calculation algorithm, the value can be set manually. Status values: • 1 = block is full • 2 = block is 0-25% free • 3 = block is 25-50% free • 4 = block is 50-75% free • 5 = block is 75-100% free The status for bitmap blocks, segment headers, and extent map blocks cannot be altered. The status for blocks in a fixed hash area cannot be altered. For index blocks, there are only two possible states: 1 = block is full and 3 = block has free space.

156.7.8 SKIP_CORRUPT_BLOCKS Procedure

This procedure enables or disables the skipping of corrupt blocks during index and table scans of the specified object.

When the object is a table, skip applies to the table and its indexes. When the object is a cluster, it applies to all of the tables in the cluster, and their respective indexes.

> **Note:**
>
> When Oracle performs an index range scan on a corrupt index after DBMS_REPAIR.SKIP_CORRUPT_BLOCKS has been set for the base table, corrupt branch blocks and root blocks are not skipped. Only corrupt non-root leaf blocks are skipped.

Syntax

```
DBMS_REPAIR.SKIP_CORRUPT_BLOCKS (
   schema_name  IN VARCHAR2,
   object_name  IN VARCHAR2,
   object_type  IN BINARY_INTEGER DEFAULT TABLE_OBJECT,
   flags        IN BINARY_INTEGER DEFAULT SKIP_FLAG);
```

Parameters

Table 156-11 SKIP_CORRUPT_BLOCKS Procedure Parameters

Parameter	Description
schema_name	Schema name of the object to be processed.
object_name	Name of the object.
object_type	Type of the object to be processed. This must be either TABLE_OBJECT (default) or CLUSTER_OBJECT. See "Constants".
flags	If SKIP_FLAG is specified, then it turns on the skip of software corrupt blocks for the object during index and table scans. If NOSKIP_FLAG is specified, then scans that encounter software corrupt blocks return an ORA-1578. See"Constants".

157
DBMS_RESCONFIG

The `DBMS_RESCONFIG` package provides an interface to operate on the resource configuration list, and to retrieve listener information for a resource.

This chapter contains the following topics:

- Overview
- Summary of DBMS_RESCONFIG Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for more information about "Resource Configuration".

157.1 DBMS_RESCONFIG Overview

The `DBMS_RESCONFIG` package contains functions and procedures to manage the resource configuration lists of individual resources and the repository.

157.2 Summary of DBMS_RESCONFIG Subprograms

This table lists and briefly describes the `DBMS_RESCONFIG` package subprograms.

Table 157-1 DBMS_RESCONFIG Package Subprograms

Subprogram	Description
ADDREPOSITORYRESCONFIG Procedure	Inserts the resource configuration specified by absolute path at the given position of the repository's configuration list
ADDRESCONFIG Procedure	Inserts the resource configuration specified by the absolute path at the given position in the target resource's configuration list
APPENDRESCONFIG Procedure	Appends the resource configuration specified by rcpath to the target resource's configuration list if it is not already included in the list
DELETEREPOSITORYRESCONFIG Procedure	Removes the configuration at the given position in the repository's configuration list.
DELETERESCONFIG Procedures	Removes the configuration at the given position in the target resource's configuration list. I
GETLISTENERS Function	Returns the list of listeners applicable for a given resource
GETREPOSITORYRESCONFIG Function	Returns the resource configuration at the specified position of the repository's configuration list

Table 157-1 (Cont.) DBMS_RESCONFIG Package Subprograms

Subprogram	Description
GETREPOSITORYRESCONFIGPATHS Function	Returns a list of resource configuration paths defined for the repository
GETRESCONFIG Function	Returns the resource configuration at the specified position of the target resource's configuration list
GETRESCONFIGPATHS Function	Returns a list of resource configuration paths defined in the target resource's configuration list
PATCHREPOSITORYRESCONFIGLIST Procedure	Removes invalid references from the repository resource configuration list, and makes the repository available

157.2.1 ADDREPOSITORYRESCONFIG Procedure

This procedure inserts the resource configuration specified by absolute path of the resource configuration at the specified position of the repository's configuration list. It shifts the element currently at that position (if any) and any subsequent elements to the right.

Syntax

```
DBMS_RESCONFIG.ADDREPOSITORYRESCONFIG(
   rcpath      IN   VARCHAR2,
   pos         IN   PLS_INTEGER := NULL);
```

Parameters

Table 157-2 ADDREPOSITORYRESCONFIG Function Parameters

Parameter	Description
`rcpath`	Absolute path of the resource configuration to be inserted. An exception is raised if `rcpath` already exists in the target's configuration list.
`pos`	Index at which the new configuration is to be inserted. If this parameter is not specified then the new configuration is appended to the end of the list. An exception is raised if the index is out of range (`pos` < 0 or `pos` > the size of the target resource's configuration list).

Usage Notes

- An error is raised if the document referenced by `rcpath` is not based on `XDBResConfig.xsd` schema.
- Users must have `XDBADMIN` role and `READ` privilege on the resource configuration to be inserted; otherwise, an error is returned.

157.2.2 ADDRESCONFIG Procedure

This procedure inserts the resource configuration specified by the absolute path of the resource configuration at the given position in the target resource's configuration list. It shifts the element currently at that position (if any) and any subsequent elements to the right.

Syntax

```
DBMS_RESCONFIG.ADDRESCONFIG(
   respath     IN   VARCHAR2,
   rcpath      IN   VARCHAR2,
   pos         IN   PLS_INTEGER := NULL);
```

Parameters

Table 157-3 ADDRESCONFIG Function Parameters

Parameter	Description
`respath`	Absolute path of the target resource
`rcpath`	Absolute path of the resource configuration to be inserted. An exception is raised if rcpath already exists in the target's configuration list.
`pos`	Index at which the new configuration is to be inserted. If this parameter is not specified then the new configuration is appended to the end of the list. An exception is raised if the index is out of range (`pos` < 0 or `pos` > the size of the target resource's configuration list).

Usage Notes

- An error is raised if the document referenced by `rcpath` is not based on `XDBResConfig.xsd` schema.
- Users must have `WRITE-CONFIG` privilege on the target resource and read privilege on the resource configuration to be inserted; otherwise, an error is returned.

157.2.3 APPENDRESCONFIG Procedure

This procedure appends the resource configuration specified by rcpath to the target resource's configuration list if it is not already included in the list.

Syntax

```
DBMS_RESCONFIG.ADDRESCONFIG(
   respath       IN   VARCHAR2,
   rcpath        IN   VARCHAR2,
   appendOption  IN   PLS_INTEGER);
```

Parameters

Table 157-4 ADDRESCONFIG Function Parameters

Parameter	Description
respath	Absolute path of the target resource

Table 157-4 (Cont.) ADDRESCONFIG Function Parameters

Parameter	Description
rcpath	Absolute path of the resource configuration to be appended at the end of the target's configuration list. If rcpath already exists in the list then nothing is appended.
appendOption	Either APPEND_RESOURCE or APPEND_RECURSIVE. If APPEND_RESOURCE is specified then only the target resource is affected. If APPEND_RECURSIVE is specified then the target resource and all its descendents will be affected.

Usage Notes

- An error is raised if the document referenced by rcpath is not based on XDBResConfig.xsd schema.
- Users must have WRITE-CONFIG privilege on all affected resources and required read privilege on the resource configuration to be inserted; otherwise, an error is returned.

157.2.4 DELETEREPOSITORYRESCONFIG Procedure

This procedure removes the configuration at the given position in the repository's configuration list. It shifts any subsequent elements to the left.

Syntax

```
DBMS_RESCONFIG.DELETEREPOSITORYRESCONFIG(
   pos            IN    PLS_INTEGER);
```

Parameters

Table 157-5 DELETEREPOSITORYRESCONFIG Function Parameters

Parameter	Description
pos	The index of the configuration to be removed. An exception is raised if the index is out of range (pos < 0 or pos >= the size of the target resource's configuration list).

Usage Notes

- Users must have XDBADMIN role to execute this.
- This statement is treated as if it is a DDL statement. This means the system will implicitly commit before and after this statement.

157.2.5 DELETERESCONFIG Procedures

This procedure removes the configuration at the given position in the target resource's configuration list. It shifts any subsequent elements to the left. Users can use the overloaded for recursive deletion.

Syntax

```
DBMS_RESCONFIG.DELETERESCONFIG(
   respath       IN   VARCHAR2,
   pos           IN   PLS_INTEGER);

DBMS_RESCONFIG.DELETERESCONFIG(
   respath        IN  VARCHAR2,
   rcpath         IN  VARCHAR2,
   deleteOption   IN  PLS_INTEGER);
```

Parameters

Table 157-6 DELETERESCONFIG Procedure Parameters

Parameter	Description
`respath`	Absolute path of the target resource
`pos`	The index of the configuration to be removed. An exception is raised if the index is out of range (`pos` < 0 or `pos` >= the size of the target resource's configuration list).
`rcpath`	Absolute path of the resource configuration to be deleted if found in list.
`deleteOption`	Either `DELETE_RESOURCE` or `DELETE_RECURSIVE`. If `DELETE_RESOURCE` is specified then only the configuration list of the target resource is affected. If `DELETE_RECURSIVE` is specified then the configuration list of the target resource and all its descendents will be affected.

Usage Notes

Users must have `WRITE-CONFIG` privilege on the target resource to execute this.

157.2.6 GETLISTENERS Function

This function returns the list of listeners applicable for a given resource.

The value returned by this function is an XML document containing the `<event-listeners>` element of the `XDBResconfig.xsd` schema. It contains all the listeners applicable to the target resource, including repository-level listeners. From the returned XML document users can use the `EXTRACT` operator to retrieve the listeners defined for a specific event.

Syntax

```
DBMS_RESCONFIG.GETLISTENERS(
   path    IN   VARCHAR2)
 RETURN XMLTYPE;
```

Parameters

Table 157-7 GETLISTENERS Function Parameters

Parameter	Description
path	Absolute path of the target resource

Usage Notes

Users must have the required access privilege on all resource configurations referenced by the repository and the target resource; otherwise, an error is returned.

157.2.7 GETREPOSITORYRESCONFIG Function

This function returns the resource configuration at the specified position of the repository's configuration list.

Syntax

```
DBMS_RESCONFIG.GETREPOSITORYRESCONFIG(
   pos    IN   PLS_INTEGER)
 RETURN XMLTYPE;
```

Parameters

Table 157-8 GETREPOSITORYRESCONFIG Function Parameters

Parameter	Description
pos	Index of element to return. An exception is raised if the index is out of range (pos < 0 or pos >= the size of the repository's configuration list).

Usage Notes

Users must have the required read privilege on the requested resource configuration; otherwise, an error is returned.

157.2.8 GETREPOSITORYRESCONFIGPATHS Function

This function returns a list of resource configuration paths defined for the repository.

Syntax

```
DBMS_RESCONFIG.GETREPOSITORYRESCONFIGPATHS
 RETURN XDB$STRING_LIST_T;
```

Usage Notes

Users must be able to access all the referenced resource configurations; otherwise, an error is returned.

157.2.9 GETRESCONFIG Function

This function returns the resource configuration at the specified position of the target resource's configuration list.

Syntax

```
DBMS_RESCONFIG.GETRESCONFIG(
   respath IN VARCHAR2,
   pos IN PLS_INTEGER)
 RETURN XMLTYPE;
```

Parameters

Table 157-9 GETRESCONFIG Function Parameters

Parameter	Description
`respath`	Absolute path of the target resource
`pos`	Index of element to return. An exception is raised if the index is out of range (`pos` < 0 or `pos` >= the size of the target resource's configuration list).

Usage Notes

Users must have the required read privilege on the requested resource configuration; otherwise, an error is returned.

157.2.10 GETRESCONFIGPATHS Function

This function returns a list of resource configuration paths defined in the target resource's configuration list.

Syntax

```
DBMS_RESCONFIG.GETRESCONFIGPATHS(
   respath IN VARCHAR2)
 RETURN XDB$STRING_LIST_T;
```

Parameters

Table 157-10 GETRESCONFIGPATHS Function Parameters

Parameter	Description
`respath`	Absolute path of the target resource

Usage Notes

Users must be able to access all the referenced resource configurations; otherwise, an error is returned.

157.2.11 PATCHREPOSITORYRESCONFIGLIST Procedure

This procedure removes invalid references from the repository resource configuration list, and makes the repository available.

Under normal circumstances, deletion of a resource configuration resource cannot be performed if it is part of the repository resource configuration list. If, for some reason, the deletion of a resource configuration resource that is part of the repository resource configuration list succeeds, then any repository operation results in a 'dangling reference' error. This procedure removes those invalid references.

This procedure must be run as `SYS`.

Syntax

```
DBMS_RESCONFIG.PATCHREPOSITORYRESCONFIGLIST;
```

158
DBMS_RESOURCE_MANAGER

The `DBMS_RESOURCE_MANAGER` package maintains plans, consumer groups, and plan directives. It also provides semantics so that you may group together changes to the plan schema.

This chapter contains the following topics:

- Deprecated Subprograms
- Security Model
- Constants
- Summary of DBMS_RESOURCE_MANAGER Subprograms

> **See Also:**
>
> For more information on using the Database Resource Manager, see *Oracle Database Administrator's Guide*

158.1 DBMS_RESOURCE_MANAGER Deprecated Subprograms

The SET_INITIAL_CONSUMER_GROUP Procedure has been deprecated with Oracle Database 11*g*.

- SET_INITIAL_CONSUMER_GROUP Procedure

> **Note:**
>
> Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

158.2 DBMS_RESOURCE_MANAGER Security Model

The invoker must have the `ADMINISTER_RESOURCE_MANAGER` system privilege to execute these procedures.

The procedures to grant and revoke this privilege are in the package DBMS_RESOURCE_MANAGER_PRIVS.

158.3 DBMS_RESOURCE_MANAGER Constants

The DBMS_RESOURCE_MANAGER package defines several constants for specifying parameter values.

These are shown in the following table.

Table 158-1 DBMS_RESOURCE_MANAGER Constants

Constant	Type	Value	Description
CLIENT_ID	VARCHAR2(30)	CLIENT_ID	Client identifier of the session
CLIENT_MACHINE	VARCHAR2(30)	CLIENT_MACHINE	Name of the computer from which the client is making the connection
CLIENT_OS_USER	VARCHAR2(30)	CLIENT_OS_USER	Operating system user name of the client that is logging in
CLIENT_PROGRAM	VARCHAR2(30)	CLIENT_PROGRAM	Name of the client program used to log in to the server
MODULE_NAME	VARCHAR2(30)	MODULE_NAME	Module name in the currently running application as set by the SET_MODULE Procedure in the DBMS_APPLICATION_INFO package, or the equivalent OCI attribute setting
MODULE_NAME_ACTION	VARCHAR2(30)	MODULE_NAME_ACTION	A combination of the current module and the action being performed as set by either of the following procedures in the DBMS_APPLICATION_INFO package, or their equivalent OCI attribute setting: • SET_MODULE Procedure • SET_ACTION Procedure The attribute is specified as the module name followed by a period (.), followed by the action name (*module_name.action_name*).
ORACLE_FUNCTION	VARCHAR2(30)	ORACLE_FUNCTION	Function the session is currently executing. Valid functions are the BACKUP, COPY, DATALOAD, and INMEMORY. BACKUP is set for sessions that are doing backup operations using RMAN. COPY is set for sessions that are doing image copies using RMAN. DATALOAD is set for sessions that are loading data using Oracle Data Pump.
ORACLE_USER	VARCHAR2(30)	ORACLE_USER	Oracle Database user name
SERVICE_MODULE	VARCHAR2(30)	SERVICE_MODULE	Combination of service and module names in this form: *service_name.module_name*

Table 158-1 (Cont.) DBMS_RESOURCE_MANAGER Constants

Constant	Type	Value	Description
SERVICE_MODULE_ACTION	VARCHAR2(30)	SERVICE_MODULE_ACTION	Combination of service name, module name, and action name, in this form: *service_name.module_name.action_name*
SERVICE_NAME	VARCHAR2(30)	SERVICE_NAME	Service name used by the client to establish a connection

158.4 Summary of DBMS_RESOURCE_MANAGER Subprograms

This table lists the DBMS_RESOURCE_MANAGER subprograms and briefly describes them.

Table 158-2 DBMS_RESOURCE_MANAGER Package Subprograms

Subprogram	Description
BEGIN_SQL_BLOCK Procedure	Indicates the start of a block of SQL statements to be treated as a group by resource manager
CALIBRATE_IO Procedure	Calibrates the I/O capabilities of storage
CLEAR_PENDING_AREA Procedure	Clears the work area for the resource manager
CREATE_CATEGORY Procedure	Creates a new resource consumer group category
CREATE_CDB_PLAN Procedure	Creates entries which define consolidation resource plans.
CREATE_CDB_PLAN_DIRECTIVE Procedure	Creates the plan directives of the consolidation resource plan
CREATE_CDB_PROFILE_DIRECTIVE Procedure	Creates the performance profile directives of the consolidation resource plan
CREATE_CONSUMER_GROUP Procedure	Creates entries which define resource consumer groups
CREATE_PENDING_AREA Procedure	Creates a work area for changes to resource manager objects
CREATE_PLAN Procedure	Creates entries which define resource plans
CREATE_PLAN_DIRECTIVE Procedure	Creates resource plan directives
CREATE_SIMPLE_PLAN Procedure	Creates a single-level resource plan containing up to eight consumer groups in one step
DELETE_CATEGORY Procedure	Deletes an existing resource consumer group category
DELETE_CDB_PLAN Procedure	Deletes the consolidation resource plan
DELETE_CDB_PLAN_DIRECTIVE Procedure	Deletes the plan directive of the consolidation resource plan
DELETE_CDB_PROFILE_DIRECTIVE Procedure	Deletes the performance profile directive of the consolidation resource plan
DELETE_CONSUMER_GROUP Procedure	Deletes entries which define resource consumer groups

Table 158-2 (Cont.) DBMS_RESOURCE_MANAGER Package Subprograms

Subprogram	Description
DELETE_PLAN Procedure	Deletes the specified plan as well as all the plan directives it refers to
DELETE_PLAN_CASCADE Procedure	Deletes the specified plan as well as all its descendants (plan directives, subplans, consumer groups)
DELETE_PLAN_DIRECTIVE Procedure	Deletes resource plan directives
DEQUEUE_PARALLEL_STATEMENT Procedure	Dequeues a parallel statement from the parallel statement queue
END_SQL_BLOCK Procedure	Indicates the end of a block of SQL statements that should be treated as a group by resource manager
SET_CONSUMER_GROUP_MAPPING Procedure	Adds, deletes, or modifies entries for the login and run-time attribute mappings
SET_CONSUMER_GROUP_MAPPING_PRI Procedure	Creates the session attribute mapping priority list
SET_INITIAL_CONSUMER_GROUP Procedure	Assigns the initial resource consumer group for a user (Caution: Deprecated Subprogram)
SUBMIT_PENDING_AREA Procedure	Submits pending changes for the resource manager
SWITCH_CONSUMER_GROUP_FOR_SESS Procedure	Changes the resource consumer group of a specific session
SWITCH_CONSUMER_GROUP_FOR_USER Procedure	Changes the resource consumer group for all sessions with a given user name
SWITCH_PLAN Procedure	Sets the current resource manager plan
UPDATE_CATEGORY Procedure	Updates an existing resource consumer group category
UPDATE_CDB_AUTOTASK_DIRECTIVE Procedure	Updates the plan directives with regard to automated maintenance tasks
UPDATE_CDB_DEFAULT_DIRECTIVE Procedure	Updates the default values for a consolidation plan
UPDATE_CDB_PLAN Procedure	Updates the consolidation resource plan
UPDATE_CDB_PLAN_DIRECTIVE Procedure	Updates the plan directives for a consolidation resource plan
UPDATE_CDB_PROFILE_DIRECTIVE Procedure	Updates the performance profile directives of the consolidation resource plan
UPDATE_CONSUMER_GROUP Procedure	Updates entries which define resource consumer groups
UPDATE_PLAN Procedure	Updates entries which define resource plans
UPDATE_PLAN_DIRECTIVE Procedure	Updates resource plan directives
VALIDATE_PENDING_AREA Procedure	Validates pending changes for the resource manager

158.4.1 BEGIN_SQL_BLOCK Procedure

This procedure, to be used with parallel statement queuing, indicates the start of a block of SQL statements that should be treated as a group by resource manager.

Syntax

```
DBMS_RESOURCE_MANAGER.BEGIN_SQL_BLOCK;
```

Usage Notes

For more information, see "Parallel Statement Queuing" and "Managing Parallel Statement Queuing with Resource Manager" in *Oracle Database VLDB and Partitioning Guide*.

158.4.2 CALIBRATE_IO Procedure

This procedure calibrates the I/O capabilities of storage. Calibration status is available from the `V$IO_CALIBRATION_STATUS` view and results for a successful calibration run are located in `DBA_RSRC_IO_CALIBRATE` table.

Syntax

```
DBMS_RESOURCE_MANAGER.CALIBRATE_IO (
   num_physical_disks      IN  PLS_INTEGER DEFAULT 1,
   max_latency             IN  PLS_INTEGER DEFAULT 20,
   max_iops                OUT PLS_INTEGER,
   max_mbps                OUT PLS_INTEGER,
   actual_latency          OUT PLS_INTEGER);
```

Parameters

Table 158-3 CALIBRATE_IO Procedure Parameters

Parameter	Description
`num_physical_disks`	Approximate number of physical disks in the database storage. This parameter is used to determine the initial I/O load for the calibration run.
`max_latency`	Maximum tolerable latency in milliseconds for database-block-sized IO requests
`max_iops`	Maximum number of I/O requests per second that can be sustained. The I/O requests are randomly-distributed, database-block-sized reads.
`max_mbps`	Maximum throughput of I/O that can be sustained, expressed in megabytes per second. The I/O requests are randomly-distributed, 1 megabyte reads.
`actual_latency`	Average latency of database-block-sized I/O requests at `max_iops` rate, expressed in milliseconds

Usage Notes

- Only users with the SYSDBA privilege can run this procedure. Qualified users must also turn on `timed_statistics`, and ensure `asynch_io` is enabled for datafiles. This can be achieved by setting `filesystemio_options` to either `ASYNCH` or `SETALL`. One can also query the `asynch_io` status by means of the following SQL statement:

  ```
  col name format a50
  SELECT name, asynch_io FROM v$datafile f,v$iostat_file i
  ```

```
  WHERE f.file#        = i.file_no
  AND   filetype_name  = 'Data File'
 /
```

- Only one calibration can be run at a time. If another calibration is initiated at the same time, it will fail.
- For an Oracle Real Application Clusters (Oracle RAC) database, the workload is simultaneously generated from all instances.
- In a multitenant container database (CDB), calibration can only be run from the CDB root (`CDB$ROOT`).
- Calibration is extremely disruptive to the database performance. It is strongly recommended to run calibration only when database users can tolerate severe deterioration to database performance.
- For optimal calibration results, no other database workloads should be running.

> **See Also:**
>
> *Oracle Database Performance Tuning Guide* for more information about calibration

Examples

Example of using I/O Calibration procedure

```
SET SERVEROUTPUT ON
DECLARE
  lat  NUMBER;
  iops INTEGER;
  mbps INTEGER;
BEGIN
-- DBMS_RESOURCE_MANAGER.CALIBRATE_IO (<DISKS>, <MAX_LATENCY>, iops, mbps, lat);
   DBMS_RESOURCE_MANAGER.CALIBRATE_IO (2, 10, iops, mbps, lat);

end;
/
```

View for I/O calibration results

```
SQL> desc V$IO_CALIBRATION_STATUS
  Name                                      Null?    Type
  ----------------------------------------- -------- ----------------------------
  STATUS                                             VARCHAR2(13)
  CALIBRATION_TIME                                   TIMESTAMP(3)

SQL> desc gv$io_calibration_status
  Name                                      Null?    Type
  ----------------------------------------- -------- ----------------------------
  INST_ID                                            NUMBER
  STATUS                                             VARCHAR2(13)
  CALIBRATION_TIME                                   TIMESTAMP(3)

Column explanation:
-------------------
STATUS:
  IN PROGRESS   : Calibration in Progress (Results from previous calibration
                  run displayed, if available)
  READY         : Results ready and available from earlier run
```

```
  NOT AVAILABLE : Calibration results not available.

CALIBRATION_TIME: End time of the last calibration run
```

DBA table that stores I/O Calibration results

```
SQL> desc DBA_RSRC_IO_CALIBRATE
  Name                                      Null?    Type
  ----------------------------------------- -------- ----------------------------
  START_TIME                                         TIMESTAMP(6)
  END_TIME                                           TIMESTAMP(6)
  MAX_IOPS                                           NUMBER
  MAX_MBPS                                           NUMBER
  MAX_PMBPS                                          NUMBER
  LATENCY                                            NUMBER
  NUM_PHYSICAL_DISKS                                 NUMBER

comment on table DBA_RSRC_IO_CALIBRATE is
'Results of the most recent I/O calibration'
/
comment on column DBA_RSRC_IO_CALIBRATE.START_TIME is
'start time of the most recent I/O calibration'
/
comment on column DBA_RSRC_IO_CALIBRATE.END_TIME is
'end time of the most recent I/O calibration'
/
comment on column DBA_RSRC_IO_CALIBRATE.MAX_IOPS is
'maximum number of data-block read requests that can be sustained per second'
/
comment on column DBA_RSRC_IO_CALIBRATE.MAX_MBPS is
'maximum megabytes per second of maximum-sized read requests that can be
sustained'
/
comment on column DBA_RSRC_IO_CALIBRATE.MAX_PMBPS is
'maximum megabytes per second of large I/O requests that
can be sustained by a single process'
/
comment on column DBA_RSRC_IO_CALIBRATE.LATENCY is
'latency for data-block read requests'
/
comment on column DBA_RSRC_IO_CALIBRATE.NUM_PHYSICAL_DISKS is
'number of physical disks in the storage subsystem (as specified by user)'
/
```

158.4.3 CLEAR_PENDING_AREA Procedure

This procedure clears pending changes for the resource manager.

Syntax

```
DBMS_RESOURCE_MANAGER.CLEAR_PENDING_AREA;
```

158.4.4 CREATE_CATEGORY Procedure

This procedure creates a new consumer group category. The primary purpose of this attribute is to support Exadata I/O Resource Manager category plans.

The view `DBA_RSRC_CATEGORIES` defines the currently defined categories. The `ADMINISTRATIVE`, `INTERACTIVE`, `BATCH`, `MAINTENANCE`, and `OTHER` categories are available.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_CATEGORY (
   category    IN    VARCHAR2,
   comment     IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 158-4 CREATE_CATEGORY Procedure Parameters

Parameter	Description
category	Name of consumer group category
comment	User comment

158.4.5 CREATE_CDB_PLAN Procedure

Creates entries which define consolidation resource plans.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_CDB_PLAN (
   plan                    IN    VARCHAR2(32),
   comment                 IN    VARCHAR2(2000) DEFAULT NULL);
```

Parameters

Table 158-5 CREATE_CDB_PLAN Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
comment	User comment

Usage Notes

This procedure can be run only from the CDB root (CDB$ROOT).

158.4.6 CREATE_CDB_PLAN_DIRECTIVE Procedure

This procedure creates the plan directives of the consolidation resource plan. Plan directives specify the resource allocation policy for pluggable databases (PDBs).

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_CDB_PLAN_DIRECTIVE (
   plan                    IN    VARCHAR2,
   pluggable_database      IN    VARCHAR2,
   comment                 IN    VARCHAR2 (2000) DEFAULT '',
   shares                  IN    NUMBER          DEFAULT NULL,
   utilization_limit       IN    NUMBER          DEFAULT NULL,
   parallel_server_limit   IN    NUMBER          DEFAULT NULL,
   memory_limit            IN    NUMBER          DEFAULT 100,
   memory_min              IN    NUMBER          DEFAULT 0);
```

Parameters

Table 158-6 CREATE_CDB_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
pluggable_database	Name of the PDB
comment	User comment
shares	Specifies the share of resource allocation for the PDB. CPU Resource Manager and Exadata I/O Resource Manager are enabled by specifying shares for each PDB. The shares parameter is also used for Parallel Statement Queuing. If no share is specified, the default is obtained from the default directive, specified through UPDATE_CDB_DEFAULT_DIRECTIVE Procedure.
utilization_limit	Specifies the maximum percentage of the CDB's CPU and Exadata I/O resources that the PDB can utilize. CPU Resource Manager and Exadata I/O Resource Manager can also be limited by setting the CPU_COUNT parameter for the PDB.
parallel_server_limit	Parallel servers that the PDB can use after which parallel statements are queued. Alternatively, you can set the parallel_servers_target at the PDB level. A PDB can set a lower limit for parallel execution servers than the limit specified in the CDB resource plan. When the PARALLEL_SERVERS_TARGET initialization parameter is set in a PDB, and parallel execution server limit is specified for a PDB in the CDB resource plan, then the lower limit is used. For example, assume that the parallel_servers_target initialization parameter is set to 100 in the CDB root and parallel_server_limit is set to 70 for hrpdb in the CDB resource plan. Also, assume that hrpdb has its parallel_servers_target initialization parameter set to 50. In this case, the limit for parallel execution servers for hrpdb is 50, because 50 is lower than the CDB resource plan limit of 70 for hrpdb. **Note:** Oracle recommends, that you use parallel_servers_target parameter instead of parallel_servers_limit in a CDB resource plan.
memory_limit	This parameter is only applicable to Oracle Exadata storage for configuring the Exadata Smart Flash Cache and Exadata PMEM Cache.
memory_min	This parameter is only applicable to Oracle Exadata storage for configuring the Exadata Smart Flash Cache and Exadata PMEM Cache.

Usage Notes

- The default value for `shares`, `utilization_limit`, and `parallel_server_limit` is `NULL`. When a user specifies `NULL`, or does not specify a value, this indicates that the default value should be used.
- This procedure can be run only from the CDB root (`CDB$ROOT`).

158.4.7 CREATE_CDB_PROFILE_DIRECTIVE Procedure

This procedure creates the performance profile directives of the consolidation resource plan. The directives specify the resource allocation policy for pluggable databases (PDBs) that use the performance profile.

For a PDB to use the new performance profile, the PDB must have the `DB_PERFORMANCE_PROFILE` initialization parameter set to the profile name.

This procedure provides an easy way to specify the directive for a large number of PDBs with the same resource requirements. Each PDB with a `DB_PERFORMANCE_PROFILE` initialization parameter set to the performance profile name inherits the settings specified by this directive, including the shares, utilization limit, and so on.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_CDB_PROFILE_DIRECTIVE (
   plan                   IN    VARCHAR2,
   profile                IN    VARCHAR2,
   comment                IN    VARCHAR2 (2000) DEFAULT '',
   shares                 IN    NUMBER          DEFAULT NULL,
   utilization_limit      IN    NUMBER          DEFAULT NULL,
   parallel_server_limit  IN    NUMBER          DEFAULT NULL,
   memory_limit           IN    NUMBER          DEFAULT 100,
   memory_min             IN    NUMBER          DEFAULT 0);
```

Parameters

Table 158-7 CREATE_CDB_PROFILE_DIRECTIVE Procedure Parameters

Parameter	Description
`plan`	Name of the consolidation plan
`profile`	Name of the performance profile
`comment`	User comment
`shares`	Specifies the share of resource allocation for PDBs that use the performance profile. CPU Resource Manager and Exadata I/O Resource Manager are enabled by specifying shares for each PDB. The `shares` parameter is also used for Parallel Statement Queuing. If no share is specified, the default is obtained from the default directive, specified through UPDATE_CDB_DEFAULT_DIRECTIVE Procedure.
`utilization_limit`	Specifies the maximum percentage of CPU Resource Manager and Exadata I/O Resource Manager that PDBs that use the performance profile can utilize.
`parallel_server_limit`	Specifies the maximum percentage of `parallel_servers_target` parallel servers that PDBs that use the performance profile can use.

Table 158-7 (Cont.) CREATE_CDB_PROFILE_DIRECTIVE Procedure Parameters

Parameter	Description
memory_limit	This parameter is only applicable to Oracle Exadata storage for configuring the Exadata Smart Flash Cache and Exadata PMEM Cache.
memory_min	This parameter is only applicable to Oracle Exadata storage for configuring the Exadata Smart Flash Cache and Exadata PMEM Cache.

158.4.8 CREATE_CONSUMER_GROUP Procedure

This procedure creates entries which define resource consumer groups.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_CONSUMER_GROUP (
   consumer_group  IN VARCHAR2,
   comment         IN VARCHAR2 DEFAULT NULL,
   cpu_mth         IN VARCHAR2 DEFAULT NULL,
   mgmt_mth        IN VARCHAR2 DEFAULT 'ROUND-ROBIN',
   category        IN VARCHAR2 DEFAULT 'OTHER');
```

Parameters

Table 158-8 CREATE_CONSUMER_GROUP Procedure Parameters

Parameter	Description
consumer_group	Name of the consumer group
comment	User comment
cpu_mth	Name of CPU resource allocation method (deprecated)
mgmt_mth	Name of CPU resource allocation method
category	Describes the category of the consumer group. The primary purpose of this attribute is to support Exadata I/O Resource Manager category plans. The view DBA_RSRC_CATEGORIES defines the currently defined categories. Categories can be modified, using the CREATE_CATEGORY Procedure, UPDATE_CATEGORY Procedure, and DELETE_CATEGORY Procedure.

158.4.9 CREATE_PENDING_AREA Procedure

This procedure makes changes to resource manager objects.

All changes to the plan schema must be done within a pending area. The pending area can be thought of as a "scratch" area for plan schema changes. The administrator creates this pending area, makes changes as necessary, possibly validates these changes, and only when the submit is completed do these changes become active.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_PENDING_AREA;
```

Usage Notes

You may, at any time while the pending area is active, view the current plan schema with your changes by selecting from the appropriate user views.

At any time, you may clear the pending area if you want to stop the current changes. You may also call the `VALIDATE` procedure to confirm whether the changes you have made are valid. You do not have to perform your changes in a given order to maintain a consistent group of entries. These checks are also implicitly done when the pending area is submitted.

> **Note:**
>
> Oracle allows "orphan" consumer groups (in other words, consumer groups that have no plan directives that refer to them). This is in anticipation that an administrator may want to create a consumer group that is not currently being used, but will be used in the future.

For resource plans, the following rules must be adhered to, and they are checked whenever the validate or submit procedures are executed:

- No plan schema may contain any loops.
- All plans and consumer groups referred to by plan directives must exist.
- All plans must have plan directives that refer to either plans or consumer groups.
- All percentages in any given level must not add up to greater than 100 for the emphasis resource allocation method.
- No plan may be deleted that is currently being used as a top plan by an active instance.
- The plan directive parameter, `parallel_degree_limit_p1`, may only appear in plan directives that refer to consumer groups (that is, not at subplans).
- There cannot be more than 28 plan directives coming from any given plan (that is, no plan can have more than 28 children).
- There cannot be more than 28 consumer groups in any active plan schema.
- Plans and consumer groups use the same namespace; therefore, no plan can have the same name as any consumer group.
- There must be a plan directive for `OTHER_GROUPS` somewhere in any active plan schema.This ensures that a session not covered by the currently active plan is allocated resources as specified by the `OTHER_GROUPS` directive.

> **Note:**
>
> These rules are not applicable for CDB resource plans.

If any of the preceding rules are broken when checked by the `VALIDATE` or `SUBMIT` procedures, then an informative error message is returned. You may then make changes to fix one or more problems and reissue the validate or submit procedures.

158.4.10 CREATE_PLAN Procedure

This procedure creates entries which define resource plans.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_PLAN (
   plan                       IN   VARCHAR2,
   comment                    IN   VARCHAR2 DEFAULT NULL,
   cpu_mth                    IN   VARCHAR2 DEFAULT NULL, -- deprecated
   active_sess_pool_mth       IN   VARCHAR2 DEFAULT 'ACTIVE_SESS_POOL_ABSOLUTE',
   parallel_degree_limit_mth  IN   VARCHAR2 DEFAULT
                                      'PARALLEL_DEGREE_LIMIT_ABSOLUTE',
   queueing_mth               IN   VARCHAR2 DEFAULT 'FIFO_TIMEOUT',
   mgmt_mth                   IN   VARCHAR2 DEFAULT 'EMPHASIS',
   sub_plan                   IN   BOOLEAN DEFAULT FALSE,
   max_iops                   IN   NUMBER DEFAULT NULL,
   max_mbps                   IN   NUMBER DEFAULT NULL);
```

Parameters

Table 158-9 CREATE_PLAN Procedure Parameters

Parameter	Description
plan	Name of the resource plan
comment	User comment
cpu_mth	Allocation method for CPU resources (deprecated)
active_sess_pool_mth	Active session pool resource allocation method. Limits the number of active sessions. All other sessions are inactive and wait in a queue to be activated. ACTIVE_SESS_POOL_ABSOLUTE is the default and only method available.
parallel_degree_limit_mth	Resource allocation method for specifying a limit on the degree of parallelism of any operation. PARALLEL_DEGREE_LIMIT_ABSOLUTE is the default and only method available.
queueing_mth	Queuing resource allocation method. Controls order in which queued inactive sessions will execute. FIFO_TIMEOUT is the default and only method available
mgmt_mth	Resource allocation method for specifying how much resources (for example, CPU or I/O) each consumer group or sub-plan gets • EMPHASIS - for multilevel plans that use percentages to specify how I/O resources are distributed among consumer groups • RATIO - for single-level plans that use ratios to specify how I/O resources are distributed
sub_plan	If TRUE, indicates that this plan is only intended for use as a sub-plan. Sub-plans are not required to have an OTHER_GROUPS directive. Default is FALSE.
max_iops	Nonoperative
max_mbps	Nonoperative

Usage Notes

If you want to use any default resource allocation method, then you do not need to specify it when creating or updating a plan.

158.4.11 CREATE_PLAN_DIRECTIVE Procedure

This procedure creates resource plan directives.

> **Note:**
>
> The parameters `max_utilization_limit` and `parallel_target_percentage` are deprecated with Oracle Database 11*g* Release 1 (11.1.0.1), and are replaced by `utilization_limit` and `parallel_server_limit`.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_PLAN_DIRECTIVE (
   plan                          IN VARCHAR2,
   group_or_subplan              IN VARCHAR2,
   comment                       IN VARCHAR2 DEFAULT NULL,
   cpu_p1                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p2                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p3                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p4                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p5                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p6                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p7                        IN NUMBER   DEFAULT NULL, -- deprecated
   cpu_p8                        IN NUMBER   DEFAULT NULL, -- deprecated
   active_sess_pool_p1           IN NUMBER   DEFAULT NULL,
   queueing_p1                   IN NUMBER   DEFAULT NULL,
   parallel_degree_limit_p1      IN NUMBER   DEFAULT NULL,
   switch_group                  IN VARCHAR2 DEFAULT NULL,
   switch_time                   IN NUMBER   DEFAULT NULL,
   switch_estimate               IN BOOLEAN  DEFAULT FALSE,
   max_est_exec_time             IN NUMBER   DEFAULT NULL,
   undo_pool                     IN NUMBER   DEFAULT NULL,
   max_idle_time                 IN NUMBER   DEFAULT NULL,
   max_idle_blocker_time         IN NUMBER   DEFAULT NULL,
   switch_time_in_call           IN NUMBER   DEFAULT NULL, -- deprecated
   mgmt_p1                       IN NUMBER   DEFAULT NULL,
   mgmt_p2                       IN NUMBER   DEFAULT NULL,
   mgmt_p3                       IN NUMBER   DEFAULT NULL,
   mgmt_p4                       IN NUMBER   DEFAULT NULL,
   mgmt_p5                       IN NUMBER   DEFAULT NULL,
   mgmt_p6                       IN NUMBER   DEFAULT NULL,
   mgmt_p7                       IN NUMBER   DEFAULT NULL,
   mgmt_p8                       IN NUMBER   DEFAULT NULL,
   switch_io_megabytes           IN NUMBER   DEFAULT NULL,
   switch_io_reqs                IN NUMBER   DEFAULT NULL,
   switch_for_call               IN BOOLEAN  DEFAULT NULL,
   max_utilization_limit         IN NUMBER   DEFAULT NULL,  -- deprecated
   parallel_target_percentage    IN NUMBER   DEFAULT NULL,  -- deprecated
   parallel_server_limit         IN NUMBER   DEFAULT NULL,
   utilization_limit             IN NUMBER   DEFAULT NULL,
   switch_io_logical             IN NUMBER   DEFAULT NULL,
   switch_elapsed_time           IN NUMBER   DEFAULT NULL,
```

```
shares                      IN NUMBER   DEFAULT NULL,
parallel_stmt_critical      IN VARCHAR2 DEFAULT NULL,
session_pga_limit           IN NUMBER   DEFAULT NULL,
pq_timeout_action           IN NUMBER   DEFAULT NULL,
parallel_queue_timeout      IN NUMBER   DEFAULT NULL,);
```

> **Note:**
>
> Oracle recommends that you use `shares` instead of `mgmt_p*`.

Parameters

Table 158-10 CREATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
`plan`	Name of the resource plan
`group_or_subplan`	Name of the consumer group or subplan
`comment`	Comment for the plan directive
`cpu_p1`	-- deprecated: use `mgmt_p1` or, even better, shares instead
`cpu_p2`	-- deprecated: use `mgmt_p2` or, even better, shares instead
`cpu_p3`	-- deprecated: use `mgmt_p3` or, even better, shares instead
`cpu_p4`	-- deprecated: use `mgmt_p4` or, even better, shares instead
`cpu_p5`	-- deprecated: use `mgmt_p5` or, even better, shares instead
`cpu_p6`	-- deprecated: use `mgmt_p6` or, even better, shares instead
`cpu_p7`	-- deprecated: use `mgmt_p7` or, even better, shares instead
`cpu_p8`	-- deprecated: use `mgmt_p8` or, even better, shares instead
`active_sess_pool_p1`	Specifies maximum number of sessions that can currently have an active call
`queueing_p1`	Specified time (in seconds) after which a call in the inactive session queue (waiting for execution) will time out. Default is `NULL`, which means unlimited.
`parallel_degree_limit_p1`	Specifies a limit on the degree of parallelism for any operation. Default is `NULL`, which means unlimited. If the value is 0, then all operations will be serial.
`switch_group`	Specifies consumer group to switch to, once a switch condition is met. If the group name is `CANCEL_SQL`, then the current call is canceled when the switch condition is met. If the group name is `KILL_SESSION`, then the session is killed when the switch condition is met. If the group name is `LOG_ONLY`, then no action is taken other than recording this event via SQL monitor.Default is `NULL`.
`switch_time`	Specifies the time on CPU (not elapsed time) that a session can execute before an action is taken. Default is `NULL`, which means unlimited. As with other switch directives, if `switch_for_call` is `TRUE`, the number of CPUs is accumulated from the start of a call. Otherwise, the number of CPUs is accumulated for the length of the session.

Table 158-10 (Cont.) CREATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
switch_estimate	If TRUE, tells Oracle to use its execution time estimate to automatically switch the consumer group of an operation before beginning its execution. This is used in conjunction with the switch_time directive. Default value is FALSE.
max_est_exec_time	Specifies the maximum execution time (in CPU seconds) allowed for a session. If the optimizer estimates that an operation will take longer than MAX_EST_EXEC_TIME, the operation is not started and ORA-07455 is issued. If the optimizer does not provide an estimate, this directive has no effect. Default is NULL, which means unlimited.
undo_pool	Limits the size in kilobytes of the undo records corresponding to uncommitted transactions by this consumer group
max_idle_time	Indicates the maximum session idle time. Default is NULL, which means unlimited.
max_idle_blocker_time	Maximum amount of time in seconds that a session can be idle while blocking another session's acquisition of a resource
switch_time_in_call	Deprecated. If this parameter is specified, switch_time is set to switch_time_in_call (in seconds) and switch_for_call is effectively set to TRUE. It is better to use switch_time and switch_for_call.
mgmt_p1	Resource allocation value for level 1 (replaces cpu_p1): • EMPHASIS - specifies the resource percentage at the first level • RATIO - specifies the weight of resource usage
mgmt_p2	Resource allocation value for level 2 (replaces cpu_p2) • EMPHASIS - specifies the resource percentage at the second level • RATIO - non-applicable
mgmt_p3	Resource allocation value for level 3 (replaces cpu_p3) • EMPHASIS - specifies the resource percentage at the third level • RATIO - non-applicable
mgmt_p4	Resource allocation value for level 4 (replaces cpu_p4) • EMPHASIS - specifies the resource percentage at the fourth level • RATIO - non-applicable
mgmt_p5	Resource allocation value for level 5 (replaces cpu_p5) • EMPHASIS - specifies the resource percentage at the fifth level • RATIO - non-applicable
mgmt_p6	Resource allocation value for level 6 (replaces cpu_p6) • EMPHASIS - specifies the resource percentage at the sixth level • RATIO - non-applicable
mgmt_p7	Resource allocation value for level 7 (replaces cpu_p7) • EMPHASIS - specifies the resource percentage at the seventh level • RATIO - non-applicable
mgmt_p8	Resource allocation value for level 8 (replaces cpu_p8) • EMPHASIS - specifies the resource percentage at the eighth level • RATIO - non-applicable

Table 158-10 (Cont.) CREATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
switch_io_megabytes	Specifies the amount of I/O (in MB) that a session can issue before an action is taken. Default is NULL, which means unlimited. As with other switch directives, if switch_for_call is TRUE, the number of CPUs is accumulated from the start of a call. Otherwise, the number of CPUs is accumulated for the length of the session.
switch_io_reqs	Specifies the number of I/O requests that a session can issue before an action is taken. Default is NULL, which means unlimited. As with other switch directives, if switch_for_call is TRUE, the number of CPUs is accumulated from the start of a call. Otherwise, the number of CPUs is accumulated for the length of the session.
switch_for_call	Specifies that if an action is taken because of the switch_time, switch_io_megabytes, switch_io_reqs, switch_io_logical or switch_elapsed_time parameters, the consumer group is restored to its original consumer group at the end of the top call. Default is NULL, which means that the original consumer group is not restored at the end of the top call.
max_utilization_limit	-- deprecated: use utilization_limit instead
parallel_target_percentage	-- deprecated: use parallel_sever_limit instead
parallel_queue_timeout	Specifies the time (in seconds) that a parallel statement may remain in its Consumer Group's parallel statement queue before it is removed and terminated with an error (ORA- 07454). **Note:** You can use the pq_timeout_action parameter to specify the action to be taken when a parallel statement is removed from the queue.
parallel_sever_limit	Specifies the maximum percentage of parallel_servers_target parallel servers that the Consumer Group can use, after which parallel statements are queued.
utilization_limit	Resource limit. Currently it includes CPU and I/O for Exadata. For CPU, this limits the CPU utilization for the consumer group. For Exadata I/O, this limits the disk utilization for the consumer group. This does not apply to parallel servers.
switch_io_logical	Number of logical IOs that will trigger the action specified by switch_group. As with other switch directives, if switch_for_call is TRUE, the number of logical IOs is accumulated from the start of a call. Otherwise, the number of logical IOs is accumulated for the length of the session.
switch_elapsed_time	Elapsed time that will trigger the action specified by switch_group. As with other switch directives, if switch_for_call is TRUE, the elapsed time is accumulated from the start of a call. Otherwise, the elapsed time is accumulated for the length of the session.

Table 158-10 (Cont.) CREATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
shares	Specifies the share of resource allocation for the consumer group. CPU Resource Manager and Exadata I/O Resource Manager are enabled by specifying shares for each consumer group. The shares parameter is also used for Parallel Statement Queuing. If CPU Resource Manager and Exadata I/O Resource Manager are enabled, then the default value is 1.
parallel_stmt_critical	If set to BYPASS_QUEUE, parallel statements from the Consumer Group are not queued, regardless of the PARALLEL_DEGREE_POLICY parameter value. If set to QUEUE, all the parallel statements from the consumer group, irrespective of the parallel_degree_policy parameter value, are eligible for queuing. Default is FALSE, which means that parallel statements are eligible for queuing, based on the parallel_degree_policy parameter value.
session_pga_limit	Maximum amount of untunable PGA (in MB) that a session in this consumer group can allocate before being terminated. NULL (default) indicates no limit. SQL operations that allocate tunable PGA (operations that can opt to use temp space) are not controlled by this limit.
pq_timeout_action	Specifies the action to be taken when a parallel statement is removed from the queue due to parallel_queue_timeout. The values are: • CANCEL — The parallel statement is terminated with error ORA-7454 • RUN — The SQL statement runs immediately, and might get downgraded if parallel servers are unavailable The default action of this parameter is CANCEL.

Usage Notes

- All parameters default to NULL.
- For max_idle_time and max_idle_blocker_time, PMON will check these limits once a minute. If it finds a session that has exceeded one of the limits, it will forcibly kill the session and clean up all its state.
- The parameter switch_for_call is mostly useful for three-tier applications where the mid-tier server is implementing session pooling. By using switch_for_call, the resource usage of one client will not affect a future client that happens to be executed on the same session.
- An error is thrown if PQ_TIMEOUT_ACTION is specified, but PARALLEL_QUEUE_TIMEOUT is not specified.
- Specifies the action to be taken when a parallel statement is removed from the queue.

158.4.12 CREATE_SIMPLE_PLAN Procedure

This procedure creates a single-level resource plan containing up to eight consumer groups in one step. You do not need to create a pending area manually before creating a resource plan,

or use the CREATE_CONSUMER_GROUP and CREATE_RESOURCE_PLAN_DIRECTIVES procedures separately.

Syntax

```
DBMS_RESOURCE_MANAGER.CREATE_SIMPLE_PLAN (
   simple_plan      IN  VARCHAR2  DEFAULT NULL,
   consumer_group1  IN  VARCHAR2  DEFAULT NULL,
   group1_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group2  IN  VARCHAR2  DEFAULT NULL,
   group2_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group3  IN  VARCHAR2  DEFAULT NULL,
   group3_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group4  IN  VARCHAR2  DEFAULT NULL,
   group4_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group5  IN  VARCHAR2  DEFAULT NULL,
   group5_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group6  IN  VARCHAR2  DEFAULT NULL,
   group6_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group7  IN  VARCHAR2  DEFAULT NULL,
   group7_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   consumer_group8  IN  VARCHAR2  DEFAULT NULL,
   group8_cpu       IN  NUMBER    DEFAULT NULL,   -- deprecated
   group1_percent   IN  NUMBER    DEFAULT NULL,
   group2_percent   IN  NUMBER    DEFAULT NULL,
   group3_percent   IN  NUMBER    DEFAULT NULL,
   group4_percent   IN  NUMBER    DEFAULT NULL,
   group5_percent   IN  NUMBER    DEFAULT NULL,
   group6_percent   IN  NUMBER    DEFAULT NULL,
   group7_percent   IN  NUMBER    DEFAULT NULL,
   group8_percent   IN  NUMBER    DEFAULT NULL);
```

Parameters

Table 158-11 CREATE_SIMPLE_PLAN Procedure Parameters

Parameter	Description
simple_plan	Name of the resource plan
consumer_group1	Name of the consumer group
group1_cpu	Percentage for group (deprecated)
consumer_group2	Name of the consumer group
group2_cpu	Percentage for group (deprecated)
consumer_group3	Name of the consumer group
group3_cpu	Percentage for group (deprecated)
consumer_group4	Name of the consumer group
group4_cpu	Percentage for group (deprecated)
consumer_group5	Name of the consumer group
group5_cpu	Percentage for group (deprecated)
consumer_group6	Name of the consumer group
group6_cpu	Percentage for group (deprecated)
consumer_group7	Name of the consumer group
group7_cpu	Percentage for group (deprecated)

Table 158-11 (Cont.) CREATE_SIMPLE_PLAN Procedure Parameters

Parameter	Description
consumer_group8	OTHER_GROUPS - all sessions that aren't mapped to a consumer group.
group8_cpu	Percentage for group (deprecated)
group1_percent	Percentage of resources allocated for this consumer group
group2_percent	Percentage of resources allocated for this consumer group
group3_percent	Percentage of resources allocated for this consumer group
group4_percent	Percentage of resources allocated for this consumer group
group5_percent	Percentage of resources allocated for this consumer group
group6_percent	Percentage of resources allocated for this consumer group
group7_percent	Percentage of resources allocated for this consumer group
group8_percent	Percentage of resources allocated to other groups

158.4.13 DELETE_CATEGORY Procedure

This procedure deletes an existing resource consumer group category.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_CATEGORY (
   category        IN    VARCHAR2);
```

Parameters

Table 158-12 DELETE_CATEGORY Procedure Parameters

Parameter	Description
category	Name of consumer group category

158.4.14 DELETE_CDB_PLAN Procedure

This procedure deletes the consolidation resource plan.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_CDB_PLAN (
   plan    IN    VARCHAR2(32)    DEFAULT NULL);
```

Parameters

Table 158-13 DELETE_CDB_PLAN Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan

Usage Notes

This procedure can be run only from the CDB root (CDB$ROOT).

158.4.15 DELETE_CDB_PLAN_DIRECTIVE Procedure

This procedure deletes the plan directives of the consolidation resource plan. Once the plan directive is deleted, the pluggable database will get the default resource allocation.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_CDB_PLAN_DIRECTIVE (
   plan                    IN    VARCHAR2(32)   DEFAULT NULL,
   pluggable_database      IN    VARCHAR2(32)   DEFAULT NULL);
```

Parameters

Table 158-14 DELETE_CDB_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
pluggable_database	Name of the pluggable database in which the plan directive is to be deleted

Usage Notes

This procedure can be run only from the CDB root (CDB$ROOT).

158.4.16 DELETE_CDB_PROFILE_DIRECTIVE Procedure

This procedure deletes the performance profile directive of the consolidation resource plan. Once the directive is deleted, the pluggable databases (PDBs) that use the performance profile use the default resource allocation.

For a PDB to use a performance profile, the PDB must have the DB_PERFORMANCE_PROFILE initialization parameter set to the performance profile name.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_CDB_PROFILE_DIRECTIVE (
   plan          IN    VARCHAR2(32)   DEFAULT NULL,
   profile       IN    VARCHAR2(32)   DEFAULT NULL);
```

Parameters

Table 158-15 DELETE_CDB_PROFILE_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
profile	Name of the performance profile directive to be deleted

Usage Notes

This procedure can be run only from the CDB root (CDB$ROOT).

158.4.17 DELETE_CONSUMER_GROUP Procedure

This procedure deletes entries which define resource consumer groups.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_CONSUMER_GROUP (
   consumer_group IN VARCHAR2);
```

Parameters

Table 158-16 DELETE_CONSUMER_GROUP Procedure Parameters

Parameters	Description
consumer_group	Name of the consumer group to be deleted

158.4.18 DELETE_PLAN Procedure

This procedure deletes the specified plan as well as all the plan directives to which it refers.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_PLAN (
   plan IN VARCHAR2);
```

Parameters

Table 158-17 DELETE_PLAN Procedure Parameters

Parameter	Description
plan	Name of the resource plan to delete

158.4.19 DELETE_PLAN_CASCADE Procedure

This procedure deletes the specified plan and all of its descendants (plan directives, subplans, consumer groups). Mandatory objects and directives are not deleted.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_PLAN_CASCADE (
   plan IN VARCHAR2);
```

Parameters

Table 158-18 DELETE_PLAN_CASCADE Procedure Parameters

Parameters	Description
plan	Name of the plan

Usage Notes

If DELETE_PLAN_CASCADE encounters any error, then it rolls back the operation, and nothing is deleted.

158.4.20 DELETE_PLAN_DIRECTIVE Procedure

This procedure deletes resource plan directives.

Syntax

```
DBMS_RESOURCE_MANAGER.DELETE_PLAN_DIRECTIVE (
   plan              IN VARCHAR2,
   group_or_subplan  IN VARCHAR2);
```

Parameters

Table 158-19 DELETE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the resource plan
group_or_subplan	Name of the group or subplan

158.4.21 DEQUEUE_PARALLEL_STATEMENT Procedure

This procedure dequeues a parallel statement from the parallel statement queue.

If the PARALLEL_DEGREE_POLICY initialization parameter is set to AUTO or ADAPTIVE, then parallel statement queuing is enabled. If a parallel statement is in the parallel statement queue, then you can use this procedure to dequeue the parallel statement so that it runs immediately.

Syntax

```
DBMS_RESOURCE_MANAGER.DEQUEUE_PARALLEL_STATEMENT (
   session_id      IN  PLS_INTEGER,
   session_serial  IN  PLS_INTEGER,
   inst_id         IN  PLS_INTEGER  DEFAULT NULL,
   sql_id          IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 158-20 DEQUEUE_PARALLEL_STATEMENT Procedure Parameters

Parameter	Description
session_id	The session id of the session running the parallel statement to be dequeued.
session_serial	The serial number of the session.
inst_id	Instance ID where the session is running. If NULL, then the current instance is used.
sql_id	The SQL ID of the session's statement to dequeue. If the session is running SQL with a different SQL ID, then the statement is not dequeued.

158.4.22 END_SQL_BLOCK Procedure

This procedure, to be used with parallel statement queuing, indicates the end of a block of SQL statements that should be treated as a group by resource manager.

Syntax

```
DBMS_RESOURCE_MANAGER.END_SQL_BLOCK;
```

Usage Notes

For more information, see "Parallel Statement Queuing" and "Managing Parallel Statement Queuing with Resource Manager" in *Oracle Database VLDB and Partitioning Guide*.

158.4.23 SET_CONSUMER_GROUP_MAPPING Procedure

This procedure adds, deletes, or modifies entries that map sessions to consumer groups, based on the session's login and runtime attributes.

Syntax

```
DBMS_RESOURCE_MANAGER.SET_CONSUMER_GROUP_MAPPING(
   attribute        IN VARCHAR2,
   value            IN VARCHAR2,
   consumer_group   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 158-21 SET_CONSUMER_GROUP_MAPPING Procedure Parameters

Parameters	Description
`attribute`	Mapping attribute to add or modify. It can be one of the Constants listed.
`value`	Attribute value to match. This includes both absolute mapping and regular expressions.
`consumer_group`	Name of the mapped consumer group, or `NULL` to delete a mapping

Usage Notes

- If no mapping exists for the given attribute and value, a mapping to the given consumer group will be created. If a mapping already exists for the given attribute and value, the mapped consumer group will be updated to the one given. If the `consumer_group` argument is `NULL`, then any mapping from the given attribute and value will be deleted.
- The subprogram supports simple regex expressions for the `value` parameter. It implements the same semantics as the SQL `'LIKE'` operator. Specifically, it uses `'%'` as amulticharacter wildcard and `'_'` as a single character wildcard. The `'\'` character can be used to escape the wildcards. Note that wildcards can only be used if the attribute is one of the following:
  - `CLIENT_OS_USER`
  - `CLIENT_PROGRAM`
  - `CLIENT_MACHINE`
  - `MODULE_NAME`

  - MODULE_NAME_ACTION
  - SERVICE_MODULE
  - SERVICE_MODULE_ACTION
- Consumer group mapping comparisons for DBMS_RESOURCE_MANAGER.CLIENT_PROGRAM are performed by stripping the @ sign and following characters from V$SESSION.PROGRAM before comparing it to the CLIENT_PROGRAM value supplied.

158.4.24 SET_CONSUMER_GROUP_MAPPING_PRI Procedure

Multiple attributes of a session can be used to map the session to a consumer group. This procedure prioritizes the attribute mappings.

Syntax

```
DBMS_RESOURCE_MANAGER.SET_CONSUMER_GROUP_MAPPING_PRI(
   explicit               IN NUMBER,
   oracle_user            IN NUMBER,
   service_name           IN NUMBER,
   client_os_user         IN NUMBER,
   client_program         IN NUMBER,
   client_machine         IN NUMBER,
   module_name            IN NUMBER,
   module_name_action     IN NUMBER,
   service_module         IN NUMBER,
   service_module_action  IN NUMBER,
   client_id              IN NUMBER DEFAULT 11);
```

Parameters

Table 158-22 SET_CONSUMER_GROUP_MAPPING_PRI Procedure Parameters

Parameters	Description
explicit	Priority of the explicit mapping
oracle_user	Priority of the Oracle user name mapping
service_name	Priority of the client service name mapping
client_os_user	Priority of the client operating system user name mapping
client_program	Priority of the client program mapping
client_machine	Priority of the client machine mapping
module_name	Priority of the application module name mapping
module_name_action	Priority of the application module name and action mapping
service_module	Priority of the service name and application module name mapping
module_name_action	Priority of the service name, application module name, and application action mapping
client_id	Client identifier

Usage Notes

- This procedure requires that you include the pseudo-attribute explicit as an argument. It must be set to 1. It indicates that explicit consumer group switches have the highest priority. You explicitly switch consumer groups with these package procedures:

  - DBMS_SESSION.SWITCH_CURRENT_CONSUMER_GROUP
  - DBMS_RESOURCE_MANAGER.SWITCH_CONSUMER_GROUP_FOR_SESS
  - DBMS_RESOURCE_MANAGER.SWITCH_CONSUMER_GROUP_FOR_USER
- Each priority value must be a unique integer from 1 to 11. Together, they establish an ordering where 1 is the highest priority and 11 is the lowest.

158.4.25 SET_INITIAL_CONSUMER_GROUP Procedure

This deprecated procedure sets the initial resource consumer group for a user.

The initial consumer group of a user is the consumer group to which any session created by that user initially belongs.

> **Note:**
>
> This procedure is deprecated in Release 11gR1. While the procedure remains available in the package, Initial Consumer Group is set by the session-to-consumer group mapping rules.

Syntax

```
DBMS_RESOURCE_MANAGER.SET_INITIAL_CONSUMER_GROUP (
   user             IN   VARCHAR2,
   consumer_group   IN   VARCHAR2);
```

Parameters

Table 158-23 SET_INITIAL_CONSUMER_GROUP Procedure Parameters

Parameters	Description
user	Name of the user
consumer_group	User's initial consumer group

Usage Notes

- The ADMINISTER_RESOURCE_MANAGER or the ALTER USER system privilege are required to be able to execute this procedure. The user, or PUBLIC, must be directly granted switch privilege to a consumer group before it can be set to be the user's initial consumer group. Switch privilege for the initial consumer group cannot come from a role granted to that user.

  > **Note:**
  >
  > These semantics are similar to those for ALTER USER DEFAULT ROLE.

- If the initial consumer group for a user has never been set, then the user's initial consumer group is automatically the consumer group: DEFAULT_CONSUMER_GROUP.

- DEFAULT_CONSUMER_GROUP has switch privileges granted to PUBLIC; therefore, all users are automatically granted switch privilege for this consumer group. Upon deletion of a consumer group, all users having the deleted group as their initial consumer group now have DEFAULT_CONSUMER_GROUP as their initial consumer group. All currently active sessions belonging to a deleted consumer group are switched to DEFAULT_CONSUMER_GROUP.

158.4.26 SUBMIT_PENDING_AREA Procedure

This procedure submits pending changes for the resource manager. It clears the pending area after validating and committing the changes (if valid).

> **Note:**
>
> A call to SUBMIT_PENDING_AREA may fail even if VALIDATE_PENDING_AREA succeeds. This may happen if a plan being deleted is loaded by an instance after a call to VALIDATE_PENDING_AREA, but before a call to SUBMIT_PENDING_AREA.

Syntax

```
DBMS_RESOURCE_MANAGER.SUBMIT_PENDING_AREA;
```

158.4.27 SWITCH_CONSUMER_GROUP_FOR_SESS Procedure

This procedure changes the resource consumer group of a specific session. It also changes the consumer group of any parallel execution servers that are related to the top user session. This procedure is RAC instance specific. You need to connect to the PDB in same RAC instance where the session to be switched is running, and then run this procedure.

Syntax

```
DBMS_RESOURCE_MANAGER.SWITCH_CONSUMER_GROUP_FOR_SESS (
   session_id      IN NUMBER,
   session_serial  IN NUMBER,
   consumer_group  IN VARCHAR2);
```

Parameters

Table 158-24 SWITCH_CONSUMER_GROUP_FOR_SESS Procedure Parameters

Parameter	Description
session_id	SID column from the view V$SESSION
session_serial	SERIAL# column from view V$SESSION.
consumer_group	Name of the consumer group to which to switch

158.4.28 SWITCH_CONSUMER_GROUP_FOR_USER Procedure

This procedure changes the resource consumer group for all sessions with a given user ID. It also changes the consumer group of any parallel execution servers that are related to the top user session.

Syntax

```
DBMS_RESOURCE_MANAGER.SWITCH_CONSUMER_GROUP_FOR_USER (
   user            IN VARCHAR2,
   consumer_group  IN VARCHAR2);
```

Parameters

Table 158-25 SWITCH_CONSUMER_GROUP_FOR_USER Procedure Parameters

Parameter	Description
user	Name of the user
consumer_group	Name of the consumer group to which to switch

Usage Notes

- The SWITCH_CONSUMER_GROUP_FOR_SESS Procedure and the SWITCH_CONSUMER_GROUP_FOR_USER procedures let you raise or lower the allocation of CPU resources of certain sessions or users. This provides a functionality similar to the nice command on UNIX.
- These procedures cause the session to be moved into the newly specified consumer group immediately.

158.4.29 SWITCH_PLAN Procedure

This procedure sets the current resource manager plan.

Syntax

```
DBMS_RESOURCE_MANAGER.SWITCH_PLAN(
   plan_name                    IN   VARCHAR2,
   sid                          IN   VARCHAR2 DEFAULT '*',
   allow_scheduler_plan_switches IN   BOOLEAN DEFAULT TRUE);
```

Parameters

Table 158-26 SWITCH_PLAN Procedure Parameters

Parameter	Description
plan_name	Name of the plan to which to switch. Passing in an empty string ('') for the plan_name, disables the resource manager
sid	The sid parameter is relevant only in an Oracle Real Application Clusters environment. This parameter lets you change the plan for a particular instance. Specify the sid of the instance where you want to change the plan. Or specify '*' if you want Oracle to change the plan for all instances.

Table 158-26 (Cont.) SWITCH_PLAN Procedure Parameters

Parameter	Description
allow_scheduler_plan_switches	FALSE - disables automated plan switches by the job scheduler at window boundaries. To reenable automated plan switches, switch_plan must be called again by the administrator with allow_scheduler_plan_switches set to TRUE. By default automated plan switches by the job scheduler are enabled.

158.4.30 UPDATE_CATEGORY Procedure

This procedure updates an existing resource consumer group category.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CATEGORY (
   category        IN    VARCHAR2,
   new_comment     IN    VARCHAR2  DEFAULT NULL);
```

Parameters

Table 158-27 UPDATE_CATEGORY Procedure Parameters

Parameter	Description
category	Name of consumer group category
new_comment	User comment

Usage Notes

To clear (reset to the directive's default value), use the value -1.

158.4.31 UPDATE_CDB_AUTOTASK_DIRECTIVE Procedure

This procedure updates the plan directives with regard to automated maintenance tasks in the CDB root (CDB$ROOT).

By default, all maintenance tasks occur directly in the PDBs themselves.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CDB_AUTOTASK_DIRECTIVE (
   plan                        IN    VARCHAR2,
   new_comment                 IN    VARCHAR2       DEFAULT NULL,
   new_shares                  IN    NUMBER         DEFAULT NULL,
   new_utilization_limit       IN    NUMBER         DEFAULT NULL,
   new_parallel_server_limit   IN    NUMBER         DEFAULT NULL,
   new_memory_limit            IN    NUMBER         DEFAULT NULL,
   new_memory_min              IN    NUMBER         DEFAULT NULL);
```

Parameters

Table 158-28 UPDATE_CDB_AUTOTASK_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
new_comment	New user comment
new_shares	Specifies the new share of resource allocation for CDB root's automated maintenance tasks
new_utilization_limit	Specifies the new maximum percentage of CPU that automated maintenance tasks in the CDB root can utilize
new_parallel_server_limit	Specifies the new maximum percentage of parallel_servers_target parallel servers that automated maintenance tasks in the CDB root are allowed to use
new_memory_limit	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.
new_memory_min	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.

Usage Notes

- By default for automated maintenance tasks, the values are
  - shares: -1
  - utilization_limit: 90
  - parallel_server_limit: 100
- The shares = -1 means that the automated maintenance tasks get an allocation of 20% of the system. If the user specifies the shares, it behaves the same properties as the other CDB plan directive functions. If the user does not change the shares or later changes it back to -1, autotask will get 20% of the system.
- This procedure can be run only from the CDB root.
- To clear (reset to the directive's default value), use the value -1.

158.4.32 UPDATE_CDB_DEFAULT_DIRECTIVE Procedure

This procedure updates the plan directives of the consolidation resource plan.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CDB_DEFAULT_DIRECTIVE (
   plan                         IN    VARCHAR2    DEFAULT NULL,
   new_comment                  IN    VARCHAR2    DEFAULT NULL,
   new_shares                   IN    NUMBER      DEFAULT NULL,
   new_utilization_limit        IN    NUMBER      DEFAULT NULL,
   new_parallel_server_limit    IN    NUMBER      DEFAULT NULL,
   new_memory_limit             IN    NUMBER      DEFAULT NULL,
   new_memory_min               IN    NUMBER      DEFAULT NULL);
```

Parameters

Table 158-29 UPDATE_CDB_DEFAULT_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
new_commnent	New user comment
new_shares	Specifies the share of resource allocation for the pluggable database. CPU Resource Manager and Exadata I/O Resource Manager are enabled by specifying shares for each PDB. The new_shares parameter is also used for Parallel Statement Queuing.
new_utilization_limit	Specifies the maximum percentage of CPU that the pluggable database can utilize.
new_parallel_server_limit	Specifies the maximum percentage of parallel_servers_target parallel servers that the pluggable database can use.
new_memory_limit	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.
new_memory_min	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.

Usage Notes

- By default, the default values are
  - new_shares: 1
  - utilization_limit: 100
  - parallel_server_limit: 100
- Note that the default values are NULL. This has the same meaning as in UPDATE_CDB_PLAN_DIRECTIVE Procedure. If the user does not specify a value, the value will not be modified.
- This procedure can be run only from the CDB root (CDB$ROOT).
- To clear (reset to the directive's default value), use the value -1.

158.4.33 UPDATE_CDB_PLAN Procedure

This procedure updates the consolidation resource plan.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CDB_PLAN (
   plan                        IN    VARCHAR2(32),
   new_comment                 IN    VARCHAR2(2000)   DEFAULT NULL);
```

Parameters

Table 158-30 UPDATE_CDB_PLAN Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
new_comment	User comment

Usage Notes

- This procedure can be run only from the CDB root (CDB$ROOT).
- To clear (reset to the directive's default value), use the value -1.

158.4.34 UPDATE_CDB_PLAN_DIRECTIVE Procedure

Updates the plan directives for a consolidation resource plan. Plan directives specify the resource allocation policy for pluggable databases (PDBs).

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CDB_PLAN_DIRECTIVE (
   plan                        IN    VARCHAR2 (30),
   pluggable_database          IN    VARCHAR2 (30)
   new_comment                 IN    VARCHAR2 (200)  DEFAULT NULL,
   new_shares                  IN    NUMBER          DEFAULT NULL,
   new_utilization_limit       IN    NUMBER          DEFAULT NULL,
   new_parallel_server_limit   IN    NUMBER          DEFAULT NULL,
   new_memory_limit            IN    NUMBER          DEFAULT NULL,
   new_memory_min              IN    NUMBER          DEFAULT NULL);
```

Parameters

Table 158-31 UPDATE_CDB_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
pluggable_database	Name of the pluggable database
new_comment	New user comment
new_shares	The share of resource allocation for the pluggable database CPU Resource Manager is enabled by specifying shares for each PDB. The shares parameter is also used for Parallel Statement Queuing. If no share is specified, the default is obtained from the default directive, specified through the UPDATE_CDB_DEFAULT_DIRECTIVE Procedure.
new_utilization_limit	The new maximum percentage of CPU that the pluggable database can utilize
new_parallel_server_limit	The new maximum percentage of parallel_servers_target parallel servers that the pluggable database can use

Table 158-31 (Cont.) UPDATE_CDB_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
new_memory_limit	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.
new_memory_ min	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.

Usage Notes

- The default value for the new_* parameters is NULL which indicates that the existing value is left unchanged. If the user does not specify one of the arguments when calling this function, the value is not modified.
- This procedure can be run only from the CDB root (CDB$ROOT).
- To clear (reset to the directive's default value), use the value -1.

158.4.35 UPDATE_CDB_PROFILE_DIRECTIVE Procedure

This procedure updates the performance profile directives of the consolidation resource plan. The directives specify the resource allocation policy for pluggable databases (PDBs) that use the performance profile.

For a PDB to use a performance profile, the PDB must have the DB_PERFORMANCE_PROFILE initialization parameter set to the performance profile name.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CDB_PROFILE_DIRECTIVE (
   plan                        IN    VARCHAR2,
   profile                     IN    VARCHAR2,
   new_comment                 IN    VARCHAR2 (2000) DEFAULT '',
   new_shares                  IN    NUMBER          DEFAULT NULL,
   new_utilization_limit       IN    NUMBER          DEFAULT NULL,
   new_parallel_server_limit   IN    NUMBER          DEFAULT NULL,
   new_memory_limit            IN    NUMBER          DEFAULT 100,
   new_memory_min              IN    NUMBER          DEFAULT 0);
```

Parameters

Table 158-32 UPDATE_CDB_PROFILE_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the consolidation plan
profile	Name of the performance profile
new_comment	New user comment
new_shares	The share of resource allocation for the PDBs that use the performance profile
new_utilization_limit	The new maximum percentage of CPU that PDBs that use the performance profile can use

Table 158-32 (Cont.) UPDATE_CDB_PROFILE_DIRECTIVE Procedure Parameters

Parameter	Description
new_parallel_server_limit	The new maximum percentage of parallel_servers_target parallel servers that PDBs that use the performance profile can use
new_memory_limit	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.
new_memory_ min	This parameter is only applicable to Oracle Exadata storage for configuring the Database Smart Flash Cache and PMEM Cache.

Usage Notes

- This procedure can be run only from the CDB root (CDB$ROOT).
- To clear (reset to the directive's default value), use the value -1.

158.4.36 UPDATE_CONSUMER_GROUP Procedure

This procedure updates entries which define resource consumer groups.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_CONSUMER_GROUP (
   consumer_group  IN VARCHAR2,
   new_comment     IN VARCHAR2 DEFAULT NULL,
   new_cpu_mth     IN VARCHAR2 DEFAULT NULL,
   new_mgmt_mth    IN VARCHAR2 DEFAULT NULL,
   new_category    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 158-33 UPDATE_CONSUMER_GROUP Procedure Parameter

Parameter	Description
consumer_group	Name of consumer group
new_comment	New user comment
new_cpu_mth	Name of new method for CPU resource allocation (deprecated)
new_mgmt_mth	Name of new method for CPU resource allocation
new_category	New consumer group category

Usage Notes

- If the parameters to the UPDATE_CONSUMER_GROUP procedure are not specified, then they remain unchanged in the data dictionary.
- To clear (reset to the directive's default value), use the value -1.

158.4.37 UPDATE_PLAN Procedure

This procedure updates entries which define resource plans.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_PLAN (
   plan                                IN VARCHAR2,
   new_comment                         IN VARCHAR2 DEFAULT NULL,
   new_cpu_mth                         IN VARCHAR2 DEFAULT NULL, -- deprecated
   new_active_sess_pool_mth            IN VARCHAR2 DEFAULT NULL,
   new_parallel_degree_limit_mth       IN VARCHAR2 DEFAULT NULL,
   new_queueing_mth                    IN VARCHAR2 DEFAULT NULL,
   new_mgmt_mth                        IN VARCHAR2 DEFAULT NULL,
   new_sub_plan                        IN BOOLEAN DEFAULT FALSE,
   new_max_iops                        IN NUMBER DEFAULT NULL,
   new_max_mbps                        IN NUMBER DEFAULT NULL);
```

Parameters

Table 158-34 UPDATE_PLAN Procedure Parameters

Parameter	Description
`plan`	Name of resource plan
`new_comment`	New user comment
`new_cpu_mth`	Name of new allocation method for CPU resources (deprecated)
`new_active_sess_pool_mth`	Name of new method for maximum active sessions
`new_parallel_degree_limit_mth`	Name of new method for degree of parallelism
`new_queueing_mth`	Specifies type of queuing policy to use with active session pool feature
`new_mgmt_mth`	Resource allocation method for specifying how much resources (for example, CPU or I/O) each consumer group or sub-plan gets • `EMPHASIS` - for multilevel plans that use percentages to specify how I/O resources are distributed among consumer groups. • `RATIO` - for single-level plans that use ratios to specify how I/O resources are distributed.
`new_sub_plan`	New setting for whether the plan is only intended for use as a sub-plan
`new_max_iops`	Nonoperative
`new_max_mbps`	Nonoperative

Usage Notes

- If the parameters to UPDATE_PLAN Procedure are not specified, then they remain unchanged in the data dictionary.
- If you want to use any default resource allocation method, then you do not need to specify it when creating or updating a plan.
- To clear (reset to the directive's default value), use the value `-1`.

158.4.38 UPDATE_PLAN_DIRECTIVE Procedure

This procedure updates resource plan directives.

> **Note:**
>
> The parameters `new_max_utilization_limit` and `new_parallel_target_percentage` are deprecated with Oracle Database 11g Release 1 (12.1.0.1), and are replaced by `new_utilization_limit` and `new_parallel_server_limit`.

Syntax

```
DBMS_RESOURCE_MANAGER.UPDATE_PLAN_DIRECTIVE (
   plan                              IN VARCHAR2,
   group_or_subplan                  IN VARCHAR2,
   new_comment                       IN VARCHAR2 DEFAULT NULL,
   new_cpu_p1                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p2                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p3                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p4                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p5                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p6                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p7                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_cpu_p8                        IN NUMBER   DEFAULT NULL, -- deprecated
   new_active_sess_pool_p1           IN NUMBER   DEFAULT NULL,
   new_queueing_p1                   IN NUMBER   DEFAULT NULL,
   new_parallel_degree_limit_p1      IN NUMBER   DEFAULT NULL,
   new_switch_group                  IN VARCHAR2 DEFAULT NULL,
   new_switch_time                   IN NUMBER   DEFAULT NULL,
   new_switch_estimate               IN BOOLEAN  DEFAULT FALSE,
   new_max_est_exec_time             IN NUMBER   DEFAULT NULL,
   new_undo_pool                     IN NUMBER   DEFAULT NULL,
   new_max_idle_time                 IN NUMBER   DEFAULT NULL,
   new_max_idle_blocker_time         IN NUMBER   DEFAULT NULL,
   switch_time_in_call               IN NUMBER   DEFAULT NULL, -- deprecated
   new_mgmt_p1                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p2                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p3                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p4                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p5                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p6                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p7                       IN NUMBER   DEFAULT NULL,
   new_mgmt_p8                       IN NUMBER   DEFAULT NULL,
   new_switch_io_megabytes           IN NUMBER   DEFAULT NULL,
   new_switch_io_reqs                IN NUMBER   DEFAULT NULL,
   new_switch_for_call               IN BOOLEAN  DEFAULT NULL,
   new_max_utilization_limit         IN NUMBER   DEFAULT NULL,
   new_parallel_target_percentage    IN NUMBER   DEFAULT NULL,
   new_parallel_queue_timeout        IN NUMBER   DEFAULT NULL,
   new_parallel_server_limit         IN NUMBER   DEFAULT NULL,
   new_utilization_limit             IN NUMBER   DEFAULT NULL,
   new_switch_io_logical             IN NUMBER   DEFAULT NULL,
   new_switch_elapsed_time           IN NUMBER   DEFAULT NULL,
   new_shares                        IN NUMBER   DEFAULT NULL,
   new_parallel_stmt_critical        IN VARCHAR2 DEFAULT NULL,
```

```
  new_session_pga_limit            IN NUMBER   DEFAULT NULL,
  new_pq_timeout_action            IN NUMBER   DEFAULT NULL);
```

Parameters

Table 158-35 UPDATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
plan	Name of the resource plan
group_or_subplan	Name of the consumer group or subplan
new_comment	Comment for the plan directive
new_cpu_p1	Deprecated - use new_mgmt_p1 instead
new_cpu_p2	Deprecated - use new_mgmt_p2 instead
new_cpu_p3	Deprecated - use new_mgmt_p3 instead
new_cpu_p4	Deprecated- use new_mgmt_p4 instead
new_cpu_p5	Deprecated - use new_mgmt_p5 instead
new_cpu_p6	Deprecated- use new_mgmt_p6 instead
new_cpu_p7	Deprecated- use new_mgmt_p7 instead
new_cpu_p8	Deprecated- use new_mgmt_p8 instead
new_active_sess_pool_p1	Specifies maximum number of concurrently active sessions for a consumer group. Default is NULL, which means unlimited.
new_queueing_p1	Specified time (in seconds) after which a job in the inactive session queue (waiting for execution) will time out. Default is NULL, which means unlimited.
new_parallel_degree_limit_p1	Specifies a limit on the degree of parallelism for any operation. Default is NULL, which means unlimited.
new_switch_group	Specifies consumer group to which this session is switched if other switch criteria are met. Default is NULL. If the group name is 'CANCEL_SQL', the current call will be canceled when other switch criteria are met. If the group name is 'KILL_SESSION', the session will be killed when other switch criteria are met.
new_switch_time	Specifies time (in CPU seconds) that a session can execute before an action is taken. Default is NULL, which means unlimited.
new_switch_estimate	If TRUE, tells Oracle to use its execution time estimate to automatically switch the consumer group of an operation before beginning its execution. Default is FALSE.
new_max_est_exec_time	Specifies the maximum execution time (in CPU seconds) allowed for a session. If the optimizer estimates that an operation will take longer than MAX_EST_EXEC_TIME, the operation is not started and ORA-07455 is issued. If the optimizer does not provide an estimate, this directive has no effect. Default is NULL, which means unlimited.
new_undo_pool	Limits the size in kilobytes of the undo records corresponding to uncommitted transactions by this consumer group
new_max_idle_time	Indicates the maximum session idle time. Default is NULL, which means unlimited.
new_max_idle_blocker_time	Maximum amount of time in seconds that a session can be idle while blocking another session's acquisition of a resource

Table 158-35 (Cont.) UPDATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
new_switch_time_in_call	Deprecated. If this parameter is specified, new_switch_time will be effectively set to new_switch_time_in_call and new_switch_for_call will be effectively set to TRUE.
new_mgmt_p1	Resource allocation value for level 1 (replaces new_cpu_p1): • EMPHASIS - specifies the resource percentage at the first level • RATIO - specifies the weight of resource usage
new_mgmt_p2	Resource allocation value for level 2 (replaces new_cpu_p2) • EMPHASIS - specifies the resource percentage at the second level • RATIO - non-applicable
new_mgmt_p3	Resource allocation value for level 3 (replaces new_cpu_p3) • EMPHASIS - specifies the resource percentage at the third level • RATIO - non-applicable
new_mgmt_p4	Resource allocation value for level 4 (replaces new_cpu_p4) • EMPHASIS - specifies the resource percentage at the fourth level • RATIO - non-applicable
new_mgmt_p5	Resource allocation value for level 5 (replaces new_cpu_p5) • EMPHASIS - specifies the resource percentage at the fifth level • RATIO - non-applicable
new_mgmt_p6	Resource allocation value for level 6 (replaces new_cpu_p6) • EMPHASIS - specifies the resource percentage at the sixth level • RATIO - non-applicable
new_mgmt_p7	Resource allocation value for level 7 (replaces new_cpu_p7) • EMPHASIS - specifies the resource percentage at the seventh level • RATIO - non-applicable
new_mgmt_p8	Resource allocation value for level 8 (replaces new_cpu_p8) • EMPHASIS - specifies the resource percentage at the eighth level • RATIO - non-applicable
new_switch_io_megabytes	Specifies the amount of I/O (in MB) that a session can issue before an action is taken. Default is NULL, which means unlimited.
new_switch_io_reqs	Specifies the number of I/O requests that a session can issue before an action is taken. Default is NULL, which means unlimited.
new_switch_for_call	Specifies that if an action is taken because of the new_switch_time, new_switch_io_megabytes, or new_switch_io_reqs parameters, the consumer group is restored to its original consumer group at the end of the top call. Default is FALSE, which means that the original consumer group is not restored at the end of the top call.
new_max_utilization_limit	Deprecated - use new_utilization_limit instead
new_parallel_target_percentage	Deprecated - use new_parallel_server_limit instead
new_parallel_server_limit	Parallel server limit. Setting this overwrites the limit for parallel server set by utilization_limit.

Table 158-35 (Cont.) UPDATE_PLAN_DIRECTIVE Procedure Parameters

Parameter	Description
new_utilization_limit	Resource limit. For CPU, this limits the CPU utilization for the consumer group. For parallel servers, this limits the parallel servers used as a percentage of parallel_servers_target.
new_switch_elapsed_time	Elapsed time that will trigger the action specified by switch_group. As with other switch directives, if new_switch_for_call is TRUE, the elapsed time is accumulated from the start of a call. Otherwise, the elapsed time is accumulated for the length of the session.
new_shares	Specifies the share of resource allocation for the pluggable database. CPU Resource Manager and Exadata I/O Resource Manager are enabled by specifying shares for each PDB. The shares parameter is also used for Parallel Statement Queuing. If CPU Resource Manager and Exadata I/O Resource Manager are enabled, then the default value is 1.
new_parallel_stmt_critical	If set to BYPASS_QUEUE, parallel statements from this consumer group are not queued. If set to QUEUE, all the parallel statements, irrespective of the parallel_degree_policy parameter value, from the consumer group get queued. Default is FALSE, which means that certain parallel statements are eligible for queuing depending upon the parallel_degree_policy parameter value.
new_session_pga_limit	Maximum amount of PGA in MB that sessions in this consumer group can allocate before being terminated. NULL (default) indicates no change.
new_parallel_queue_timeout	Specifies the time (in seconds) that a parallel statement may remain in its Consumer Group's parallel statement queue before it is removed. The default action of this parameter is ERROR. This action can be altered using the new_pq_timeout_action parameter.
new_pq_timeout_action	Specifies the action to be taken when a parallel statement is removed from the queue due to new_parallel_queue_timeout. The values are: • CANCEL — The SQL statement is terminated with error ORA-7454 • RUN — The SQL statement runs immediately, and might get downgraded if parallel servers are unavailable

Usage Notes

- If the parameters for UPDATE_PLAN_DIRECTIVE are left unspecified, then they remain unchanged in the data dictionary.
- For new_max_idle_time and new_max_idle_blocker_time, PMON will check these limits once a minute. If it finds a session that has exceeded one of the limits, it will forcibly kill the session and clean up all its state.
- The parameter new_switch_time_in_call is mostly useful for three-tier applications where the mid-tier server is implementing session pooling. By turning on new_switch_time_in_call, the resource usage of one client will not affect the consumer group of a future client that happens to be executed on the same session.
- To clear (reset to the directive's default value), use the value -1.

158.4.39 VALIDATE_PENDING_AREA Procedure

This procedure validates pending changes for the resource manager.

Syntax

```
DBMS_RESOURCE_MANAGER.VALIDATE_PENDING_AREA;
```

159
DBMS_RESOURCE_MANAGER_PRIVS

The `DBMS_RESOURCE_MANAGER_PRIVS` package maintains privileges associated with the Resource Manager.

This chapter contains the following topics:

- Summary of DBMS_RESOURCE_MANAGER_PRIVS Subprograms

> **See Also:**
>
> For more information on using the Database Resource Manager, see*Oracle Database Administrator's Guide*

159.1 Summary of DBMS_RESOURCE_MANAGER_PRIVS Subprograms

This table lists the `DBMS_RESOURCE_MANAGER_PRIVS` subprograms and briefly describes them.

Table 159-1 DBMS_RESOURCE_MANAGER_PRIVS Package Subprograms

Subprogram	Description
GRANT_SWITCH_CONSUMER_GROUP Procedure	Grants the privilege to switch to resource consumer groups
GRANT_SYSTEM_PRIVILEGE Procedure	Performs a grant of a system privilege
REVOKE_SWITCH_CONSUMER_GROUP Procedure	Revokes the privilege to switch to resource consumer groups.
REVOKE_SYSTEM_PRIVILEGE Procedure	Performs a revoke of a system privilege

159.1.1 GRANT_SWITCH_CONSUMER_GROUP Procedure

This procedure grants the privilege to switch to a resource consumer group.

Syntax

```
DBMS_RESOURCE_MANAGER_PRIVS.GRANT_SWITCH_CONSUMER_GROUP (
   grantee_name   IN VARCHAR2,
   consumer_group IN VARCHAR2,
   grant_option   IN BOOLEAN);
```

Parameters

Table 159-2 GRANT_SWITCH_CONSUMER_GROUP Procedure Parameters

Parameter	Description
grantee_name	Name of the user or role to whom privilege is to be granted.
consumer_group	Name of consumer group.
grant_option	TRUE if grantee should be allowed to grant access, FALSE otherwise.

Usage Notes

If you grant permission to switch to a particular consumer group to a user, then that user can immediately switch their current consumer group to the new consumer group.

If you grant permission to switch to a particular consumer group to a role, then any users who have been granted that role and have enabled that role can immediately switch their current consumer group to the new consumer group.

If you grant permission to switch to a particular consumer group to PUBLIC, then any user can switch to that consumer group.

If the grant_option parameter is TRUE, then users granted switch privilege for the consumer group may also grant switch privileges for that consumer group to others.

In order to set the initial consumer group of a user, you must grant the switch privilege for that group to the user.

> **See Also:**
>
> DBMS_RESOURCE_MANAGER

Examples

```
BEGIN
DBMS_RESOURCE_MANAGER_PRIVS.GRANT_SWITCH_CONSUMER_GROUP (
 'scott', 'mail_maintenance_group', true);
DBMS_RESOURCE_MANAGER.CREATE_PENDING_AREA();
DBMS_RESOURCE_MANAGER.set_consumer_group_mapping(
  dbms_resource_manager.oracle_user, 'scott','mail_maintenance_group');
DBMS_RESOURCE_MANAGER.SUBMIT_PENDING_AREA();
END;
/
```

159.1.2 GRANT_SYSTEM_PRIVILEGE Procedure

This procedure performs a grant of a system privilege to a user or role.

Syntax

```
DBMS_RESOURCE_MANAGER_PRIVS.GRANT_SYSTEM_PRIVILEGE (
   grantee_name    IN VARCHAR2,
   privilege_name  IN VARCHAR2 DEFAULT 'ADMINISTER_RESOURCE_MANAGER',
   admin_option    IN BOOLEAN);
```

Parameters

Table 159-3 GRANT_SYSTEM_PRIVILEGE Procedure Parameters

Parameter	Description
grantee_name	Name of the user or role to whom privilege is to be granted.
privilege_name	Name of the privilege to be granted.
admin_option	TRUE if the grant is with admin_option, FALSE otherwise.

Usage Notes

Currently, Oracle provides only one system privilege for the Resource Manager: ADMINISTER_RESOURCE_MANAGER. Database administrators have this system privilege with the ADMIN option. The grantee and the revokee can either be a user or a role. Users that have been granted the system privilege with the ADMIN option can also grant this privilege to others.

Examples

The following call grants this privilege to a user called scott without the ADMIN option:

```
BEGIN
DBMS_RESOURCE_MANAGER_PRIVS.GRANT_SYSTEM_PRIVILEGE (
 grantee_name => 'scott',
 privilege_name => 'ADMINISTER_RESOURCE_MANAGER',
 admin_option => FALSE);
END;
/
```

159.1.3 REVOKE_SWITCH_CONSUMER_GROUP Procedure

This procedure revokes the privilege to switch to a resource consumer group.

Syntax

```
DBMS_RESOURCE_MANAGER_PRIVS.REVOKE_SWITCH_CONSUMER_GROUP (
   revokee_name   IN VARCHAR2,
   consumer_group IN VARCHAR2);
```

Parameters

Table 159-4 REVOKE_SWITCH_CONSUMER_GROUP Procedure Parameter

Parameter	Description
revokee_name	Name of user/role from which to revoke access.
consumer_group	Name of consumer group.

Usage Notes

If you revoke a user's switch privilege for a particular consumer group, then any subsequent attempts by that user to switch to that consumer group will fail.

If you revoke the initial consumer group from a user, then that user will automatically be part of the DEFAULT_CONSUMER_GROUP consumer group when logging in.

If you revoke the switch privilege for a consumer group from a role, then any users who only had switch privilege for the consumer group through that role will not be able to switch to that consumer group.

If you revoke the switch privilege for a consumer group from `PUBLIC`, then any users who could previously only use the consumer group through `PUBLIC` will not be able to switch to that consumer group.

Examples

The following example revokes the privileges to switch to `mail_maintenance_group` from Scott:

```
BEGIN
DBMS_RESOURCE_MANAGER_PRIVS.REVOKE_SWITCH_CONSUMER_GROUP (
 'scott', 'mail_maintenance_group');
END;
/
```

159.1.4 REVOKE_SYSTEM_PRIVILEGE Procedure

This procedure performs a revoke of a system privilege from a user or role.

Syntax

```
DBMS_RESOURCE_MANAGER_PRIVS.REVOKE_SYSTEM_PRIVILEGE (
   revokee_name   IN VARCHAR2,
   privilege_name IN VARCHAR2 DEFAULT 'ADMINISTER_RESOURCE_MANAGER');
```

Parameters

Table 159-5 REVOKE_SYSTEM_PRIVILEGE Procedure Parameters

Parameter	Description
`revokee_name`	Name of the user or role from whom privilege is to be revoked.
`privilege_name`	Name of the privilege to be revoked.

Examples

The following call revokes the `ADMINISTER_RESOURCE_MANAGER` from user scott:

```
BEGIN
DBMS_RESOURCE_MANAGER_PRIVS.REVOKE_SYSTEM_PRIVILEGE ('scott');
END;
/
```

160
DBMS_RESULT_CACHE

The DBMS_RESULT_CACHE package provides an interface to allow the DBA to administer that part of the shared pool that is used by the SQL result cache and the PL/SQL function result cache.

Both these caches use the same infrastructure. Therefore, for example, DBMS_RESULT_CACHE.BYPASS determines whether both caches are bypassed or both caches are used, and DBMS_RESULT_CACHE.FLUSH flushes both all the cached results for SQL queries and all the cached results for PL/SQL functions.

This chapter contains the following topics:

- Security Model
- Constants
- Summary of DBMS_RESULT_CACHE Subprograms

> **See Also:**
>
> - *Oracle Database Performance Tuning Guide*, for more information about "Result Cache Concepts"
> - *Oracle Database PL/SQL Language Reference*, for more information about PL/SQL function Result Cache "Using the Cross-Session PL/SQL Function Result Cache"
> - *Database PL/SQL Language Reference*, for more information about Result Cache Management.

160.1 DBMS_RESULT_CACHE Security Model

Only database administrators should be granted the EXECUTE privilege for this package.

160.2 DBMS_RESULT_CACHE Constants

The DBMS_RESULT_CACHE package defines several constants for specifying parameter values.

The following table describes these constants.

Table 160-1 DBMS_RESULT_CACHE Constants

Constant	Definition	Description
STATUS_BYPS	CONSTANT VARCHAR2(10) := 'BYPASS';	Cache has been made temporarilyunavailable.
STATUS_CORR	CONSTANT VARCHAR2(10) := 'CORRUPT';	The result cache is in an unusable state.

Table 160-1 (Cont.) DBMS_RESULT_CACHE Constants

Constant	Definition	Description
STATUS_DISA	CONSTANT VARCHAR2(10) := 'DISABLED';	Cache is not available.
STATUS_ENAB	CONSTANT VARCHAR2(10) := 'ENABLED';	Cache is available.
STATUS_SYNC	CONSTANT VARCHAR2(10) := 'SYNC';	Cache is available, but synchronizing with Oracle RAC nodes.

160.3 Summary of DBMS_RESULT_CACHE Subprograms

This table lists the DBMS_RESULT_CACHE subprograms and briefly describes them.

Table 160-2 DBMS_RESULT_CACHE Package Subprograms

Subprogram	Description
BLOCK_LIST Function	Returns all the block listed cache ids of a local instance.
BLOCK_LIST_ADD Procedure	Adds a cache_id to the block list.
BLOCK_LIST_CLEAR Procedure	Removes all cache_ids from the block list.
BLOCK_LIST_REMOVE Procedure	Removes the cache_id from the block list.
BYPASS Procedure	Sets the bypass mode for the result cache.
FLUSH Function & Procedure	Attempts to remove all the objects from the result cache, and depending on the arguments retains or releases the memory and retains or clears the statistics.
IGNORE_LIST Function	This table function returns a row for each entry that's been added into the ignore list.
IGNORE_LIST_ADD Procedure	This procedure adds to the ignore list a table specified by object number or by owner and table name.
IGNORE_LIST_CLEAR Procedure	This procedure drops the entire ignore list.
IGNORE_LIST_REMOVE Procedure	This procedure removes a table specified by object number or by owner and table name.
INVALIDATE Functions & Procedures	Invalidates all the result-set objects that are dependent upon the specified dependency object.
INVALIDATE_OBJECT Functions & Procedures	Invalidates the specified result-set object(s)
MEMORY_REPORT Procedure	Produces the memory usage report for the result cache.
OBJECT_BLOCK_LIST Function	Returns all the object numbers currently in the block list
OBJECT_BLOCK_LIST_ADD Procedure	Adds objects to the object block list
OBJECT_BLOCK_LIST_CLEAR Procedure	Clears the entire block list
OBJECT_BLOCK_LIST_REMOVE Procedure	Removes objects from the object block list.
STATUS Function	Checks the status of the result cache.

160.3.1 BLACK_LIST Function

This procedure is deprecated with Oracle Database 23ai. Use the BLOCK_LIST Function instead.

160.3.2 BLACK_LIST_ADD Procedure

This procedure is deprecated with Oracle Database 23ai. Use the BLOCK_LIST_ADD Procedure instead.

160.3.3 BLACK_LIST_CLEAR Procedure

This procedure is deprecated with Oracle Database 23ai. Use the BLOCK_LIST_CLEAR Procedure instead.

160.3.4 BLACK_LIST_REMOVE Procedure

This procedure is deprecated with Oracle Database 23ai. Use the BLOCK_LIST_REMOVE Procedure instead.

160.3.5 BYPASS Procedure

This procedure sets the bypass mode for the Result Cache.

It sets one of the following bypass modes:

- When bypass mode is turned on, it implies that cached results are no longer used and that no new results are saved in the cache.
- When bypass mode is turned off, the cache resumes normal operation.

Syntax

```
DBMS_RESULT_CACHE.BYPASS (
   bypass_mode    IN   BOOLEAN,
   session        IN   BOOLEAN);
```

Parameters

Table 160-3 BYPASS Procedure Parameters

Parameter	Description
`bypass_mode`	• `TRUE` => Result Cache usage is bypassed • `FALSE` => Result Cache usage is turned on
`session`	• `TRUE` => Applies to current session • `FALSE` (default) => Applies to all sessions

Usage Notes

This operation is database instance specific.

Examples

This operation can be used when there is a need to hot patch PL/SQL code in a running system. If a code-patch is applied to a PL/SQL module on which a result cached function directly or transitively depends, then the cached results associated with the result cache function are not automatically flushed (if the instance is not restarted/bounced). This must be manually achieved.

To ensure correctness during the patching process follow these steps:

1. Place the result cache in bypass mode, and flush existing result.

```
BEGIN
   DBMS_RESULT_CACHE.BYPASS(TRUE);
   DBMS_RESULT_CACHE.FLUSH;
END;
/
```

   This step must be performed on each instance if in a Oracle Real Application Clusters environment.

2. Apply the PL/SQL code patches.
3. Resume use of the result cache, by turning off the cache bypass mode.

```
BEGIN
   DBMS_RESULT_CACHE.BYPASS(FALSE);
END;
/
```

   This step must be performed on each instance if in a Oracle Real Application Clusters environment.

160.3.6 FLUSH Function & Procedure

This function and procedure attempts to remove all the objects from the Result Cache, and depending on the arguments retains or releases the memory and retains or clears the statistics.

Syntax

```
DBMS_RESULT_CACHE.FLUSH (
   retainMem  IN  BOOLEAN DEFAULT FALSE,
   retainSta  IN  BOOLEAN DEFAULT FALSE)
  RETURN BOOLEAN;

DBMS_RESULT_CACHE.FLUSH (
   retainMem  IN  BOOLEAN DEFAULT FALSE,
   retainSta  IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 160-4 FLUSH Function & Procedure Parameters

Parameter	Description
retainMem	• TRUE => retains the free memory in the cache • FALSE (default) => releases the free memory to the system

Table 160-4 (Cont.) FLUSH Function & Procedure Parameters

Parameter	Description
retainSta	• TRUE => retains the existing cache statistics • FALSE (default) => clears the existing cache statistics

Return Values

TRUE if successful in removing all the objects.

160.3.7 IGNORE_LIST Function

This table function returns a row for each entry that's been added into the ignore list.

Objects in the ignore list are never considered for auto block listing. This allows users to override auto block listing.

Syntax

```
type ign_recT is record( objNo NUMBER);
type ign_tabT is table of ign_recT;
DBMS_RESULT_CACHE.IGNORE_LIST ()
   RETURN ign_tabT pipelined;
```

160.3.8 IGNORE_LIST_ADD Procedure

This procedure adds to the ignore list a table specified by object number or by owner and table name.

Objects in the ignore list are never considered for auto block listing. This allows users to override auto block listing.

Syntax

```
DBMS_RESULT_CACHE.IGNORE_LIST_ADD (
   objNo    IN  NATURALN);
```

Syntax

```
DBMS_RESULT_CACHE.IGNORE_LIST_ADD (
   owner    IN  VARCHAR2,
   name     IN  VARCHAR2);
```

Parameters

Table 160-5 IGNORE_LIST_ADD Procedure Parameters

Parameter	Description
objNo	The objNo is added to the ignore list.

Table 160-5 (Cont.) IGNORE_LIST_ADD Procedure Parameters

Parameter	Description
name	The table specified by owner and name are added to the ignore list.
owner	The table specified by owner and name are added to the ignore list.

160.3.9 IGNORE_LIST_CLEAR Procedure

This procedure drops the entire ignore list.

Objects in the ignore list are never considered for auto block listing. This allows users to override auto block listing.

Syntax

```
DBMS_RESULT_CACHE.IGNORE_LIST_CLEAR;
```

160.3.10 IGNORE_LIST_REMOVE Procedure

This procedure removes a table specified by object number or by owner and table name.

Syntax

```
DBMS_RESULT_CACHE.IGNORE_LIST_REMOVE (
   objNo    IN  NATURALN);
```

Syntax

```
DBMS_RESULT_CACHE.IGNORE_LIST_REMOVE (
   owner    IN  VARCHAR2,
   name     IN  VARCHAR2);
```

Parameters

Table 160-6 IGNORE_LIST_REMOVE Procedure Parameters

Parameter	Description
objNo	The objNo is removed from the ignore list.
name	The table specified by owner and name is removed to the ignore list.
owner	The table specified by owner and name is removed to the ignore list.

160.3.11 INVALIDATE Functions & Procedures

This function and procedure invalidates all the result-set objects that dependent upon the specified dependency object.

Syntax

```
DBMS_RESULT_CACHE.INVALIDATE (
   owner        IN  VARCHAR2,
   name         IN  VARCHAR2)
 RETURN NUMBER;

DBMS_RESULT_CACHE.INVALIDATE (
   owner       IN  VARCHAR2,
   name        IN  VARCHAR2);

DBMS_RESULT_CACHE.INVALIDATE (
  object_id    IN BINARY_INTEGER)
 RETURN NUMBER;

DBMS_RESULT_CACHE.INVALIDATE (
   object_id    IN BINARY_INTEGER);
```

Parameters

Table 160-7 INVALIDATE Function & Procedure Parameters

Parameter	Description
owner	Schema name
name	Object name
object_id	Dictionary object number

Return Values

The number of objects invalidated.

160.3.12 INVALIDATE_OBJECT Functions & Procedures

This function and procedure invalidates the specified result-set object(s).

Syntax

```
DBMS_RESULT_CACHE.INVALIDATE_OBJECT (
   id          IN  BINARY_INTEGER)
 RETURN NUMBER;

DBMS_RESULT_CACHE.INVALIDATE_OBJECT (
   id          IN  BINARY_INTEGER);

DBMS_RESULT_CACHE.INVALIDATE_OBJECT (
  cache_id     IN  VARCHAR2)
 RETURN NUMBER;

DBMS_RESULT_CACHE.INVALIDATE_OBJECT (
   cache_id   IN  VARCHAR2);
```

Parameters

Table 160-8 INVALIDATE_OBJECT Function & Procedure Parameters

Parameter	Description
id	Address of the cache object in the Result Cache
cache_id	Result cache identifier of a SQL cursor or PL/SQL function.

Return Values

The number of objects invalidated.

160.3.13 MEMORY_REPORT Procedure

This procedure produces the memory usage report for the Result Cache.

Syntax

```
DBMS_RESULT_CACHE.MEMORY_REPORT (
   detailed   IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 160-9 MEMORY_REPORT Procedure Parameters

Parameter	Description
detailed	• TRUE => produces a more detailed report • FALSE (default) => produces the standard report

Usage Notes

Invoking this procedure from SQL*Plus requires that the serveroutput be turned on.

Examples

```
SET SERVEROUTPUT ON

EXECUTE DBMS_RESULT_CACHE.MEMORY_REPORT;
```

160.3.14 OBJECT_BLACK_LIST Function

This function is deprecated with Oracle Database 23ai. Use the OBJECT_BLOCK_LIST Function instead.

160.3.15 OBJECT_BLACK_LIST_ADD Procedure

This procedure is deprecated with Oracle Database 23ai. Use the OBJECT_BLOCK_LIST_ADD Procedure instead.

160.3.16 OBJECT_BLACK_LIST_CLEAR Procedure

This procedure is deprecated with Oracle Database 23ai. Use the OBJECT_BLOCK_LIST_CLEAR Procedure instead.

160.3.17 OBJECT_BLACK_LIST_REMOVE Procedure

This procedure is removed with Oracle Database 23ai. Use OBJECT_BLOCK_LIST_REMOVE Procedure instead.

160.3.18 STATUS Function

This function checks the status of the Result Cache.

Syntax

```
DBMS_RESULT_CACHE.STATUS
   RETURN VARCHAR2;
```

> **Note:**
>
> For more information on the constants, see DBMS_RESULT_CACHE Constants.

160.3.19 BLOCK_LIST Function

This table function returns all the block-listed cache ids of a local instance.

Syntax

```
DBMS_RESULT_CACHE.BLOCK_LIST
   RETURN BL_TABT;
```

160.3.20 BLOCK_LIST_ADD Procedure

This procedure adds a cache_id to the block list.

Syntax

```
DBMS_RESULT_CACHE.BLOCK_LIST_ADD (
   cache_id   IN  VARCHAR2,
   global     IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 160-10 BLOCK_LIST_ADD Procedure Parameters

Parameter	Description
`cache_id`	The `cache_id` is added to the block list.
`global`	`TRUE` — applies to all caches in a RAC cluster. `FALSE` (default) — applies only to the local instance cache.

160.3.21 BLOCK_LIST_CLEAR Procedure

This procedure removes all cache_ids from the block list.

Syntax

```
DBMS_RESULT_CACHE.BLOCK_LIST_CLEAR (
   global    IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 160-11 BLOCK_LIST_CLEAR Procedure Parameters

Parameter	Description
`global`	`TRUE` — applies to all caches in a RAC cluster. `FALSE` (default) — applies only to the local instance cache.

160.3.22 BLOCK_LIST_REMOVE Procedure

This procedures removes the cache_id from the block list.

Syntax

```
DBMS_RESULT_CACHE.BLOCK_LIST_REMOVE (
   cache_id   IN  VARCHAR2,
   global     IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 160-12 BLOCK_LIST_REMOVE Procedure Parameters

Parameter	Description
`cache_id`	The `cache_id` is removed from the block list.
`global`	`TRUE` — applies to all caches in a RAC cluster. `FALSE` (default) — applies only to the local instance cache.

160.3.23 OBJECT_BLOCK_LIST Function

This table function returns all the object numbers currently in the block list.

Syntax

```
DBMS_RESULT_CACHE.OBJECT_BLOCK_LIST ()
   RETURN NATURALN;
```

160.3.24 OBJECT_BLOCK_LIST_ADD Procedure

This procedures adds objects to the object block list. These procedures only have to be run once in the cluster as they take effect everywhere at once. However, they do not persist between cluster shutdowns and startups.

Syntax

```
DBMS_RESULT_CACHE.OBJECT_BLOCK_LIST_ADD (
   objNo       IN  BINARY_INTEGER);

DBMS_RESULT_CACHE.OBJECT_BLOCK_LIST_ADD (
   owner       IN  VARCHAR2,
   name        IN  VARCHAR2);
```

Parameters

Table 160-13 OBJECT_BLOCK_LIST_ADD Procedure Parameters

Parameter	Description
objNo	The object number.
owner	The owner of the object.
name	The name of the table.

160.3.25 OBJECT_BLOCK_LIST_CLEAR Procedure

This procedure clears the entire block list.

Syntax

```
DBMS_RESULT_CACHE.OBJECT_BLOCK_LIST_CLEAR ();
```

160.3.26 OBJECT_BLOCK_LIST_REMOVE Procedure

This procedures removes objects from the object block list.

Syntax

```
DBMS_RESULT_CACHE.OBJECT_BLOCK_LIST_REMOVE (
   objNo       IN  BINARY_INTEGER);

DBMS_RESULT_CACHE.OBJECT_BLOCK_LIST_REMOVE (
   owner       IN  VARCHAR2,
   name        IN  VARCHAR2);
```

Parameters

Table 160-14 OBJECT_BLOCK_LIST_REMOVE Procedure Parameters

Parameter	Description
objNo	The object number.
owner	The owner of the object.
name	The name of the table.

161
DBMS_RESUMABLE

With the DBMS_RESUMABLE package, you can suspend large operations that run out of space or reach space limits after executing for a long time, fix the problem, and make the statement resume execution. In this way you can write applications without worrying about running into space-related errors.

This chapter contains the following topics:

- Operational Notes
- Summary of DBMS_RESUMABLE Subprograms

161.1 DBMS_RESUMABLE Operational Notes

When you suspend a statement, you should log the suspension in the alert log. You should also register a procedure to be executed when the statement is suspended. Using a view, you can monitor the progress of the statement and indicate whether the statement is currently executing or suspended.

Suspending a statement automatically results in suspending the transaction. Thus all transactional resources are held during a statement suspend and resume. When the error condition disappears, the suspended statement automatically resumes execution. A resumable space allocation can be suspended and resumed multiple times during execution.

A suspension timeout interval is associated with resumable space allocations. A resumable space allocation that is suspended for the timeout interval (the default is two hours) wakes up and returns an exception to the user. A suspended statement may be forced to throw an exception using the DBMS_RESUMABLE.ABORT() procedure.

161.2 Summary of DBMS_RESUMABLE Subprograms

This table lists the DBMS_RESUMABLE subprograms and briefly describes them.

Table 161-1 ***DBMS_RESUMABLE Package Subprograms***

Subprogram	Description
ABORT Procedure	Aborts a suspended resumable space allocation
GET_SESSION_TIMEOUT Function	Returns the current timeout value of the resumable space allocations for a session with session_id
GET_TIMEOUT Function	Returns the current timeout value of resumable space allocations for the current session
SET_SESSION_TIMEOUT Procedure	Sets the timeout of resumable space allocations for a session with session_id
SET_TIMEOUT Procedure	Sets the timeout of resumable space allocations for the current session
SPACE_ERROR_INFO Function	Looks for space-related errors in the error stack, otherwise returning FALSE

161.2.1 ABORT Procedure

This procedure aborts a suspended resumable space allocation.

The parameter `session_id` is the session ID in which the statement is executed. For a parallel DML/DDL, `session_id` is any session ID that participates in the parallel DML/DDL. This operation is guaranteed to succeed. The procedure can be called either inside or outside of the `AFTER SUSPEND` trigger.

Syntax

```
DBMS_RESUMABLE.ABORT (
   session_id  IN NUMBER);
```

Parameters

Table 161-2 ABORT Procedure Parameters

Parameter	Description
`session_id`	The session identifier of the resumable space allocation.

Usage Notes

To call an `ABORT` procedure, you must be the owner of the session with `session_id`, have `ALTER SYSTEM` privileges, or be a DBA.

161.2.2 GET_SESSION_TIMEOUT Function

This function returns the current timeout value of resumable space allocations for a session with `session_id`.

Syntax

```
DBMS_RESUMABLE.GET_SESSION_TIMEOUT (
   session_id  IN NUMBER)
RETURN NUMBER;
```

Parameters

Table 161-3 GET_SESSION_TIMEOUT Function Parameters

Parameter	Description
`session_id`	The session identifier of the resumable space allocation.

Return Values

Table 161-4 GET_SESSION_TIMEOUT Function Return Values

Return Value	Description
`NUMBER`	The current timeout value of resumable space allocations for a session with `session_id`. The timeout is returned in seconds.

Usage Notes

If `session_id` does not exist, the `GET_SESSION_TIMEOUT` function returns -1.

161.2.3 GET_TIMEOUT Function

This function returns the current timeout value of resumable space allocations for the current session.

Syntax

```
DBMS_RESUMABLE.GET_TIMEOUT
 RETURN NUMBER;
```

Return Values

Table 161-5 GET_TIMEOUT Function Return Values

Return Value	Description
`NUMBER`	The current timeout value of resumable space allocations for the current session. The returned value is in seconds.

Usage Notes

If the current session is not resumable enabled, the `GET_TIMEOUT` function returns -1.

161.2.4 SET_SESSION_TIMEOUT Procedure

This procedure sets the timeout of resumable space allocations for a session with `session_id`.

The new timeout setting applies to the session immediately. If `session_id` does not exist, no operation occurs.

Syntax

```
DBMS_RESUMABLE.SET_SESSION_TIMEOUT (
   session_id  IN NUMBER,
   timeout     IN NUMBER);
```

Parameters

Table 161-6 SET_SESSION_TIMEOUT Procedure Parameters

Parameter	Description
`session_id`	The session identifier of the resumable space allocation.
`timeout`	The timeout of the resumable space allocation.

161.2.5 SET_TIMEOUT Procedure

This procedure sets the timeout of resumable space allocations for the current session. The new timeout setting applies to the session immediately.

Syntax

```
DBMS_RESUMABLE.SET_TIMEOUT (
   timeout  IN NUMBER);
```

Parameters

Table 161-7 SET_TIMEOUT Procedure Parameters

Parameter	Description
timeout	The timeout of the resumable space allocation.

161.2.6 SPACE_ERROR_INFO Function

This function looks for space-related errors in the error stack.

If it cannot find a space related error, it will return `FALSE`. Otherwise, `TRUE` is returned and information about the particular object that causes the space error is returned.

Syntax

```
DBMS_RESUMABLE.SPACE_ERROR_INFO
   error_type        OUT VARCHAR2,
   object_type       OUT VARCHAR2,
   object_owner      OUT VARCHAR2,
   table_space_name  OUT VARCHAR2,
   object_name       OUT VARCHAR2,
   sub_object_name   OUT VARCHAR2)
RETURN BOOLEAN;
```

Parameters

Table 161-8 SPACE_ERROR_INFO Function Parameters

Parameter	Description
error_type	The space error type. It will be one of the following: • NO MORE SPACE • MAX EXTENTS REACHED • SPACE QUOTA EXCEEDED

Table 161-8 (Cont.) SPACE_ERROR_INFO Function Parameters

Parameter	Description
`object_type`	The object type. It will be one of the following: • TABLE • INDEX • CLUSTER • TABLE SPACE • ROLLBACK SEGMENT • UNDO SEGMENT • LOB SEGMENT • TEMP SEGMENT • INDEX PARTITION • TABLE PARTITION • LOB PARTITION • TABLE SUBPARTITION • INDEX SUBPARTITION • LOB SUBPARTITION The type can also be `NULL` if it does not apply.
`object_owner`	The owner of the object. `NULL` if it cannot be determined.
`table_space_name`	The table space where the object resides. `NULL` if it cannot be determined.
`object_name`	The name of rollback segment, temp segment, table, index, or cluster.
`sub_object_name`	The partition name or sub-partition name of `LOB, TABLE,` or `INDEX.` `NULL` if it cannot be determined.

162
DBMS_RLS

The DBMS_RLS package contains the fine-grained access control administrative interface, which is used to implement Virtual Private Database (VPD).

DBMS_RLS is available with the Enterprise Edition only.

> **See Also:**
>
> *Oracle Database Security Guide*for usage information about DBMS_RLS

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operational Notes
- Rules and Limits
- Summary of DBMS_RLS Subprograms

162.1 DBMS_RLS Overview

The functionality to support fine-grained access control is based on dynamic predicates, where security rules are not embedded in views, but are acquired at the statement parse time, when the base table or view is referenced in a DML statement.

A dynamic predicate for a table, view, or synonym is generated by a PL/SQL function, which is associated with a security policy through a PL/SQL interface. For example:

```
DBMS_RLS.ADD_POLICY (
   'hr', 'employees', 'emp_policy', 'hr', 'emp_sec', 'select', 'user_ctx', 'time');
```

Whenever the EMPLOYEES table, under the HR schema, is referenced in a query or subquery (SELECT), the server calls the EMP_SEC function (under the HR schema). This function returns a predicate specific to the current schema for the EMP_POLICY policy. The policy function may generate the predicates based on the session environment variables available during the function call. These variables usually appear in the form of application contexts. The policy can specify any combination of security-relevant columns and of these statement types: INDEX, SELECT, INSERT, UPDATE, or DELETE.

The server then produces a transient view with the text:

```
SELECT * FROM hr.employees WHERE P1
```

Here, P1 (for example, where SAL > 10000, or even a subquery) is the predicate returned from the EMP_SEC function. The server treats the EMPLOYEES table as a view and does the view

expansion just like the ordinary view, except that the view text is taken from the transient view instead of the data dictionary.

If the predicate contains subqueries, then the owner (definer) of the policy function is used to resolve objects within the subqueries and checks security for those objects. In other words, users who have access privilege to the policy-protected objects do not need to know anything about the policy. They do not need to be granted object privileges for any underlying security policy. Furthermore, the users do not require `EXECUTE` privilege on the policy function, because the server makes the call with the function definer's right.

> **Note:**
>
> The transient view can preserve the updatability of the parent object because it is derived from a single table or view with predicate only; that is, no `JOIN`, `ORDER BY`, `GROUP BY`, and so on.

`DBMS_RLS` also provides the interface to drop or enable security policies. For example, you can drop or enable the `EMP_POLICY` with the following PL/SQL statements:

```
DBMS_RLS.DROP_POLICY('hr', 'employees', 'emp_policy');
DBMS_RLS.ENABLE_POLICY('hr', 'employees', 'emp_policy', TRUE);
```

162.2 DBMS_RLS Security Model

A security check is performed when the transient view is created with a subquery. The schema owning the policy function, which generates the dynamic predicate, is the transient view's definer for security check and object lookup.

162.3 DBMS_RLS Constants

The `DBMS_RLS` package includes constants that can be used for specifying parameter values.

Table 162-1 DBMS_RLS Constants

Constant	Type	Value	Description
`ADD_ATTRIBUTE_ASSOCIATION`	`BINARY_INTEGER`	1	Used with `DBMS_RLS.ALTER_POLICY` and `DBMS_RLS_ALTER_GROUPED_POLICY`: adds the specified namespace and attribute to the policy or grouped policy
`REMOVE_ATTRIBUTE_ASSOCIATION`	`BINARY_INTEGER`	2	Used with `DBMS_RLS.ALTER_POLICY` and `DBMS_RLS.ALTER_GROUPED_POLICY`: removes the specified namespace and attribute to the policy or grouped policy.

162.4 DBMS_RLS Operational Notes

The DBMS_RLS procedures cause current DML transactions, if any, to commit before the operation. However, the procedures do not cause a commit first if they are inside a DDL event trigger. With DDL transactions, the DBMS_RLS procedures are part of the DDL transaction.

For example, you may create a trigger for CREATE TABLE. Inside the trigger, you may add a column through ALTER TABLE, and you can add a policy through DBMS_RLS. All these operations are in the same transaction as CREATE TABLE, even though each one is a DDL statement. The CREATE TABLE succeeds only if the trigger is completed successfully.

Views of current cursors and corresponding predicates are available from V$VPD_POLICIES.

A synonym can reference only a view or a table.

162.5 DBMS_RLS Rules and Limits

Using long identifiers is supported for VPD. The maximum length for arguments such as object_schema, object_name, and policy_name, which apply to objects (table names, policy names, and subprogram names) and views is 128 bytes.

162.6 Summary of DBMS_RLS Subprograms

This table lists and briefly describes the subprograms available in DBMS_RLS.

Table 162-2 DBMS_RLS Package Subprograms

Subprogram	Description
ADD_GROUPED_POLICY Procedure	Adds a policy associated with a policy group
ADD_POLICY Procedure	Adds a fine-grained access control policy to a table, view, or synonym
ADD_POLICY_CONTEXT Procedure	Adds the context for the active application
ALTER_POLICY Procedure	Associates an application context attribute with VPD policies
ALTER_GROUPED_POLICY Procedure	Adds application context related changes
CREATE_POLICY_GROUP Procedure	Creates a policy group
DELETE_POLICY_GROUP Procedure	Deletes a policy group
DISABLE_GROUPED_POLICY Procedure	Disables a row-level group security policy
DROP_GROUPED_POLICY Procedure	Drops a policy associated with a policy group
DROP_POLICY Procedure	Drops a fine-grained access control policy from a table, view, or synonym
DROP_POLICY_CONTEXT Procedure	Drops a driving context from the object so that it will have one less driving context
ENABLE_GROUPED_POLICY Procedure	Enables or disables a row-level group security policy
ENABLE_POLICY Procedure	Enables or disables a fine-grained access control policy
REFRESH_GROUPED_POLICY Procedure	Reparses the SQL statements associated with a refreshed policy

Table 162-2 (Cont.) DBMS_RLS Package Subprograms

Subprogram	Description
REFRESH_POLICY Procedure	Causes all the cached statements associated with the policy to be reparsed

162.6.1 ADD_GROUPED_POLICY Procedure

This procedure adds a policy associated with a policy group.

Syntax

```
DBMS_RLS.ADD_GROUPED_POLICY(
   object_schema            IN  VARCHAR2        DEFAULT NULL,
   object_name              IN  VARCHAR2,
   policy_group             IN  VARCHAR2        DEFAULT 'SYS_DEFAULT',
   policy_name              IN  VARCHAR2,
   function_schema          IN  VARCHAR2        DEFAULT NULL,
   policy_function          IN  VARCHAR2,
   statement_types          IN  VARCHAR2        DEFAULT NULL,
   update_check             IN  BOOLEAN         DEFAULT FALSE,
   enable                   IN  BOOLEAN         DEFAULT TRUE,
   static_policy            IN  BOOLEAN         DEFAULT FALSE,
   policy_type              IN  BINARY_INTEGER  DEFAULT NULL,
   long_predicate           IN  BOOLEAN         DEFAULT FALSE,
   sec_relevant_cols        IN  VARCHAR2,
   sec_relevant_cols_opt    IN  BINARY_INTEGER  DEFAULT NULL,
   namespace                IN  VARCHAR2        DEFAULT NULL,
   attribute                IN  VARCHAR2        DEFAULT NULL);
```

Parameters

Table 162-3 ADD_GROUPED_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym to which the policy is added
`policy_group`	Name of the policy group to which the policy belongs
`policy_name`	Name of the policy; must be unique for the same table or view
`function_schema`	Schema owning the policy function. If no `function_schema` is specified or is `NULL`, then the current schema is used.
`policy_function`	Name of the function that generates a predicate for the policy. If the function is defined within a package, the name of the package must be present.
`statement_types`	Statement types to which the policy applies. It can be any combination of `INDEX`, `SELECT`, `UPDATE`, or `DELETE`. The default is to apply to all of these types except `INSERT` and `INDEX`.

Table 162-3 (Cont.) ADD_GROUPED_POLICY Procedure Parameters

Parameter	Description
update_check	For INSERT and UPDATE statements only, setting update_check to TRUE causes the server to check the policy against the value after INSERT or UPDATE. The check applies only to the security relevant columns that are included in the policy definition. In other words, the INSERT or UPDATE operation will fail only if the security relevant column that is defined in the policy is added or updated in the INSERT or UPDATE statement.
enable	Indicates if the policy is enable when it is added. The default is TRUE.
static_policy	Default is FALSE. If it is set to TRUE, the server assumes that the policy function for the static policy produces the same predicate string for anyone accessing the object, except for SYS or the privilege user who has the EXEMPT ACCESS POLICY privilege.
policy_type	Default is NULL, which means policy_type is decided by the value of static_policy. The available policy types are listed in Table 162-5. Specifying any of these policy types overrides the value of static_policy.
long_predicate	Default is FALSE, which means the policy function can return a predicate with a length of up to 4000 bytes. TRUE means the predicate text string length can be up to 32K bytes.Policies existing prior to the availability of this parameter retain a 32K limit.
sec_relevant_cols	Enables column-level Virtual Private Database (VPD), which enforces security policies when a column containing sensitive information is referenced in a query. Applies to tables and views, but not to synonyms. Specify a list of comma- or space-separated valid column names of the policy-protected object. The policy is enforced only if a specified column is referenced (or, for an abstract datatype column, its attributes are referenced) in the user SQL statement or its underlying view definition. Default is all the user-defined columns for the object.
namespace	Name which determines the application context namespace
attribute	Attribute which determines the application context attribute name

Usage Notes

- This procedure adds a policy to the specified table, view, or synonym and associates the policy with the specified policy group.
- The policy group must have been created by using the CREATE_POLICY_GROUP Procedure.
- The policy name must be unique within a policy group for a specific object.
- Policies from the default policy group, SYS_DEFAULT, are always executed regardless of the active policy group; however, fine-grained access control policies do not apply to users with EXEMPT ACCESS POLICY system privilege.

162.6.2 ADD_POLICY Procedure

This procedure adds a fine-grained access control policy to a table, view, or synonym.

The procedure causes the current transaction, if any, to commit before the operation is carried out. However, this does not cause a commit first if it is inside a DDL event trigger.

> **See Also:**
>
> Operational Notes

A `COMMIT` is also performed at the end of the operation.

Syntax

```
DBMS_RLS.ADD_POLICY (
   object_schema            IN  VARCHAR2       DEFAULT NULL,
   object_name              IN  VARCHAR2,
   policy_name              IN  VARCHAR2,
   function_schema          IN  VARCHAR2       DEFAULT NULL,
   policy_function          IN  VARCHAR2,
   statement_types          IN  VARCHAR2       DEFAULT NULL,
   update_check             IN  BOOLEAN        DEFAULT FALSE,
   enable                   IN  BOOLEAN        DEFAULT TRUE,
   static_policy            IN  BOOLEAN        DEFAULT FALSE,
   policy_type              IN  BINARY_INTEGER DEFAULT NULL,
   long_predicate           IN  BOOLEAN        DEFAULT FALSE,
   sec_relevant_cols        IN  VARCHAR2       DEFAULT NULL,
   sec_relevant_cols_opt    IN  BINARY_INTEGER DEFAULT NULL,
   namespace                IN  VARCHAR2       DEFAULT NULL,
   attribute                IN  VARCHAR2       DEFAULT NULL);
```

Parameters

Table 162-4 ADD_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of table, view, or synonym to which the policy is added.
`policy_name`	Name of policy to be added. It must be unique for the same table or view. Do not enter special characters such as spaces or commas. If you want to use special characters for the policy name, then enclose the name in quotation marks.
`function_schema`	Schema owning the policy function. If no `function_schema` is specified or is `NULL`, then the current schema is used.
`policy_function`	Name of a function which generates a predicate for the policy. If the function is defined within a package, then the name of the package must be present.
`statement_types`	Statement types to which the policy applies. It can be any combination of `INDEX`, `SELECT`, `UPDATE`, or `DELETE`. The default is to apply to all of these types except `INSERT` and `INDEX` .
`update_check`	Optional argument for the `INSERT` or `UPDATE` statement type. The default is `FALSE`. If you plan to use the `INSERT` statement type, then you must set `update_check` to `TRUE`. Otherwise, an `ORA-28104 input value for string is not valid` error is generated. The check applies only to the security relevant columns that are included in the policy definition. In other words, the `INSERT` or `UPDATE` operation will fail only if the security relevant column that is defined in the policy is added or updated in the `INSERT` or `UPDATE` statement.

Table 162-4 (Cont.) ADD_POLICY Procedure Parameters

Parameter	Description
enable	Indicates if the policy is enabled when it is added. The default is TRUE.
static_policy	The default is FALSE. If it is set to TRUE, the server assumes that the policy function for the static policy produces the same predicate string for anyone accessing the object, except for SYS or the privileged user who has the EXEMPT ACCESS POLICY privilege.
policy_type	Default is NULL, which means policy_type is decided by the value of static_policy. The available policy types are listed in Table 162-5. Specifying any of these policy types overrides the value of static_policy.
long_predicate	Default is FALSE, which means the policy function can return a predicate with a length of up to 4000 bytes. TRUE means the predicate text string length can be up to 32K bytes.Policies existing prior to the availability of this parameter retain a 32K limit.
sec_relevant_cols	Enables column-level Virtual Private Database (VPD), which enforces security policies when a column containing sensitive information is referenced in a query. Applies to tables and views, but not to synonyms. Specify a list of comma- or space-separated valid column names of the policy-protected object. The policy is enforced only if a specified column is referenced (or, for an abstract datatype column, its attributes are referenced) in the user SQL statement or its underlying view definition. Default is all the user-defined columns for the object.
sec_relevant_cols_opt	Use with sec_relevant_cols to display all rows for column-level VPD filtered queries (SELECT only), but where sensitive columns appear as NULL. Default is set to NULL, which allows the filtering defined with sec_relevant_cols to take effect. Set to dbms_rls.ALL_ROWS to display all rows, but with sensitive column values, which are filtered by sec_relevant_cols, displayed as NULL. See Usage Notes for restrictions and additional information about this option.
namespace	Name which determines the application context namespace
attribute	Attribute which determines the application context attribute name

Table 162-5 DBMS_RLS.ADD_POLICY Policy Types

Policy Type	Description
STATIC	Predicate is assumed to be the same regardless of the runtime environment. Static policy functions are executed once and then cached in SGA. Statements accessing the same object do not reexecute the policy function. However, each execution of the same cursor could produce a different row set even for the same predicate because the predicate may filter the data differently based on attributes such as SYS_CONTEXT or SYSDATE. Applies to only one object.
SHARED_STATIC	Same as STATIC except that the server first looks for a cached predicate generated by the same policy function of the same policy type. Shared across multiple objects.
CONTEXT_SENSITIVE	Server re-evaluates the policy function at statement execution time if it detects context changes since the last use of the cursor. For session pooling where multiple clients share a database session, the middle tier must reset context during client switches. Note that the server does not cache the value returned by the function for this policy type; it always executes the policy function on statement parsing. Applies to only one object.

Table 162-5 (Cont.) DBMS_RLS.ADD_POLICY Policy Types

Policy Type	Description
SHARED_CONTEXT_SENSITIVE	Same as CONTEXT_SENSITIVE except that the server first looks for a cached predicate generated by the same policy function of the same policy type within the same database session. If the predicate is found in the session memory, the policy function is not reexecuted and the cached value is valid until session private application context changes occur. Shared across multiple objects.
DYNAMIC	The default policy type. Server assumes the predicate may be affected by any system or session environment at any time, and so always reexecutes the policy function upon each statement parsing and execution. Applies to only one object.

Usage Notes

- SYS is free of any security policy.
- The policy functions are called by the server. Following is the interface for the function:

```
FUNCTION policy_function (object_schema IN VARCHAR2, object_name VARCHAR2)
    RETURN VARCHAR2
--- object_schema is the schema owning the table or view.
--- object_name is the name of table, view, or synonym to which the policy applies.
```

- The policy functions must have the purity level of WNDS (write no database state).

> **See Also:**
>
> The *Oracle Database Development Guide* has more details about the RESTRICT_REFERENCES pragma.

- Predicates generated from different VPD policies for the same object have the combined effect of a conjunction (ANDed) of all the predicates.
- The security check and object lookup are performed against the owner of the policy function for objects in the subqueries of the dynamic predicates.
- If the function returns a zero length predicate, then it is interpreted as no restriction being applied to the current user for the policy.
- When a table alias is required (for example, parent object is a type table) in the predicate, the name of the table or view itself must be used as the name of the alias. The server constructs the transient view as something like

```
"select c1, c2, ... from tab tab where <predicate>"
```

- Validity of the function is checked at runtime for ease of installation and other dependency issues during import and export.
- Column-level VPD column masking behavior (specified with sec_relevant_cols_opt => dbms_rls.ALL_ROWS) is fundamentally different from all other VPD policies, which return only a subset of rows. Instead the column masking behavior returns all rows specified by the user's query, but the sensitive column values display as NULL. The restrictions for this option are as follows:
  - Only applies to SELECT statements

  - Unlike regular VPD predicates, the masking condition that is generated by the policy function must be a simple boolean expression.
  - If your application performs calculations, or does not expect `NULL` values, then you should use the default behavior of column-level VPD, which is specified with the `sec_relevant_cols` parameter.
  - If you use `UPDATE AS SELECT` with this option, then only the values in the columns you are allowed to see will be updated.
  - This option may prevent some rows from displaying. For example:

    ```
    SELECT * FROM employees
    WHERE salary = 10
    ```

    This query may not return rows if the `salary` column returns a `NULL` value because the column masking option has been set.

- When you add a VPD policy to a synonym, it causes all the dependent objects of the synonym, including policy functions that reference the synonym, to be marked `INVALID`.
- You cannot associate a global application context with a context sensitive policy or a context shared sensitive policy.
- The maximum number of policies that can be created for a single object is 255.

Examples

As the first of two examples, the following creates a policy that applies to the `hr.employee` table. This is a column-level VPD policy that will be enforced only if a `SELECT` or an `INDEX` statement refers to the `salary`, `birthdate`, or `SSN` columns of the table explicitly, or implicitly through a view. It is also a `CONTEXT_SENSITIVE` policy, so the server will invoke the policy function `hr.hrfun` at parse time. The namespace and attribute application context parameters restrict the policy evaluation only when the application context values change. During execution, it will only invoke the function if there has been any session private context change since the last use of the statement cursor. The predicate generated by the policy function must not exceed 4000 bytes, the default length limit, since the `long_predicate` parameter is omitted from the call.

```
BEGIN
   DBMS_RLS.ADD_POLICY(
      object_schema      => 'hr',
      object_name        => 'employee',
      policy_name        => 'hr_policy',
      function_schema    => 'hr',
      policy_function    => 'hrfun',
      statement_types    => 'select,index',
      policy_type        => DBMS_RLS.CONTEXT_SENSITIVE,
      sec_relevant_cols  => 'salary,birthdate,ssn',
      namespace          => 'empno_ctx',
      attribute          => 'emp_id');
END;
```

As the second example, the following command creates another policy that applies to the same object for hosting, so users can access only data based on their subscriber ID. Since it is defined as a `SHARED_STATIC` policy type, the server will first try to find the predicate in the SGA cache. The server will only invoke the policy function, `subfun`, if that search fails.

```
BEGIN
   DBMS_RLS.ADD_POLICY(
      object_schema      => 'hr',
      object_name        => 'employee',
```

```
      policy_name        => 'hosting_policy',
      function_schema    => 'hr',
      policy_function    => 'subfun',
      policy_type        => dbms_rls.SHARED_STATIC);
END;
```

162.6.3 ADD_POLICY_CONTEXT Procedure

This procedure adds the context for the active application.

Syntax

```
DBMS_RLS.ADD_POLICY_CONTEXT (
   object_schema    IN VARCHAR2 NULL,
   object_name      IN VARCHAR2,
   namespace        IN VARCHAR2,
   attribute        IN VARCHAR2);
```

Parameters

Table 162-6 ADD_POLICY_CONTEXT Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym to which the policy is added.
`namespace`	Name which determines the application context namespace
`attribute`	Attribute which determines the application context attribute name

Usage Notes

Note the following:

- This procedure indicates the application context that drives the enforcement of policies; this is the context that determines which application is running.
- The driving context can be session or global.
- At execution time, the server retrieves the name of the active policy group from the value of this context.
- There must be at least one driving context defined for each object that has fine-grained access control policies; otherwise, all policies for the object will be executed.
- Adding multiple context to the same object will cause policies from multiple policy groups to be enforced.
- If the driving context is `NULL`, policies from all policy groups are used.
- If the driving context is a policy group with policies, all enabled policies from that policy group will be applied, along with all policies from the `SYS_DEFAULT` policy group.
- To add a policy to table `HR.EMPLOYEES` in group `access_control_group`, the following command is issued:

  ```
  DBMS_RLS.ADD_GROUPED_POLICY('hr','employees','access_control_group','policy1','SYS',
  'HR.ACCESS');
  ```

162.6.4 ALTER_POLICY Procedure

This procedure associates an application context attribute with VPD policies.

Syntax

```
DBMS_RLS.ALTER_POLICY (
   object_schema    IN VARCHAR2 DEFAULT NULL,
   object_name      IN VARCHAR2,
   policy_name      IN VARCHAR2,
   alter_option     IN NUMBER,
   namespace        IN VARCHAR2 DEFAULT NULL,
   attribute        IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 162-7 ALTER_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym to which the policy is added
`policy_name`	Name of the policy, unique for the same table or view
`alter_option`	Used to determine whether the application context is being added or removed from an Oracle Virtual Private Database policy
`namespace`	Name which determines the application context namespace
`attribute`	Attribute which determines the application context attribute name

Usage Notes

Note the following:

- This procedure associates an application context namespace and application context attribute to context sensitive and shared context sensitive policies only. Specifying application context namespace and application context attribute for `DYNAMIC`, `STATIC` or `SHARED_STATIC` policies will result in an error. If `namespace` is specified, `attribute` should also be specified for the procedure call.
- You cannot associate a global application context with a context sensitive policy or a context shared sensitive policy.
- Invocations of `ALTER_POLICY` which modify a shared context sensitive VPD policy have an effect on all shared context sensitive VPD policies that have the same VPD policy function.
- The driving context can be session or global.
- At execution time, the server retrieves the name of the active policy group from the value of this context.
- There must be at least one driving context defined for each object that has fine-grained access control policies; otherwise, all policies for the object will be executed.
- Adding multiple context to the same object will cause policies from multiple policy groups to be enforced.
- If the driving context is `NULL`, policies from all policy groups are used.

- If the driving context is a policy group with policies, all enabled policies from that policy group will be applied, along with all policies from the `SYS_DEFAULT` policy group.
- To add a policy to table `hr.employees` in group `access_control_group`, the following command is issued:

```
DBMS_RLS.ADD_GROUPED_POLICY(
   'hr','employees','access_control_group','policy1','SYS', 'HR.ACCESS');
```

162.6.5 ALTER_GROUPED_POLICY Procedure

This procedure adds application context related changes.

Syntax

```
DBMS_RLS.ALTER_GROUPED_POLICY (
   object_schema    IN VARCHAR2 DEFAULT NULL,
   object_name      IN VARCHAR2,
   policy_group     IN VARCHAR2 DEFAULT SYS_DEFAULT,
   policy_name      IN VARCHAR2,
   alter_option     IN NUMBER,
   namespace        IN VARCHAR2 DEFAULT NULL,
   attribute        IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 162-8 ALTER_GROUPED_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym to which the policy is added
`policy_group`	Name of the policy group to which this policy belongs; must be unique for the same table or view
`policy_name`	Name of the policy, unique for the same table or view
`alter_option`	Used to determine whether the application context is being added or removed from the Oracle Virtual Private Database policy
`namespace`	Name that determines the application context namespace
`attribute`	Attribute determines the application context attribute name

Usage Notes

Note the following:

- This procedure will associate an application context namespace and application context attribute to context sensitive and shared context sensitive policies only. Specifying application context namespace and application context attribute for `DYNAMIC`, `STATIC` or `SHARED_STATIC` policies will result in an error. If `namespace` is specified, `attribute` should also be specified for the procedure call.
- You cannot associate a global application context with a context sensitive policy or a context shared sensitive policy.
- Invocations of `ALTER_GROUPED_POLICY` which modify a shared context sensitive VPD policy have an effect on all shared context sensitive VPD policies that have the same VPD policy function.

- The driving context can be session or global.
- At execution time, the server retrieves the name of the active policy group from the value of this context.
- There must be at least one driving context defined for each object that has fine-grained access control policies; otherwise, all policies for the object will be executed.
- Adding multiple context to the same object will cause policies from multiple policy groups to be enforced.
- If the driving context is `NULL`, policies from all policy groups are used.
- If the driving context is a policy group with policies, all enabled policies from that policy group will be applied, along with all policies from the `SYS_DEFAULT` policy group.
- To add a policy to table `hr.employees` in group `access_control_group`, the following command is issued:

```
DBMS_RLS.ADD_GROUPED_POLICY (
   'hr','employees','access_control_group','policy1','SYS', 'HR.ACCESS');
```

162.6.6 CREATE_POLICY_GROUP Procedure

This procedure creates a policy group.

Syntax

```
DBMS_RLS.CREATE_POLICY_GROUP (
   object_schema    IN VARCHAR2 NULL,
   object_name      IN VARCHAR2,
   policy_group     IN VARCHAR2);
```

Parameters

Table 162-9 CREATE_POLICY_GROUP Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym to which the policy is added
`policy_group`	Name of the policy group that the policy belongs to

Usage Notes

The group must be unique for each table or view.

162.6.7 DELETE_POLICY_GROUP Procedure

This procedure deletes a policy group.

Syntax

```
DBMS_RLS.DELETE_POLICY_GROUP (
  object_schema    IN VARCHAR2 NULL,
  object_name      IN VARCHAR2,
  policy_group     IN VARCHAR2);
```

Parameters

Table 162-10 DELETE_POLICY_GROUP Procedure Parameters

Parameter	Description
object_schema	Schema containing the table, view, or synonym. If no object_schema is specified or is NULL, then the current schema is used.
object_name	Name of the table, view, or synonym to which the policy is added
policy_group	Name of the policy group that the policy belongs to

Usage Notes

Note the following:

- This procedure deletes a policy group for the specified table, view, or synonym.
- No policy can be in the policy group.

162.6.8 DISABLE_GROUPED_POLICY Procedure

This procedure disables a row-level group security policy.

Syntax

```
DBMS_RLS.DISABLE_GROUPED_POLICY (
   object_schema   IN VARCHAR2 NULL,
   object_name     IN VARCHAR2,
   group_name      IN VARCHAR2,
   policy_name     IN VARCHAR2);
```

Parameters

Table 162-11 DISABLE_GROUPED_POLICY Procedure Parameters

Parameter	Description
object_schema	Schema containing the table, view, or synonym
object_name	Name of the table, view, or synonym with which the policy is associated
group_name	Name of the group of the policy
policy_name	Name of the policy to be enabled or disabled

Usage Notes

- The procedure causes the current transaction, if any, to commit before the operation is carried out.
- A commit is performed at the end of the operation.
- A policy is disabled when this procedure is executed or when the ENABLE_GROUPED_POLICY procedure is executed with "enable" set to FALSE.

162.6.9 DROP_GROUPED_POLICY Procedure

This procedure drops a policy associated with a policy group.

Syntax

```
DBMS_RLS.DROP_GROUPED_POLICY (
   object_schema    IN VARCHAR2 NULL,
   object_name      IN VARCHAR2,
   policy_group     IN VARCHAR2 'SYS_DEFAULT',
   policy_name      IN VARCHAR2);
```

Parameters

Table 162-12 DROP_GROUPED_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym to which the policy is dropped
`policy_group`	Name of the policy group to which the policy belongs
`policy_name`	Name of the policy

162.6.10 DROP_POLICY Procedure

This procedure drops a fine-grained access control policy from a table, view, or synonym.

The procedure causes the current transaction, if any, to commit before the operation is carried out. However, this does not cause a commit first if it is inside a DDL event trigger.

> **See Also:**
>
> Operational Notes

A COMMIT is also performed at the end of the operation.

Syntax

```
DBMS_RLS.DROP_POLICY (
   object_schema    IN VARCHAR2 NULL,
   object_name      IN VARCHAR2,
   policy_name      IN VARCHAR2);
```

Parameters

Table 162-13 DROP_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view or synonym. If no `object_schema` is specified, or `NULL` is provided, then the current user's schema is assumed.

Table 162-13 (Cont.) DROP_POLICY Procedure Parameters

Parameter	Description
object_name	Name of the table, view, or synonym for which the policy is dropped
policy_name	Name of policy to be dropped from table, view, or synonym

Usage Notes

- When you drop a VPD policy from a synonym, it causes all the dependent objects of the synonym, including policy functions that reference the synonym, to be marked INVALID.

162.6.11 DROP_POLICY_CONTEXT Procedure

This procedure drops a driving context from the object so that it will have one less driving context.

Syntax

```
DBMS_RLS.DROP_POLICY_CONTEXT (
   object_schema    IN VARCHAR2 NULL,
   object_name      IN VARCHAR2,
   namespace        IN VARCHAR2,
   attribute        IN VARCHAR2);
```

Parameters

Table 162-14 DROP_POLICY_CONTEXT Procedure Parameters

Parameter	Description
object_schema	Schema containing the table, view, or synonym. If no object_schema is specified or is NULL, then the current schema is used.
object_name	Name of the table, view, or synonym to which the policy is dropped
namespace	Namespace of the driving context
attribute	Attribute of the driving context

162.6.12 ENABLE_GROUPED_POLICY Procedure

This procedure enables or disables a row-level group security policy.

Syntax

```
DBMS_RLS.ENABLE_GROUPED_POLICY (
   object_schema    IN VARCHAR2 NULL,
   object_name      IN VARCHAR2,
   group_name       IN VARCHAR2,
   policy_name      IN VARCHAR2,
   enable           IN BOOLEAN TRUE);
```

Parameters

Table 162-15 ENABLE_GROUPED_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of the table, view, or synonym with which the policy is associated
`group_name`	Name of the group of the policy
`policy_name`	Name of the policy to be enabled or disabled
`enable`	`TRUE` enables the policy; `FALSE` disables the policy

Usage Notes

- The procedure causes the current transaction, if any, to commit before the operation is carried out.
- A commit is performed at the end of the operation.
- A policy is enabled when it is created.

162.6.13 ENABLE_POLICY Procedure

This procedure enables or disables a fine-grained access control policy. A policy is enabled when it is created.

The procedure causes the current transaction, if any, to commit before the operation is carried out. However, this does not cause a commit first if it is inside a DDL event trigger.

> **See Also:**
>
> Operational Notes

A `COMMIT` is also performed at the end of the operation.

Syntax

```
DBMS_RLS.ENABLE_POLICY (
   object_schema IN VARCHAR2 NULL,
   object_name   IN VARCHAR2,
   policy_name   IN VARCHAR2,
   enable        IN BOOLEAN TRUE);
```

Parameters

Table 162-16 ENABLE_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.

Table 162-16 (Cont.) ENABLE_POLICY Procedure Parameters

Parameter	Description
object_name	Name of table, view, or synonym with which the policy is associated
policy_name	Name of policy to be enabled or disabled
enable	TRUE to enable the policy, FALSE to disable the policy

162.6.14 REFRESH_GROUPED_POLICY Procedure

This procedure reparses the SQL statements associated with a refreshed policy.

Syntax

```
DBMS_RLS.REFRESH_GROUPED_POLICY (
   object_schema   IN VARCHAR2 NULL,
   object_name     IN VARCHAR2 NULL,
   group_name      IN VARCHAR2 NULL,
   policy_name     IN VARCHAR2 NULL);
```

Parameters

Table 162-17 REFRESH_GROUPED_POLICY Procedure Parameters

Parameter	Description
object_schema	Schema containing the table, view, or synonym. If no object_schema is specified or is NULL, then the current schema is used.
object_name	Name of the table, view, or synonym with which the policy is associated
group_name	Name of the group of the policy
policy_name	Name of the policy

Usage Notes

- This procedure causes all the cached statements associated with the policy to be reparsed. This guarantees that the latest change to the policy has immediate effect after the procedure is executed.
- The procedure causes the current transaction, if any, to commit before the operation is carried out.
- A commit is performed at the end of the operation.
- The procedure returns an error if it tries to refresh a disabled policy.
- The procedure removes the cached results of context and shared sensitive VPD policies.

162.6.15 REFRESH_POLICY Procedure

This procedure causes all the cached statements associated with the policy to be reparsed. This guarantees that the latest change to this policy will have immediate effect after the procedure is executed.

The procedure causes the current transaction, if any, to commit before the operation is carried out. However, this does not cause a commit first if it is inside a DDL event trigger.

> **See Also:**
>
> Operational Notes

A `COMMIT` is also performed at the end of the operation.

Syntax

```
DBMS_RLS.REFRESH_POLICY (
   object_schema IN VARCHAR2 NULL,
   object_name   IN VARCHAR2 NULL,
   policy_name   IN VARCHAR2 NULL);
```

Parameters

Table 162-18 REFRESH_POLICY Procedure Parameters

Parameter	Description
`object_schema`	Schema containing the table, view, or synonym. If no `object_schema` is specified or is `NULL`, then the current schema is used.
`object_name`	Name of table, view, or synonym with which the policy is associated
`policy_name`	Name of policy to be refreshed

Usage Notes

- The procedure returns an error if it tries to refresh a disabled policy.
- The procedure removes the cached results of context and shared sensitive VPD policies.

163
DBMS_ROLLING

The `DBMS_ROLLING` PL/SQL package is used to implement the Rolling Upgrade Using Active Data Guard feature.

The Active Data Guard feature of Oracle Data Guard streamlines the process of upgrading Oracle Database software in a Data Guard configuration in a rolling fashion. The Rolling Upgrade Using Active Data Guard feature requires a license for the Oracle Active Data Guard option. You can use this feautre for database release upgrades starting with the first patchset of Oracle Database 12c.

Starting with Oracle Database 23ai, Transaction Guard works during `DBMS_ROLLING` operations to ensure continuous application functions during switchover, issued by `DBMS_ROLLING` to Transient Logical Standby. Transaction Guard returns the commit outcome of the current in-flight transaction when an error or outage occurs. Applications embed the Transaction Guard APIs in their error handling procedures to ensure that work continues without any in-flight work lost or duplicate submissions after an outage.

Additionally, you can use this feature immediately for other database maintenance tasks. The database where maintenance is performed must be operating at a minimum of Oracle Database 12c Release 1 (12.1). Such maintenance tasks include:

- Adding partitioning to non-partitioned tables
- Changing BasicFiles LOBs to SecureFiles LOBs
- Changing `XMLType` stored as `CLOB` to `XMLtype` stored as binary XML
- Altering tables to be OLTP-compressed

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_ROLLING Subprograms

> **See Also:**
>
> *Oracle Data Guard Concepts and Administration* for information about using `DBMS_ROLLING` to perform a rolling upgrade

163.1 DBMS_ROLLING Overview

The `DBMS_ROLLING` PL/SQL package provides procedures that you can use to perform any change throughout a Data Guard configuration in a rolling fashion.
These procedures include a rolling upgrade of the Oracle Database software. Although the focus of this document is rolling upgrade operations, the content is applicable to the deployment of any rolling changes.

Starting with Oracle Database 23ai, Transaction Guard works during DBMS_ROLLING operations to ensure continuous application functions during switchover, issued by DBMS_ROLLING to Transient Logical Standby.

All the procedures are executed at the current primary database, which eliminates the potential confusion of moving between remote databases to perform various operations related to the rolling upgrade. If necessary, all the procedures can be called again to resume the rolling upgrade after an error or interruption. (The upgrade script must still be run at the standby.)

The package also provides a procedure that allows you to return a Data Guard configuration back to its original, pre-upgrade state in the event users wish to abandon the rolling upgrade.

The actual execution of a rolling upgrade has been reduced to three steps (excluding the upgrade of the Oracle Database software itself and the on-disk setup of the new Oracle Database software). The number of steps remains the same regardless of the size of the Data Guard configuration.

Conceptually, for the purposes of the DBMS_ROLLING package, you divide your Data Guard configuration into two groups: the leading group and the trailing group. The databases in the leading group undergo the upgrade operation (or any other change that you are deploying) first. The databases in the trailing group undergo the upgrade of the Oracle Database software (or any other change that you are deploying) only after the switchover operation. This insulates them from the upgrade and gives you time to evaluate the effect of the change in the leading group databases.

Each group has a master database: the future primary database as specified in the DBMS_ROLLING.INIT_PLAN procedure is the master of the leading group, called Leading Group Master (LGM), while the original primary database is the master of the trailing group called Trailing Group Master (TGM). You can configure databases to protect the LGM and the TGM. Standbys designated to protect the LGM are referred to as Leading Group Standbys (LGS). Standbys designated to protect the TGM are referred to as Trailing Group Standbys (TGS). These terms are used throughout this documentation.

If you perform a rolling upgrade using a transient logical standby database (including using DBMS_ROLLING), then you must manually load audit records from the operating system spillover location before the switchover. Before you run DBMS_ROLLING.SWITCHOVER, run DBMS_AUDIT_MGMT.LOAD_UNIFIED_AUDIT_FILES on the trailing group master (current primary database).

163.2 DBMS_ROLLING Security Model

The DBMS_ROLLING package is available to users who have been granted the DBA role.

163.3 Summary of DBMS_ROLLING Subprograms

This table lists and briefly describes the DBMS_ROLLING package subprograms.

Table 163-1 DBMS_ROLLING Package Subprograms

Subprogram	Description
INIT_PLAN Procedure	Initializes a rolling operation plan with system-generated default values.
DESTROY_PLAN Procedure	Destroys any existing rolling operation plan, its parameters, and all resources associated with the rolling operation.

Table 163-1 (Cont.) DBMS_ROLLING Package Subprograms

Subprogram	Description
BUILD_PLAN Procedure	Validates plan parameters and creates or modifies a rolling operation plan.
SET_PARAMETER Procedure	Modifies a rolling operation parameter.
START_PLAN Procedure	Starts the rolling operation.
SWITCHOVER Procedure	Performs a switchover between the current primary database and the transient logical standby database.
FINISH_PLAN Procedure	Finalizes the rolling operation.
ROLLBACK_PLAN Procedure	Completely rolls back the rolling operation.

163.3.1 INIT_PLAN Procedure

This procedure initializes a rolling operation plan with system-generated default values.

Syntax

```
DBMS_ROLLING.INIT_PLAN (
   future_primary   IN VARCHAR2);
```

Parameters

Table 163-2 *INIT_PLAN Procedure Parameters*

Parameter	Description
future_primary	DB_UNIQUE_NAME of the future primary (also known as the Leading Group Master (LGM))

Exceptions

- ORA-45400: operation not permitted on current database
- ORA-45401: upgrade plan is already active
- ORA-45402: LOG_ARCHIVE_CONFIG must contain the DG_CONFIG attribute
- ORA-45403: database %s must be specified in DG_CONFIG
- ORA-45411: operation requires additional arguments
- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- This procedure communicates with all databases defined in the log_archive_config init.ora parameter and validates that each database is a valid participant in the rolling upgrade. Valid participants other than the primary database must have a database role of either physical standby, DG, or logical standby.
- The designated future primary must be a physical standby or ADG.

163.3.2 DESTROY_PLAN Procedure

This procedure destroys any existing upgrade plan, its parameters, and all resources associated with a rolling operation.

Syntax

```
DBMS_ROLLING.DESTROY_PLAN ();
```

Parameters

This procedure has no parameters.

Exceptions

- ORA-45422: operation requires existing plan
- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- When a rolling operation is complete, this procedure can be called to completely purge all states related to a rolling operation.
- This procedure must also be called after a ROLLBACK_PLAN to purge the metadata.

163.3.3 BUILD_PLAN Procedure

This procedure validates plan parameters and creates or modifies a rolling operation plan.

A successfully constructed plan is required in order to perform a rolling operation. This procedure must return successfully before the START_PLAN procedure can be called to start the rolling operation. Parameter changes made after a plan has been created may require calling the BUILD_PLAN procedure to modify the existing plan. The DBA_ROLLING_EVENTS view will indicate if any invocation of the SET_PARAMETER procedure requires a plan rebuild. Failure to rebuild the plan will result in an ORA-45416 error when attempting to resume the rolling operation.

Syntax

```
DBMS_ROLLING.BUILD_PLAN ();
```

Parameters

This procedure has no parameters.

Exceptions

- ORA-45400: operation not permitted on current database
- ORA-45403: database %s must be specified in the DG_CONFIG
- ORA-45414: could not connect to a remote database
- ORA-45419: DB_UNIQUE_NAME parameter must be specified
- ORA-45433: failover was detected on an unsupported database
- ORA-45434: multiple failovers of the same type detected

- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- This procedure connects to databases specified as plan parameters. These instances must be mounted or open, and must be reachable via the network.

163.3.4 SET_PARAMETER Procedure

This procedure modifies a rolling operation parameter.

Starting with Oracle Database 20c, a new parameter called BLOCK_UNSUPPORTED is added to the DBMS_ROLLING.SET_PARAMETER procedure. By default, BLOCK_UNSUPPORTED is set to 1, indicating that operations performed on tables that are unsupported by Transient Logical Standby will be blocked on the primary database. If set to 0, then the DBMS_ROLLING package does not block operations on unsupported tables. Those tables will not be maintained by Transient Logical Standby, and will diverge from the primary database.

Syntax

```
DBMS_ROLLING.SET_PARAMETER (
   scope             IN VARCHAR2 DEFAULT NULL,
   name              IN NUMBER,
   value             IN VARCHAR2);
```

Parameters

Table 163-3 *SET_PARAMETER Procedure Parameters*

Parameter	Description
scope	Parameter scope. It can either be NULL for global parameters, or the DB_UNIQUE_NAME of a specific database for local parameters.
name	The DBMS_ROLLING constant for a given parameter.
value	New value for the parameter or NULL to revert to a default value.

Exceptions

- ORA-45400: operation not permitted on current database
- ORA-45408: parameter name is unknown
- ORA-45409: parameter value is invalid or out of bounds
- ORA-45410: parameter may not be modified
- ORA-45411: operation requires additional arguments
- ORA-45412: parameter scope argument is unknown
- ORA-45413: parameter has no default value
- ORA-45414: could not connect to a remote database
- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- Changes to a parameter value may require a call to the DBMS_ROLLING.BUILD_PLAN procedure to modify the existing plan. Users should check the DBA_ROLLING_EVENTS view after setting a parameter to determine if a rebuild is necessary.
- Table 163-4 lists all the available parameters and their descriptions. The parameter names and values described in the table are all of type VARCHAR2.
- The MINVAL and MAXVAL columns in the DBA_ROLLING_PARAMETERS view identify the valid range of values for a parameter. The view does not contain any parameters until the DBMS_ROLLING.INIT_PLAN procedure has been successfully invoked.

Table 163-4 Valid Values for DBMS_ROLLING.SET_PARAMETER Procedure

Parameter Name	Global?	Description	Default
ACTIVE_SESSIONS_TIMEOUT	Yes	The maximum amount of time in seconds to enforce ACTIVE_SESSIONS_WAIT before halting the rolling upgrade. This parameter is only valid if ACTIVE_SESSIONS_WAIT is set to 1.	3600
ACTIVE_SESSIONS_WAIT	Yes	Whether the switchover operation will wait for active sessions to finish. If set to 1, the SWITCHOVER procedure waits for active sessions to compete. If set to 0, the SWITCHOVER procedure kills active sessions to expedite the switchover.	0
BACKUP_CONTROLFILE	Yes	File name of the backup control file that is created during a rolling upgrade.	rolling_change_backup.f
BLOCK_UNSUPPORTED	Yes	This parameter can be set for Rolling Upgrade to Block Unsupported Operations on the source database. Possible values are 0 and 1. • 1: Turns blocking on. This is the default value. • 0: Turns blocking off. Scope is not considered for this parameter.	1
DGBROKER	Yes	Use Data Guard broker for managing apply, recovery, and log archive destinations.	1 if broker is enabled, 0 otherwise.
DICTIONARY_LOAD_TIMEOUT	Yes	The maximum amount of time in seconds to enforce DICTIONARY_LOAD_WAIT before halting the rolling upgrade. This parameter is only valid if DICTIONARY_LOAD_WAIT is set to 1.	3600
DICTIONARY_LOAD_WAIT	Yes	Whether the instantiation of the transient logical standby will include a wait for the complete loading of the data dictionary snapshot in redo. If set to 1, then the START_PLAN procedure will not return until the dictionary has been completely loaded. If set to 0, then the START_PLAN procedure will only verify that the loading of the dictionary has started.	0
DICTIONARY_PLS_WAIT_INIT	Yes	The time in seconds to wait in between attempts to quiesce PL/SQL activity in order to write the data dictionary to redo.	300
DICTIONARY_PLS_WAIT_TIMEOUT	Yes	The maximum amount of time in seconds to attempt to quiesce PL/SQL activity in order to write the data dictionary to redo.	3600
EVENT_RECORDS	Yes	The maximum number of records to permit in DBA_ROLLING_EVENTS	10000

Table 163-4 (Cont.) Valid Values for DBMS_ROLLING.SET_PARAMETER Procedure

Parameter Name	Global?	Description	Default
FAILOVER	Yes	Automatically attempt to adjust the upgrade plan as a result of a failover event. This parameter resets its value to 0 upon completion of a subsequent call to BUILD_PLAN.	0
GRP_PREFIX	Yes	Execution of procedures in DBMS_ROLLING results in a number of Guaranteed Restore Points (GRP) taken in various databases participating in the Data Guard configuration. All such GRPs have the same prefix in their names. You can use this parameter to override the default prefix.	DBMSRU
IGNORE_BUILD_WARNINGS	Yes	Ignore warnings which would otherwise raise exceptions during execution of the BUILD_PLAN procedure.	1
IGNORE_LAST_ERROR	Yes	Ignore last encountered error upon startup of next rolling operation. This parameter resets its value to 0 upon invocation of a procedure call which resumes the rolling upgrade.	0
LAD_ENABLED_TIMEOUT	Yes	The maximum time in seconds to wait for a recently enabled log archive destination to reach a VALID state.	600
LOG_LEVEL	Yes	Logging level for the DBS_ROLLING PL/SQL package. A value of INFO results in the logging of errors and relevant non-fatal warnings. A value of FULL results in the logging of all events.	INFO
MEMBER	No	The upgrade group in which the specified database is a member. A value of LEADING indicates that the standby is a member of the leading upgrade group. As such, it is a standby of the Leading Group Master (LGM). The LGM is the database which is converted into the transient logical standby, and which becomes the new primary after the switchover. A value of TRAILING indicates that the standby is a member of the trailing upgrade group. As such, it is a standby of the Trailing Group Master (TGM). The TGM is the original primary database.	LEADING
READY_LGM_LAG_TIME	Yes	The apply lag time in seconds associated with the READY_LGM_LAG_WAIT parameter.	600
READY_LGM_LAG_TIMEOUT	Yes	The maximum amount of time in seconds to enforce READY_LGM_LAG_WAIT before halting the rolling upgrade. This parameter is only valid if READY_LGM_LAG_WAIT is set to 1.	60
READY_LGM_LAG_WAIT	Yes	Whether the START_PLAN procedure will wait for the apply lag on the leading group master to fall below READY_LGM_LAG_TIME seconds before returning control back to the user. If set to 1, the wait is performed. If set to 0, the wait is not performed.	0
SWITCH_LGM_LAG_TIME	Yes	The apply lag time in seconds associated with the SWITCH_LGM_LAG_WAIT parameter.	600
SWITCH_LGM_LAG_TIMEOUT	Yes	The maximum amount of time in seconds to enforce SWITCH_LGM_LAG_WAIT before halting the rolling upgrade. This parameter is only valid if SWITCH_LGM_LAG_WAIT is set to 1.	60
SWITCH_LGM_LAG_WAIT	Yes	Whether the SWITCHOVER procedure will wait for the apply lag on the leading group master to fall below SWITCH_LGM_LAG_TIME seconds before initiating the switchover. If set to 1, the wait is performed. If set to 0, the wait is not performed.	1

Table 163-4 (Cont.) Valid Values for DBMS_ROLLING.SET_PARAMETER Procedure

Parameter Name	Global?	Description	Default
SWITCH_LGS_LAG_TIME	Yes	The apply lag time in seconds associated with the SWITCH_LGS_LAG_WAIT parameter.	60
SWITCH_LGS_LAG_TIMEOUT	Yes	The maximum amount of time in seconds to enforce SWITCH_LGS_LAG_WAIT before halting the rolling upgrade. This parameter is only valid if SWITCH_LGS_LAG_WAIT is set to 1.	60
SWITCH_LGS_LAG_WAIT	Yes	Whether the SWITCHOVER procedure will wait for the apply lag on the leading group standbys to fall below SWITCH_LGS_LAG_TIME seconds before initiating the switchover. If set to 1, the wait is performed. If set to 0, the wait is not performed.	0
UPDATED_LGS_TIMEOUT	Yes	The maximum amount of time in seconds to enforce UPDATED_LGS_WAIT before halting the rolling upgrade. This parameter is only valid if UPDATED_LGS_WAIT is set to 1.	10800
UPDATED_LGS_WAIT	Yes	Whether the SWITCHOVER procedure will wait for the leading group standbys to complete recovery of all upgrade redo before initiating the switchover. If set to 1, the wait is performed. If set to 0, the wait is not performed	1
UPDATED_TGS_TIMEOUT	Yes	The maximum amount of time in seconds to enforce UPDATED_TGS_WAIT before halting the rolling upgrade. This parameter is only valid if UPDATED_TGS_WAIT is set to 1.	10800
UPDATED_TGS_WAIT	Yes	Whether the FINISH_PLAN procedure will wait for the trailing group standbys to complete recovery of all upgrade redo before returning control to the user. If set to 1, the wait is performed. If set to 0, the wait is not performed.	1

163.3.5 START_PLAN Procedure

This procedure starts the rolling operation. This procedure must be executed on the primary database to formally start the rolling operation.

When the START_PLAN procedure is complete, the future_primary parameter in the INIT_PLAN procedure will be converted into a fully configured transient logical standby database.

Syntax

```
DBMS_ROLLING.START_PLAN ();
```

Parameters

This procedure has no parameters.

Exceptions

- ORA-45400: operation not permitted on current database
- ORA-45414: could not connect to a remote database
- ORA-45415: instruction execution failure
- ORA-45416: operation cannot start until plan rebuild
- ORA-45417: operation not permitted since current phase was not %s

- ORA-45422: operation requires existing plan
- ORA-45426: managed recovery process was not running
- ORA-45427: logical standby Redo Apply process was not running
- ORA-45428: database was not in expected database role
- ORA-45435: managed recovery process was running
- ORA-45436: logical standby Redo Apply process was running
- ORA-45438: database is not in mounted mode
- ORA-45439: database is not in open read/write mode
- ORA-45486: database update progress is inconsistent
- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- A rolling operation plan must have previously been generated through the BUILD_PLAN procedure.

163.3.6 SWITCHOVER Procedure

This procedure performs a switchover between the current primary database and the transient logical standby database.

At the successful completion of the procedure, the LGM assumes the primary role for the Data Guard configuration.

Syntax

```
DBMS_ROLLING.SWITCHOVER ();
```

Parameters

This procedure has no parameters.

Exceptions

- ORA-16224: database guard is enabled
- ORA-45400: operation not permitted on current database
- ORA-45414: could not connect to a remote database
- ORA-45415: instruction execution failure
- ORA-45416: operation cannot start until plan rebuild
- ORA-45417: operation not permitted since current phase was not %s
- ORA-45422: operation requires existing plan
- ORA-45426: managed recovery process was not running
- ORA-45427: logical standby Redo Apply process was not running
- ORA-45428: database was not in expected database role
- ORA-45435: managed recovery process was running
- ORA-45436: logical standby Redo Apply process was running

- ORA-45438: database is not in mounted mode
- ORA-45439: database is not in open read/write mode
- ORA-45486: database update progress is inconsistent
- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- This procedure can only be called after you have manually upgraded the transient logical standby and opened it on the higher Oracle Database version.
- Once the future primary has been upgraded, logical apply should be restarted and allowed to catch up with the primary since the primary will have been open for business and generating redo which the standby must now apply. Failing to do so will cause the switchover to take a long time since it will have to apply all of that redo.

163.3.7 FINISH_PLAN Procedure

This procedure finalizes the rolling operation.

It configures the former primary as a physical standby of the new primary by flashing it back to an earlier taken GRP (guaranteed restore point), converting its role to physical standby, and then starting managed recovery which will recover all redo generated by the new primary including upgrade redo and any other redo generated before or after the upgrade.

Syntax

```
DBMS_ROLLING.FINISH_PLAN ();
```

Parameters

This procedure has no parameters.

Exceptions

- ORA-45400: operation not permitted on current database
- ORA-45414: could not connect to a remote database
- ORA-45415: instruction execution failure
- ORA-45416: operation cannot start until plan rebuild
- ORA-45417: operation not permitted since current phase was not %s
- ORA-45422: operation requires existing plan
- ORA-45426: managed recovery process was not running
- ORA-45427: logical standby Redo Apply process was not running
- ORA-45428: database was not in expected database role
- ORA-45435: managed recovery process was running
- ORA-45436: logical standby Redo Apply process was running
- ORA-45438: database is not in mounted mode
- ORA-45439: database is not in open read/write mode
- ORA-45486: database update progress is inconsistent

- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- This procedure can only be called after you have remounted the former primary and remaining physical standbys on the higher Oracle Database version.

163.3.8 ROLLBACK_PLAN Procedure

This procedure rolls back the configuration-wide rolling operation.

Once completed, all of the databases in the leading group become physical standbys of the original primary database. This procedure can only be called if the configuration has not yet gone through a switchover operation since the START_PLAN procedure was invoked.

Syntax

```
DBMS_ROLLING.ROLLBACK_PLAN;
```

Parameters

This procedure has no parameters.

Exceptions

- ORA-45400: operation not permitted on current database
- ORA-45414: could not connect to a remote database
- ORA-45415: instruction execution failure
- ORA-45441: no databases eligible for rollback
- ORA-45442: rollback is not permitted after a role change
- ORA-65040: operation not allowed from within a pluggable database

Usage Notes

- You must manually restart media recovery on the lower Oracle Database version if the upgrade of the transient logical standby has already been performed.

164
DBMS_ROWID

The DBMS_ROWID package lets you create ROWIDs and obtain information about ROWIDs from PL/SQL programs and SQL statements. You can find the data block number, the object number, and other ROWID components without writing code to interpret the base-64 character external ROWID. DBMS_ROWID is intended for upgrading from Oracle database version 7 to Oracle database version 8.X.

> **Note:**
>
> DBMS_ROWID is not to be used with universal ROWIDs (UROWIDs).

This chapter contains the following topics:

- Security Model
- Types
- Exceptions
- Operational Notes
- Examples
- Summary of DBMS_ROWID Subprograms

164.1 DBMS_ROWID Security Model

This package runs with the privileges of calling user, rather than the package owner SYS.

164.2 DBMS_ROWID Types

There are four DBMS_ROWID types.

These are:

- Extension and restriction types
- Verification types
- Object types
- Conversion types

Extension and Restriction Type

The types are as follows:

- RESTRICTED—restricted ROWID
- EXTENDED—extended ROWID

For example:

```
rowid_type_restricted constant integer := 0;
rowid_type_extended   constant integer := 1;
```

> **Note:**
>
> Extended ROWIDs are only used in Oracle database version 8.X*i* and higher.

Verification Types

Table 164-1 Verification Types

Result	Description
VALID	Valid ROWID
INVALID	Invalid ROWID

For example:

```
rowid_is_valid   constant integer := 0;
rowid_is_invalid constant integer := 1;
```

Object Types

Table 164-2 Object Types

Result	Description
UNDEFINED	Object Number not defined (for restricted ROWIDs)

For example:

```
rowid_object_undefined constant integer := 0;
```

Conversion Types

Table 164-3 Conversion Types

Result	Description
INTERNAL	Convert to/from column of ROWID type
EXTERNAL	Convert to/from string format

For example:

```
rowid_convert_internal constant integer := 0;
rowid_convert_external constant integer := 1;
```

164.3 DBMS_ROWID Exceptions

This table describes the Exceptions raised by DBMS_ROWID subprograms.

Table 164-4 Exceptions

Exception	Description
ROWID_INVALID	Invalid rowid format
ROWID_BAD_BLOCK	Block is beyond end of file

For example:

```
ROWID_INVALID exception;
   pragma exception_init(ROWID_INVALID, -1410);

ROWID_BAD_BLOCK exception;
   pragma exception_init(ROWID_BAD_BLOCK, -28516);
```

164.4 DBMS_ROWID Operational Notes

These operation notes apply to DBMS_ROWID.

- Some of the functions in this package take a single parameter, such as a ROWID. This can be a character or a PL/SQL ROWID, either restricted or extended, as required.
- You can call the DBMS_ROWID functions and procedures from PL/SQL code, and you can also use the functions in SQL statements.

> **Note:**
>
> ROWID_INFO is a procedure. It can only be used in PL/SQL code.

- You can use functions from the DBMS_ROWID package just like built-in SQL functions; in other words, you can use them wherever you can use an expression. In this example, the ROWID_BLOCK_NUMBER function is used to return just the block number of a single row in the EMP table:

```
SELECT DBMS_ROWID.ROWID_BLOCK_NUMBER(rowid)
   FROM emp
   WHERE ename = 'KING';
```

- If Oracle returns the error "ORA:452, 0, 'Subprogram '%s' violates its associated pragma' for pragma restrict_references, it could mean the violation is due to:
  - A problem with the current procedure or function
  - Calling a procedure or function without a pragma or due to calling one with a less restrictive pragma
  - Calling a package procedure or function that touches the initialization code in a package or that sets the default values

164.5 DBMS_ROWID Examples

This example returns the ROWID for a row in the EMP table, extracts the data object number from the ROWID, using the ROWID_OBJECT function in the DBMS_ROWID package, then displays the object number:

```
DECLARE
  object_no   INTEGER;
  row_id      ROWID;
  ...
BEGIN
  SELECT ROWID INTO row_id FROM emp
    WHERE empno = 7499;
  object_no := DBMS_ROWID.ROWID_OBJECT(row_id);
  DBMS_OUTPUT.PUT_LINE('The obj. # is '|| object_no);
  ...
```

164.6 Summary of DBMS_ROWID Subprograms

This table lists the DBMS_ROWID subprograms and briefly describes them.

Table 164-5 DBMS_ROWID Package Subprograms

Subprogram	Description
ROWID_BLOCK_NUMBER Function	Returns the block number of a ROWID
ROWID_CREATE Function	Creates a ROWID, for testing only
ROWID_INFO Procedure	Returns the type and components of a ROWID
ROWID_OBJECT Function	Returns the object number of the extended ROWID
ROWID_RELATIVE_FNO Function	Returns the file number of a ROWID
ROWID_ROW_NUMBER Function	Returns the row number
ROWID_TO_ABSOLUTE_FNO Function	Returns the absolute file number associated with the ROWID for a row in a specific table
ROWID_TO_EXTENDED Function	Converts a ROWID from restricted format to extended
ROWID_TO_RESTRICTED Function	Converts an extended ROWID to restricted format
ROWID_TYPE Function	Returns the ROWID type: 0 is restricted, 1 is extended
ROWID_VERIFY Function	Checks if a ROWID can be correctly extended by the ROWID_TO_EXTENDED function

164.6.1 ROWID_BLOCK_NUMBER Function

This function returns the database block number for the input ROWID.

Syntax

```
DBMS_ROWID.ROWID_BLOCK_NUMBER (
   row_id      IN   ROWID,
   ts_type_in  IN   VARCHAR2 DEFAULT 'SMALLFILE')
  RETURN NUMBER;
```

Pragmas

```
pragma RESTRICT_REFERENCES(rowid_block_number,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-6 ROWID_BLOCK_NUMBER Function Parameters

Parameter	Description
row_id	ROWID to be interpreted
ts_type_in	The type of the tablespace (bigfile/smallfile) to which the row belongs

Examples

The example SQL statement selects the block number from a ROWID and inserts it into another table:

```
INSERT INTO T2 (SELECT dbms_rowid.rowid_block_number(ROWID, 'BIGFILE')
   FROM some_table
   WHERE key_value = 42);
```

164.6.2 ROWID_CREATE Function

This function lets you create a ROWID, given the component parts as parameters.

This is useful for testing ROWID operations, because only the Oracle Server can create a valid ROWID that points to data in a database.

Syntax

```
DBMS_ROWID.ROWID_CREATE (
   rowid_type    IN NUMBER,
   object_number IN NUMBER,
   relative_fno  IN NUMBER,
   block_number  IN NUMBER,
   row_number    IN NUMBER)
  RETURN ROWID;
```

Pragmas

```
pragma RESTRICT_REFERENCES(rowid_create,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-7 ROWID_CREATE Function Parameters

Parameter	Description
rowid_type	Type (restricted or extended) Set the rowid_type parameter to 0 for a restricted ROWID. Set it to 1 to create an extended ROWID. If you specify rowid_type as 0, then the required object_number parameter is ignored, and ROWID_CREATE returns a restricted ROWID.
object_number	Data object number (rowid_object_undefined for restricted)
relative_fno	Relative file number

Table 164-7 (Cont.) ROWID_CREATE Function Parameters

Parameter	Description
block_number	Block number in this file
row_number	Returns row number in this block

Examples

Create a dummy extended ROWID:

```
my_rowid := DBMS_ROWID.ROWID_CREATE(1, 9999, 12, 1000, 13);
```

Find out what the rowid_object function returns:

```
obj_number := DBMS_ROWID.ROWID_OBJECT(my_rowid);
```

The variable obj_number now contains 9999.

164.6.3 ROWID_INFO Procedure

This procedure returns information about a ROWID, including its type (restricted or extended), and the components of the ROWID.

This is a procedure, and it cannot be used in a SQL statement.

Syntax

```
DBMS_ROWID.ROWID_INFO (
   rowid_in          IN   ROWID,
   rowid_type        OUT  NUMBER,
   object_number     OUT  NUMBER,
   relative_fno      OUT  NUMBER,
   block_number      OUT  NUMBER,
   row_number        OUT  NUMBER);
```

Pragmas

```
pragma RESTRICT_REFERENCES(rowid_info,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-8 ROWID_INFO Procedure Parameters

Parameter	Description
rowid_in	ROWID to be interpreted. This determines if the ROWID is a restricted (0) or extended (1) ROWID.
rowid_type	Returns type (restricted/extended)
object_number	Returns data object number (rowid_object_undefined for restricted)
relative_fno	Returns relative file number
block_number	Returns block number in this file
row_number	Returns row number in this block

> **See Also:**
>
> "ROWID_TYPE Function"

Examples

This example reads back the values for the ROWID that you created in the ROWID_CREATE:

```
DBMS_ROWID.ROWID_INFO (
   my_rowid, rid_type, obj_num, file_num, block_num, row_num, 'BIGFILE');
```

164.6.4 ROWID_OBJECT Function

This function returns the data object number for an extended ROWID.

The function returns zero if the input ROWID is a restricted ROWID.

Syntax

```
DBMS_ROWID.ROWID_OBJECT (
   rowid_id IN ROWID)
  RETURN NUMBER;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_object,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-9 ROWID_OBJECT Function Parameters

Parameter	Description
row_id	ROWID to be interpreted

> **Note:**
>
> The ROWID_OBJECT_UNDEFINED constant is returned for restricted ROWIDs.

Examples

```
SELECT dbms_rowid.rowid_object(ROWID)
   FROM emp
   WHERE empno = 7499;
```

164.6.5 ROWID_RELATIVE_FNO Function

This function returns the relative file number of the `ROWID` specified as the `IN` parameter. (The file number is relative to the tablespace.)

Syntax

```
DBMS_ROWID.ROWID_RELATIVE_FNO (
   rowid_id      IN   ROWID,
   ts_type_in    IN   VARCHAR2 DEFAULT 'SMALLFILE')
  RETURN NUMBER;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_relative_fno,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-10 ROWID_RELATIVE_FNO Function Parameters

Parameter	Description
`row_id`	`ROWID` to be interpreted
`ts_type_in`	Type of the tablespace (bigfile/smallfile) to which the row belongs

Examples

The example PL/SQL code fragment returns the relative file number:

```
DECLARE
   file_number    INTEGER;
   rowid_val      ROWID;
BEGIN
   SELECT ROWID INTO rowid_val
     FROM dept
     WHERE loc = 'Boston';
   file_number :=
     dbms_rowid.rowid_relative_fno(rowid_val, 'SMALLFILE');
    ...
```

164.6.6 ROWID_ROW_NUMBER Function

This function extracts the row number from the `ROWID IN` parameter.

Syntax

```
DBMS_ROWID.ROWID_ROW_NUMBER (
   row_id IN ROWID)
  RETURN NUMBER;
```

Pragmas

```
   PRAGMA RESTRICT_REFERENCES(rowid_row_number,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-11 ROWID_ROW_NUMBER Function Parameters

Parameter	Description
row_id	ROWID to be interpreted.

Examples

Select a row number:

```
SELECT dbms_rowid.rowid_row_number(ROWID)
   FROM emp
   WHERE ename = 'ALLEN';
```

164.6.7 ROWID_TO_ABSOLUTE_FNO Function

This function extracts the absolute file number from a ROWID, where the file number is absolute for a row in a given schema and table.

The schema name and the name of the schema object (such as a table name) are provided as IN parameters for this function.

Syntax

```
DBMS_ROWID.ROWID_TO_ABSOLUTE_FNO (
   row_id      IN ROWID,
   schema_name IN VARCHAR2,
   object_name IN VARCHAR2)
  RETURN NUMBER;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_to_absolute_fno,WNDS,WNPS,RNPS);
```

Parameters

Table 164-12 ROWID_TO_ABSOLUTE_FNO Function Parameters

Parameter	Description
row_id	ROWID to be interpreted
schema_name	Name of the schema which contains the table
object_name	Table name

Examples

```
DECLARE
   abs_fno        INTEGER;
   rowid_val      CHAR(18);
   object_name    VARCHAR2(20) := 'EMP';
BEGIN
    SELECT ROWID INTO rowid_val
     FROM emp
     WHERE empno = 9999;
```

```
abs_fno := dbms_rowid.rowid_to_absolute_fno(
rowid_val, 'SCOTT', object_name);
```

> **Note:**
>
> For partitioned objects, the name must be a table name, not a partition or a sub/partition name.

164.6.8 ROWID_TO_EXTENDED Function

This function translates a restricted `ROWID` that addresses a row in a schema and table that you specify to the extended `ROWID` format.

Later, it may be removed from this package into a different place.

Syntax

```
DBMS_ROWID.ROWID_TO_EXTENDED (
   old_rowid       IN ROWID,
   schema_name     IN VARCHAR2,
   object_name     IN VARCHAR2,
   conversion_type IN INTEGER)
  RETURN ROWID;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_to_extended,WNDS,WNPS,RNPS);
```

Parameters

Table 164-13 ROWID_TO_EXTENDED Function Parameters

Parameter	Description
`old_rowid`	`ROWID` to be converted
`schema_name`	Name of the schema which contains the table (optional)
`object_name`	Table name (optional).
`conversion_type`	The following constants are defined: `ROWID_CONVERT_INTERNAL` (:=0) `ROWID_CONVERT_EXTERNAL` (:=1)

Return Values

`ROWID_TO_EXTENDED` returns the `ROWID` in the extended character format. If the input `ROWID` is `NULL`, then the function returns `NULL`. If a zero-valued `ROWID` is supplied (00000000.0000.0000), then a zero-valued restricted `ROWID` is returned.

Examples

Assume that there is a table called `RIDS` in the schema `SCOTT`, and that the table contains a column `ROWID_COL` that holds `ROWIDs` (restricted), and a column `TABLE_COL` that point to other tables in the `SCOTT` schema. You can convert the `ROWIDs` to extended format with the statement:

```
UPDATE SCOTT.RIDS
   SET rowid_col =
   dbms_rowid.rowid_to_extended (
      rowid_col, 'SCOTT", TABLE_COL, 0);
```

Usage Notes

- If the schema and object names are provided as `IN` parameters, then this function verifies `SELECT` authority on the table named, and converts the restricted `ROWID` provided to an extended `ROWID`, using the data object number of the table. That `ROWID_TO_EXTENDED` returns a value, however, does not guarantee that the converted `ROWID` actually references a valid row in the table, either at the time that the function is called, or when the extended `ROWID` is actually used.
- If the schema and object name are not provided (are passed as `NULL`), then this function attempts to fetch the page specified by the restricted `ROWID` provided. It treats the file number stored in this `ROWID` as the absolute file number. This can cause problems if the file has been dropped, and its number has been reused prior to the migration. If the fetched page belongs to a valid table, then the data object number of this table is used in converting to an extended `ROWID` value. This is very inefficient, and Oracle recommends doing this only as a last resort, when the target table is not known. The user must still know the correct table name at the time of using the converted value.
- If an extended `ROWID` value is supplied, the data object number in the input extended `ROWID` is verified against the data object number computed from the table name parameter. If the two numbers do not match, the `INVALID_ROWID` exception is raised. If they do match, the input `ROWID` is returned.
- `ROWID_TO_EXTENDED` cannot be used with partition tables.

> **See Also:**
>
> The ROWID_VERIFY Function has a method to determine if a given `ROWID` can be converted to the extended format.

164.6.9 ROWID_TO_RESTRICTED Function

This function converts an extended `ROWID` into restricted `ROWID` format.

Syntax

```
DBMS_ROWID.ROWID_TO_RESTRICTED (
   old_rowid       IN ROWID,
   conversion_type IN INTEGER)
  RETURN ROWID;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_to_restricted,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-14 ROWID_TO_RESTRICTED Function Parameters

Parameter	Description
old_rowid	ROWID to be converted
conversion_type	The following constants are defined: ROWID_CONVERT_INTERNAL (:=0) ROWID_CONVERT_EXTERNAL (:=1)

164.6.10 ROWID_TYPE Function

This function returns 0 if the ROWID is a restricted ROWID, and 1 if it is extended.

Syntax

```
DBMS_ROWID.ROWID_TYPE (
   rowid_id IN ROWID)
  RETURN NUMBER;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_type,WNDS,RNDS,WNPS,RNPS);
```

Parameters

Table 164-15 ROWID_TYPE Function Parameters

Parameter	Description
row_id	ROWID to be interpreted

Examples

```
IF DBMS_ROWID.ROWID_TYPE(my_rowid) = 1 THEN
   my_obj_num := DBMS_ROWID.ROWID_OBJECT(my_rowid);
```

164.6.11 ROWID_VERIFY Function

This function verifies the ROWID.

It returns 0 if the input restricted ROWID can be converted to extended format, given the input schema name and table name, and it returns 1 if the conversion is not possible.

> **Note:**
>
> You can use this function in a WHERE clause of a SQL statement, as shown in the example.

Syntax

```
DBMS_ROWID.ROWID_VERIFY (
   rowid_in        IN ROWID,
   schema_name     IN VARCHAR2,
   object_name     IN VARCHAR2,
   conversion_type IN INTEGER
  RETURN NUMBER;
```

Pragmas

```
   pragma RESTRICT_REFERENCES(rowid_verify,WNDS,WNPS,RNPS);
```

Parameters

Table 164-16 ROWID_VERIFY Function Parameters

Parameter	Description
rowid_in	ROWID to be verified
schema_name	Name of the schema which contains the table
object_name	Table name
conversion_type	The following constants are defined: ROWID_CONVERT_INTERNAL (:=0) ROWID_CONVERT_EXTERNAL (:=1)

Examples

Considering the schema in the example for the ROWID_TO_EXTENDED function, you can use the following statement to find bad ROWIDs prior to conversion. This enables you to fix them beforehand.

```
SELECT ROWID, rowid_col
   FROM SCOTT.RIDS
   WHERE dbms_rowid.rowid_verify(rowid_col, NULL, NULL, 0) =1;
```

See Also:

UTL_RAW , UTL_REF

165
DBMS_RULE

The `DBMS_RULE` package contains subprograms that enable the evaluation of a rule set for a specified event.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_RULE Subprograms

165.1 DBMS_RULE Overview

This package contains subprograms that enable the evaluation of a rule set for a specified event.

> **See Also:**
>
> - Rule TYPEs for more information about the types used with the `DBMS_RULE` package
> - DBMS_RULE_ADM

165.2 DBMS_RULE Security Model

`PUBLIC` is granted `EXECUTE` privilege on this package.

> **See Also:**
>
> *Oracle Database Security Guide* for more information about user group `PUBLIC`

165.3 Summary of DBMS_RULE Subprograms

This table lists the `DBMS_RULE` subprograms and briefly describes them.

Table 165-1 DBMS_RULE Package Subprograms

Subprogram	Description
CLOSE_ITERATOR Procedure	Closes an open iterator
EVALUATE Procedure	Evaluates the rules in the specified rule set that use the evaluation context specified

Table 165-1 (Cont.) DBMS_RULE Package Subprograms

Subprogram	Description
EVALUATE_EXPRESSION Procedure	Evaluates an expression under the logged in user in a session
EVALUATE_EXPRESSION_ITERATOR Procedure	Finds the relevant datapoints and pass `re$value_list` into evaluation interface
EVALUATE_RULE Procedure	Evaluates the condition defined in the Rule
EVALUATE_RULE_ITERATOR Procedure	Finds the relevant datapoints and pass `re$value_list` into evaluation interface
GET_NEXT_HIT Function	Returns the next rule that evaluated to `TRUE` from a true rules iterator, or returns the next rule that evaluated to `MAYBE` from a maybe rules iterator; returns `NULL` if there are no more rules that evaluated to `TRUE` or `MAYBE`.
IS_FAST Procedure	Returns `TRUE` if the expression can be evaluated fast. An expression can be evaluated fast if the engine does not need to run any internal SQL and does not need to go to PL/SQL layer in case there are any PL/SQL functions referred.
GET_NEXT_RESULT Function	Iterates over result from the expression given `result_val_iterator`.

165.3.1 CLOSE_ITERATOR Procedure

This procedure closes an open iterator.

Syntax

```
DBMS_RULE.CLOSE_ITERATOR(
   iterator  IN  BINARY_INTEGER);
```

Parameter

Table 165-2 CLOSE_ITERATOR Procedure Parameter

Parameter	Description
`iterator`	Iterator to be closed

Usage Notes

This procedure requires an open iterator that was returned by an earlier call to `DBMS_RULE.EVALUATE` in the same session. The user who runs this procedure does not require any privileges on the rule set being evaluated.

Closing an iterator frees resources, such as memory, associated with the iterator. Therefore, Oracle recommends that you close an iterator when it is no longer needed.

> **See Also:**
>
> EVALUATE Procedure

165.3.2 EVALUATE Procedure

This procedure evaluates the rules in the specified rule set that use the evaluation context specified for a specified event.

This procedure is overloaded. The `true_rules` and `maybe_rules` parameters are mutually exclusive with the `true_rules_iterator` and `maybe_rules_iterator` parameters. In addition, the procedure with the `true_rules` and `maybe_rules` parameters includes the `stop_on_first_hit` parameter, but the other procedure does not.

Syntax

```
DBMS_RULE.EVALUATE(
   rule_set_name          IN     VARCHAR2,
   evaluation_context     IN     VARCHAR2,
   event_context          IN     SYS.RE$NV_LIST               DEFAULT NULL,
   table_values           IN     SYS.RE$TABLE_VALUE_LIST      DEFAULT NULL,
   column_values          IN     SYS.RE$COLUMN_VALUE_LIST     DEFAULT NULL,
   variable_values        IN     SYS.RE$VARIABLE_VALUE_LIST   DEFAULT NULL,
   attribute_values       IN     SYS.RE$ATTRIBUTE_VALUE_LIST  DEFAULT NULL,

  skip_rules           IN   SYS.RE$RULE_NAME_LIST        DEFAULT NULL,

  dop                  IN   NUMBER,

   result_cache           IN     BOOLEAN                      DEFAULT FALSE,
   stop_on_first_hit      IN     BOOLEAN                      DEFAULT FALSE,
   simple_rules_only      IN     BOOLEAN                      DEFAULT FALSE,
   true_rules             OUT    SYS.RE$RULE_HIT_LIST,
   maybe_rules            OUT    SYS.RE$RULE_HIT_LIST);

DBMS_RULE.EVALUATE(
   rule_set_name          IN     VARCHAR2,
   evaluation_context     IN     VARCHAR2,
   event_context          IN     SYS.RE$NV_LIST               DEFAULT NULL,
   table_values           IN     SYS.RE$TABLE_VALUE_LIST      DEFAULT NULL,
   column_values          IN     SYS.RE$COLUMN_VALUE_LIST     DEFAULT NULL,
   variable_values        IN     SYS.RE$VARIABLE_VALUE_LIST   DEFAULT NULL,
   attribute_values       IN     SYS.RE$ATTRIBUTE_VALUE_LIST  DEFAULT NULL,

  skip_rules            IN   SYS.RE$RULE_NAME_LIST        DEFAULT NULL,

  dop                   IN   NUMBER,

   simple_rules_only      IN     BOOLEAN                      DEFAULT FALSE,
   true_rules_iterator    OUT    BINARY_INTEGER,
   maybe_rules_iterator   OUT    BINARY_INTEGER);
```

Parameters

Table 165-3 EVALUATE Procedure Parameters

Parameter	Description
`rule_set_name`	Name of the rule set in the form `[`*`schema_name`*`.]`*`rule_set_name`*. For example, to evaluate all of the rules in a rule set named `hr_rules` in the `hr` schema, enter `hr.hr_rules` for this parameter. If the schema is not specified, then the schema of the current user is used.

Table 165-3 (Cont.) EVALUATE Procedure Parameters

Parameter	Description
evaluation_context	An evaluation context name in the form [*schema_name.*]*evaluation_context_name*. If the schema is not specified, then the name of the current user is used. Only rules that use the specified evaluation context are evaluated.
event_context	A list of name-value pairs that identify events that cause evaluation
table_values	Contains the data for table rows using the table aliases specified when the evaluation context was created. Each table alias in the list must be unique.
column_values	Contains the partial data for table rows. It must not contain column values for tables, whose values are already specified in table_values.
variable_values	A list containing the data for variables. The only way for an explicit variable value to be known is to specify its value in this list. If an implicit variable value is not specified in the list, then the function used to obtain the value of the implicit variable is invoked. If an implicit variable value is specified in the list, then this value is used and the function is not invoked.
attribute_values	Contains the partial data for variables. It must not contain attribute values for variables whose values are already specified in variable_values.
stop_on_first_hit	If TRUE, then the rules engine stops evaluation as soon as it finds a TRUE rule. If TRUE and there are no TRUE rules, then the rules engine stops evaluation as soon as it finds a rule that may evaluate to TRUE given more data. If FALSE, then the rules engine continues to evaluate rules even after it finds a TRUE rule.
simple_rules_only	If TRUE, then only those rules that are simple enough to be evaluated fast (without issuing SQL) are considered for evaluation. If FALSE, then evaluates all rules.
true_rules	Receives the output of the EVALUATE procedure into a varray of RE$RULE_HIT_LIST type. If no rules evaluate to TRUE, then true_rules is empty. If at least one rule evaluates to TRUE and stop_on_first_hit is TRUE, then true_rules contains one rule that evaluates to TRUE. If stop_on_first_hit is FALSE, then true_rules contains all rules that evaluate to TRUE.
maybe_rules	If all rules can be evaluated completely, without requiring any additional data, then maybe_rules is empty. If stop_on_first_hit is TRUE, then if there is at least one rule that may evaluate to TRUE given more data, and no rules evaluate to TRUE, then maybe_rules contains one rule that may evaluate to TRUE. If stop_on_first_hit is FALSE, then maybe_rules contains all rules that may evaluate to TRUE given more data.
true_rules_iterator	Contains the iterator for accessing rules that are TRUE

Table 165-3 (Cont.) EVALUATE Procedure Parameters

Parameter	Description
maybe_rules_iterator	Contains the iterator for accessing rules that may be TRUE given additional data or the ability to issue SQL
skip_rules	List of rules to skip within this evaluation.
dop	Degree of parallelism
result_cache	If TRUE, Result Cache will be created. If evalate procedure is called with either true_rules_iterator or maybe_rules_iterator, then result_cache is not enabled.

Usage Notes

> **Note:**
>
> Rules in the rule set that use an evaluation context different from the one specified are not considered for evaluation.

The rules in the rule set are evaluated using the data specified for table_values, column_values, variable_values, and attribute_values. These values must refer to tables and variables in the specified evaluation context. Otherwise, an error is raised.

The caller may specify, using stop_on_first_hit, if evaluation must stop as soon as the first TRUE rule or the first MAYBE rule (if there are no TRUE rules) is found.

The caller may also specify, using simple_rules_only, if only rules that are simple enough to be evaluated fast (which means without SQL) should be considered for evaluation. This makes evaluation faster, but causes rules that cannot be evaluated without SQL to be returned as MAYBE rules.

Partial evaluation is supported. The EVALUATE procedure can be called with data for only some of the tables, columns, variables, or attributes. In such a case, rules that cannot be evaluated because of a lack of data are returned as MAYBE rules, unless they can be determined to be TRUE or FALSE based on the values of one or more simple expressions within the rule. For example, given a value of 1 for attribute "a.b" of variable "x", a rule with the following rule condition can be returned as TRUE, without a value for table "tab":

```
(:x.a.b = 1) or (tab.c > 10)
```

The results of an evaluation are the following:

- TRUE rules, which is the list of rules that evaluate to TRUE based on the given data. These rules are returned either in the OUT parameter true_rules, which returns all of the rules that evaluate to TRUE, or in the OUT parameter true_rules_iterator, which returns each rule that evaluates to TRUE one at a time.
- MAYBE rules, which is the list of rules that could not be evaluated for one of the following reasons:
  - The rule refers to data that was unavailable. For example, a variable attribute "x.a.b" is specified, but no value is specified for the variable "x", the attribute "a", or the attribute "a.b".

  - The rule is not simple enough to be evaluated fast (without SQL) and `simple_rules_only` is specified as `TRUE`, or partial data is available.

  Maybe rules are returned either in the `OUT` parameter `maybe_rules`, which returns all of the rules that evaluate to `MAYBE`, or in the `OUT` parameter `maybe_rules_iterator`, which returns each rule that evaluates to `MAYBE` one at a time.

The caller may specify whether the procedure returns all of the rules that evaluate to `TRUE` and `MAYBE` for the event or an iterator for rules that evaluate to `TRUE` and `MAYBE`. A true rules iterator enables the client to fetch each rule that evaluates to `TRUE` one at a time, and a maybe rules iterator enables the client to fetch each rule that evaluates to `MAYBE` one at a time.

If you use an iterator, then you use the `GET_NEXT_HIT` function in the `DBMS_RULE` package to retrieve the next rule that evaluates to `TRUE` or `MAYBE` from an iterator. Oracle recommends that you close an iterator if it is no longer needed to free resources, such as memory, used by the iterator. An iterator can be closed in the following ways:

- The `CLOSE_ITERATOR` procedure in the `DBMS_RULE` package is run with the iterator specified.
- The iterator returns `NULL` because no more rules evaluate to `TRUE` or `MAYBE`.
- The session in which the iterator is running ends.

To run the `DBMS_RULE.EVALUATE` procedure, a user must meet at least one of the following requirements:

- Have `EXECUTE_ON_RULE_SET` privilege on the rule set
- Have `EXECUTE_ANY_RULE_SET` system privilege
- Be the rule set owner

> **Note:**
>
> The rules engine does not invoke any actions. An action context can be returned with each returned rule, but the client of the rules engine must invoke any necessary actions.

> **See Also:**
>
> - Rule TYPEs for more information about the types used with the `DBMS_RULE` package
> - GET_NEXT_HIT Function
> - CLOSE_ITERATOR Procedure

165.3.3 EVALUATE_EXPRESSION Procedure

This procedure allows user to evaluate an expression under the logged in user in a session.

Any re-execute of the same expression with same table alias and variable type will result in reusing the same compiled context. With fixed compile cache size, its possible of aging....

Syntax

```
DBMS_RULE.EVALUATE_EXPRESSION(
   rule_expression          IN         VARCHAR2,
   table_aliases            IN         SYS.RE$TABLE_ALIAS_LIST:= NULL,
   variable_types           IN         SYS.RE$VARIABLE_TYPE_LIST:= NULL,
   table_values             IN         SYS.RE$TABLE_VALUE_LIST:= NULL,
   column_values            IN         SYS.RE$COLUMN_VALUE_LIST:=NULL,
   variable_values          IN         SYS.RE$VARIABLE_VALUE_LIST:=NULL,
   attribute_values         IN         SYS.RE$ATTRIBUTE_VALUE_LIST:=NULL,
   cache                    IN         BOOLEAN DEFAULT FALSE,
   result_val               OUT        BOOLEAN);
```

Parameters

Table 165-4 EVALUATE_EXPRESSION Procedure Parameters

Parameter	Description
rule_expression	Contains an expression string.
table_alias	Contains alias of tables referred in the expression string.
variable_types	Contains type definitions of variables used in expression.
table_values	Contains ROWID of table row for expression evaluation.
column_values	Contains values of columns referred in the expression.
variable_values	Contains values of variables referred in the expression.
attribute_values	Contains values of attributes referred in the expression.
cache	If TRUE, Result Cache will be created.
result_val	Result of the evaluation.

165.3.4 EVALUATE_EXPRESSION_ITERATOR Procedure

This is an user visible interface. Because PL/SQL based callbacks can be expensive, we provide an array based approach. The client program is assumed to find the relevant datapoints and pass re$value_list into evaluation interface. The expression evaluation engine is expected to walk through this list and evaluate expression for each datapoint (re$value_list) element.

Syntax

```
DBMS_RULE.EVALUATE_EXPRESSION_ITERATOR(
       rule_expression        IN             varchar2,
       table_aliases          IN             sys.re$table_alias_list:= NULL,
       variable_types         IN             sys.re$variable_type_list:= NULL,
       values                 IN             sys.re$value_list,
       cache                  IN             boolean DEFAULT FALSE,
       result_val_iter_id     OUT            BINARY_INTEGER)
```

Parameters

Table 165-5 EVALUATE_EXPRESSION_ITERATOR Procedure Parameter

Parameter	Description
rule_expression	Contains an expression string.
table_alias	Alias of tables referred in the above expression string.
variable_types	Type definitions of variables used in expression.
values	List of datapoint values for evaluation.
cache	If TRUE, Result Cache will be created.
result_val_iter_id	Contains iterator for result of array of values sent using value.

165.3.5 EVALUATE_RULE Procedure

The Rule Evaluation API expects that CREATE_RULE procedure has been called with an legitimate EVALUATION_CONTEXT prior. This API will evaluate the condition defined in the Rule.

Syntax

```
DBMS_RULE.EVALUATE_RULE(
  rule_name            IN   VARCHAR2,
  event_context        IN   SYS.RE$NV_LIST               DEFAULT NULL,
  table_values         IN   SYS.RE$TABLE_VALUE_LIST      DEFAULT NULL,
  column_values        IN   SYS.RE$COLUMN_VALUE_LIST     DEFAULT NULL,
  variable_values      IN   SYS.RE$VARIABLE_VALUE_LIST   DEFAULT NULL,
  attribute_values     IN   SYS.RE$ATTRIBUTE_VALUE_LIST  DEFAULT NULL,
  cache                IN   BOOLEAN DEFAULT FALSE,
  result_val           OUT  BOOLEAN);
```

Parameters

Table 165-6 EVALUATE_RULE Procedure Parameter

Parameter	Description
rule_name	Name of the rule previously create using CREATE_RULE procedure.
event_context	A list of name-value pairs that identify events that cause evaluation.
table_values	ROWID of table row for expression evaluation.
column_values	Values of columns referred in the expression
variable_values	Values of variables referred in expression
attribute_values	Values of attributes referred in expression
cache	If TRUE, Result Cache will be created.
result_val	Result of the evaluation

165.3.6 EVALUATE_RULE_ITERATOR Procedure

This is an iterative interface. The client program is assumed to find the relevant datapoints and pass `re$value_list` into evaluation interface.

Evaluation engine is expected to walk through this list and evaluate expression for each datapoint (`re$value_list`) element. User can use `DBMS_RULE.GET_NEXT_RESULT` procedure to iterate through the result list.

Syntax

```
DBMS_RULE.EVALUATE_RULE_ITERATOR)
  rule_name             IN       VARCHAR2,
  event_context         IN       SYS.RE$NV_LIST DEFAULT NULL,
  values                IN       SYS.RE$VALUE_LIST,
  cache                 IN       BOOLEAN DEFAULT FALSE,
  result_val_iter_id    OUT      BINARY_INTEGER);
```

Parameters

Table 165-7 EVALUATE_RULE_ITERATOR Procedure Parameter

Parameter	Description
`rule_name`	Name of the rule previously create using `CREATE_RULE` procedure.
`event_context`	A list of name-value pairs that identify events that cause evaluation
`values`	List of datapoint values for evaluation.
`cache`	If `TRUE`, Result Cache will be created.
`result_val_iter_id`	Contains iterator for result of array of values sent using `values`

165.3.7 GET_NEXT_HIT Function

This function returns the next rule that evaluated to `TRUE` from a true rules iterator, or returns the next rule that evaluated to `MAYBE` from a maybe rules iterator. The function returns `NULL` if there are no more rules that evaluated to `TRUE` or `MAYBE`.

Syntax

```
DBMS_RULE.GET_NEXT_HIT(
   iterator  IN  BINARY_INTEGER)
 RETURN SYS.RE$RULE_HIT;
```

Parameter

Table 165-8 GET_NEXT_HIT Function Parameter

Parameter	Description
`iterator`	The iterator from which the rule that evaluated to `TRUE` or `MAYBE` is retrieved

Usage Notes

This procedure requires an open iterator that was returned by an earlier call to `DBMS_RULE.EVALUATE` in the same session. The user who runs this procedure does not require any privileges on the rule set being evaluated.

When an iterator returns `NULL`, it is closed automatically. If an open iterator is no longer needed, then use the `CLOSE_ITERATOR` procedure in the `DBMS_RULE` package to close it.

> **Note:**
>
> This function raises an error if the rule set being evaluated was modified after the call to the `DBMS_RULE.EVALUATE` procedure that returned the iterator. Modifications to a rule set include added rules to the rule set, changing existing rules in the rule set, dropping rules from the rule set, and dropping the rule set.

> **See Also:**
>
> - Rule TYPEs for more information about the types used with the `DBMS_RULE` package
> - EVALUATE Procedure
> - CLOSE_ITERATOR Procedure

165.3.8 GET_NEXT_RESULT Function

This function iterates over result from the expression given in `RESULT_VAL_ITERATOR`. It returns the expression at iterator evaluated to `TRUE` or `FALSE`.

Syntax

```
DBMS_RULE.GET_NEXT_RESULT (
 result_val_iterator_id    IN   BINARY_INTEGER)
 RETURN  BOOLEAN;
```

Parameter

Table 165-9 GET_NEXT_RESULT Function Parameter

Parameter	Description
`result_val_iterator_id`	Iterator returned from `EVALUATE_EXPRESSION_ITERATOR`

165.3.9 IS_FAST Procedure

Given an expression, of either rule or Independent Expression, this procedure will return `TRUE` if the expression can be evaluated as fast. An expression can be evaluated as fast if the engine

does not need to run any internal SQL and does not need to go to PL/SQL layer in case there are any PL/SQL functions referred.

Syntax

```
DBMS_RULE.IS_FAST(
   expression            IN       VARCHAR2,
   table_aliases         IN       SYS.RE$TABLE_ALIAS_LIST:= NULL,
   variable_types        IN       SYS.RE$VARIABLE_TYPE_LIST:= NULL,
   result_val            OUT      BOOLEAN);
```

Parameter

Table 165-10 IS_FAST Procedure Parameter

Parameter	Description
expression	Expression to check
table_aliases	Alias of tables referred in the above expression string
variable_types	Type definitions of variables used in expression
result_val	If the expression can be evaluated as fast

166
DBMS_RULE_ADM

The `DBMS_RULE_ADM` package provides the subprograms for creating and managing rules, rule sets, and rule evaluation contexts.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_RULE_ADM Subprograms

166.1 DBMS_RULE_ADM Overview

This package provides the subprograms for creating and managing rules, rule sets, and rule evaluation contexts.

> **See Also:**
>
> - Rule TYPEs for more information about the types used with the `DBMS_RULE_ADM` package
> - DBMS_RULE

166.2 DBMS_RULE_ADM Security Model

`PUBLIC` is granted `EXECUTE` privilege on this package.

> **See Also:**
>
> *Oracle Database Security Guide* for more information about user group `PUBLIC`

166.3 Summary of DBMS_RULE_ADM Subprograms

This table lists the `DBMS_RULE_ADM` subprograms and briefly describes them.

Table 166-1 DBMS_RULE_ADM Package Subprograms

Subprogram	Description
ADD_RULE Procedure	Adds the specified rule to the specified rule set
ALTER_EVALUATION_CONTEXT Procedure	Alters a rule evaluation context

Table 166-1 (Cont.) DBMS_RULE_ADM Package Subprograms

Subprogram	Description
ALTER_RULE Procedure	Changes one or more aspects of the specified rule
CREATE_EVALUATION_CONTEXT Procedure	Creates a rule evaluation context
CREATE_RULE Procedure	Creates a rule with the specified name
CREATE_RULE_SET Procedure	Creates a rule set with the specified name
DROP_EVALUATION_CONTEXT Procedure	Drops the rule evaluation context with the specified name
DROP_RULE Procedure	Drops the rule with the specified name
DROP_RULE_SET Procedure	Drops the rule set with the specified name
GRANT_OBJECT_PRIVILEGE Procedure	Grants the specified object privilege on the specified object to the specified user or role
GRANT_SCHEMA_PRIVILEGE Procedure	Grants the specified schema privilege on the specified schema to the specified user or role.
GRANT_SYSTEM_PRIVILEGE Procedure	Grants the specified system privilege to the specified user or role
REMOVE_RULE Procedure	Removes the specified rule from the specified rule set
REVOKE_OBJECT_PRIVILEGE Procedure	Revokes the specified object privilege on the specified object from the specified user or role
REVOKE_SCHEMA_PRIVILEGE Procedure	Revokes the specified schema privilege on the specified schema from the specified user or role.
REVOKE_SYSTEM_PRIVILEGE Procedure	Revokes the specified system privilege from the specified user or role

Note:

All subprograms commit unless specified otherwise.

166.3.1 ADD_RULE Procedure

This procedure adds the specified rule to the specified rule set.

Syntax

```
DBMS_RULE_ADM.ADD_RULE(
   rule_name           IN  VARCHAR2,
   rule_set_name       IN  VARCHAR2,
   evaluation_context  IN  VARCHAR2   DEFAULT NULL,
   rule_comment        IN  VARCHAR2   DEFAULT NULL);
```

Parameters

Table 166-2 ADD_RULE Procedure Parameters

Parameter	Description
rule_name	The name of the rule you are adding to the rule set, specified as [*schema_name.*]*rule_name*. For example, to add a rule named all_a in the hr schema, enter hr.all_a for this parameter. If the schema is not specified, then the current user is the default.
rule_set_name	The name of the rule set to which you are adding the rule, specified as [*schema_name.*]*rule_set_name*. For example, to add the rule to a rule set named apply_rules in the hr schema, enter hr.apply_rules for this parameter. If the schema is not specified, then the current user is the default.
evaluation_context	An evaluation context name in the form [*schema_name.*]*evaluation_context_name*. If the schema is not specified, then the current user is the default. Only specify an evaluation context if the rule itself does not have an evaluation context and you do not want to use the rule set's evaluation context for the rule.
rule_comment	Optional description, which can contain the reason for adding the rule to the rule set

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Have ALTER_ON_RULE_SET privilege on the rule set
- Have ALTER_ANY_RULE_SET system privilege
- Be the owner of the rule set

Also, the rule set owner must meet at least one of the following requirements:

- Have EXECUTE_ON_RULE privilege on the rule
- Have EXECUTE_ANY_RULE system privilege
- Be the rule owner

If the rule has no evaluation context and no evaluation context is specified when you run this procedure, then the rule uses the evaluation context associated with the rule set. In such a case, the rule owner must have the necessary privileges on all the base objects accessed by the rule using the evaluation context.

If an evaluation context is specified, then the rule set owner must meet at least one of the following requirements:

- Have EXECUTE_ON_EVALUATION_CONTEXT privilege on the evaluation context
- Have EXECUTE_ANY_EVALUATION_CONTEXT system privilege, and the owner of the evaluation context must not be SYS
- Be the evaluation context owner

Also, the rule owner must have the necessary privileges on all the base objects accessed by the rule using the evaluation context.

166.3.2 ALTER_EVALUATION_CONTEXT Procedure

This procedure alters a rule evaluation context. A rule evaluation context defines external data that can be referenced in rule conditions. The external data can either exist as variables or as table data.

Syntax

```
DBMS_RULE_ADM.ALTER_EVALUATION_CONTEXT(
  evaluation_context_name      IN  VARCHAR2,
  table_aliases                IN  SYS.RE$TABLE_ALIAS_LIST    DEFAULT NULL,
  remove_table_aliases         IN  BOOLEAN                    DEFAULT FALSE,
  variable_types               IN  SYS.RE$VARIABLE_TYPE_LIST  DEFAULT NULL,
  remove_variable_types        IN  BOOLEAN                    DEFAULT FALSE,
  evaluation_function          IN  VARCHAR2                   DEFAULT NULL,
  remove_evaluation_function   IN  BOOLEAN                    DEFAULT FALSE,
  evaluation_context_comment   IN  VARCHAR2                   DEFAULT NULL,
  remove_eval_context_comment  IN  BOOLEAN                    DEFAULT FALSE);
```

Parameters

Table 166-3 ALTER_EVALUATION_CONTEXT Procedure Parameters

Parameter	Description
evaluation_context_name	The name of the evaluation context you are altering, specified as [*schema_name*.]*evaluation_context_name*. For example, to alter an evaluation context named dept_eval_context in the hr schema, enter hr.dept_eval_context for this parameter. If the schema is not specified, then the current user is the default.
table_aliases	If NULL and remove_table_aliases is FALSE, then the procedure retains the existing table aliases. If NULL and remove_table_aliases is TRUE, then the procedure removes the existing table aliases. If non-NULL, then the procedure replaces the existing table aliases for the evaluation context with the specified table aliases. Table aliases specify the tables in an evaluation context. The table aliases can be used to reference tables in rule conditions.
remove_table_aliases	If TRUE and table_aliases is NULL, then the procedure removes the existing table aliases for the evaluation context. If TRUE and table_aliases is non-NULL, then the procedure raises an error. If FALSE, then the procedure does not remove table aliases.
variable_types	If NULL and remove_variable_types is FALSE, then the procedure retains the variable types. If NULL and remove_variable_types is TRUE, then the procedure removes the existing variable types. If non-NULL, then the procedure replaces the existing variable types for the evaluation context with the specified variable types.

Table 166-3 (Cont.) ALTER_EVALUATION_CONTEXT Procedure Parameters

Parameter	Description
remove_variable_types	If TRUE and variable_types is NULL, then the procedure removes the existing variable types for the evaluation context. If TRUE and variable_types is non-NULL, then the procedure raises an error. If FALSE, then the procedure does not remove the variable types.
evaluation_function	If NULL and remove_evaluation_function is FALSE, then the procedure retains the existing evaluation function. If NULL and remove_evaluation_function is TRUE, then the procedure removes the existing evaluation function. If non-NULL, then the procedure replaces the existing evaluation function for the evaluation context with the specified evaluation function. An evaluation function is an optional function that will be called to evaluate rules that use the evaluation context. It must have the same form as the DBMS_RULE.EVALUATE procedure. If the schema is not specified, then the current user is the default. See CREATE_EVALUATION_CONTEXT Procedurefor more information about evaluation functions.
remove_evaluation_function	If TRUE and evaluation_function is NULL, then the procedure removes the existing evaluation function for the evaluation context. If TRUE and evaluation_function is non-NULL, then the procedure raises an error. If FALSE, then the procedure does not remove the evaluation function.
evaluation_context_comment	If NULL and remove_eval_context_comment is FALSE, then the procedure retains the existing evaluation context comment. If NULL and remove_evaluation_function is TRUE, then the procedure removes the existing evaluation context comment. If non-NULL, then the procedure replaces the existing comment for the evaluation context with the specified comment. An evaluation context comment is an optional description of the rule evaluation context.
remove_eval_context_comment	If TRUE and evaluation_context_comment is NULL, then the procedure removes the existing comment for the evaluation context. If TRUE and evaluation_context_comment is non-NULL, then the procedure raises an error. If FALSE, then the procedure does not remove the evaluation context comment.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the evaluation context being altered

- Have ALL_ON_EVALUATION_CONTEXT or ALTER_ON_EVALUATION_CONTEXT object privilege on an evaluation context owned by another user
- Have ALTER_ANY_EVALUATION_CONTEXT system privilege

> **See Also:**
>
> Rule TYPEs for more information about the types used with the DBMS_RULE_ADM package

166.3.3 ALTER_RULE Procedure

This procedure changes one or more aspects of the specified rule.

Syntax

```
DBMS_RULE_ADM.ALTER_RULE(
   rule_name                  IN  VARCHAR2,
   condition                  IN  VARCHAR2         DEFAULT NULL,
   evaluation_context         IN  VARCHAR2         DEFAULT NULL,
   remove_evaluation_context  IN  BOOLEAN          DEFAULT FALSE,
   action_context             IN  SYS.RE$NV_LIST   DEFAULT NULL,
   remove_action_context      IN  BOOLEAN          DEFAULT FALSE,
   rule_comment               IN  VARCHAR2         DEFAULT NULL,
   remove_rule_comment        IN  BOOLEAN          DEFAULT FALSE);
```

Parameters

Table 166-4 ALTER_RULE Procedure Parameters

Parameter	Description
rule_name	The name of the rule you are altering, specified as [*schema_name.*]*rule_name*. For example, to alter a rule named all_a in the hr schema, enter hr.all_a for this parameter. If the schema is not specified, then the current user is the default.
condition	The condition to be associated with the rule. If non-NULL, then the procedure replaces the existing condition of the rule with the specified condition.
evaluation_context	An evaluation context name in the form [*schema_name.*]*evaluation_context_name*. If the schema is not specified, then the current user is the default. If non-NULL, then the procedure replaces the existing evaluation context of the rule with the specified evaluation context.
remove_evaluation_context	If TRUE, then the procedure sets the evaluation context for the rule to NULL, which effectively removes the evaluation context from the rule. If FALSE, then the procedure retains any evaluation context for the specified rule. If the evaluation_context parameter is non-NULL, then this parameter should be set to FALSE.

Table 166-4 (Cont.) ALTER_RULE Procedure Parameters

Parameter	Description
action_context	If non-NULL, then the procedure changes the action context associated with the rule. A rule action context is information associated with a rule that is interpreted by the client of the rules engine when the rule is evaluated.
remove_action_context	If TRUE, then the procedure sets the action context for the rule to NULL, which effectively removes the action context from the rule. If FALSE, then the procedure retains any action context for the specified rule. If the action_context parameter is non-NULL, then this parameter should be set to FALSE.
rule_comment	If non-NULL, then the existing comment of the rule is replaced by the specified comment.
remove_rule_comment	If TRUE, then the procedure sets the comment for the rule to NULL, which effectively removes the comment from the rule. If FALSE, then the procedure retains any comment for the specified rule. If the rule_comment parameter is non-NULL, then this parameter should be set to FALSE.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Have ALTER_ON_RULE privilege on the rule
- Have ALTER_ANY_RULE system privilege
- Be the owner of the rule being altered

If an evaluation context is specified, then the rule owner must meet at least one of the following requirements:

- Have EXECUTE_ON_EVALUATION_CONTEXT privilege on the evaluation context
- Have EXECUTE_ANY_EVALUATION_CONTEXT system privilege, and the owner of the evaluation context must not be SYS
- Be the evaluation context owner

Also, the rule owner must have the necessary privileges on all the base objects accessed by the rule using the evaluation context.

> **See Also:**
>
> Rule TYPEs for more information about the types used with the DBMS_RULE_ADM package

166.3.4 CREATE_EVALUATION_CONTEXT Procedure

This procedure creates a rule evaluation context. A rule evaluation context defines external data that can be referenced in rule conditions. The external data can either exist as variables or as table data.

Syntax

```
DBMS_RULE_ADM.CREATE_EVALUATION_CONTEXT(
  evaluation_context_name      IN  VARCHAR2,
  table_aliases                IN  SYS.RE$TABLE_ALIAS_LIST    DEFAULT NULL,
  variable_types               IN  SYS.RE$VARIABLE_TYPE_LIST  DEFAULT NULL,
  evaluation_function          IN  VARCHAR2                   DEFAULT NULL,
  evaluation_context_comment   IN  VARCHAR2                   DEFAULT NULL);
```

Parameters

Table 166-5 CREATE_EVALUATION_CONTEXT Procedure Parameters

Parameter	Description
evaluation_context_name	The name of the evaluation context you are creating, specified as [*schema_name.*]*evaluation_context_name*. For example, to create an evaluation context named dept_eval_context in the hr schema, enter hr.dept_eval_context for this parameter. If the schema is not specified, then the current user is the default.
table_aliases	Table aliases that specify the tables in an evaluation context. The table aliases can be used to reference tables in rule conditions.
variable_types	A list of variables for the evaluation context
evaluation_function	An optional function that will be called to evaluate rules using the evaluation context. It must have the same form as the DBMS_RULE.EVALUATE procedure. If the schema is not specified, then the current user is the default. See "Usage Notes" for more information about the evaluation function.
evaluation_context_comment	An optional description of the rule evaluation context.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the evaluation context being created and have CREATE_EVALUATION_CONTEXT_OBJ system privilege
- Have CREATE_ANY_EVALUATION_CONTEXT system privilege

> **See Also:**
>
> Rule TYPEs for more information about the types used with the DBMS_RULE_ADM package

The evaluation function must have the following signature:

```
FUNCTION evaluation_function_name(
  rule_set_name       IN   VARCHAR2,
  evaluation_context  IN   VARCHAR2,
  event_context       IN   SYS.RE$NV_LIST               DEFAULT NULL,
  table_values        IN   SYS.RE$TABLE_VALUE_LIST      DEFAULT NULL,
  column_values       IN   SYS.RE$COLUMN_VALUE_LIST     DEFAULT NULL,
  variable_values     IN   SYS.RE$VARIABLE_VALUE_LIST   DEFAULT NULL,
  attribute_values    IN   SYS.RE$ATTRIBUTE_VALUE_LIST  DEFAULT NULL,
  stop_on_first_hit   IN   BOOLEAN                      DEFAULT FALSE,
  simple_rules_only   IN   BOOLEAN                      DEFAULT FALSE,
  true_rules          OUT  SYS.RE$RULE_HIT_LIST,
  maybe_rules         OUT  SYS.RE$RULE_HIT_LIST);
RETURN BINARY_INTEGER;
```

> **Note:**
>
> Each parameter is required and must have the specified datatype. However, you can change the names of the parameters.

The return value of the function must be one of the following:

- `DBMS_RULE_ADM.EVALUATION_SUCCESS`: The user specified evaluation function completed the rule set evaluation successfully. The rules engine returns the results of the evaluation obtained by the evaluation function to the rules engine client using the `DBMS_RULE.EVALUATE` procedure.
- `DBMS_RULE_ADM.EVALUATION_CONTINUE`: The rules engine evaluates the rule set as if there were no evaluation function. The evaluation function is not used, and any results returned by the evaluation function are ignored.
- `DBMS_RULE_ADM.EVALUATION_FAILURE`: The user specified evaluation function failed. Rule set evaluation stops, and an error is raised.

166.3.5 CREATE_RULE Procedure

This procedure creates a rule.

Syntax

```
DBMS_RULE_ADM.CREATE_RULE(
   rule_name           IN  VARCHAR2,
   condition           IN  VARCHAR2,
   evaluation_context  IN  VARCHAR2        DEFAULT NULL,
   action_context      IN  SYS.RE$NV_LIST  DEFAULT NULL,
   rule_comment        IN  VARCHAR2        DEFAULT NULL);
```

Parameters

Table 166-6 CREATE_RULE Procedure Parameters

Parameter	Description
rule_name	The name of the rule you are creating, specified as [*schema_name.*]*rule_name*. For example, to create a rule named all_a in the hr schema, enter hr.all_a for this parameter. If the schema is not specified, then the current user is the default.
condition	The condition to be associated with the rule. A condition evaluates to TRUE or FALSE and can be any condition allowed in the WHERE clause of a SELECT statement. For example, the following is a valid rule condition: `department_id = 30` Ensure that the proper case is used for text in rule conditions. **Note:** Do not include the word "WHERE" in the condition.
evaluation_context	An optional evaluation context name in the form [*schema_name.*]*evaluation_context_name*, which is associated with the rule. If the schema is not specified, then the current user is the default. If evaluation_context is not specified, then the rule inherits the evaluation context from its rule set.
action_context	The action context associated with the rule. A rule action context is information associated with a rule that is interpreted by the client of the rules engine when the rule is evaluated.
rule_comment	An optional description of the rule

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the rule being created and have the CREATE_RULE_OBJ system privilege
- Have CREATE_ANY_RULE system privilege

If an evaluation context is specified, then the rule owner must meet at least one of the following requirements:

- Have EXECUTE_ON_EVALUATION_CONTEXT privilege on the evaluation context
- Have EXECUTE_ANY_EVALUATION_CONTEXT system privilege, and the owner of the evaluation context must not be SYS.
- Be the evaluation context owner

Also, the rule owner must have the necessary privileges on all the base objects accessed by the rule using the evaluation context.

> **See Also:**
>
> Rule TYPEs for more information about the types used with the DBMS_RULE_ADM package

166.3.6 CREATE_RULE_SET Procedure

This procedure creates a rule set.

Syntax

```
DBMS_RULE_ADM.CREATE_RULE_SET(
   rule_set_name       IN  VARCHAR2,
   evaluation_context  IN  VARCHAR2  DEFAULT NULL,
   rule_set_comment    IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 166-7 CREATE_RULE_SET Procedure Parameters

Parameter	Description
rule_set_name	The name of the rule set you are creating, specified as [*schema_name.*]*rule_set_name*. For example, to create a rule set named apply_rules in the hr schema, enter hr.apply_rules for this parameter. If the schema is not specified, then the current user is the default.
evaluation_context	An optional evaluation context name in the form [*schema_name.*]*evaluation_context_name*, which applies to all rules in the rule set that are not associated with an evaluation context explicitly. If the schema is not specified, then the current user is the default.
rule_set_comment	An optional description of the rule set

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the rule set being created and have CREATE_RULE_SET_OBJ system privilege
- Have CREATE_ANY_RULE_SET system privilege

If an evaluation context is specified, then the rule set owner must meet at least one of the following requirements:

- Have EXECUTE_ON_EVALUATION_CONTEXT privilege on the evaluation context
- Have EXECUTE_ANY_EVALUATION_CONTEXT system privilege, and the owner of the evaluation context must not be SYS
- Be the evaluation context owner

166.3.7 DROP_EVALUATION_CONTEXT Procedure

This procedure drops a rule evaluation context.

Syntax

```
DBMS_RULE_ADM.DROP_EVALUATION_CONTEXT(
   evaluation_context_name  IN  VARCHAR2,
   force                    IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-8 DROP_EVALUATION_CONTEXT Procedure Parameters

Parameter	Description
evaluation_context_name	The name of the evaluation context you are dropping, specified as [*schema_name.*]*evaluation_context_name*. For example, to drop an evaluation context named dept_eval_context in the hr schema, enter hr.dept_eval_context for this parameter. If the schema is not specified, then the current user is the default.
force	If TRUE, then the procedure removes the rule evaluation context from all rules and rule sets that use it. If FALSE and no rules or rule sets use the rule evaluation context, then the procedure drops the rule evaluation context. If FALSE and one or more rules or rule sets use the rule evaluation context, then the procedure raises an exception. **Caution:** Setting force to TRUE can result in rules and rule sets that do not have an evaluation context. If neither a rule nor the rule set it is in has an evaluation context, and no evaluation context was specified for the rule by the ADD_RULE procedure, then the rule cannot be evaluated.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the evaluation context
- Have DROP_ANY_EVALUATION_CONTEXT system privilege

166.3.8 DROP_RULE Procedure

This procedure drops a rule.

Syntax

```
DBMS_RULE_ADM.DROP_RULE(
   rule_name  IN  VARCHAR2,
   force      IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-9 DROP_RULE Procedure Parameters

Parameter	Description
rule_name	The name of the rule you are dropping, specified as [*schema_name.*]*rule_name*. For example, to drop a rule named all_a in the hr schema, enter hr.all_a for this parameter. If the schema is not specified, then the current user is the default.

Table 166-9 (Cont.) DROP_RULE Procedure Parameters

Parameter	Description
force	If TRUE, then the procedure removes the rule from all rule sets that contain it. If FALSE and no rule sets contain the rule, then the procedure drops the rule. If FALSE and one or more rule sets contain the rule, then the procedure raises an exception.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the rule
- Have DROP_ANY_RULE system privilege

> **Note:**
>
> – To remove a rule from a rule set without dropping the rule from the database, use the REMOVE_RULE procedure.
>
> – The rule evaluation context associated with the rule, if any, is not dropped when you run this procedure.

166.3.9 DROP_RULE_SET Procedure

This procedure drops a rule set.

Syntax

```
DBMS_RULE_ADM.DROP_RULE_SET(
   rule_set_name  IN  VARCHAR2,
   delete_rules   IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-10 DROP_RULE_SET Procedure Parameters

Parameter	Description
rule_set_name	The name of the rule set you are dropping, specified as *[schema_name.]rule_set_name*. For example, to drop a rule set named apply_rules in the hr schema, enter hr.apply_rules for this parameter. If the schema is not specified, then the current user is the default.
delete_rules	If TRUE, then the procedure drops any rules that are in the rule set. If any of the rules in the rule set are also in another rule set, then these rules are not dropped. If FALSE, then the procedure does not drop the rules in the rule set.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Have `DROP_ANY_RULE_SET` system privilege
- Be the owner of the rule set

> **Note:**
>
> The rule evaluation context associated with the rule set, if any, is not dropped when you run this procedure.

166.3.10 GRANT_OBJECT_PRIVILEGE Procedure

This procedure grants the specified object privilege on the specified object to the specified user or role. If a user owns the object, then the user automatically is granted all privileges on the object, with grant option.

Syntax

```
DBMS_RULE_ADM.GRANT_OBJECT_PRIVILEGE(
   privilege     IN  BINARY_INTEGER,
   object_name   IN  VARCHAR2,
   grantee       IN  VARCHAR2,
   grant_option  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-11 GRANT_OBJECT_PRIVILEGE Procedure Parameters

Parameter	Description
`privilege`	The name of the object privilege to grant to the grantee on the object. See "Usage Notes" for the available object privileges.
`object_name`	The name of the object for which you are granting the privilege to the grantee, specified as `[schema_name.]object_name`. For example, to grant the privilege on a rule set named `apply_rules` in the `hr` schema, enter `hr.apply_rules` for this parameter. If the schema is not specified, then the current user is the default. The object must be an existing rule, rule set, or evaluation context.
`grantee`	The name of the user or role for which the privilege is granted. The specified user cannot be the owner of the object.
`grant_option`	If `TRUE`, then the specified user or users granted the specified privilege can grant this privilege to others. If `FALSE`, then the specified user or users granted the specified privilege cannot grant this privilege to others.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Be the owner of the object on which the privilege is granted
- Have the same privilege as the privilege being granted with the grant option

In addition, if the object is a rule set, then the user must have `EXECUTE` privilege on all the rules in the rule set with grant option or must own the rules in the rule set.

Table 166-12 lists the object privileges.

Table 166-12 Object Privileges for Evaluation Contexts, Rules, and Rule Sets

Privilege	Description
SYS.DBMS_RULE_ADM.ALL_ON_EVALUATION_CONTEXT	Alter and execute a particular evaluation context in another user's schema
SYS.DBMS_RULE_ADM.ALL_ON_RULE	Alter and execute a particular rule in another user's schema
SYS.DBMS_RULE_ADM.ALL_ON_RULE_SET	Alter and execute a particular rule set in another user's schema
SYS.DBMS_RULE_ADM.ALTER_ON_EVALUATION_CONTEXT	Alter a particular evaluation context in another user's schema
SYS.DBMS_RULE_ADM.ALTER_ON_RULE	Alter a particular rule in another user's schema
SYS.DBMS_RULE_ADM.ALTER_ON_RULE_SET	Alter a particular rule set in another user's schema
SYS.DBMS_RULE_ADM.EXECUTE_ON_EVALUATION_CONTEXT	Execute a particular evaluation context in another user's schema
SYS.DBMS_RULE_ADM.EXECUTE_ON_RULE	Execute a particular rule in another user's schema
SYS.DBMS_RULE_ADM.EXECUTE_ON_RULE_SET	Execute a particular rule set in another user's schema

Examples

For example, to grant the `HR` user the privilege to alter a rule named `hr_dml` in the `strmadmin` schema, enter the following:

```
BEGIN
  DBMS_RULE_ADM.GRANT_OBJECT_PRIVILEGE(
    privilege    => SYS.DBMS_RULE_ADM.ALTER_ON_RULE,
    object_name  => 'strmadmin.hr_dml',
    grantee      => 'hr',
    grant_option => FALSE);
END;
/
```

166.3.11 GRANT_SCHEMA_PRIVILEGE Procedure

This procedure grants the specified schema privilege on the specified schema to the specified user or role.

Syntax

```
DBMS_RULE_ADM.GRANT_SCHEMA_PRIVILEGE(
   privilege     IN  BINARY_INTEGER,
   schema        IN  VARCHAR2,
   grantee       IN  VARCHAR2,
   grant_option  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-13 GRANT_SCHEMA_PRIVILEGE Procedure Parameters

Parameter	Description
privilege	The name of the schema privilege to grant to the grantee.
schema	The name of the schema on which schema privilege is to be granted.
grantee	The name of the user or role to which the schema privilege is to be granted.
grant_option	If TRUE, then the specified user or users granted the specified privilege can grant the system privilege to others. If FALSE, then the specified user or users granted the specified privilege cannot grant the system privilege to others.

Usage Notes

Table 166-16 lists the system privileges.

Table 166-14 System Privileges for Evaluation Contexts, Rules, and Rule Sets

Privilege	Description
SYS.DBMS_RULE_ADM.ALTER_ANY_EVALUATION_CONTEXT	Alter any evaluation context owned by any user
SYS.DBMS_RULE_ADM.ALTER_ANY_RULE	Alter any rule owned by any user
SYS.DBMS_RULE_ADM.ALTER_ANY_RULE_SET	Alter any rule set owned by any user
SYS.DBMS_RULE_ADM.CREATE_ANY_EVALUATION_CONTEXT	Create a new evaluation context in any schema
SYS.DBMS_RULE_ADM.CREATE_EVALUATION_CONTEXT_OBJ	Create a new evaluation context in the grantee's schema
SYS.DBMS_RULE_ADM.CREATE_ANY_RULE	Create a new rule in any schema
SYS.DBMS_RULE_ADM.CREATE_RULE_OBJ	Create a new rule in the grantee's schema
SYS.DBMS_RULE_ADM.CREATE_ANY_RULE_SET	Create a new rule set in any schema
SYS.DBMS_RULE_ADM.CREATE_RULE_SET_OBJ	Create a new rule set in the grantee's schema
SYS.DBMS_RULE_ADM.DROP_ANY_EVALUATION_CONTEXT	Drop any evaluation context in any schema
SYS.DBMS_RULE_ADM.DROP_ANY_RULE	Drop any rule in any schema
SYS.DBMS_RULE_ADM.DROP_ANY_RULE_SET	Drop any rule set in any schema
SYS.DBMS_RULE_ADM.EXECUTE_ANY_EVALUATION_CONTEXT	Execute any evaluation context owned by any user
SYS.DBMS_RULE_ADM.EXECUTE_ANY_RULE	Execute any rule owned by any user
SYS.DBMS_RULE_ADM.EXECUTE_ANY_RULE_SET	Execute any rule set owned by any user

166.3.12 GRANT_SYSTEM_PRIVILEGE Procedure

This procedure grant the specified system privilege to the specified user or role.

Syntax

```
DBMS_RULE_ADM.GRANT_SYSTEM_PRIVILEGE(
   privilege     IN  BINARY_INTEGER,
```

```
  grantee       IN  VARCHAR2,
  grant_option  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-15 GRANT_SYSTEM_PRIVILEGE Procedure Parameters

Parameter	Description
privilege	The name of the system privilege to grant to the grantee.
grantee	The name of the user or role for which the privilege is granted
grant_option	If TRUE, then the specified user or users granted the specified privilege can grant the system privilege to others. If FALSE, then the specified user or users granted the specified privilege cannot grant the system privilege to others.

Usage Notes

Table 166-16 lists the system privileges.

Table 166-16 System Privileges for Evaluation Contexts, Rules, and Rule Sets

Privilege	Description
SYS.DBMS_RULE_ADM.ALTER_ANY_EVALUATION_CONTEXT	Alter any evaluation context owned by any user
SYS.DBMS_RULE_ADM.ALTER_ANY_RULE	Alter any rule owned by any user
SYS.DBMS_RULE_ADM.ALTER_ANY_RULE_SET	Alter any rule set owned by any user
SYS.DBMS_RULE_ADM.CREATE_ANY_EVALUATION_CONTEXT	Create a new evaluation context in any schema
SYS.DBMS_RULE_ADM.CREATE_EVALUATION_CONTEXT_OBJ	Create a new evaluation context in the grantee's schema
SYS.DBMS_RULE_ADM.CREATE_ANY_RULE	Create a new rule in any schema
SYS.DBMS_RULE_ADM.CREATE_RULE_OBJ	Create a new rule in the grantee's schema
SYS.DBMS_RULE_ADM.CREATE_ANY_RULE_SET	Create a new rule set in any schema
SYS.DBMS_RULE_ADM.CREATE_RULE_SET_OBJ	Create a new rule set in the grantee's schema
SYS.DBMS_RULE_ADM.DROP_ANY_EVALUATION_CONTEXT	Drop any evaluation context in any schema
SYS.DBMS_RULE_ADM.DROP_ANY_RULE	Drop any rule in any schema
SYS.DBMS_RULE_ADM.DROP_ANY_RULE_SET	Drop any rule set in any schema
SYS.DBMS_RULE_ADM.EXECUTE_ANY_EVALUATION_CONTEXT	Execute any evaluation context owned by any user
SYS.DBMS_RULE_ADM.EXECUTE_ANY_RULE	Execute any rule owned by any user
SYS.DBMS_RULE_ADM.EXECUTE_ANY_RULE_SET	Execute any rule set owned by any user

For example, to grant the strmadmin user the privilege to create a rule set in any schema, enter the following:

```
BEGIN
  DBMS_RULE_ADM.GRANT_SYSTEM_PRIVILEGE(
    privilege    => SYS.DBMS_RULE_ADM.CREATE_ANY_RULE_SET,
    grantee      => 'strmadmin',
    grant_option => FALSE);
```

```
END;
/
```

166.3.13 REMOVE_RULE Procedure

This procedure removes the specified rule from the specified rule set.

Syntax

```
DBMS_RULE_ADM.REMOVE_RULE(
   rule_name                IN  VARCHAR2,
   rule_set_name            IN  VARCHAR2,
   evaluation_context       IN  VARCHAR2  DEFAULT NULL,
   all_evaluation_contexts  IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 166-17 REMOVE_RULE Procedure Parameters

Parameter	Description
rule_name	The name of the rule you are removing from the rule set, specified as [*schema_name.*]*rule_name*. For example, to remove a rule named all_a in the hr schema, enter hr.all_a for this parameter. If the schema is not specified, then the current user is the default.
rule_set_name	The name of the rule set from which you are removing the rule, specified as [*schema_name.*]*rule_set_name*. For example, to remove the rule from a rule set named apply_rules in the hr schema, enter hr.apply_rules for this parameter. If the schema is not specified, then the current user is the default.
evaluation_context_name	The name of the evaluation context associated with the rule you are removing, specified as [*schema_name.*]*evaluation_context_name*. For example, to specify an evaluation context named dept_eval_context in the hr schema, enter hr.dept_eval_context for this parameter. If the schema is not specified, then the current user is the default. If an evaluation context was specified for the rule you are removing when you added the rule to the rule set using the ADD_RULE procedure, then specify the same evaluation context. If you added the same rule more than once with different evaluation contexts, then specify the rule with the evaluation context you want to remove. If you specify an evaluation context that is not associated with the rule, then the procedure raises an error. Specify NULL if you did not specify an evaluation context when you added the rule to the rule set. If you specify NULL and there are one or more evaluation contexts associated with the rule, then the procedure raises an error.
all_evaluation_contexts	If TRUE, then the procedure removes the rule from the rule set with all of its associated evaluation contexts. If FALSE, then the procedure only removes the rule with the specified evaluation context. This parameter is relevant only if the same rule is added more than once to the rule set with different evaluation contexts.

Usage Notes

To run this procedure, a user must meet at least one of the following requirements:

- Have `ALTER_ON_RULE_SET` privilege on the rule set
- Have `ALTER_ANY_RULE_SET` system privilege
- Be the owner of the rule set

> **Note:**
>
> This procedure does not drop a rule from the database. To drop a rule from the database, use the `DROP_RULE` procedure.

166.3.14 REVOKE_OBJECT_PRIVILEGE Procedure

This procedure revokes the specified object privilege on the specified object from the specified user or role.

Syntax

```
DBMS_RULE_ADM.REVOKE_OBJECT_PRIVILEGE(
   privilege    IN  BINARY_INTEGER,
   object_name  IN  VARCHAR2,
   revokee      IN  VARCHAR2);
```

Parameters

Table 166-18 REVOKE_OBJECT_PRIVILEGE Procedure Parameters

Parameter	Description
`privilege`	The name of the object privilege on the object to revoke from the revokee. See GRANT_OBJECT_PRIVILEGE Procedure for a list of the object privileges.
`object_name`	The name of the object for which you are revoking the privilege from the revokee, specified as `[`*`schema_name`*`.]`*`object_name`*. For example, to revoke an object privilege on a rule set named `apply_rules` in the `hr` schema, enter `hr.apply_rules` for this parameter. If the schema is not specified, then the current user is the default. The object must be an existing rule, rule set, or evaluation context.
`revokee`	The name of the user or role from which the privilege is revoked. The user who owns the object cannot be specified.

166.3.15 REVOKE_SCHEMA_PRIVILEGE Procedure

This procedure revokes the specified schema privilege on the specified schema from the specified user or role.

Syntax

```
DBMS_RULE_ADM.REVOKE_SCHEMA_PRIVILEGE(
   privilege  IN  BINARY_INTEGER,
```

```
   schema      IN  VARCHAR2,
   revokee     IN  VARCHAR2);
```

Parameters

Table 166-19 REVOKE_SCHEMA_PRIVILEGE Procedure Parameters

Parameter	Description
privilege	The name of the schema privilege to revoke from the revokee. See GRANT_SCHEMA_PRIVILEGE Procedure for a list of the schema privileges.
schema	The name of the schema from which the schema privilege is to be revoked.
revokee	The name of the user or role from which the schema privilege is to be revoked.

166.3.16 REVOKE_SYSTEM_PRIVILEGE Procedure

This procedure revokes the specified system privilege from the specified user or role.

Syntax

```
DBMS_RULE_ADM.REVOKE_SYSTEM_PRIVILEGE(
   privilege   IN  BINARY_INTEGER,
   revokee     IN  VARCHAR2);
```

Parameters

Table 166-20 REVOKE_SYSTEM_PRIVILEGE Procedure Parameters

Parameter	Description
privilege	The name of the system privilege to revoke from the revokee. See GRANT_SYSTEM_PRIVILEGE Procedure for a list of the system privileges.
revokee	The name of the user or role from which the privilege is revoked

167
DBMS_SAGA_ADM

The DBMS_SAGA_ADM package provides a collection of saga administration functions and procedures to define and manage saga participants, coordinators and brokers.

This chapter contains the following topics:

- DBMS_SAGA_ADM Overview
- DBMS_SAGA_ADM Security Model
- Summary of DBMS_SAGA_ADM Subprograms

167.1 DBMS_SAGA_ADM Overview

The SYS.DBMS_SAGA_ADM package provides the saga administration APIs to manage sagas.

The saga administration APIs implemented using the DBMS_SAGA_ADM package enables the administrators to define and manage saga participants, coordinators and brokers. The package APIs are described here for your reference.

167.2 DBMS_SAGA_ADM Security Model

A participant must invoke the procedures in DBMS_SAGA_ADM package from inside its schema and PDB. Invocation of APIs in the DBMS_SAGA_ADM package requires the SAGA_ADM_ROLE role.

167.3 Summary of DBMS_SAGA_ADM Subprograms

This table lists and briefly describes the DBMS_SAGA_ADM package subprograms.

Table 167-1 DBMS_SAGA_ADM Package Subprograms

Subprogram	Description
ADD_COORDINATOR Procedure	Adds a saga coordinator to the saga framework
DROP_COORDINATOR Procedure	Drops the given coordinator from the saga framework.
ADD_PARTICIPANT Procedure	Adds a participant to the local PDB and the broker
DROP_PARTICIPANT Procedure	Drops the given participant on the local PDB and the broker if there are no pending sagas and pending saga messages for the participant.
ADD_BROKER Procedure	Adds a broker to the saga framework and creates a JMS topic for the broker.
DROP_BROKER Procedure	Drops a broker and the associated JMS topic.
REGISTER_SAGA_CALLBACK Procedure	Allows users to create or modify a callback package if created using add_participant procedure.

167.3.1 ADD_BROKER Procedure

This procedure adds a broker to the saga framework.

As a part of adding a broker, a JMS topic is created for the broker. Both incoming and outgoing message propagation channels from participants to the broker use the broker's JMS topic.

Syntax

```
add_broker(
 broker_name                IN varchar2,
 broker_schema              IN varchar2 default sys_context('userenv', 'CURRENT_USER'),
 storage_clause             IN varchar2 default NULL,
 queue_partitions           IN number DEFAULT 1,
 version                    IN number default 1
 );
```

Parameters

Table 167-2 ADD_BROKER Procedure Parameters

Parameter	Description
broker_name	A unique broker name.
broker_schema	The schema for the microservice participant being added. If no value is provided, CURRENT_USER is used.
storage_clause	The storage parameter is included in the CREATE_TRANSACTIONAL_EVENT_QUEUE statement when the inbound and outbound JMS topics are created. The storage_clause argument can take any text that you can use in a standard DBMS_AQADM.CREATE_TRANSACTIONAL_EVENT_QUEUE storage_clause argument.
queue_partitions	The number of partitions for the Saga entity IN/OUT queues. This allows for higher parallelism and throughput for the Saga entity. DEFAULT is 1.
version	Saga version

167.3.2 DROP_BROKER Procedure

This procedure drops a broker and the associated JMS topic.

A broker can only be dropped if there are no registered participants and pending messages for the broker.

Syntax

```
drop_broker(broker_name      IN varchar2)
```

Parameters

Table 167-3 DROP_BROKER Procedure Parameters

Parameter	Description
broker_name	The broker name to be dropped.

167.3.3 ADD_COORDINATOR Procedure

This procedure adds a Saga coordinator to the Saga framework. A Saga coordinator acts as a transaction manager for Sagas. The Saga coordinator maintains the state of a Saga across various participants.

Syntax

```
add_coordinator(
 coordinator_name       IN varchar2,
 coordinator_schema     IN varchar2 DEFAULT sys_context('USERENV', 'CURRENT_USER'),
 storage_clause         IN varchar2 DEFAULT NULL,
 dblink_to_broker       IN varchar2 DEFAULT NULL,
 mailbox_schema         IN varchar2,
 broker_name            IN varchar2,
 dblink_to_coordinator  IN varchar2 DEFAULT NULL,
 queue_partitions       IN number DEFAULT 1,
 listener_count         IN number DEFAULT -1,
 version                IN number default 1);
```

Parameters

Table 167-4 ADD_COORDINATOR Procedure Parameters

Parameter	Description
coordinator_name	A unique name for the Saga coordinator.
coordinator_schema	The schema for the Saga coordinator being added. The parameter uses CURRENT_USER if no value is specified.
storage_clause	The storage parameter is included in the CREATE_TRANSACTIONAL_EVENT_QUEUE statement when the inbound and outbound JMS topics are created. The storage_clause argument can take any text that you can use in a standard DBMS_AQADM.CREATE_TRANSACTIONAL_EVENT_QUEUE storage_clause argument.
dblink_to_broker	Pre-created database link to connect from the coordinator PDB to the broker PDB.
mailbox_schema	The schema for the mailbox owner at the broker.
broker_name	The name of the broker.
dblink_to_coordinator	Pre-created database link to connect from the broker PDB to the coordinator PDB.
queue_partitions	The number of partitions for the Saga entity IN/OUT queues. This allows for higher parallelism and throughput for the Saga entity. DEFAULT is 1.

> **Note:**
>
> The number of queue partitions should be equal to the partition count specified while performing the add_broker() operation.

Table 167-4 (Cont.) ADD_COORDINATOR Procedure Parameters

Parameter	Description
listener_count	The number of concurrent Saga listener processes consuming inbound messages for the Saga coordinator. listener_count can have one of the following values: • DBMS_SAGA_ADM.AQ_NTFN, which is also the DEFAULT. In this case, the Saga coordinator processes messages using the AQ notification mechanism. • DBMS_SAGA_ADM.AUTO_NTFN- In this case, the Saga coordinator processes messages using job queue processes that dequeue from the coordinator queue. The job automatically adds and removes jobs depending on the queue depth. • If any positive number (N) is specified, then N number of fixed jobs are created to process the messages for the Saga coordinator.
version	Saga version

167.3.4 DROP_COORDINATOR Procedure

This procedure drops the given coordinator and disables message propagation between the coordinator and the broker. A coordinator can only be dropped if there are no participants associated with the coordinator.

Syntax

```
drop_coordinator(coordinator_name  IN varchar2)
```

Parameters

Table 167-5 DROP_COORDINATOR Procedure Parameters

Parameter	Description
coordinator_name	The name of the saga coordinator that is being dropped.

167.3.5 ADD_PARTICIPANT Procedure

This procedure adds a Saga participant to the database and the broker.

A participant name is unique to the broker where the participant is added. The ADD_PARTICIPANT procedure creates system-defined inbound and outbound JMS topics for the participant. A JMS topic is also created at the broker's side if the topic does not exist. Message propagation relationships are set up from:

- The outbound JMS topic of the participant to the JMS topic of the broker.
- The JMS topic of the broker to the inbound JMS topic of the participant.

Syntax

```
add_participant(
 participant_name           IN varchar2,
 participant_schema         IN varchar2 DEFAULT sys_context('USERENV', 'CURRENT_USER'),
 storage_clause             IN varchar2 DEFAULT NULL,
 coordinator_name           IN varchar2 DEFAULT NULL,
```

```
dblink_to_broker           IN varchar2,
mailbox_schema             IN varchar2,
broker_name                IN varchar2,
callback_schema            IN varchar2 DEFAULT sys_context('USERENV','CURRENT_USER'),
callback_package           IN varchar2,
dblink_to_participant      IN varchar2,
queue_partitions           IN number DEFAULT 1,
version                    IN number DEFAULT 1
);
```

Parameters

Table 167-6 ADD_PARTICIPANT Procedure Parameters

Parameter	Description
participant_name	A unique name for the saga participant.
participant_schema	The schema for the saga participant being added. The parameter uses CURRENT_USER if no value is specified.
storage_clause	The storage parameter is included in the CREATE_TRANSACTIONAL_EVENT_QUEUE statement when the inbound and outbound JMS topics are created. The storage_clause argument can take any text that you can use in a standard DBMS_AQADM.CREATE_TRANSACTIONAL_EVENT_QUEUE storage_clause argument.
coordinator_name	The name of the saga coordinator. If no value is provided and create_coordinator is TRUE, a system generated name is assumed.
dblink_to_broker	The pre-created database link to connect from the participant PDB to broker PDB.
mailbox_schema	The schema for the mailbox owner at the broker's side.
broker_name	The name of the broker. If no value is provided, the mailbox_schema name is assumed.
callback_schema	The schema of the user-defined callback package. If no value is provided, CURRENT_USER is used.
callback_package	The name of the user-defined callback package. The callback package isolates application developers from the internal details of the saga infrastructure, and enables the developers to focus on business logic.
dblink_to_participant	The pre-created database link to connect from broker PDB to the participant PDB.
queue_partitions	The number of partitions for the Saga entity IN/OUT queues. This allows for higher parallelism and throughput for the Saga entity. DEFAULT is 1. **Note:** The number of queue partitions should be equal to the partition count specified while performing the add_broker() operation.
version	Saga version

Usage Notes

The `ADD_PARTICIPANT` procedure needs coordinator name as an argument, hence the `ADD_COORDINATOR` procedure must be used before the `ADD_PARTICIPANT` procedure.

167.3.6 DROP_PARTICIPANT Procedure

This procedure drops the given participant on the local PDB and the broker.

The `DROP_PARTICIPANT` procedure drops the given participant on the local PDB and the broker if there are no pending sagas and pending saga messages for the participant. The `DROP_PARTICIPANT` procedure also disables message propagation between the participant and the broker.

Syntax

```
drop_participant(participant_name  IN varchar2)
```

Parameters

Table 167-7 DROP_PARTICIPANT Procedure Parameters

Parameter	Description
`participant_name`	The name of the participant microservice that is being dropped.

167.3.7 REGISTER_SAGA_CALLBACK Procedure

This procedure enables users to create or modify a callback package if the package is created using ADD_PARTICIPANT procedure.

If the callback package was not added before, the REGISTER_SAGA_CALLBACK procedure sets up the required subscribers and the notification callback. While doing an add participant the user has the option of specifying the callback package. If they want to defer specifying the callback package or if they want to change the callback package they can use the REGISTER_SAGA_CALLBACK procedure. If the callback already exists for a participant, then it is simply replaced. The REGISTER_SAGA_CALLBACK procedure is useful in cases where the user is not sure about whether to implement PL/SQL or Java, and wants to decide about it later, or the user wants to switch implementations (for example, from a Java client to a PL/SQL based client).

Syntax

```
register_saga_callback(participant_name IN varchar2,
 callback_schema        IN varchar2 default sys_context('USERENV','CURRENT_USER'),
 callback_package       IN varchar2);
```

Parameters

Table 167-8 REGISTER_SAGA_CALLBACK Procedure Parameters

Parameter	Description
`participant_name`	The name of the participant already created.

Table 167-8 (Cont.) REGISTER_SAGA_CALLBACK Procedure Parameters

Parameter	Description
`callback_schema`	The schema for the user-defined callback package. If no value is provided, `CURRENT_USER` is used.
`callback_package`	The name of the user-defined callback package. The callback package isolates application developers from the internal details of the saga infrastructure, and enables them to focus on business logic.

168
DBMS_SAGA

The DBMS_SAGA package provides a collection of saga functions and procedures to initiate and finalize sagas.

This chapter contains the following topics:

- DBMS_SAGA Overview
- DBMS_SAGA Security Model
- Summary of DBMS_SAGA Subprograms

168.1 DBMS_SAGA Overview

The DBMS_SAGA package enables you (developers) to use PL/SQL and develop packaged microservices applications in the database without requiring a mid-tier.

The DBMS_SAGA package provides the PL/SQL interfaces to allow client programs to interact with database sagas.

If you want to implement microservices using a mid-tier to communicate with a database, it is recommended that you use the AQJMS extensions to implement the saga functionality.

> **See Also:**
>
> Managing a Saga Using JMS Interface for more information about the AQJMS extensions.

168.2 DBMS_SAGA Security Model

The DBMS_SAGA package requires the SAGA_PARTICPANT role to participate in sagas.

168.3 Summary of DBMS_SAGA Subprograms

This table lists and briefly describes the DBMS_SAGA package subprograms.

Table 168-1 DBMS_SAGA Package Subprograms

Subprogram	Description
BEGIN_SAGA Function	Creates and returns a new saga_id_t (GUID). The new Saga ID is inserted into the Saga$ dictionary table.
GET_SAGA_ID Function	Gets the Saga ID for a Saga associated with a database session.
SET_SAGA_ID Procedure	Sets the Saga ID for a database session to the saga_id provided.
COMMIT_SAGA Procedure	Commits the Saga identified by the saga_id parameter.

Table 168-1 (Cont.) DBMS_SAGA Package Subprograms

Subprogram	Description
ROLLBACK_SAGA Procedure	Rolls back the Saga identified by the `saga_id` parameter.
IS_INCOMPLETE Function	Returns `TRUE` if the Saga corresponding to the provided `saga_id` is considered incomplete.
SET_INCOMPLETE Procedure	Marks the Saga identified by the provided `saga_id` as incomplete.
SEND_REQUEST Procedure	Enrolls a participant to an initiated Saga.

168.3.1 BEGIN_SAGA Function

This function creates and returns a new saga GUID: `saga_id_t`. The new saga ID is inserted into the `saga$` dictionary table.

Syntax

```
function begin_saga(initiator_name IN varchar2,
          timeout   IN number default NULL,
          version   IN number default 1)
 return saga_id_t;
```

Parameters

Table 168-2 BEGIN_SAGA Procedure Parameters

Parameter	Description
`initiator_name`	The name of the saga initiator.
`timeout`	The timeout value specified as number of seconds. This value overrides the value of the `max_saga_duration` parameter.

168.3.2 GET_SAGA_ID Function

This function returns the saga, if any, associated with the current database session.

Syntax

```
function get_saga_id() return saga_id_t;
```

168.3.3 SET_SAGA_ID Procedure

This procedure sets the saga ID for a database session to the `saga_id` provided.

Syntax

```
procedure set_saga_id(saga_id  IN saga_id_t);
```

Parameters

Table 168-3 SET_SAGA Procedure Parameters

Parameter	Description
saga_id	The saga identifier of the type saga_id_t.

168.3.4 COMMIT_SAGA Procedure

This procedure commits the saga that is identified with the saga_id parameter.

The force option is used to forcefully commit a saga participant transaction for pending or incomplete sagas, For the COMMIT_SAGA() call on the participant database, force=TRUE only commits the saga branch operation on the participant database. The default value of force is TRUE.

Syntax

```
procedure commit_saga(saga_participant IN VARCHAR2,
 saga_id IN saga_id_t,
 force   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 168-4 COMMIT_SAGA Procedure Parameters

Parameter	Description
saga_participant	The saga participant (participant_name).
saga_id	The saga identifier of the type saga_id_t.
force	The flag to indicate whether the commit saga operation is initiated for a saga branch.

168.3.5 ROLLBACK_SAGA Procedure

This procedure aborts the saga corresponding to the saga_id parameter.

The force option can be used to forcefully roll back a saga participant transaction for pending or incomplete sagas, For the ROLLBACK_SAGA() call on the participant database, force=TRUE only rolls back the saga branch operation on the participant database. The default value of force is TRUE.

Syntax

```
procedure rollback_saga (saga_participant IN VARCHAR2,
 saga_id IN saga_id_t,
 force   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 168-5 ROLLBACK_SAGA Procedure Parameters

Parameter	Description
saga_participant	The saga participant (participant_name).
saga_id	The saga identifier of the type saga_id_t.
force	The flag to indicate whether the rollback saga operation is initiated for a participant transaction. The default value is TRUE.

168.3.6 IS_INCOMPLETE Function

This function returns TRUE if the saga corresponding to the provided saga_id is considered incomplete.

A saga is considered incomplete for the following reasons:

- If the saga exceeds the duration identified by the max_saga_duration parameter or the duration established by the begin_saga() call.
- If the saga participant fails to finalize the saga (commit or rollback).

Syntax

```
function is_incomplete (saga_id  IN saga_id_t)    returns boolean;
```

Parameters

Table 168-6 IS_COMPLETE Function Parameters

Parameter	Description
saga_id	The saga identifier of the type saga_id_t.

168.3.7 SET_INCOMPLETE Procedure

This procedure marks the saga corresponding to the provided saga_id as incomplete.

This is an administrative interface to mark certain sagas as incomplete such that these sagas can be flagged for manual intervention.

Syntax

```
procedure set_incomplete (saga_id  IN saga_id_t);
```

Parameters

Table 168-7 SET_INCOMPLETE Procedure Parameters

Parameter	Description
saga_id	The saga identifier of the type saga_id_t.

168.3.8 SEND_REQUEST Procedure

This procedure can be used by the client to enroll a participant in the initiated saga.

This procedure abstracts the creation of the saga message, conversion of JSON payloads, and enqueuing the message to the participant. The application developer can manually perform this step as well.

Syntax

```
procedure send_request(saga_id IN saga_id_t,
 recipient    IN VARCHAR2,
 payload      IN JSON DEFAULT NULL,
 saga_version IN NUMBER DEFAULT 1,
 saga_spare   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 168-8 SEND_REQUEST Procedure Parameters

Parameter	Description
`saga_id`	The saga identifier that is derived using the `BEGIN_SAGA()` function.
`recipient`	The participant to enroll in this saga.
`payload`	The JSON payload to call request() of the participant being enrolled.
`saga_version`	The version of the saga framework (default is 1).
`saga_spare`	The currently unused field.

169
DBMS_SCHEDULER

The `DBMS_SCHEDULER` package provides a collection of scheduling functions and procedures that can be called from any PL/SQL program.

This chapter contains the following topics:

- Deprecated Subprograms
- Security Model
- Rules and Limits
- Operational Notes
- Data Structures
- Summary of DBMS_SCHEDULER Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more information regarding how to use `DBMS_SCHEDULER`

169.1 DBMS_SCHEDULER Deprecated Subprograms

Oracle recommends that you do not use deprecated subprograms in new applications. Support for deprecated features is for backward compatibility only

The following subprograms are deprecated with Oracle Database 12*c* Release 1 (12.1):

- CREATE_CREDENTIAL Procedure
- DROP_CREDENTIAL Procedure

169.2 DBMS_SCHEDULER Security Model

The `DBMS_SCHEDULER` package ignores privileges granted on scheduler objects, such as jobs or chains, through roles. Object privileges must be granted directly to the user.

169.3 DBMS_SCHEDULER Rules and Limits

These rules apply when using the `DBMS_SCHEDULER` package.

- Only `SYS` can perform actions on objects in the `SYS` schema.
- Several of the procedures accept comma-delimited lists of object names. If you provide a list of names, then the Scheduler stops executing the list at the first object that returns an error. Therefore, the Scheduler does not perform the tasks needed for the remaining objects on the list.

For example, consider the statement `DBMS_SCHEDULER.STOP_JOB ('job1, job2, job3, sys.jobclass1, sys.jobclass2, sys.jobclass3');`

If `job3` cannot be stopped, then the jobs that follow it, `jobclass1`, `jobclass2`, and `jobclass3` cannot be stopped. The jobs that preceded `job3`, `job1` and `job2`, are stopped.

- Performing an action on an object that does not exist returns a PL/SQL exception stating that the object does not exist.

169.4 DBMS_SCHEDULER Operational Notes

The Scheduler uses a rich **calendaring syntax** to enable you to define repeating schedules, such as "every Tuesday and Friday at 4:00 p.m." or "the second Wednesday of every month." This calendaring syntax is used in calendaring expressions in the `repeat_interval` argument of a number of package subprograms. Evaluating a calendaring expression results in a set of discrete timestamps.

See *Oracle Database Administrator's Guide* for examples of the calendaring syntax.

Calendaring Syntax

This section starts with the calendaring syntax. It is followed by descriptions of various parts of the syntax.

In the calendaring syntax, * means 0 or more.

```
repeat_interval = regular_schedule | combined_schedule

regular_schedule = frequency_clause
[";" interval_clause] [";" bymonth_clause] [";" byweekno_clause]
[";" byyearday_clause] [";" bydate_clause] [";" bymonthday_clause]
[";" byday_clause] [";" byhour_clause] [";" byminute_clause]
[";" bysecond_clause] [";" bysetpos_clause] [";" include_clause]
[";" exclude_clause] [";" intersect_clause][";" periods_clause]
[";" byperiod_clause]


frequency_clause = "FREQ" "=" ( predefined_frequency | user_defined_frequency )
predefined_frequency = "YEARLY" | "MONTHLY" | "WEEKLY" | "DAILY" |
   "HOURLY" | "MINUTELY" | "SECONDLY"
user_defined_frequency = named_schedule

interval_clause = "INTERVAL" "=" intervalnum
   intervalnum = 1 through 99
bymonth_clause = "BYMONTH" "=" monthlist
   monthlist = month ( "," month)*
   month = numeric_month | char_month
   numeric_month = 1 | 2 | 3 ...  12
   char_month = "JAN" | "FEB" | "MAR" | "APR" | "MAY" | "JUN" |
   "JUL" | "AUG" | "SEP" | "OCT" | "NOV" | "DEC"
byweekno_clause = "BYWEEKNO" "=" weeknumber_list
   weeknumber_list = weeknumber ( "," weeknumber)*
   weeknumber = [minus] weekno
   weekno = 1 through 53
byyearday_clause = "BYYEARDAY" "=" yearday_list
   yearday_list = yearday ( "," yearday)*
   yearday = [minus] yeardaynum
   yeardaynum = 1 through 366
bydate_clause = "BYDATE" "=" date_list
   date_list = date ( "," date)*
   date = [YYYY]MMDD [ offset | span ]
```

```
bymonthday_clause = "BYMONTHDAY" "=" monthday_list
   monthday_list = monthday ( "," monthday)*
   monthday = [minus] monthdaynum
   monthdaynum = 1 through 31
byday_clause = "BYDAY" "=" byday_list
   byday_list = byday ( "," byday)*
   byday = [weekdaynum] day
   weekdaynum = [minus] daynum
   daynum = 1 through 53 /* if frequency is yearly */
   daynum = 1 through 5  /* if frequency is monthly */
   day = "MON" | "TUE" | "WED" | "THU" | "FRI" | "SAT" | "SUN"
BYTIME clause:  BYTIME=[hour_minute_second_list|minute_second_list]
   hour_minute_second_list: hh24mmss, .., hh24mmss
   minute_second_list: mmss, .. mmss
byhour_clause = "BYHOUR" "=" hour_list
   hour_list = hour ( "," hour)*
   hour = 0 through 23
byminute_clause = "BYMINUTE" "=" minute_list
   minute_list = minute ( "," minute)*
   minute = 0 through 59
bysecond_clause = "BYSECOND" "=" second_list
   second_list = second ( "," second)*
   second = 0 through 59
bysetpos_clause = "BYSETPOS" "=" setpos_list
   setpos_list = setpos ("," setpos)*
   setpos = [minus] setpos_num
   setpos_num = 1 through 9999

include_clause = "INCLUDE" "=" schedule_list
exclude_clause = "EXCLUDE" "=" schedule_list
intersect_clause = "INTERSECT" "=" schedule_list
schedule_list = schedule_clause ("," schedule_clause)*
schedule_clause = named_schedule [ offset ]
named_schedule = [schema "."] schedule
periods_clause = "PERIODS" "=" periodnum
byperiod_clause = "BYPERIOD" "=" period_list
period_list = periodnum ("," periodnum)*
periodnum = 1 through 100

offset = ("+" | "-") ["OFFSET:"] duration_val
span = ("+" | "-" | "^") "SPAN:" duration_val
duration_val = dur-weeks | dur_days
dur_weeks = numofweeks "W"
dur_days = numofdays "D"
numofweeks = 1 through 53
numofdays = 1 through 376
minus = "-"

combined_schedule = schedule_list
```

Table 169-1 Values for repeat_interval

Name	Description
FREQ	This specifies the type of recurrence. It must be specified. The possible predefined frequency values are YEARLY, MONTHLY, WEEKLY, DAILY, HOURLY, MINUTELY, and SECONDLY. Alternatively, specifies an existing schedule to use as a user-defined frequency.
INTERVAL	This specifies a positive integer representing how often the recurrence repeats. The default is 1, which means every second for secondly, every day for daily, and so on. The maximum value is 99.

Table 169-1 (Cont.) Values for repeat_interval

Name	Description
BYMONTH	This specifies which month or months you want the job to execute in. You can use numbers such as 1 for January and 3 for March, as well as three-letter abbreviations such as FEB for February and JUL for July.
BYWEEKNO	This specifies the week of the year as a number. It follows ISO-8601, which defines the week as starting with Monday and ending with Sunday; and the first week of a year as the first week, which is mostly within the Gregorian year. The first week is equivalent to the following two variants: the week that contains the first Thursday of the Gregorian year; and the week containing January 4th. The ISO-8601 week numbers are integers from 1 to 52 or 53; parts of week 1 may be in the previous calendar year; parts of week 52 may be in the following calendar year; and if a year has a week 53, parts of it must be in the following calendar year. As an example, in the year 1998, the ISO week 1 began on Monday December 29th, 1997; and the last ISO week (week 53) ended on Sunday January 3rd, 1999. So December 29th, 1997, is in the ISO week 1998-01, and January 1st, 1999, is in the ISO week 1998-53. `byweekno` is only valid for YEARLY. Examples of invalid specifications are "`FREQ=YEARLY; BYWEEKNO=1; BYMONTH=12`" and "`FREQ=YEARLY;BYWEEKNO=53;BYMONTH=1`".
BYYEARDAY	This specifies the day of the year as a number. Valid values are 1 to 366. An example is 69, which is March 10 (31 for January, 28 for February, and 10 for March). 69 evaluates to March 10 for non-leap years and March 9 in leap years. -2 will always evaluate to December 30th independent of whether it is a leap year.

BYDATE	This specifies a list of dates, where each date is of the form `[YYYY]MMDD`. A list of consecutive dates can be generated by using the `SPAN` modifier, and a date can be adjusted with the `OFFSET` modifier. An example of a simple `BYDATE` clause follows:
	`BYDATE=0115,0315,0615,0915,1215,20060115`
	The following `SPAN` example is equivalent to `BYDATE=0110,0111,0112,0113,0114`, which is a span of 5 days starting at 1/10:
	`BYDATE=0110+SPAN:5D`
	The plus sign in front of the `SPAN` keyword indicates a span starting at the supplied date. The minus sign indicates a span ending at the supplied date, and the "^" sign indicates a span of *n* days or weeks centered around the supplied date. If *n* is an even number, it is adjusted up to the next odd number.
	Offsets adjust the supplied date by adding or subtracting *n* days or weeks. `BYDATE=0205-OFFSET:2W` is equivalent to `BYDATE=0205-14D` (the `OFFSET:` keyword is optional), which is also equivalent to `BYDATE=0122`.

BYMONTHDAY	This specifies the day of the month as a number. Valid values are 1 to 31. An example is 10, which means the 10th day of the selected month. You can use the minus sign (-) to count backward from the last day, so, for example, `BYMONTHDAY=-1` means the last day of the month and `BYMONTHDAY=-2` means the next to last day of the month.
BYDAY	This specifies the day of the week from Monday to Sunday in the form MON, TUE, and so on. Using numbers, you can specify the 26th Friday of the year, if using a YEARLY frequency, or the 4th THU of the month, using a MONTHLY frequency. Using the minus sign, you can say the second to last Friday of the month. For example, `-1 FRI` is the last Friday of the month.
BYHOUR	This specifies the hour on which the job is to run. Valid values are 0 to 23. As an example, 10 means 10 a.m.

Table 169-1 (Cont.) Values for repeat_interval

Name	Description
BYMINUTE	This specifies the minute on which the job is to run. Valid values are 0 to 59. As an example, 45 means 45 minutes past the chosen hour.
BYSECOND	This specifies the second on which the job is to run. Valid values are 0 to 59. As an example, 30 means 30 seconds past the chosen minute.
BYSETPOS	This selects one or more items, by position, in the list of timestamps that result after the whole calendaring expression is evaluated. It is useful for requirements such as running a job on the last workday of the month. Rather than attempting to express this with the other BY clauses, you can code the calendaring expression to evaluate to a list of every workday of the month, and then add the BYSETPOS clause to select only the last item of that list. Assuming that workdays are Monday through Friday, the syntax would then be: `FREQ=MONTHLY; BYDAY=MON,TUE,WED,THU,FRI; BYSETPOS=-1` Valid values are 1 through 9999. A negative number selects an item from the end of the list (-1 is the last item, -2 is the next to last item, and so on) and a positive number selects from the front of the list. The BYSETPOS clause is always evaluated last. BYSETPOS is only supported with the MONTHLY and YEARLY frequencies. The BYSETPOS clause is applied to the list of timestamps once per frequency period. For example, when the frequency is defined as MONTHLY, the Scheduler determines all valid timestamps for the month, orders that list, and then applies the BYSETPOS clause. The Scheduler then moves on to the next month and repeats the procedure. Assuming a start date of Jun 10, 2004, the example evaluates to: Jun 30, Jul 30, Aug 31, Sep 30, Oct 29, and so on.
INCLUDE	This includes one or more named schedules in the calendaring expression. That is, the set of timestamps defined by each included named schedule is added to the results of the calendaring expression. If an identical timestamp is contributed by both an included schedule and the calendaring expression, it is included in the resulting set of timestamps only once. The named schedules must have been defined with the CREATE_SCHEDULE procedure. This clause only works on a full day and therefore cannot be used with BYHOUR, BYMIN, and BYSECOND.
EXCLUDE	This excludes one or more named schedules from the calendaring expression. That is, the set of timestamps defined by each excluded named schedule is removed from the results of the calendaring expression. The named schedules must have been defined with the CREATE_SCHEDULE procedure. This clause only works on a full day and therefore cannot be used with BYHOUR, BYMIN, and BYSECOND.

Table 169-1 (Cont.) Values for repeat_interval

Name	Description

INTERSECT	This specifies an intersection between the calendaring expression results and the set of timestamps defined by one or more named schedules. Only the timestamps that appear both in the calendaring expression and in one of the named schedules are included in the resulting set of timestamps.
	For example, assume that the named schedule `last_sat` indicates the last Saturday in every month, and that for the year 2005, the only months where the last day of the month is also a Saturday are April and December. Assume also that the named schedule `end_qtr` indicates the last day of each quarter in 2005:
	`3/31/2005, 6/30/2005, 9/30/2005, 12/31/2005`
	These calendaring expressions result in the dates that follow:
	`3/31/2005, 4/30/2005, 6/30/2005, 9/30/2005, 12/31/2005`
	`FREQ=MONTHLY; BYMONTHDAY=-1; INTERSECT=last_sat,end_qtr`
	In this example, the terms `FREQ=MONTHLY; BYMONTHDAY=-1` indicate the last day of each month.
	This clause only works on a full day and therefore cannot be used with `BYHOUR`, `BYMIN`, and `BYSECOND`.

PERIODS	This identifies the number of periods that together form one cycle of a user-defined frequency. It is used in the `repeat_interval` expression of the schedule that defines the user-defined frequency. It is mandatory when the `repeat_interval` expression in the main schedule contains a `BYPERIOD` clause. The following example defines the quarters of a fiscal year. `FREQ=YEARLY;BYDATE=0301,0601,0901,1201;PERIODS=4`
BYPERIOD	This selects periods from a user-defined frequency. For example, if a main schedule names a user-defined frequency schedule that defines the fiscal quarters shown in the previous example, the clause `BYPERIOD=2,4` in the main schedule selects the 2nd and 4th fiscal quarters.

Combining Schedules

There are two ways to combine schedules:

- Using a combined schedule expression, which is a list of individual schedules

  For example, to create a schedule for all company holidays, you provide a list of individual schedules, where each schedule in the list defines a single holiday. The Scheduler evaluates each individual schedule, and then returns a union of the timestamps returned by each individual schedule.

- Embedding other schedules into the main schedule using `include`, `exclude`, and `intersect` clauses

  With this method, the embedded schedules inherit certain attributes from the main schedule.

  – Timestamps generated by the `INCLUDE` clause that fall into periods that are skipped by the main schedule are ignored. This is the case when the main schedule skips periods due to the `INTERVAL` clause, the `BYPERIOD` clause, or the `BYMONTH` clause for `freq=monthly`.

- Days that are added by the INCLUDE clause follow the hourly/minutely/secondly execution pattern of the main schedule.
- When the INCLUDE clause is present, no date-specific defaults are retrieved from the start date (but time-specific defaults can be). (See "Start Dates and Repeat Intervals", later in this section.) For example, a repeat_interval of FREQ=MONTHLY;INCLUDE=HOLIDAY executes only on holidays and not on the month/day defaults retrieved from the start date.

The following is an example:

```
BEGIN
dbms_scheduler.create_schedule('embed_sched', repeat_interval =>
  'FREQ=YEARLY;BYDATE=0130,0220,0725');
dbms_scheduler.create_schedule('main_sched', repeat_interval =>
  'FREQ=MONTHLY;INTERVAL=2;BYMONTHDAY=15;BYHOUR=9,17;INCLUDE=embed_sched');
END;
/
```

In this example, the dates 1/30, 2/20, and 7/25 are added to the main schedule. However, the Scheduler does not include dates that fall in months that are skipped by the INTERVAL clause. If the start date of the main schedule is 1/1/2005, then 2/20 is not added. On the dates that are added, the embedded schedule follows the execution pattern of the main schedule: jobs are executed at 9:00 a.m. and 5:00 p.m. on 1/30 and 7/25. If the embedded schedule does not itself have a start date, it inherits the start date from the main schedule.

User-Defined Frequencies

Instead of using predefined frequencies like DAILY, WEEKLY, MONTHLY, and so on, you can create your own frequencies by creating a schedule that returns the start date of each period. For example, the following repeat_interval expression is used in a schedule named fiscal_year that defines the start of each quarter in a fiscal year:

```
FREQ=YEARLY;BYDATE=0301,0601,0901,1201;PERIODS=4
```

To return the last Wednesday of every quarter, you create a schedule (the "main schedule") that uses the fiscal_year schedule as a user-defined frequency:

```
FREQ=fiscal_year;BYDAY=-1WED
```

Periods in a user-defined frequency do not have to be equal in length. In the main schedule, the BYSETPOS clause and numbered weekdays are recalculated based on the size of each period. To select dates in specific periods, you must use the BYPERIOD clause in the main schedule. To enable this, the schedule that is used as the user-defined frequency must include a PERIODS clause, and it must set its start date appropriately. The first date returned by this schedule is used as the starting point of period 1.

As another example, assuming work days are Monday through Friday, to get the last work day of the 2nd and 4th quarters of the fiscal year, the repeat_interval clause in the main schedule is the following:

```
FREQ=fiscal_year;BYDAY=MON,TUE,WED,THU,FRI;BYPERIOD=2,4;BYSETPOS=-1
```

Start Dates and Repeat Intervals

The Scheduler retrieves the date and time from the job or schedule start date and incorporates them as defaults into the repeat_interval. For example, if the specified frequency is yearly and there is no BYMONTH or BYMONTHDAY clause in the repeat interval, then the month and day that the job runs on are retrieved from the start date. Similarly, if frequency is monthly but there is no BYMONTHDAY clause in the repeat interval, then the day of the month that the job runs on is

retrieved from the start date. If present, BYHOUR, BYMINUTE, and BYSECOND defaults are also retrieved from the start date, and used if those clauses are not specified. Note that if the INCLUDE, EXCLUDE, or INTERSECT clauses are present, no date-related defaults are retrieved from the start date, but time-related defaults are.The following are some examples:

```
start_date:      4/15/05 9:00:00
repeat_interval: freq=yearly
```

is expanded internally to:

```
freq=yearly;bymonth=4;bymonthday=15;byhour=9;byminute=0;bysecond=0
```

The preceding schedule executes on 04/15/05 9:00:00, 04/15/06 9:00:00, 04/15/07 9:00:00, and so on.

For the next example, assume that schedule S1 has a repeat_interval of FREQ=YEARLY;BYDATE=0701.

```
start_date:      01/20/05 9:00:00
repeat_interval: freq=yearly;include=S1
```

is expanded internally to:

```
freq=yearly;byhour=9;byminute=0;bysecond=0;include=S1
```

Because an INCLUDE clause is present, date-related information is not retrieved from the start date. However, time-specific information is, so the preceding schedule executes on 07/01/05 9:00:00, 07/01/06 9:00:00, 07/01/08 9:00:00, and so on.

General Rules

When using a calendaring expression, consider the following rules:

- For a regular schedule (as opposed to a combined schedule), the calendar string must start with the frequency clause. All other clauses are optional and can be put in any order.
- All clauses are separated by a semicolon, and each clause can be present at most once, with the exception of the include, exclude, and intersect clauses.
- Spaces are allowed between syntax elements and the strings are case-insensitive.
- The list of values for a specific BY clause do not need to be ordered.
- When not enough BY clauses are present to determine what the next date is, this information is retrieved from the start date. For example, "FREQ=YEARLY" with a start date of 02/15/2003 becomes "FREQ=YEARLY;BYMONTH=FEB; BYMONTHDAY=15", which means every year on the 15th of February.

  "FREQ=YEARLY;BYMONTH=JAN,JUL" with start date 01/21/2003 becomes "FREQ=YEARLY;BYMONTH=JAN,JUL;BYMONTHDAY=21", which means every year on January 21 and July 21.
- The byweekno clause is only allowed if the frequency is YEARLY. It cannot be used with other frequencies. When it is present, it will return all days in that week number. If you want to limit it to specific days within the week, you have to add a BYDAY clause. For example, "FREQ=YEARLY;BYWEEKNO=2" with a start date of 01/01/2003 will return:

  ```
  01/06/2003, 01/07/2003, 01/08/2003, 01/09/2003, 01/10/2003, 01/11/2003, 01/12/2003,
  01/05/2004, 01/06/2004, 01/07/2004, .... and so on.
  ```

Note that when the `byweekno` clause is used, it is possible that the dates returned are from a year other than the current year. For example, if returning dates for the year 2004 and the calendar string is `"FREQ=YEARLY;BYWEEKNO=1,53"` for the specified week numbers in 2004, it will return the dates:

```
12/29/03, 12/30/03, 12/31/03, 01/01/04, 01/02/04, 01/03/04, 01/04/04, 12/27/04,
12/28/04, 12/29/04, 12/30/04, 12/31/04, 01/01/05, 01/02/05
```

- For those `BY` clauses that do not have a consistent range of values, you can count backward by putting a "-" in front of the numeric value. For example, specifying `BYMONTHDAY=31` will not give you the last day of every month, because not every month has 31 days. Instead, `BYMONTHDAY=-1` will give you the last day of the month.

  This is not supported for `BY` clauses that are fixed in size. In other words, `BYMONTH`, `BYHOUR`, `BYMINUTE`, and `BYSECOND` are not supported.

- The basic values for the `BYDAY` clause are the days of the week. When the frequency is `YEARLY`, or `MONTHLY`, you are allowed to specify a positive or negative number in front of each day of the week. In the case of `YEARLY`, `BYDAY=40MON`, indicates the 40th Monday of the year. In the case of `MONTHLY, BYDAY=-2SAT`, indicates the second to last Saturday of the month.

  Note that positive or negative numbers in front of the weekdays are not supported for other frequencies and that in the case of yearly, the number ranges from -53 ... -1, 1 ... 53, whereas for the monthly frequency it is limited to -5 ... -1, 1... 5.

  If no number is present in front of the weekday it specifies, every occurrence of that weekday in the specified frequency.

- The first day of the week is Monday.

- Repeating jobs with frequencies smaller than daily follow their frequencies exactly across daylight savings adjustments. For example, suppose that a job is scheduled to repeat every 3 hours, the clock is moved forward from 1:00 a.m. to 2:00 a.m., and the last time the job ran was midnight. Its next scheduled time will be 4:00 a.m. Thus, the 3 hour period between subsequent job runs is retained. The same applies when the clock is moved back. This behavior is not the case for repeating jobs that have frequencies of daily or larger. For example, if a repeating job is supposed to be executed on a daily basis at midnight, it will continue to run at midnight if the clock is moved forward or backward. When the execution time of such a daily (or larger frequency) job happens to fall inside a window where the clock is moved forward, the job executes at the end of the window.

- The calendaring syntax does not allow you to specify a time zone. Instead the Scheduler retrieves the time zone from the `start_date` argument. If jobs must follow daylight savings adjustments, then you must specify a region name for the time zone of the `start_date`. For example specifying the `start_date` time zone as `'US/Eastern'` in New York ensures that daylight saving adjustments are automatically applied. If instead, the time zone of the `start_date` is set to an absolute offset, such as `'-5:00'`, then daylight savings adjustments are not followed and your job execution is off by an hour for half the year.

- When `start_date` is `NULL`, the Scheduler determines the time zone for the repeat interval as follows:

  1. It checks whether or not the session time zone is a region name. The session time zone can be set by either:

     – Issuing an `ALTER SESSION` statement, for example:

     ```
     SQL> ALTER SESSION SET time_zone = 'Asia/Shanghai';
     ```

     – Setting the `ORA_SDTZ` environment variable.

2. If the session time zone is an absolute offset instead of a region name, the Scheduler uses the value of the `DEFAULT_TIMEZONE` Scheduler attribute. For more information, see the SET_SCHEDULER_ATTRIBUTE Procedure.
3. If the `DEFAULT_TIMEZONE` attribute is `NULL`, the Scheduler uses the time zone of `systimestamp` when the job or window is enabled.

BYSETPOS Clause Rules

The following are rules for the `BYSETPOS` clause.

- The `BYSETPOS` clause is the last clause to be evaluated. It is processed after all other `BY` clauses and the `INCLUDE`, `EXCLUDE` and `INTERSECT` clauses have been evaluated.
- The `INTERVAL` clause does not change the size of the period to which the `BYSETPOS` clause is applied. For example, when the frequency is set to monthly and interval is set to 3, the list of timestamps to which `BYSETPOS` is applied is generated from a month, not a quarter. The only impact of the `INTERVAL` clause is to cause months to be skipped. However, you can still select the second to last workday of the quarter like this:

  ```
  FREQ=MONTHLY;INTERVAL=3;BYDAY=MON,TUE,WED,THU,FRI;BYSETPOS=-2
  ```

  provided that you set the start date in the right month. This example returns the next to last workday of a month, and repeats once a quarter.
- To get consistent results, the set to which `BYSETPOS` is applied is determined from the beginning of the frequency period independently of when the evaluation occurs. Whether the Scheduler evaluates

  ```
  FREQ=MONTHLY;BYDAY=MON,TUE,FRI;BYSETPOS=1,3
  ```

  on 01/01/2004 or 01/15/2004, in both cases the expression evaluates to Friday 01/02/2004, and Tuesday 01/06/2004. The only difference is that when the expression is evaluated on 01/15/2004, the Scheduler determines that there are no matches in January because the timestamps found are in the past, and it moves on to the matches in the next month, February.

BYDATE Clause Rules

The following are rules for the `BYDATE` clause.

- If dates in the `BYDATE` clause do not have their optional year component, the job runs on those dates every year.
- The job execution times on the included dates are derived from the `BY` clauses in the calendaring expression. For example, if `repeat_interval` is defined as

  ```
  freq=daily;byhour=8,13,18;byminute=0;bysecond=0;bydate=0502,0922
  ```

  then the execution times on 05/02 and 09/22 are 8:00 a.m., 1:00 p.m., and 6:00 p.m.

EXCLUDE Clause Rules

Excluded dates without a time component are 24 hour periods. All timestamps that fall on an excluded date are removed. In the following example, `jan_fifteen` is a named schedule that resolves to the single date of 01/15:

```
freq=monthly;bymonthday=15,30;byhour=8,13,18;byminute=0;bysecond=0;
     exclude=jan_fifteenth
```

In this case, all three instances of the job are removed for 01/15.

OFFSET Rules

You can adjust the dates of individual named schedules by adding positive offsets to them. For example, to execute `JOB2` exactly 15 days after every occurrence of `JOB1`, add `+OFFSET:15D` to the schedule of `JOB1`, as follows:

```
BEGIN
dbms_scheduler.create_schedule('job2_schedule', repeat_interval =>
  'job1_schedule+OFFSET:15D');
END;
/
```

Note that negative offsets to named schedules are not supported.

Example 169-1 Putting It All Together

This example demonstrates the use of user-defined frequencies, spans, offsets, and the `BYSETPOS` and `INCLUDE` clauses. (Note that the `OFFSET:` keyword is optional in an offset clause.)

Many companies in the retail industry share the same fiscal year. The fiscal year starts on the Sunday closest to February 1st, and subsequent quarters start exactly 13 weeks later. The fiscal year schedule for the retail industry can be defined as the following:

```
begin
 dbms_scheduler.create_schedule('year_start', repeat_interval=>
       'FREQ=YEARLY;BYDATE=0201^SPAN:1W;BYDAY=SUN');
 dbms_scheduler.create_schedule('retail_fiscal_year',
        to_timestamp_tz('15-JAN-2005 12:00:00','DD-MON-YYYY HH24:MI:SS'),
         'year_start,year_start+13w,year_start+26w,year_start+39w;periods=4');
end;
/
```

The following schedule can be used to execute a job on the 5th day off in the 2nd and the 4th quarters of the retail industry. This assumes that Saturday and Sunday are off days as well as the days in the existing `holiday` schedule.

```
begin
 dbms_scheduler.create_schedule('fifth_day_off', repeat_interval=>
  'FREQ=retail_fiscal_year;BYDAY=SAT,SUN;INCLUDE=holiday;
    BYPERIOD=2,4;BYSETPOS=5');
end;
/
```

169.5 DBMS_SCHEDULER Data Structures

The `DBMS_SCHEDULER` package defines `OBJECT` types and `TABLE` types.

OBJECT Types

- JOBARG Object Type
- JOB_DEFINITION Object Type
- JOBATTR Object Type
- SCHEDULER$_STEP_TYPE Object Type
- SCHEDULER$_EVENT_INFO Object Type
- SCHEDULER_FILEWATCHER_RESULT Object Type
- SCHEDULER_FILEWATCHER_REQUEST Object Type

TABLE Types

- JOBARG_ARRAY Table Type
- JOB_DEFINITION_ARRAY Table Type
- JOBATTR_ARRAY Table Type
- SCHEDULER$_STEP_TYPE_LIST Table Type

169.5.1 DBMS_SCHEDULER JOBARG Object Type

This type is used by the `JOB` and `JOBATTR` object types. It represents a job argument in a batch of job arguments.

Syntax

```
TYPE jobarg IS OBJECT (
   arg_position         NUMBER,
   arg_text_value       VARCHAR2(4000),
   arg_anydata_value    ANYDATA,
   arg_operation        VARCHAR2(5));
```

Attributes

Table 169-2 JOBARG Object Type Attributes

Attribute	Description
`arg_position`	Position of the argument
`arg_text_value`	Value of the argument if the type is `VARCHAR2`
`arg_anydata_value`	Value of the argument if the type is `AnyData`
`arg_operation`	Type of the operation: • `SET` • `RESET`

JOBARG Constructor Function

This constructor function constructs a job argument. It is overloaded to construct job arguments with different types of values.

Syntax

Constructs a job argument with a text value.

```
constructor function jobarg (
   arg_position        IN POSITIVEN,
   arg_value           IN VARCHAR2)
   RETURN SELF AS RESULT;
```

Constructs a job argument with an `AnyData` value.

```
constructor function jobarg (
   arg_position        IN POSITIVEN,
   arg_value           IN ANYDATA)
   RETURN SELF AS RESULT;
```

Constructs a job argument with a `NULL` value.

```
constructor function jobarg (
   arg_position        IN POSITIVEN,
   arg_reset           IN BOOLEAN DEFAULT FALSE)
   RETURN SELF AS RESULT;
```

Parameters

Table 169-3 JOBARG Constructor Function Parameters

Parameter	Description
arg_position	Position of the argument
arg_value	Value of the argument
arg_reset	If arg_reset is TRUE, then the argument at that position is reset. Setting arg_reset to FALSE (which is the default) will create an argument with a NULL value.

JOBARG_ARRAY Table Type

Syntax

```
TYPE jobarg_array IS TABLE OF jobarg;
```

169.5.2 JOBARG_ARRAY Table Type

The jobarg_array type is a table of jobarg.

Syntax

```
TYPE jobarg_array IS TABLE OF jobarg;
```

169.5.3 DBMS_SCHEDULER JOB_DEFINITION Object Type

This type is used by the CREATE_JOBS procedure and represents a job in a batch of jobs.

Syntax

```
TYPE job_definition IS OBJECT (
  job_name                       VARCHAR2(100),
  job_class                      VARCHAR2(32),
  job_style                      VARCHAR2(11),
  program_name                   VARCHAR2(100),
  job_action                     VARCHAR2(4000),
  job_type                       VARCHAR2(20),
  schedule_name                  VARCHAR2(65),
  repeat_interval                VARCHAR2(4000),
  schedule_limit                 INTERVAL DAY TO SECOND,
  start_date                     TIMESTAMP WITH TIME ZONE,
  end_date                       TIMESTAMP WITH TIME ZONE,
  event_condition                VARCHAR2(4000),
  queue_spec                     VARCHAR2(100),
  number_of_arguments            NUMBER,
  arguments                      SYS.JOBARG_ARRAY,
  job_priority                   NUMBER,
  job_weight                     NUMBER,
  max_run_duration               INTERVAL DAY TO SECOND,
  max_runs                       NUMBER,
  max_failures                   NUMBER,
```

```
  logging_level                 NUMBER,
  restartable                   VARCHAR2(5),
  stop_on_window_close          VARCHAR2(5),
  raise_events                  NUMBER,
  comments                      VARCHAR2(240),
  auto_drop                     VARCHAR2(5),
  enabled                       VARCHAR2(5),
  follow_default_timezone       VARCHAR2(5),
  parallel_instances            VARCHAR2(5),
  aq_job                        VARCHAR2(5),
  instance_id                   NUMBER,
  credential_name               VARCHAR2(65),
  destination                   VARCHAR2(4000),
  database_role                 VARCHAR2(20),
  allow_runs_in_restricted_mode VARCHAR2(5);
  restart_on_recovery           BOOLEAN;
  restart_on_failure            BOOLEAN;)
```

Object Attributes

Table 169-4 provides brief descriptions of the attributes of the `JOB_DEFINITION` object type. For more complete information about these attributes, see the "CREATE_JOB Procedure" and the "SET_ATTRIBUTE Procedure".

Table 169-4 JOB_DEFINITION Object Types

Attribute	Description
`job_name`	Name of the job
`job_class`	Name of the job class
`job_style`	Style of the job: • `REGULAR` • `LIGHTWEIGHT` • `IN_MEMORY_RUNTIME` • `IN_MEMORY_FULL`
`program_name`	Name of the program that the job runs
`job_action`	Inline action of the job. This is either the code for an anonymous PL/SQL block or the name of a stored procedure, external executable, or chain.
`job_type`	Job action type (`'PLSQL_BLOCK'`, `'STORED_PROCEDURE'`, `'EXECUTABLE'`, `'CHAIN'`, `'EXTERNAL_SCRIPT'`, `'SQL_SCRIPT'`, and `'BACKUP_SCRIPT'`)
`schedule_name`	Name of the schedule that specifies when the job has to execute
`repeat_interval`	Inline time-based schedule
`schedule_limit`	Maximum delay time between scheduled and actual job start before a job run is canceled
`start_date`	Start date and time of the job
`end_date`	End date and time of the job
`event_condition`	Event condition for event-based jobs
`queue_spec`	File watcher name or queue specification for event-based jobs
`number_of_arguments`	Number of job arguments
`arguments`	Array of job arguments
`job priority`	Job priority

Table 169-4 (Cont.) JOB_DEFINITION Object Types

Attribute	Description
job_weight	*** Deprecated in Oracle Database 11*g*R2. Do not change the value of this attribute from the default, which is 1. Weight of the job for parallel execution.
max_run_duration	Maximum run duration of the job
max_runs	Maximum number of runs before the job is marked as completed
max_failures	Maximum number of failures tolerated before the job is marked as broken
logging_level	Job logging level
restartable	Indicates whether the job is restartable (TRUE) or not (FALSE)
stop_on_window_close	Indicates whether the job is stopped when the window that it runs in ends (TRUE) or not (FALSE). Equivalent to the stop_on_window_close job attribute described in the SET_ATTRIBUTE Procedure.
raise_events	State changes that raise events
comments	Comments on the job
auto_drop	If TRUE (the default), indicates that the job should be dropped once completed
enabled	Indicates whether the job should be enabled immediately after creating it (TRUE) or not (FALSE)
follow_default_timezone	If TRUE and if the job start_date is null, then when the default_timezone scheduler attribute is changed, the Scheduler recomputes the next run date and time for this job so that it is in accordance with the new time zone.
parallel_instances	For event-based jobs only. If TRUE, on the arrival of the specified event, the Scheduler creates a new lightweight job to handle that event, so multiple instances of the same event-based job can run in parallel. If FALSE, then an event is discarded if it is raised while the job that handles it is already running,
aq_job	For internal use only
instance_id	The instance ID of the instance that the job must run on For in-memory full jobs, the instance_id value determines in which instance to stop the job; if left NULL, the job is stopped in all instances.
credential_name	The credential to use for a single destination or the default credential for a group of destinations
destination	The name of a single external destination or database destination, or a group name of type external destination or database destination
database_role	In an Oracle Data Guard environment, the database role ('PRIMARY' or 'LOGICAL STANDBY') for which the job runs
allow_runs_in_restricted_mode	If TRUE, the job is permitted to run when the database is in restricted mode, provided that the job owner is permitted to log in during this mode

Table 169-4 (Cont.) JOB_DEFINITION Object Types

Attribute	Description
restart_on_recovery	If set to TRUE for a job and the job is stopped by a database shutdown, then the job is restarted when the database is recovered. If set to FALSE, and the job is stopped by a database shutdown, then the job is marked as stopped when the database is recovered.
restart_on_failure	If set to TRUE for a job and the job fails due to an application error, then the job is retried using the normal Scheduler retry mechanism (after 1 second, after 10 seconds, after 100 seconds, and so on, up to a maximum of 6 times). If all 6 retries fail (after about 30 hours), then the job is marked FAILED. If set to FALSE (the default), a failed job is immediately marked FAILED.

JOB_DEFINITION Constructor Function

This constructor function constructs a job_definition object.

Syntax

```
constructor function job_definition (
    job_name                IN     VARCHAR2,
    job_style               IN     VARCHAR2 DEFAULT 'REGULAR',
    program_name            IN     VARCHAR2 DEFAULT NULL,
    job_action              IN     VARCHAR2 DEFAULT NULL,
    job_type                IN     VARCHAR2 DEFAULT NULL,
    schedule_name           IN     VARCHAR2 DEFAULT NULL,
    repeat_interval         IN     VARCHAR2 DEFAULT NULL,
    event_condition         IN     VARCHAR2 DEFAULT NULL,
    queue_spec              IN     VARCHAR2 DEFAULT NULL,
    start_date              IN     TIMESTAMP WITH TIME ZONE DEFAULT NULL,
    end_date                IN     TIMESTAMP WITH TIME ZONE DEFAULT NULL,
    number_of_arguments     IN     NATURAL DEFAULT NULL,
    arguments               IN     SYS.JOBARG_ARRAY DEFAULT NULL,
    job_class               IN     VARCHAR2 DEFAULT 'DEFAULT_JOB_CLASS',
    schedule_limit          IN     INTERVAL DAY TO SECOND DEFAULT NULL,
    job_priority            IN     NATURAL DEFAULT NULL,
    job_weight              IN     NATURAL DEFAULT NULL,
    max_run_duration        IN     INTERVAL DAY TO SECOND DEFAULT NULL,
    max_runs                IN     NATURAL DEFAULT NULL,
    max_failures            IN     NATURAL DEFAULT NULL,
    logging_level           IN     NATURALN DEFAULT 64,
    restartable             IN     BOOLEAN DEFAULT FALSE,
    stop_on_window_close    IN     BOOLEAN DEFAULT FALSE,
    raise_events            IN     NATURAL DEFAULT NULL,
    comments                IN     VARCHAR2 DEFAULT NULL,
    auto_drop               IN     BOOLEAN DEFAULT TRUE,
    enabled                 IN     BOOLEAN DEFAULT FALSE,
    follow_default_timezone IN     BOOLEAN DEFAULT FALSE,
    parallel_instances      IN     BOOLEAN DEFAULT FALSE,
    aq_job                  IN     BOOLEAN DEFAULT FALSE,
    instance_id             IN     NATURAL DEFAULT NULL,
    credential_name         IN     VARCHAR2 DEFAULT NULL,
    destination             IN     VARCHAR2 DEFAULT NULL,
    database_role           IN     VARCHAR2 DEFAULT NULL,
    allow_runs_in_restricted_mode IN BOOLEAN DEFAULT FALSE)
    RETURN SELF AS RESULT;
```

JOB_DEFINITION_ARRAY Table Type

Syntax

```
TYPE job_definition_array IS TABLE OF job_definition;
```

169.5.4 JOB_DEFINITION_ARRAY Table Type

The type `job_definition_array` is a table of `job_definition`.

Syntax

```
TYPE job_definition_array IS TABLE OF job_definition;
```

169.5.5 JOBATTR Object Type

This type is used by the `SET_JOB_ATTRIBUTES` procedure and represents a job attribute in a batch of job attributes.

Syntax

```
TYPE jobattr IS OBJECT (
   job_name             VARCHAR2(100),
   attr_name            VARCHAR2(30),
   char_value           VARCHAR2(4000),
   char_value2          VARCHAR2(4000),
   args_value           JOBARG_ARRAY,
   num_value            NUMBER,
   timestamp_value      TIMESTAMP(6) WITH TIME ZONE,
   interval_value       INTERVAL DAY(2) TO SECOND(6));
```

Attributes

Table 169-5 JOBATTR Object Type Attributes

Attribute	Description
`job_name`	Name of the job
`attr_name`	Name of the attribute
`char_value`	Value of the argument if the type is `VARCHAR2`
`char_value2`	Second `VARCHAR2` attribute value
`args_value`	Value of the argument if the type is a `JOBARG` array
`num_value`	Value of the argument if the type is `NUMBER`
`timestamp_value`	Value of the argument if the type is `TIMESTAMP WITH TIME ZONE`
`interval_value`	Value of the argument if the type is `INTERVAL DAY TO SECOND`

JOBATTR Constructor Function

This constructor function constructs a job attribute. It is overloaded to create attribute values of the following types: `VARCHAR2`, `NUMBER`, `TIMESTAMP WITH TIME ZONE`, `INTERVAL DAY TO SECOND`, and an array of `JOBARG` types.

Syntax

```
constructor function jobattr (
   job_name            IN VARCHAR2,
   attr_name           IN VARCHAR2,
   attr_value          IN VARCHAR2,
   attr_value2         IN VARCHAR2 DEFAULT NULL)
   RETURN SELF AS RESULT;

constructor function jobattr (
   job_name            IN VARCHAR2,
   attr_name           IN VARCHAR2,
   attr_value          IN [NUMBER, BOOLEAN,
                           TIMESTAMP WITH TIME ZONE,
                           INTERVAL DAY TO SECOND, JOBARG_ARRAY])
   RETURN SELF AS RESULT;

constructor function jobattr (
   job_name            IN VARCHAR2,
   attr_name           IN VARCHAR2)
   RETURN SELF AS RESULT;
```

Parameters

Table 169-6 JOBATTR Constructor Function Parameters

Parameter	Description
job_name	Name of the job
attr_name	Name of the argument
attr_value	Value of the argument
attr_value2	Most attributes have only one value associated with them, but some can have two. The attr_value2 argument is for this optional second value.

JOBATTR Table Type

Syntax

```
TYPE jobattr_array IS TABLE OF jobattr;
```

169.5.6 JOBATTR_ARRAY Table Type

The type jobattr_array is a table of jobattr.

Syntax

```
TYPE jobattr_array IS TABLE OF jobattr;
```

169.5.7 SCHEDULER$_STEP_TYPE Object Type

This type is used by RUN_CHAIN to return a list of chain steps with an initial state.

Syntax

```
TYPE scheduler$_step_type IS OBJECT (
   step_name  VARCHAR2(32),
   step_type  VARCHAR2(32));
```

Attributes

Table 169-7 SCHEDULER$_STEP_TYPE Object Type Attributes

Attribute	Description
step_name	Name of the step
step_type	State of the step

169.5.8 SCHEDULER$_STEP_TYPE_LIST Table Type

This type is a table of scheduler$_step_type.

Syntax

```
TYPE scheduler$_step_type_list IS TABLE OF scheduler$_step_type;
```

169.5.9 SCHEDULER$_EVENT_INFO Object Type

This the datatype of the Scheduler event queue SYS.SCHEDULER$_EVENT_QUEUE, from which your application consumes job state events raised by the Scheduler.

It is a secure queue owned by SYS.

Syntax

```
TYPE SCHEDULER$_EVENT_INFO IS OBJECT (
  event_type         VARCHAR2(4000),
  object_owner       VARCHAR2(4000),
  object_name        VARCHAR2(4000),
  event_timestamp    TIMESTAMP WITH TIME ZONE,
  error_code         NUMBER,
  error_msg          VARCHAR2(4000),
  event_status       NUMBER,
  log_id             NUMBER,
  run_count          NUMBER,
  failure_count      NUMBER,
  retry_count        NUMBER,
  spare1             NUMBER,
  spare2             NUMBER,
  spare3             VARCHAR2(4000),
  spare4             VARCHAR2(4000),
  spare5             TIMESTAMP WITH TIME ZONE,
  spare6             TIMESTAMP WITH TIME ZONE,
  spare7             RAW(2000),
  spare8             RAW(2000));
```

Attributes

Table 169-8 SCHEDULER_EVENT_INFO Object Type Attributes

Attribute	Description
event_type	One of "JOB_STARTED", "JOB_SUCCEEDED", "JOB_FAILED", "JOB_BROKEN", "JOB_COMPLETED", "JOB_STOPPED", "JOB_SCH_LIM_REACHED", "JOB_DISABLED", "JOB_CHAIN_STALLED", "JOB_OVER_MAX_DUR". For descriptions of these event types, see Table 169-85.
object_owner	Owner of the job that raised the event
object_name	Name of the job that raised the event
event_timestamp	Time at which the event occurred
error_code	Applicable only when an error is thrown during job execution. Contains the top-level error code.
error_msg	Applicable only when an error is thrown during job execution. Contains the entire error stack.
event_status	Adds further qualification to the event type. If event_type is "JOB_STARTED," status 1 indicates that it is a normal start, and status 2 indicates that it is a retry. If event_type is "JOB_FAILED," status 4 indicates that it was a failure due to an error that was thrown during job execution, and status 8 indicates that it was an unusual termination of some kind. If event_type is "JOB_STOPPED," status 16 indicates that it was a normal stop, and status 32 indicates that it was a stop with the FORCE option set to TRUE.
log_id	Points to the ID in the scheduler job log from which additional information can be obtained. Note that there need not always be a log entry corresponding to an event. In such cases, log_id is NULL.
run_count	Run count for the job when the event was raised.
failure_count	Failure count for the job when the event was raised.
retry_count	Retry count for the job when the event was raised.
spare1 – spare8	Not currently in use.

169.5.10 SCHEDULER_FILEWATCHER_RESULT Object Type

This is the datatype of a file arrival event message.

You access the event message as a parameter of an event-based job (or a parameter of a program referenced by an event-based job). The message contains information needed to locate and process a file that arrived on a local or remote system.

Syntax

```
TYPE scheduler_filewatcher_result IS OBJECT (
  destination          VARCHAR2(4000),
  directory_path       VARCHAR2(4000),
  actual_file_name     VARCHAR2(4000),
  file_size            NUMBER,
  file_timestamp       TIMESTAMP WITH TIME ZONE,
  ts_ms_from_epoch     NUMBER,
  matching_requests    SYS.SCHEDULER_FILEWATCHER_REQ_LIST);
```

Attributes

Table 169-9 SCHEDULER_FILEWATCHER_RESULT Object Type Attributes

Attribute	Description
destination	Destination at which the file was found, expressed as a host name or IP address.
directory_path	Absolute path of directory in which the file was found.
actual_file_name	Actual name of the file that was found. If the file name specified in the file watcher did not contain wildcards, then this is the same as the name specified in the file watcher.
file_size	Size of the file that was found, in bytes.
file_timestamp	Timestamp assigned to the file when the file watcher considered the file found, based on the minimum file size and steady state duration attributes.
ts_ms_from_epoch	For internal use only.
matching_requests	List of matching requests. This is a TABLE of type objects SCHEDULER_FILEWATCHER_REQUEST. Each matching request corresponds to a file watcher whose destination, directory_path, and file_name attributes matched the arrived file. See "SCHEDULER_FILEWATCHER_REQUEST Object Type".

169.5.11 SCHEDULER_FILEWATCHER_REQUEST Object Type

This type is returned in the matching_requests attribute of the SCHEDULER_FILEWATCHER_RESULT Object Type. Its attributes are similar to the attributes of a file watcher.

Syntax

```
TYPE scheduler_filewatcher_request IS OBJECT (
  owner                 VARCHAR2(4000),
  name                  VARCHAR2(4000),
  requested_path_name   VARCHAR2(4000),
  requested_file_name   VARCHAR2(4000),
  credential_owner      VARCHAR2(4000),
  credential_name       VARCHAR2(4000),
  min_file_size         NUMBER,
  steady_state_dur      NUMBER);
```

Attributes

Table 169-10 SCHEDULER_FILEWATCHER_REQUEST Object Type Attributes

Attribute	Description
owner	Owner of the matched file watcher.
name	Name of the matched file watcher.
requested_path_name	Value of the directory_path attribute of the matched file watcher.
requested_file_name	Value of the file_name attribute of the matched file watcher.
credential_owner	Owner of the credential referenced by the matched file watcher.

Table 169-10 (Cont.) SCHEDULER_FILEWATCHER_REQUEST Object Type Attributes

Attribute	Description
`credential_name`	Name of the credential referenced by the matched file watcher.
`min_file_size`	Value of the `min_file_size` attribute of the matched file watcher.
`steady_state_dur`	Value of the `steady_state_duration` attribute of the matched file watcher.

Related Topics

- SCHEDULER_FILEWATCHER_RESULT Object Type
  This is the datatype of a file arrival event message.

169.6 Summary of DBMS_SCHEDULER Subprograms

This table lists the `DBMS_SCHEDULER` subprograms and briefly describes them.

Table 169-11 DBMS_SCHEDULER Package Subprograms

Subprogram	Description
ADD_EVENT_QUEUE_SUBSCRIBER Procedure	Adds a user as a subscriber to the Scheduler event queue `SYS.SCHEDULER$_EVENT_QUEUE`
ADD_GROUP_MEMBER Procedure	Adds one or more members to an existing group
ADD_JOB_EMAIL_NOTIFICATION Procedure	Adds e-mail notifications for a job for a list of recipients and a list of job state events
ADD_TO_INCOMPATIBILITY Procedure	Adds jobs or programs to an existing incompatibility definition
ALTER_CHAIN Procedure	Alters specified steps of a chain
ALTER_RUNNING_CHAIN Procedure	Alters specified steps of a running chain
CLOSE_WINDOW Procedure	Closes an open window prematurely
COPY_JOB Procedure	Copies an existing job
CREATE_CHAIN Procedure	Creates a chain, which is a named series of programs that are linked together for a combined objective
CREATE_CREDENTIAL Procedure	Creates a credential
CREATE_DATABASE_DESTINATION Procedure	Creates a database destination for use with remote database jobs
CREATE_EVENT_SCHEDULE Procedure	Creates an event schedule, which is a schedule that starts a job based on the detection of an event
CREATE_FILE_WATCHER Procedure	Creates a file watcher, which is a Scheduler object that defines the location, name, and other properties of a file whose arrival on a system causes the Scheduler to start a job

Table 169-11 (Cont.) DBMS_SCHEDULER Package Subprograms

Subprogram	Description
CREATE_GROUP Procedure	Creates a group
CREATE_INCOMPATIBILITY Procedure	Creates an incompatibility definition
CREATE_JOB Procedure	Creates a single job
CREATE_JOB_CLASS Procedure	Creates a job class, which provides a way to group jobs for resource allocation and prioritization
CREATE_JOBS Procedure	Creates multiple jobs
CREATE_PROGRAM Procedure	Creates a program
CREATE_RESOURCE Procedure	Specifies resources used by jobs or creates a new resource
CREATE_SCHEDULE Procedure	Creates a schedule
CREATE_WINDOW Procedure	Creates a window, which provides a way to automatically activate different resource plans at different times
DEFINE_ANYDATA_ARGUMENT Procedure	Defines a program argument whose value is of a complex type and must be passed encapsulated in an `AnyData` object
DEFINE_CHAIN_EVENT_STEP Procedure	Adds or replaces a chain step and associates it with an event schedule or inline event. See also: `DEFINE_CHAIN_STEP`.
DEFINE_CHAIN_RULE Procedure	Adds a rule to an existing chain
DEFINE_CHAIN_STEP Procedure	Defines a chain step, which can be a program or another (nested) chain. See also: `DEFINE_CHAIN_EVENT_STEP`.
DEFINE_METADATA_ARGUMENT Procedure	Defines a special metadata argument for the program. You can retrieve specific metadata through this argument.
DEFINE_PROGRAM_ARGUMENT Procedure	Defines a program argument whose value can be passed as a string literal to the program
DISABLE Procedure	Disables a program, job, chain, window, database destination, external destination, file watcher, group, or incompatibilty
DROP_AGENT_DESTINATION Procedure	Drops one or more external destinations. Use only when the preferred method of dropping external destinations—unregistering the Scheduler agent with the database—fails.
DROP_CHAIN Procedure	Drops an existing chain
DROP_CHAIN_RULE Procedure	Removes a rule from an existing chain
DROP_CHAIN_STEP Procedure	Drops a chain step
DROP_CREDENTIAL Procedure	Drops a credential

Table 169-11 (Cont.) DBMS_SCHEDULER Package Subprograms

Subprogram	Description
DROP_DATABASE_DESTINATION Procedure	Drops one or more database destinations
DROP_FILE_WATCHER Procedure	Drops one or more file watchers
DROP_GROUP Procedure	Drops one or more groups
DROP_INCOMPATIBILITY Procedure	Drops an existing incompatibility definition
DROP_JOB Procedure	Drops a job or all jobs in a job class
DROP_JOB_CLASS Procedure	Drops a job class
DROP_PROGRAM Procedure	Drops a program
DROP_PROGRAM_ARGUMENT Procedure	Drops a program argument
DROP_SCHEDULE Procedure	Drops a schedule
DROP_WINDOW Procedure	Drops a window
DUMP_IN_MEMORY_TRACE Procedure	Dumps the scheduler in-memory trace buffer of the specified process state object address into the current trace file of the requester process.
ENABLE Procedure	Enables a program, job, chain, window, database destination, external destination, file watcher, or group
END_DETACHED_JOB_RUN Procedure	Ends a running detached job
EVALUATE_CALENDAR_STRING Procedure	Evaluates the calendar string and tells you what the next execution date of a job or window will be
EVALUATE_RUNNING_CHAIN Procedure	Forces reevaluation of the rules of a running chain to trigger any rules for conditions that have been satisfied
GENERATE_JOB_NAME Function	Generates a unique name for a job. This enables you to identify jobs by adding a prefix, so, for example, Sally's jobs would be named `sally1`, `sally2`, and so on
GET_AGENT_INFO Function	Returns job information specific to an agent, such as how many are running and so on, depending on the attribute selected
GET_AGENT_VERSION Function	Returns the version string of a Scheduler agent that is registered with the database and is currently running
GET_ATTRIBUTE Procedure	Retrieves the value of an attribute of an object
GET_FILE Procedure	Retrieves a file from a host
GET_SCHEDULER_ATTRIBUTE Procedure	Retrieves the value of a Scheduler attribute
OPEN_WINDOW Procedure	Opens a window prematurely. The window is opened immediately for the duration

Table 169-11 (Cont.) DBMS_SCHEDULER Package Subprograms

Subprogram	Description
PURGE_LOG Procedure	Purges specific rows from the job and window logs
PUT_FILE Procedure	Saves a file to one or more hosts
REMOVE_EVENT_QUEUE_SUBSCRIBER Procedure	Unsubscribes a user from the Scheduler event queue `SYS.SCHEDULER$_EVENT_QUEUE`
REMOVE_FROM_INCOMPATIBILITY Procedure	Removes jobs or programs from an incompatibility definition
REMOVE_GROUP_MEMBER Procedure	Removes one or more members from a group
REMOVE_JOB_EMAIL_NOTIFICATION Procedure	Removes e-mail notifications for a job
RESET_JOB_ARGUMENT_VALUE Procedure	Resets the current value assigned to an argument defined with the associated program
RUN_CHAIN Procedure	Immediately runs a chain by creating a run-once job
RUN_JOB Procedure	Runs a job immediately
SET_AGENT_REGISTRATION_PASS Procedure	Sets the agent registration password for a database
SET_ATTRIBUTE Procedure	Changes an attribute of a job, schedule, or other Scheduler object
SET_ATTRIBUTE_NULL Procedure	Changes an attribute of an object to `NULL`
SET_JOB_ANYDATA_VALUE Procedure	Sets the value of a job argument encapsulated in an `AnyData` object
SET_JOB_ARGUMENT_VALUE Procedure	Sets the value of a job argument
SET_JOB_ATTRIBUTES Procedure	Sets the value of a job attribute
SET_RESOURCE_CONSTRAINT Procedure	Specifies the resources used by jobs
SET_SCHEDULER_ATTRIBUTE Procedure	Sets the value of a Scheduler attribute
STOP_JOB Procedure	Stops a currently running job or all jobs in a job class

169.6.1 ADD_EVENT_QUEUE_SUBSCRIBER Procedure

This procedure adds a user as a subscriber to the Scheduler event queue `SYS.SCHEDULER$_EVENT_QUEUE`, and grants the user permission to dequeue from this queue using the designated agent.

Syntax

```
DBMS_SCHEDULER.ADD_EVENT_QUEUE_SUBSCRIBER (
   subscriber_name         IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-12 ADD_EVENT_QUEUE_SUBSCRIBER Procedure Parameters

Parameter	Description
subscriber_name	Name of the Oracle Advanced Queuing (AQ) agent to be used to subscribe to the Scheduler event queue. If NULL, an agent is created and assigned the user name of the calling user.

Usage Notes

The subscription is rule-based. The rule permits the user to see only events raised by jobs that the user owns, and filters out all other messages. If an AQ agent with the same name already exists, an error is raised.

169.6.2 ADD_GROUP_MEMBER Procedure

This procedure adds one or more members to an existing group.

Syntax

```
DBMS_SCHEDULER.ADD_GROUP_MEMBER (
   group_name              IN VARCHAR2,
   member                  IN VARCHAR2);
```

Parameters

Table 169-13 ADD_GROUP_MEMBER Procedure Parameters

Parameter	Description
group_name	The name of the group.
member	A comma-separated list of members to add to the group. Members must match the group type. A group of the same type can be a member. The Scheduler immediately expands the included group name into its list of members. An error is returned if any of the members do not exist. A member that is already in the group is skipped, and no error is generated. The keyword LOCAL can be included as a member for database destination or external destination groups. See the "CREATE_GROUP Procedure" for information about this keyword.

Usage Notes

The following users may add members to a group:

- The group owner
- A user that has been granted the ALTER object privilege on the group
- A user with the CREATE ANY JOB system privilege

You must have the MANAGE SCHEDULER privilege to add a member to a group of type WINDOW.

> **See Also:**
>
> "CREATE_GROUP Procedure"

169.6.3 ADD_JOB_EMAIL_NOTIFICATION Procedure

This procedure adds e-mail notifications for a job. E-mails are then sent to the specified list of recipients whenever any of the specified job state events is raised.

Syntax

```
DBMS_SCHEDULER.ADD_JOB_EMAIL_NOTIFICATION (
    job_name             IN VARCHAR2,
    recipients           IN VARCHAR2,
    sender               IN VARCHAR2 DEFAULT NULL,
    subject              IN VARCHAR2 DEFAULT DBMS_SCHEDULER.DEFAULT_NOTIFICATION_SUBJECT,
    body                 IN VARCHAR2 DEFAULT DBMS_SCHEDULER.DEFAULT_NOTIFICATION_BODY,
    events               IN VARCHAR2 DEFAULT 'JOB_FAILED,JOB_BROKEN,JOB_SCH_LIM_REACHED,
                            JOB_CHAIN_STALLED,JOB_OVER_MAX_DUR',
    filter_condition     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-14 ADD_JOB_EMAIL_NOTIFICATION Procedure Parameters

Parameter	Description
job_name	Name of the job that e-mail notifications are added for. Cannot be NULL.
recipients	Comma-separated list of e-mail addresses to send notifications to. E-mail notifications for all listed events are sent to all recipients. Cannot be NULL.
sender	e-mail address to use as the sender address (the From: address) in the e-mail header. If NULL or omitted, the e-mail address specified in the Scheduler attribute email_sender is used. See *Oracle Database Administrator's Guide* for more information on this Scheduler attribute.
subject	The subject to use in the e-mail header. Table 169-15 describes the variables that you can include within this parameter. The Scheduler assigns values to these variables before sending the notification. If subject is omitted, the default subject is used. The default subject is the following text, where text enclosed in the '%' character represents a variable: 'Oracle Scheduler Job Notification - %job_owner%.%job_name%.%job_subname% %event_type%'

Table 169-14 (Cont.) ADD_JOB_EMAIL_NOTIFICATION Procedure Parameters

Parameter	Description
body	The body of the e-mail message. Table 169-15 describes the variables that you can include within this parameter. The Scheduler assigns values to these variables before sending the notification. If body is omitted, the default body is used. The default body is the following text, where text enclosed in the '%' character represents a variable: 'Job: %job_owner%.%job_name%.%job_subname% Event: %event_type% Date: %event_timestamp% Log id: %log_id% Job class: %job_class_name% Run count: %run_count% Failure count: %failure_count% Retry count: %retry_count% Error code: %error_code% Error message: %error_message%'
events	Comma-separate list of job state events to send e-mail notifications for. Cannot be NULL. A notification is sent to all recipients if any of the listed events is raised. Table 169-85 lists the valid events for this parameter. If events is omitted, notifications are sent for the following default events: JOB_FAILED,JOB_BROKEN,JOB_SCH_LIM_REACHED,JOB_CHAIN_STALLED,JOB_OVER_MAX_DUR
filter_condition	Used to filter events to send e-mail notifications for. If NULL, all occurrences of the specified events cause e-mail notifications to be sent. filter_condition must be a boolean SQL WHERE clause that may refer to the :event bind variable. This bind variable is automatically bound to an object of type SCHEDULER$_EVENT_INFO that represents the raised event. For example, to send an e-mail notification only when the error number in an event is 600 or 700, use the following filter_condition: :event.error_code=600 or :event.error_code=700 See "SCHEDULER$_EVENT_INFO Object Type".

Table 169-15 lists the variables that you can use in the subject and body arguments.

Table 169-15 Variables Used in the SUBJECT and BODY Parameters

Variable	Comment
%job_owner%	Schema in which job was created
%job_name%	Name of the job that e-mail notifications are added for
%job_subname%	Present for event-based jobs with the parallel_instances attribute set and for chain steps
%event_type%	Valid values are listed in Table 169-85
%event_timestamp%	Time at which the event occurred
%log_id%	Refers to the LOG_ID column in views *_SCHEDULER_JOB_LOG and *_SCHEDULER_JOB_RUN_DETAILS
%error_code%	Number of the error code.

Table 169-15 (Cont.) Variables Used in the SUBJECT and BODY Parameters

Variable	Comment
%error_message%	The text of the error message
%run_count%	Run count for the job when the event was raised
%failure_count%	Failure count for the job when the event was raised
%retry_count%	Retry count for the job when the event was raised

Usage Notes

You can call ADD_JOB_EMAIL_NOTIFICATION once for each different set of notifications that you want to configure for a particular job. For example, you may want to send notifications for the JOB_FAILED, JOB_BROKEN, JOB_SCH_LIM_REACHED, and JOB_CHAIN_STALLED events to the principle DBA and all senior DBAs, but send a notification for the JOB_OVER_MAX_DUR event only to the principle DBA.

This procedure succeeds only if the Scheduler attribute email_server is set to a valid SMTP server. See *Oracle Database Administrator's Guide* for more information.

To call this procedure, you must be the job owner or have the CREATE ANY JOB system privilege or have the ALTER object privilege on the job.

169.6.4 ADD_TO_INCOMPATIBILITY Procedure

This procedure adds jobs or programs to an existing incompatibility definition.

Syntax

```
DBMS_SCHEDULER.ADD_TO_INCOMPATIBILITY (
   incompatibility_name    IN VARCHAR2,
   object_name             IN VARCHAR2);
```

Parameters

Table 169-16 ADD_TO_INCOMPATIBILITY Procedure Parameters

Parameter	Description
incompatibility_name	The name of the incompatibility definition.
object_name	One or more (comma-separated) programs or jobs

Usage Notes

This procedure does not raise an error if any specified objects already exist in the incompatibility definition.

> **See Also:**
>
> Using Incompatibility Definitions in *Oracle Database Administrator's Guide*

169.6.5 ALTER_CHAIN Procedure

This procedure alters an attribute of the specified steps of a chain. This affects all future runs of the specified steps, both in the currently running chain job and in future runs of the same chain job or other chain jobs that point to the chain.

Syntax

Alters the value of a boolean attribute of one or more steps:

```
DBMS_SCHEDULER.ALTER_CHAIN (
   chain_name              IN VARCHAR2,
   step_name               IN VARCHAR2,
   attribute               IN VARCHAR2,
   value                   IN BOOLEAN);
```

Alters the value of a character attribute of one or more steps:

```
DBMS_SCHEDULER.ALTER_CHAIN (
   chain_name              IN VARCHAR2,
   step_name               IN VARCHAR2,
   attribute               IN VARCHAR2,
   char_value              IN VARCHAR2);
```

Parameters

Table 169-17 ALTER_CHAIN Procedure Parameters

Parameter	Description
chain_name	The name of the chain to alter
step_name	The name of the step or a comma-separated list of steps to alter. This cannot be NULL.

Table 169-17 (Cont.) ALTER_CHAIN Procedure Parameters

Parameter	Description
attribute	The attribute of the steps to change. Must be one of the following: • 'PAUSE' If set to TRUE for a step, after the step has run, its state changes to PAUSED (and the completed attribute remains FALSE). If PAUSE is reset to FALSE for a paused chain step (using ALTER_RUNNING_CHAIN), the state is set to its completion state (SUCCEEDED, FAILED, or STOPPED) and the completed attribute is set to TRUE. Setting PAUSE has no effect on steps that have already run. This allows execution of a chain to be suspended after the execution of certain steps. • 'PAUSED_BEFORE' If set to TRUE for a step and if any of the rule conditions that start the step are true, then its state changes to PAUSED and the step does not run. If PAUSE_BEFORE is reset to FALSE for a chain step that has paused before starting (using ALTER_RUNNING_CHAIN), then the step starts running if any of the rule conditions that start the step are true. Setting PAUSE_BEFORE has no effect on steps that are running or have already run. This allows execution of a chain to be suspended before the execution of certain steps. • 'SKIP' If set to TRUE for a step, when the step condition is met, instead of being run, the step is treated as if it has immediately succeeded. Setting SKIP to TRUE has no effect for a step that is running, scheduled to run after a delay, or has already run. If SKIP is set TRUE for a step that PAUSE is also set for, when the step condition is met, the step immediately changes to state PAUSED. • 'RESTART_ON_FAILURE' If set to TRUE for a step and the step fails due to an application error, then the step is retried using the normal Scheduler retry mechanism (after 1 second, after 10 seconds, after 100 seconds, and so on, up to a maximum of 6 times). If all 6 retries fail (after about 30 hours), then the chain step is marked FAILED. If set to FALSE (the default), a failed chain step is immediately marked FAILED. • 'RESTART_ON_RECOVERY' If set to TRUE for a step and the step is stopped by a database shutdown, then the step is restarted when the database is recovered. If set to FALSE, and the step is stopped by a database shutdown, then the step is marked as stopped when the database is recovered and the chain continues. • 'DESTINATION_NAME' The name of an existing database destination or external destination. You can view external destination names in the view ALL_SCHEDULER_EXTERNAL_DESTS, and database destination names in the views *_SCHEDULER_DB_DESTS. You cannot specify a destination group for this attribute. This parameter is NULL by default. • 'CREDENTIAL_NAME' The credential to use when running this step. NULL by default.
value	The value to set for the attribute (for a boolean attribute).
char_value	The value to set for the attribute (for a character attribute).

Usage Notes

Altering a chain requires `ALTER` privileges on the chain either by being the owner of the chain, or by having the `ALTER` object privilege on the chain or by having the `CREATE ANY JOB` system privilege.

169.6.6 ALTER_RUNNING_CHAIN Procedure

This procedure alters an attribute of the specified steps of a chain. This affects only steps of the instance of the chain for the specified running chain job.

Syntax

```
DBMS_SCHEDULER.ALTER_RUNNING_CHAIN (
   job_name                IN VARCHAR2,
   step_name               IN VARCHAR2,
   attribute               IN VARCHAR2,
   value                   IN {BOOLEAN|VARCHAR2});
```

Parameters

Table 169-18 ALTER_RUNNING_CHAIN Procedure Parameters

Parameter	Description
`job_name`	The name of the job that is running the chain
`step_name`	The name of the step or a comma-separated list of steps to alter. If this is set to `NULL` and attribute is `PAUSE` or `SKIP`, then all steps of the running chain are altered.
attribute	The attribute of the steps to change. Valid values are: • 'PAUSE' If the `PAUSE` attribute is set `TRUE` for a step, then after the step runs, its state changes to `PAUSED` (and the `completed` attribute remains false). If `PAUSE` is reset to `FALSE` for a paused chain step (using `ALTER_RUNNING_CHAIN`), the state is set to completion (`SUCCEEDED`, `FAILED`, or `STOPPED`) and the `completed` attribute is set to `TRUE`. Setting `PAUSE` has no effect on steps that have already run. This allows execution of a chain to be suspended after the execution of certain steps. If `step_name` is set to `NULL`, `PAUSE` is set to `TRUE` for all steps of this running chain. • 'PAUSE_BEFORE' If set to `TRUE` for a step that has not yet run and if any of the rule conditions that start the step are true, then its state changes to `PAUSED` and the step does not run. If `PAUSE_BEFORE` is reset to `FALSE` for a chain step that has paused before starting, then the step starts running if any of the rule conditions that start the step are true. Setting `PAUSE_BEFORE` has no effect on steps that are running or have already run. This allows execution of a chain to be suspended before the execution of certain steps. If `step_name` is set to `NULL`, then `PAUSE_BEFORE` is set to the specified value for all steps of this running chain.

Table 169-18 (Cont.) ALTER_RUNNING_CHAIN Procedure Parameters

Parameter	Description
attribute CONTINUED	• 'SKIP' If the SKIP attribute is set to TRUE for a step, when the step condition is met, instead of being run, the step is treated as if it has immediately succeeded. Setting SKIP to TRUE has no effect for a step that is running, scheduled to run after a delay, or has already run. If step_name is set to NULL, SKIP is set TRUE for all steps of this running chain. If SKIP is set TRUE for a step that PAUSE is also set for, when the step condition is met the step immediately changes to state PAUSED. • 'RESTART_ON_FAILURE' If set to TRUE for a step and the step fails due to an application error, then the step is retried using the normal Scheduler retry mechanism (after 1 second, after 10 seconds, after 100 seconds, and so on, up to a maximum of 6 times). If all 6 retries fail (after about 30 hours), then the chain step is marked FAILED. If set to FALSE (the default), a failed chain step is immediately marked FAILED. • 'RESTART_ON_RECOVERY' If the RESTART_ON_RECOVERY attribute is set to TRUE for a step, then if the step is stopped by a database shutdown, it is restarted when the database is recovered. If set to FALSE, then if the step is stopped by a database shutdown, the step is marked as stopped when the database is recovered and the chain continues. • 'STATE' This changes the state of the steps. The state can only be changed if the step is not running. The state can only be changed to one of the following: 'NOT_STARTED', 'SUCCEEDED', 'FAILED *error_code*' If the state is being changed to FAILED, an error code must be included (this must be a positive integer).
value	The value to set for the attribute. Valid values are: TRUE, FALSE, 'NOT_STARTED', 'SUCCEEDED', or 'FAILED *error_code*'

Usage Notes

Altering a running chain requires you to have alter privileges on the job that is running (either as the owner, or as a user with ALTER privileges on the job or the CREATE ANY JOB system privilege).

When trying to update a step defined with a nested chain, it is necessary to specify the job_name as <SCHEMA>.<JOB_NAME>.<STEP_NAME_IN_TOP_LEVEL_CHAIN> to be able to make reference to the steps inside the subchain.

169.6.7 CLOSE_WINDOW Procedure

This procedure closes an open window prematurely. A closed window means that it is no longer in effect. When a window is closed, the Scheduler switches the resource plan to the one that is in effect outside the window, or in the case of overlapping windows, to another window.

Syntax

```
DBMS_SCHEDULER.CLOSE_WINDOW (
   window_name             IN VARCHAR2);
```

Parameters

Table 169-19 CLOSE_WINDOW Procedure Parameters

Parameter	Description
window_name	The name of the window

Usage Notes

If you try to close a window that does not exist or is not open, an error is generated.

A job that is running does not stop when the window it is running in closes, unless the attribute stop_on_window_close is set to TRUE for the job. However, the resources allocated to the job can change if the resource plan changes.

When a running job has a group of type WINDOW as its schedule, the job is not stopped when its window is closed if another window in the same window group becomes active. This is the case even if the job has the attribute stop_on_window_close set to TRUE.

Closing a window requires the MANAGE SCHEDULER privilege.

169.6.8 COPY_JOB Procedure

This procedure copies all attributes of an existing job to a new job. The new job is created disabled, while the state of the existing job is unaltered.

Syntax

```
DBMS_SCHEDULER.COPY_JOB (
   old_job                 IN VARCHAR2,
   new_job                 IN VARCHAR2);
```

Parameters

Table 169-20 COPY_JOB Procedure Parameters

Parameter	Description
old_job	The name of the existing job
new_job	The name of the new job

Usage Notes

To copy a job, you must have privileges to create a job in the schema of the new job (the `CREATE JOB` system privilege if it is in your own schema, otherwise, the `CREATE ANY JOB` system privilege). If the old job is not in the your own schema, then you must also have `ALTER` privileges on the old job or the `CREATE ANY JOB` system privilege.

169.6.9 CREATE_CHAIN Procedure

This procedure creates a new chain. The chain name can be optionally qualified with a schema name (for example, `myschema.myname`).

A chain is always created as disabled and must be enabled with the ENABLE Procedure before it can be used.

Syntax

```
DBMS_SCHEDULER.CREATE_CHAIN (
   chain_name              IN VARCHAR2,
   rule_set_name           IN VARCHAR2 DEFAULT NULL,
   evaluation_interval     IN INTERVAL DAY TO SECOND DEFAULT NULL,
   comments                IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-21 CREATE_CHAIN Procedure Parameters

Parameter	Description
`chain_name`	The name to assign to the new chain, which can optionally be qualified with a schema. This must be unique in the SQL namespace, therefore, there cannot already be a table or other object with this name and schema.
`rule_set_name`	In the normal case, no rule set should be passed in. The Scheduler automatically creates a rule set and associated empty evaluation context. You then use `DEFINE_CHAIN_RULE` to add rules and `DROP_CHAIN_RULE` to remove them. Advanced users can create a rule set that describes their chain dependencies and pass it in here. This allows greater flexibility in defining rules. For example, conditions can refer to external variables, and tables can be exposed through the evaluation context. If you pass in a rule set, you must ensure that it is in the format of a chain rule set. (For example, all steps must be listed as variables in the evaluation context). If no rule set is passed in, the rule set created is of the form `SCHED_RULESET${N}` and the evaluation context created is of the form `SCHED_EVCTX${N}`
`evaluation_interval`	If this is `NULL`, reevaluation of the rules of a running chain are performed only when the job starts and when a step completes. A non-`NULL` value causes rule evaluations to also occur periodically at the specified interval. Because evaluation may be CPU-intensive, this should be conservatively set to the highest possible value or left at `NULL` if possible. `evaluation_interval` cannot be less than a minute or greater than a day.
`comments`	An optional comment describing the purpose of the chain

Usage Notes

To create a chain in your own schema, you must have the `CREATE JOB` system privilege. To create a chain in a different schema you must have the `CREATE ANY JOB` system privilege. If you do not provide a `rule_set_name`, a rule set and evaluation context is created in the schema

that the chain is being created in, so you must have the privileges required to create these objects. See the `DBMS_RULE_ADM.CREATE_RULE_SET` and `DBMS_RULE_ADM.CREATE_EVALUATION_CONTEXT` procedures for more information.

169.6.10 CREATE_CREDENTIAL Procedure

This deprecated procedure creates a stored username/password pair. Credentials are assigned to jobs so that they can authenticate with a local or remote host operating system or a remote Oracle database.

> **Note:**
>
> This procedure is deprecated with Oracle Database 12*c* Release 1 (12.1). While the procedure remains available in this package, for reasons of backward compatibility, Oracle recommends using the alternative enhanced functionality provided in the DBMS_CREDENTIAL package, specifically the CREATE_CREDENTIAL Procedure.

Syntax

```
DBMS_SCHEDULER.CREATE_CREDENTIAL (
   credential_name          IN VARCHAR2,
   username                 IN VARCHAR2,
   password                 IN VARCHAR2,
   database_role            IN VARCHAR2 DEFAULT NULL,
   windows_domain           IN VARCHAR2 DEFAULT NULL,
   comments                 IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-22 CREATE_CREDENTIAL Procedure Parameters

Parameter	Description
`credential_name`	The name to assign to the credential. It can optionally be prefixed with a schema name. It cannot be set to `NULL`. It is converted to uppercase unless enclosed in double quotation marks.
`username`	The user name for logging into to the host operating system or remote Oracle database. This cannot be set to `NULL` and is case-sensitive. It cannot contain double quotes or spaces. Maximum length is 64.
`password`	The password for the user name. This cannot be set to `NULL` and is case sensitive. The password is stored obfuscated and is not displayed in the Scheduler dictionary views. Maximum length is 128.
`database_role`	The value of the `database_role` attribute is used as the system privilege for logging into a remote database to run a remote database job. Valid values are: `SYSDBA` and `SYSOPER`
`windows_domain`	For a Windows remote executable target, this is the domain that the specified user belongs to. The domain is converted to uppercase automatically. Maximum length is 64.
`comments`	A text string that can be used to describe the credential. Scheduler does not use this parameter. Maximum length is 240.

Usage Notes

Credentials reside in a particular schema and can be created by any user with the `CREATE JOB` system privilege. To create a credential in a schema other than your own, you must have the `CREATE ANY JOB` privilege.

169.6.11 CREATE_DATABASE_DESTINATION Procedure

This procedure creates a database destination. A database destination represents an Oracle database on which remote database jobs run.

The host that the remote database resides on must have a running Scheduler agent that is registered with the database that this procedure is called from.

Syntax

```
DBMS_SCHEDULER.CREATE_DATABASE_DESTINATION (
   destination_name        IN VARCHAR2,
   agent                   IN VARCHAR2,
   tns_name                IN VARCHAR2,
   comments                IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-23 CREATE_DATABASE_DESTINATION Procedure Parameters

Parameter	Description
`destination_name`	The name to assign to the database destination. It can optionally be prefixed with a schema name. Cannot be `NULL`. It is converted to uppercase unless enclosed in double quotation marks.
`agent`	The external destination name of the Scheduler agent to connect. Equivalent to an agent name. The external destination must already exist. The external destination representing an agent is created automatically on a database instance when the agent registers with that instance. An agent's name is specified in its agent configuration file. If it is not specified, it defaults to the first part (before the first period) of the name of the host it resides on.
`tns_name`	An Oracle Net connect identifier that is resolved to the Oracle database instance being connected to. The exact syntax depends on the Oracle Net configuration.The connect identifier can be a complete Oracle Net connect descriptor (network address and database service name) or a *net service name*, which is an alias for a connect descriptor. The alias must be resolved in the tnsnames.ora file on the local computer. The maximum size for `tns_name` is 2000 characters. If `tns_name` is `NULL`, the agent connects to the default Oracle database on its host. You specify the default database by assigning values to the `ORACLE_HOME` and `ORACLE_SID` parameters in the agent configuration file, schagent.conf, located in the agent home directory. See *Oracle Database Net Services Administrator's Guide* for more information on connect identifiers.
`comments`	A text string that describes the database destination. Scheduler does not use this argument.

Usage Notes

Database destinations reside in a particular schema and can be created by any user with the `CREATE JOB` system privilege. To create a database destination in a schema other than your own, you must have the `CREATE ANY JOB` privilege.

169.6.12 CREATE_EVENT_SCHEDULE Procedure

This procedure creates an event schedule, which is used to start a job when a particular event is raised.

Syntax

```
DBMS_SCHEDULER.CREATE_EVENT_SCHEDULE (
   schedule_name           IN VARCHAR2,
   start_date              IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   event_condition         IN VARCHAR2 DEFAULT NULL,
   queue_spec              IN VARCHAR2,
   end_date                IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   comments                IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-24 CREATE_EVENT_SCHEDULE Parameters

Parameter	Description
`schedule_name`	The name to assign to the schedule. The name must be unique in the SQL namespace. For example, a schedule cannot have the same name as a table in a schema. If no name is specified, then an error occurs.
`start_date`	This attribute specifies the date and time that this schedule becomes valid. Occurrences of the event before this date and time are ignored in the context of this schedule.
`event_condition`	This is a conditional expression based on the columns of the event source queue table. The expression must have the syntax of an Advanced Queuing rule. Accordingly, you can include user data properties in the expression, provided that the message payload is an object type, and that you prefix object attributes in the expression with `tab.user_data`. For more information on rules, see the `DBMS_AQADM.ADD_SUBSCRIBER` procedure.
`queue_spec`	This argument specifies either a file watcher name or the queue into which events that start this particular job are enqueued (the **source queue**). If the source queue is a secure queue, the `queue_spec` argument is a string containing a pair of values of the form *queue_name*, *agent name*. For non-secure queues, only the queue name need be provided. If a fully qualified queue name is not provided, the queue is assumed to be in the job owner's schema. In the case of secure queues, the agent name provided should belong to a valid agent that is currently subscribed to the queue.
`end_date`	The date and time after which jobs do not run and windows do not open. An event schedule that has no `end_date` is valid forever. `end_date` must be after the `start_date`. If it is not, then an error is generated when the schedule is created.
`comments`	This attribute specifies an optional comment about the schedule. By default, this attribute is `NULL`.

Usage Notes

You must have the `CREATE JOB` privilege to create a schedule in your own schema or the `CREATE ANY JOB` privilege to create a schedule in someone else's schema by specifying `schema.schedule_name`. Once a schedule has been created, it can be used by other users. The schedule is created with access to `PUBLIC`. Therefore, there is no need to explicitly grant access to the schedule.

> **See Also:**
>
> "CREATE_FILE_WATCHER Procedure"

169.6.13 CREATE_FILE_WATCHER Procedure

This procedure creates a file watcher, which is a Scheduler object that defines the location, name, and other properties of a file whose arrival on a system causes the Scheduler to start a job. After you create a file watcher, you reference it in an event-based job or event schedule.

Syntax

```
DBMS_SCHEDULER.CREATE_FILE_WATCHER (
   file_watcher_name            IN VARCHAR2,
   directory_path               IN VARCHAR2,
   file_name                    IN VARCHAR2,
   credential_name              IN VARCHAR2,
   destination                  IN VARCHAR2  DEFAULT NULL,
   min_file_size                IN PLS_INTEGER DEFAULT 0,
   steady_state_duration        IN INTERVAL DAY TO SECOND DEFAULT NULL,
   comments                     IN VARCHAR2 DEFAULT NULL,
   enabled                      IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 169-25 CREATE_FILE_WATCHER Parameters

Parameter	Description
`file_watcher_name`	The name to assign to the file watcher. The name must be unique in the SQL namespace. For example, a file watcher cannot have the same name as a table in a schema. This can optionally be prefixed with a schema name. Cannot be `NULL`.
`directory_path`	Directory in which the file is expected to arrive. The single wildcard '?' at the beginning of the path denotes the Oracle home path. For example, '?/rdbms/log' denotes the rdbms/log subdirectory of the Oracle home directory.
`file_name`	Name of the file to look for. Two wildcards are permitted anywhere in the file name: '?' denotes any single character, and '*' denotes zero or more characters. This attribute cannot be `NULL`.
`credential_name`	Name of a valid credential object. The file watcher uses the credential to authenticate itself with the host operating system to access the watched-for file. The file watcher owner must have `EXECUTE` privileges on the credential. Cannot be `NULL`.

Table 169-25 (Cont.) CREATE_FILE_WATCHER Parameters

Parameter	Description
destination	Name of an external destination. You create an external destination by registering a remote Scheduler agent with the database. See the view ALL_SCHEDULER_EXTERNAL_DESTS for valid external destination names. If this parameter is NULL, the file watcher is created on the local host.
min_file_size	Minimum size in bytes that the file must be before the file watcher considers the file found. Default is 0.
steady_state_duration	Minimum time interval that the file must remain unchanged before the file watcher considers the file found. Cannot exceed one hour. If NULL, an internal value is used. The minimum value is 10 seconds. Oracle recommends similar steady_state_duration values for all file watchers for efficient file watcher job operation. Also, the repeat interval of the file watcher schedule must be equal or greater than the steady_state_duration value.
comments	Optional comment.
enabled	If TRUE (the default), the file watcher is enabled.

Usage Notes

You must have the CREATE JOB system privilege to create a file watcher in your own schema. You require the CREATE ANY JOB system privilege to create a file watcher in a schema different from your own (except the SYS schema, which is disallowed).

169.6.14 CREATE_GROUP Procedure

This procedure creates a group. Groups contain members, which you can specify when you create the group or at a later time. There are three types of groups: window groups, database destination groups, and external destination groups.

You can use a group name in other DBMS_SCHEDULER package procedures to specify a list of objects. For example, to specify multiple destinations for a remote database job, you provide a group name for the DESTINATION_NAME parameter of the job.

Syntax

```
DBMS_SCHEDULER.CREATE_GROUP (
   group_name           IN VARCHAR2,
   group_type           IN VARCHAR2,
   member               IN VARCHAR2 DEFAULT NULL,
   comments             IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-26 CREATE_GROUP Procedure Parameters

Parameter	Description
group_name	The name to assign to the group. It can optionally be prefixed with a schema name. It cannot be NULL. It is converted to uppercase unless enclosed in double quotation marks.

Table 169-26 (Cont.) CREATE_GROUP Procedure Parameters

Parameter	Description
group_type	The type of members in the group. All members must be of the same type. Possible types are: • 'DB_DEST' Database destination: Members are database destinations, for running remote database jobs. • 'EXTERNAL_DEST External destination: Members are external destinations, for running remote external jobs. • 'WINDOW' Members are Scheduler windows. You must have the MANAGE SCHEDULER privilege to create a group of this type. Members in database destination and external destination groups have the following format: [[*schema.*]*credential*@][*schema.*]*destination* where: • *credential* is the name of an existing credential. • *destination* is the name of an existing database destination or external destination. The credential portion of a destination member is optional. If omitted, the job using this destination member uses its default credential. Members in window groups are window names. Because all Scheduler windows reside in the SYS schema, you do not specify a schema name for windows.
member	Optional comma-separated list of group members. The default is NULL. If NULL, use the ADD_GROUP_MEMBER procedure to add members. You can also use ADD_GROUP_MEMBER to add additional members at a later time. The keyword LOCAL can be used as a member in database destination groups and external destination groups. • In database destination groups, LOCAL represents the source database on which the job is created. It cannot be preceded with a credential. • In external destination groups, LOCAL represents the host on which the source database resides. It can be optionally preceded with a credential name. If no credential is provided, jobs that use this group as their destination must have a default credential.
comments	A text string that describes the group. Scheduler does not use this argument.

Usage Notes

Groups reside in a particular schema and can be created by any user with the CREATE JOB system privilege. To create a group in a schema other than your own, you must have the CREATE ANY JOB privilege. The group name must be unique among all Scheduler objects.

You can grant the SELECT or READ privilege on a group so that other users can reference the group when creating jobs or schedules. To enable other users to modify a group, you can grant the ALTER privilege on the group.

Each group member must be unique within the group. For destination groups, the credential/destination name pairs must be unique within the group. An error is generated if any of the

group members do not exist. For destination groups, both the credential and destination portions of a member must exist.

Another group of the same type can be a group member. The Scheduler immediately expands the included group name into its list of members.

Groups are created enabled, but you can disable them.

Example

The following PL/SQL block creates a group named `production_dest1`, whose members are database destinations for a collection of production databases.

```
BEGIN
  DBMS_SCHEDULER.CREATE_GROUP(
    GROUP_NAME    => 'production_dest1',
    GROUP_TYPE    => 'DB_DEST',
    MEMBER        => 'LOCAL, oracle_cred@prodhost1, prodhost2',
    COMMENTS      => 'All sector1 production machines');
END;
```

169.6.15 CREATE_INCOMPATIBILITY Procedure

This procedure creates an incompatibility definition.

Syntax

```
DBMS_SCHEDULER.CREATE_INCOMPATIBILITY (
   incompatibility_name    IN VARCHAR2,
   object_name             IN VARCHAR2,
   constraint_level        IN VARCHAR2 DEFAULT 'JOB_LEVEL',
   enabled                 IN BOOLEAN DEFAULT TRUE,
   comments                IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-27 CREATE_INCOMPATIBILITY Procedure Parameters

Parameter	Description
`incompatibility_name`	The name of the incompatibility definition.
`object_name`	One or more (comma-separated) programs or jobs.
`constraint_level`	One or more (comma-separated) programs or jobs.
`enabled`	Specifies whether the constraint is initially enabled (`true`) or not enabled (`false`).
`comments`	Optional descriptive comment.

Usage Notes

If `object_name` contains multiple (comma-separated) values, they must be either all programs or all jobs that are incompatible with each other (that is, they cannot be run at the same time). For jobs, the list must consist of two or more jobs, and `constraint_level` must be 'JOB_LEVEL'. For programs, `constraint_level` can be either `'JOB_LEVEL'` or `'PROGRAM_LEVEL'`. When set to the default value `'JOB_LEVEL'`, only a single job that is based on the program (or programs) mentioned in `object_name` can run at the same time. When

`constraint_level` is set to `'PROGRAM_LEVEL'`, the programs are incompatible, but the jobs based on the same program are not incompatible.

For example, if the value of `object_name` is `'P1,P2,P3'` and `constraint_level` is `'PROGRAM_LEVEL'`, many jobs based on P1 can be running at the same time, but if any P1 based job is running, none based on P2 or P3 can be running. Or, similarly, many jobs based on P3 can be running at the same time, but none based on P1 or P2. If `constraint_level` is set to `'JOB_LEVEL'`, then only a single job out of all the jobs based on programs P1, P2 and P3 can be running at a time.

> **See Also:**
>
> Using Incompatibility Definitionsin *Oracle Database Administrator's Guide*

169.6.16 CREATE_JOB Procedure

This procedure creates a single job.

If you create the job as enabled by setting the `enabled` attribute to `TRUE`, the Scheduler automatically runs the job according to its schedule. If you create the job disabled, the job does not run until you enable it with the SET_ATTRIBUTE Procedure.

The procedure is overloaded. The different functionality of each form of syntax is presented along with the syntax declaration.

Syntax

Creates a job in a single call without using an existing program or schedule:

```
DBMS_SCHEDULER.CREATE_JOB (
   job_name             IN VARCHAR2,
   job_type             IN VARCHAR2,
   job_action           IN VARCHAR2,
   number_of_arguments  IN PLS_INTEGER              DEFAULT 0,
   start_date           IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   repeat_interval      IN VARCHAR2                 DEFAULT NULL,
   end_date             IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   job_class            IN VARCHAR2                 DEFAULT 'DEFAULT_JOB_CLASS',
   enabled              IN BOOLEAN                  DEFAULT FALSE,
   auto_drop            IN BOOLEAN                  DEFAULT TRUE,
   comments             IN VARCHAR2                 DEFAULT NULL,
   credential_name      IN VARCHAR2                 DEFAULT NULL,
   destination_name     IN VARCHAR2                 DEFAULT NULL);
```

Creates a job using a named schedule object and a named program object:

```
DBMS_SCHEDULER.CREATE_JOB (
   job_name               IN VARCHAR2,
   program_name           IN VARCHAR2,
   schedule_name          IN VARCHAR2,
   job_class              IN VARCHAR2              DEFAULT 'DEFAULT_JOB_CLASS',
   enabled                IN BOOLEAN               DEFAULT FALSE,
   auto_drop              IN BOOLEAN               DEFAULT TRUE,
   comments               IN VARCHAR2              DEFAULT NULL,
   job_style              IN VARCHAR2              DEFAULT 'REGULAR',
   credential_name        IN VARCHAR2              DEFAULT NULL,
   destination_name       IN VARCHAR2              DEFAULT NULL);
```

Creates a job using a named program object and an inlined schedule:

```
DBMS_SCHEDULER.CREATE_JOB (
   job_name             IN VARCHAR2,
   program_name         IN VARCHAR2,
   start_date           IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   repeat_interval      IN VARCHAR2                 DEFAULT NULL,
   end_date             IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   job_class            IN VARCHAR2                 DEFAULT 'DEFAULT_JOB_CLASS',
   enabled              IN BOOLEAN                  DEFAULT FALSE,
   auto_drop            IN BOOLEAN                  DEFAULT TRUE,
   comments             IN VARCHAR2                 DEFAULT NULL,
   job_style            IN VARCHAR2                 DEFAULT 'REGULAR',
   credential_name      IN VARCHAR2                 DEFAULT NULL,
   destination_name     IN VARCHAR2                 DEFAULT NULL);
```

Creates a job using a named schedule object and an inlined program:

```
DBMS_SCHEDULER.CREATE_JOB (
   job_name                IN VARCHAR2,
   schedule_name           IN VARCHAR2,
   job_type                IN VARCHAR2,
   job_action              IN VARCHAR2,
   number_of_arguments     IN PLS_INTEGER       DEFAULT 0,
   job_class               IN VARCHAR2          DEFAULT 'DEFAULT_JOB_CLASS',
   enabled                 IN BOOLEAN           DEFAULT FALSE,
   auto_drop               IN BOOLEAN           DEFAULT TRUE,
   comments                IN VARCHAR2          DEFAULT NULL,
   credential_name         IN VARCHAR2          DEFAULT NULL,
   destination_name        IN VARCHAR2          DEFAULT NULL);
```

Creates a job using an inlined program and an event:

```
DBMS_SCHEDULER.CREATE_JOB (
   job_name                IN VARCHAR2,
   job_type                IN VARCHAR2,
   job_action              IN VARCHAR2,
   number_of_arguments     IN PLS_INTEGER       DEFAULT 0,
   start_date              IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   event_condition         IN VARCHAR2          DEFAULT NULL,
   queue_spec              IN VARCHAR2,
   end_date                IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   job_class               IN VARCHAR2          DEFAULT 'DEFAULT_JOB_CLASS',
   enabled                 IN BOOLEAN           DEFAULT FALSE,
   auto_drop               IN BOOLEAN           DEFAULT TRUE,
   comments                IN VARCHAR2          DEFAULT NULL,
   credential_name         IN VARCHAR2          DEFAULT NULL,
   destination_name        IN VARCHAR2          DEFAULT NULL);
```

Creates a job using a named program object and an event:

```
DBMS_SCHEDULER.CREATE_JOB (
   job_name                IN VARCHAR2,
   program_name            IN VARCHAR2,
   start_date              IN TIMESTAMP WITH TIME ZONE,
   event_condition         IN VARCHAR2,
   queue_spec              IN VARCHAR2,
   end_date                IN TIMESTAMP WITH TIME ZONE,
   job_class               IN VARCHAR2          DEFAULT 'DEFAULT_JOB_CLASS',
   enabled                 IN BOOLEAN           DEFAULT FALSE,
   auto_drop               IN BOOLEAN           DEFAULT TRUE,
   comments                IN VARCHAR2          DEFAULT NULL,
   job_style               IN VARCHAR2          DEFAULT 'REGULAR',
```

```
   credential_name         IN VARCHAR2         DEFAULT NULL,
   destination_name        IN VARCHAR2         DEFAULT NULL);
```

Parameters

Table 169-28 *CREATE_JOB Procedure Parameters*

Parameter	Description
job_name	The name to assign to the job. The name must be unique in the SQL namespace. For example, a job cannot have the same name as a table in a schema. If the job being created will reside in another schema, it must be qualified with the schema name. If job_name is not specified, an error is generated. If you want to have a name generated by the Scheduler, you can use the GENERATE_JOB_NAME procedure to generate a name and then use the output in the CREATE_JOB procedure. The GENERATE_JOB_NAME procedure generates a number from a sequence, which is the job name. You can prefix the number with a string. The job name will then be the string with the number from the sequence appended to it. See "GENERATE_JOB_NAME Function" for more information.

Table 169-28 (Cont.) *CREATE_JOB Procedure Parameters*

Parameter	Description
job_type	This attribute specifies the type of job that you are creating. If it is not specified, an error is generated. See job_action in the next row for related information. The supported values are: • 'PLSQL_BLOCK' This specifies that the job is an anonymous PL/SQL block. Job or program arguments are not supported when the job or program type is PLSQL_BLOCK. In this case, the number of arguments must be 0. • 'STORED_PROCEDURE' This specifies that the job is a PL/SQL or Java stored procedure, or an external C subprogram. Only procedures, not functions with return values, are supported. • 'EXECUTABLE' This specifies that the job is going to be run outside the database using an external executable. External jobs are anything that can be executed from the command line of the operating system. Anydata arguments are not supported with a job or program type of EXECUTABLE. The job owner must have the CREATE EXTERNAL JOB system privilege before the job can be enabled or run. • 'CHAIN' This specifies that the job is a chain. Arguments are not supported for a chain, so number_of_arguments must be 0. • 'EXTERNAL_SCRIPT' This specifies that the job is an external script that uses the command shell of the computer running the job. For Windows this is cmd.exe and for UNIX based systems the sh shell, unless a different interpreter is specified by prefixing the first line of the script with #!. • 'SQL_SCRIPT' This specifies that the job is a SQL*Plus script. The job must point to a credential that contains a valid operating system username and password. The SQL*Plus script is run by the SQL*Plus executable. The job may point to a connect credential that contains a database credential. If so, this credential is used to connect to the database before running the SQL*Plus script. Note that if you choose to use connect credential, you must use set_attribute to specify the Connect_Credential_Name attribute. If you do not have connect credential, you must include an explicit SQL*Plus connect statement providing a valid database userid / password. The job owner must have the CREATE EXTERNAL JOB system privilege. • 'BACKUP_SCRIPT' This specifies that the job is an RMAN backup script. The script runs a connect statement that uses either a password or OS authentication before it executes any target commands. The job points to a credential that contains a valid operating system username and password. The RMAN session runs under this operating system user. The Scheduler uses the RMAN executable from the current Oracle home to run the script and throws an error if this is missing. The job owner must have the CREATE EXTERNAL JOB system privilege.

Table 169-28 (Cont.) *CREATE_JOB Procedure Parameters*

Parameter	Description
job_action	This attribute specifies the action of the job. If job_action is not specified for an inline program, then an error is generated when creating the job. The job action is executed inside an autonomous transaction, and all autonomous transaction guidelines and restrictions apply. For example, online DDL operations are not allowed inside an autonomous transaction, and therefore cannot be used in the job action. The following actions are possible: • For a PL/SQL block: The action is to execute PL/SQL code. These blocks must end with a semicolon. For example, my_proc(); or BEGIN my_proc(); END; or DECLARE arg pls_integer:= 10; BEGIN my_proc2(arg); END;. Note that the Scheduler wraps job_action in its own block and passes the following to PL/SQL for execution: DECLARE ... BEGIN job_action END; This is done to declare some internal Scheduler variables. You can include any Scheduler metadata attribute except event_message in your PL/SQL code. You use the attribute name as you use any other PL/SQL identifier, and the Scheduler assigns it a value. See Table 169-40 for details on available metadata attributes. • For a stored procedure: The action is the name of the stored procedure. You have to specify the schema if the procedure resides in another schema than the job. If case sensitivity is needed, enclose the schema name and the store procedure name in double quotes. For example, job_action_action=>'"*Schema*"."*Procedure*"'. PL/SQL procedures with INOUT or OUT arguments are not supported as job_action when the job or program type is STORED_PROCEDURE. • For an executable: The action is the name of the external executable, including the full path name, but excluding any command-line arguments. If the action starts with a single question mark ('?'), the question mark is replaced by the path to the Oracle home directory for a local job or to the Scheduler agent home for a remote job. If the action contains an at-sign ('@') and the job is local, the at-sign is replaced with the SID of the current Oracle instance. NOTE: Shell script syntax is not supported, only syntax for the name of and path to an executable is supported. • For a chain: The action is the name of a Scheduler chain object. You must specify the schema of the chain if it resides in a different schema than the job. • For an external script: The job_action must be either the path to an operating system script or an inline operating system script. If the job_action is a path to a script, then the script must reside on every computer that the job runs on. The job_action may contain calls to SQL*Plus or RMAN executables directly, without having to specify its full path, given that they are stored on their default location for every computer that runs the job. The job can only have arguments that are strings or that can be cast to strings. These arguments are passed positionally when the script is called. The job must point to a credential that contains a valid operating system username and password. • For a SQL script:

Table 169-28 (Cont.) *CREATE_JOB Procedure Parameters*

Parameter	Description
	The `job_action` must be either the path to a SQL*Plus script or an inline SQL*Plus script. If the `job_action` is a path to a script, then the script must reside on every computer that the job runs on. The job can only have arguments that are strings or that can be cast to strings. These arguments are passed positionally when the script is called. If the arguments are named, they are also bound to named variables in the SQL*Plus session. • For a backup script: The `job_action` is either the path to a RMAN script or an inline RMAN script. If the `program_action` is a path to a script, then the script must reside on every computer that the program runs on. The job can only have arguments that are strings or that can be cast to strings. These arguments are passed positionally when the script is called.
`number_of_arguments`	This attribute specifies the number of arguments that the job expects. The range is 0-255, with the default being 0.
`program_name`	The name of the program associated with this job. If the program is of type `EXECUTABLE`, the job owner must have the `CREATE EXTERNAL JOB` system privilege before the job can be enabled or run.
`start_date`	This attribute specifies the first date and time on which this job is scheduled to start. If `start_date` and `repeat_interval` are left null, then the job is scheduled to run as soon as the job is enabled. For repeating jobs that use a calendaring expression to specify the repeat interval, `start_date` is used as a reference date. The first time the job runs is the first match of the calendaring expression that is on or after the current date and time. The Scheduler cannot guarantee that a job executes on an exact time because the system may be overloaded and thus resources unavailable.
`event_condition`	This is a conditional expression based on the columns of the event source queue table. The expression must have the syntax of an Advanced Queuing rule. Accordingly, you can include user data properties in the expression provided that the message payload is an object type, and that you prefix object attributes in the expression with `tab.user_data`. For more information on rules, see the `DBMS_AQADM.ADD_SUBSCRIBER` procedure.
`queue_spec`	This argument specifies either of the following: • The source queue where events that start this particular job are enqueued. If it is secure, then the `queue_spec` argument is a pair of values of the form *`queue_name, agent name`*. If it is not secure, then only the queue name need be provided. If a fully qualified queue name is not provided, the queue is assumed to be in the job owner's schema. In the case of secure queues, the agent name provided should belong to a valid agent that is currently subscribed to the queue. • A file watcher name. For more information on this option, see *Oracle Database Administrator's Guide*.
`repeat_interval`	This attribute specifies how often the job repeats. You can specify the repeat interval by using calendaring or PL/SQL expressions. The expression specified is evaluated to determine the next time the job should run. If `repeat_interval` is not specified, the job runs only once at the specified start date. See "Calendaring Syntax" for further information.
`schedule_name`	The name of the schedule, window, or window group associated with this job.
`job_class`	The class this job is associated with.

Table 169-28 (Cont.) *CREATE_JOB Procedure Parameters*

Parameter	Description
end_date	This attribute specifies the date and time after which the job expires and is no longer run. After the end_date, if auto_drop is TRUE, the job is dropped. If auto_drop is FALSE, the job is disabled and the STATE of the job is set to COMPLETED. If no value for end_date is specified, the job repeats forever unless max_runs or max_failures is set, in which case the job stops when either value is reached. The value for end_date must be after the value for start_date. If end_date is less than start_date, then an error will be generated. If end_date is the same as start_date, then the job will not execute and no error will be generated.
comments	This attribute specifies a comment about the job. By default, this attribute is NULL.
job_style	Style of the job being created. This argument can have one of the following values: • 'REGULAR' creates a regular job. This is the default. • 'LIGHTWEIGHT' creates a lightweight job. This value is permitted only when the job references a program object. Use lightweight jobs when you have many short-duration jobs that run frequently. Under certain circumstances, using lightweight jobs can deliver a small performance gain. • 'IN_MEMORY_RUNTIME' creates an in-memory runtime job. These jobs are based on lightweight job structures, so the same rules and restrictions apply; however, they further boost performance by keeping an in-memory cache, so they minimize disk access for pre-run and post-run actions. • 'IN_MEMORY_FULL' creates an in-memory full job. In-memory full jobs require a program and cannot have a schedule or repeat interval. They run automatically when the job is enabled, and after running they are discarded. They keep all the job information in memory and are not backed up on disk, meaning that they are lost when the instance is rebooted. They are designed to run actions that must be performed immediately with the least amount of overhead possible.
credential_name	The default credential to use with the job. Applicable only to remote database jobs, remote external jobs, local external jobs, script jobs, and event-based jobs that process file arrival events. The credential must exist. For local database jobs, it must be NULL. For local external jobs only, if this attribute is NULL (the default), then a preferred (default) credential is selected. See *Oracle Database Administrator's Guide* for information about preferred credentials for local external jobs. See also: "CREATE_CREDENTIAL Procedure"

Table 169-28 (Cont.) *CREATE_JOB Procedure Parameters*

Parameter	Description
destination_name	The database destination or external destination for the job. Use for remote database jobs and remote external jobs only. Must be NULL for jobs running on the local database or for local external jobs (executables). This attribute can be a single destination name or the name of a group of type 'EXTERNAL_DEST' or 'DB_DEST'. The single destination or group must already exist. The following applies to this attribute: • If it is a database destination, it must have been created by the CREATE_DATABASE_DESTINATION Procedure. • If it is an external destination, it must have been implicitly created by registering a remote Scheduler agent with the local database. • If it is a group, each member of the group must exist, and the job must run on all destinations named in the group. See "CREATE_GROUP Procedure". destination_name cannot reference a destination group when: • The job type is 'CHAIN' • The job style is 'LIGHTWEIGHT', 'IN_MEMORY_RUNTIME', or 'IN_MEMORY_FULL', If the credential_name argument of CREATE_JOB is NULL, each destination must be preceded by a credential, in the following format: *credential.destination* The credential must already exist. If the credential_name argument is provided, then it serves as the default credential for every destination that is not preceded by a credential. You can query the views *_SCHEDULER_DB_DESTS and ALL_SCHEDULER_EXTERNAL DESTS for existing destinations and *_SCHEDULER_GROUP_MEMBERS for existing groups and their members. *** destination job attribute is deprecated in Oracle Database 11*g*R2 and superseded by destination_name.
enabled	This attribute specifies whether the job is created enabled or not. The possible settings are TRUE or FALSE. By default, this attribute is set to FALSE and, therefore, the job is created as disabled. A disabled job means that the metadata about the job has been captured, and the job exists as a database object. However, the Scheduler ignores the job and the job coordinator does not pick it for processing. In order for the job coordinator to process the job, the job must be enabled. You can enable a job by setting this argument to TRUE or by using the ENABLE procedure.
auto_drop	This flag, if TRUE, causes a job to be automatically dropped after it has completed or has been automatically disabled. A job is considered completed if: • Its end date (or the end date of the job schedule) has passed. Note that a job with a Window schedule will not be auto-dropped when the window closes, because this is not considered to be the end of the Window. • It has run max_runs number of times. max_runs must be set with SET_ATTRIBUTE. • It is not a repeating job and has run once. A job is disabled when it has failed max_failures times. max_failures is also set with SET_ATTRIBUTE. If this flag is set to FALSE, the jobs are not dropped and their metadata is kept until the job is explicitly dropped with the DROP_JOB procedure. By default, jobs are created with auto_drop set to TRUE.

Usage Notes

Jobs are created as disabled by default. You must explicitly enable them so that they will become active and scheduled. Before enabling a job, ensure that all program arguments, if any, are defined, either by defining default values in the program object or by supplying values with the job.

The `JOB_QUEUE_PROCESSES` initialization parameter specifies the maximum number of processes that can be created for the execution of jobs. Beginning with Oracle Database 11*g* Release 2, `JOB_QUEUE_PROCESSES` applies to `DBMS_SCHEDULER` jobs. Setting this parameter to 0 disables `DBMS_SCHEDULER` jobs.

To create a job in your own schema, you need to have the `CREATE JOB` privilege. A user with the `CREATE ANY JOB` privilege can create a job in any schema. If the job being created will reside in another schema, the job name must be qualified with the schema name. For a job of type `EXECUTABLE` (or for a job that points to a program of type `EXECUTABLE`), the job owner must have the `CREATE EXTERNAL JOB` system privilege before the job can be enabled or run.

Associating a job with a particular class or program requires `EXECUTE` privileges for that class or program.

Not all possible job attributes can be set with `CREATE_JOB`. Some must be set after the job is created. For example, job arguments must be set with the SET_JOB_ARGUMENT_VALUE Procedure or the SET_JOB_ANYDATA_VALUE Procedure. Other job attributes, such as `job_priority` and `max_runs`, are set with the SET_ATTRIBUTE Procedure.

To create multiple jobs efficiently, use the `CREATE_JOBS` procedure.

> **Note:**
>
> The Scheduler runs event-based jobs for each occurrence of an event that matches the event condition of the job. However, events that occur while the job is already running are ignored; the event gets consumed, but does not trigger another run of the job.

169.6.17 CREATE_JOB_CLASS Procedure

This procedure creates a job class. Job classes are created in the `SYS` schema.

Syntax

```
DBMS_SCHEDULER.CREATE_JOB_CLASS (
   job_class_name            IN VARCHAR2,
   resource_consumer_group   IN VARCHAR2 DEFAULT NULL,
   service                   IN VARCHAR2 DEFAULT NULL,
   logging_level             IN PLS_INTEGER
                                DEFAULT DBMS_SCHEDULER.LOGGING_RUNS,
   log_history               IN PLS_INTEGER DEFAULT NULL,
   comments                  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-29 CREATE_JOB_CLASS Procedure Parameters

Parameter	Description
job_class_name	The name to assign to the job class. Job classes can only be created in the SYS schema. This attribute specifies the name of the job class and uniquely identifies the job class. The name must be unique in the SQL namespace. For example, a job class cannot have the same name as a table in a schema.
resource_consumer_group	This attribute specifies the resource consumer group that his class is associated with. A resource consumer group is a set of synchronous or asynchronous sessions that are grouped together based on their processing needs. A job class has a many-to-one relationship with a resource consumer group. The resource consumer group that the job class associates with determines the resources that are allocated to the job class. If a resource consumer group is dropped, job classes associated with it are then associated with the default resource consumer group. If no resource consumer group is specified, job classes are associated with the default resource consumer group. If the specified resource consumer group does not exist when creating the job class, an error occurs.
service	This attribute specifies the database service that the jobs in this class have affinity to. In an Oracle RAC environment, this means that the jobs in this class only run on those database instances that are assigned to the specific service. Note that a service can be mapped to a resource consumer group, so you can also control resources allocated to jobs by specifying a service. See DBMS_RESOURCE_MANAGER.SET_CONSUMER_GROUP_MAPPING for details. If both the resource_consumer_group and service attributes are specified, and if the service is mapped to a resource consumer group, the resource_consumer_group attribute takes precedence. If no service is specified, the job class belongs to the default service, which means it has no service affinity and any one of the database instances within the cluster might run the job. If the service that a job class belongs to is dropped, the job class will then belong to the default service. If the specified service does not exist when creating the job class, then an error occurs.
logging_level	This attribute specifies how much information is logged. The possible options are: • DBMS_SCHEDULER.LOGGING_OFF No logging is performed for any jobs in this class. • DBMS_SCHEDULER.LOGGING_RUNS The Scheduler writes detailed information to the job log for all runs of each job in this class. This is the default. • DBMS_SCHEDULER.LOGGING_FAILED_RUNS The Scheduler logs only jobs that failed in this class. • DBMS_SCHEDULER.LOGGING_FULL In addition to recording every run of a job, the Scheduler records all operations performed on all jobs in this class. Every time a job is created, enabled, disabled, altered (with SET_ATTRIBUTE), stopped, and so, an entry is recorded in the log.

Table 169-29 (Cont.) CREATE_JOB_CLASS Procedure Parameters

Parameter	Description
log_history	This attribute controls the number of days that job log entries for jobs in this class are retained. It helps prevent the job log from growing indiscriminately. The range of valid values is 0 through1000000. If set to 0, no history is kept. If NULL (the default), retention days are set by the log_history Scheduler attribute (set with SET_SCHEDULER_ATTRIBUTE).
comments	This attribute is for an optional comment about the job class. By default, this attribute is NULL.

Usage Notes

For users to create jobs that belong to a job class, the job owner must have EXECUTE privileges on the job class. Therefore, after the job class has been created, EXECUTE privileges must be granted on the job class so that users create jobs belonging to that class. You can also grant the EXECUTE privilege to a role.

Creating a job class requires the MANAGE SCHEDULER system privilege.

169.6.18 CREATE_JOBS Procedure

This procedure creates multiple jobs and sets the values of their arguments in a single call.

Syntax

```
DBMS_SCHEDULER.CREATE_JOBS (
   jobdef_array      IN SYS.JOB_DEFINITION_ARRAY,
   commit_semantics  IN VARCHAR2 DEFAULT 'STOP_ON_FIRST_ERROR');
```

Parameters

Table 169-30 CREATE_JOBS Procedure Parameters

Parameter	Description
jobdef_array	The array of job definitions. See "Data Structures" for a description of the JOB_DEFINITION_ARRAY and JOB_DEFINITION datatypes.
commit_semantics	The commit semantics. The following types are supported: • STOP_ON_FIRST_ERROR returns on the first error. Previous successfully created jobs are committed to disk. This is the default. • TRANSACTIONAL returns on the first error and everything that happened before that error is rolled back. • ABSORB_ERRORS tries to absorb any errors and attempts to create the rest of the jobs on the list. It commits all successfully created jobs. If errors occur, you can query the view SCHEDULER_BATCH_ERRORS for details.

Usage Notes

This procedure creates many jobs in the context of a single transaction. To realize the desired performance gains, the jobs being created must be grouped in batches of sufficient size. Calling CREATE_JOBS with a small array size may not be much faster than calling CREATE_JOB once for each job.

You cannot use this procedure to create multiple-destination jobs. That is, the `destination` attribute of the `job_definition` object cannot reference a destination group.

Examples

See *Oracle Database Administrator's Guide*.

169.6.19 CREATE_PROGRAM Procedure

This procedure creates a program.

Syntax

```
DBMS_SCHEDULER.CREATE_PROGRAM (
   program_name             IN VARCHAR2,
   program_type             IN VARCHAR2,
   program_action           IN VARCHAR2,
   number_of_arguments      IN PLS_INTEGER DEFAULT 0,
   enabled                  IN BOOLEAN DEFAULT FALSE,
   comments                 IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-31 CREATE_PROGRAM Procedure Parameters

Parameter	Description
`program_name`	The name to assign to the program. The name must be unique in the SQL namespace. For example, a program cannot have the same name as a table in a schema. If no name is specified, then an error occurs.

Table 169-31 (Cont.) CREATE_PROGRAM Procedure Parameters

Parameter	Description
program_type	This attribute specifies the type of program you are creating. If it is not specified then you get an error. These are the supported values for program_type: • 'PLSQL_BLOCK' This specifies that the program is a PL/SQL block. Job or program arguments are not supported when the job or program type is PLSQL_BLOCK. In this case, the number of arguments must be 0. • 'STORED_PROCEDURE' This specifies that the program is a PL/SQL or Java stored procedure, or an external C subprogram. Only procedures, not functions with return values, are supported. PL/SQL procedures with INOUT or OUT arguments are not supported. • 'EXECUTABLE' This specifies that the job is going to be run outside the database using an external executable. External programs imply anything that can be executed from the operating system command line. AnyData arguments are not supported with job or program type EXECUTABLE. • 'EXTERNAL_SCRIPT' This specifies that the job is an external script that uses the command shell of the computer running the job. For Windows this is cmd.exe and for UNIX based systems the sh shell, unless a different interpreter is specified by prefixing the first line of the script with #!. • 'SQL_SCRIPT' This specifies that the program is a SQL*Plus script. A job using this program must point to a credential that contains a valid operating system username and password. The SQL*Plus script is run by SQL*Plus executable. The job using this program may point to a connect credential that contains a database credential. If so, this credential is used to connect to the database before running the SQL*Plus script. Note that if you choose to use connect credential, you must use set_attribute to specify the Connect_Credential_Name attribute. If you do not have connect credential, you must include an explicit SQL*Plus connect statement providing a valid database userid / password. • 'BACKUP_SCRIPT' This specifies that the program is an RMAN backup script. The script runs a connect statement that uses either a password or OS authentication before it executes any target commands. The Scheduler uses the RMAN executable from the current Oracle home to run the script and throws an error if this is missing.

Table 169-31 (Cont.) CREATE_PROGRAM Procedure Parameters

Parameter	Description
program_action	This attribute specifies the action of the program. If program_action is not specified, an error is generated. The following actions are possible: • For a PL/SQL block, the action is to execute PL/SQL code. These blocks must end with a semicolon. For example, my_proc(); or BEGIN my_proc(); END; or DECLARE arg pls_integer:= 10; BEGIN my_proc2(arg); END;. Note that the Scheduler wraps job_action in its own block and passes the following to PL/SQL for execution: DECLARE ... BEGIN job_action END; This is done to declare some internal Scheduler variables. You can include any Scheduler metadata attribute except event_message in your PL/SQL code. You use the attribute name as you use any other PL/SQL identifier, and the Scheduler assigns it a value. See Table 169-40 for details on available metadata attributes. If it is an anonymous block, special Scheduler metadata may be accessed using the following variable names: job_name, job_owner, job_start, window_start, window_end. For more information, see the "DEFINE_METADATA_ARGUMENT Procedure". • For a stored procedure, the action is the name of the stored procedure. You have to specify the schema if the procedure resides in a schema other than the job. If case sensitivity is needed, enclose the schema name and the store procedure name in double quotes. For example, program_action=>'"*Schema*"."*Procedure*"'. • For an executable, the action is the name of the external executable, including the full path name, but excluding any command-line arguments. If the action starts with a single question mark ('?'), the question mark is replaced by the path to the Oracle home directory for a local job or to the Scheduler agent home for a remote job. If the action contains an at sign ('@') and the job is local, the at sign is replaced with the SID of the current Oracle instance. • For an external script, the action must be either the path to an operating system script or an inline operating system script. If the program_action is a path to a script, then the script must reside on every computer that the program runs on. The program_action may contain calls to SQL*Plus or RMAN executables directly, without having to specify its full path, given that they are stored on their default location for every computer that runs the job. The job can only have arguments that are strings or that can be cast to strings. These arguments are passed positionally when the script is called. The program points to a credential that contains a valid operating system username and password. • For a SQL script, the action must be either the path to a SQL*Plus script or an inline SQL*Plus script. If the program_action is a path to a script, then the script must reside on every computer that the program runs on. The job can only have arguments that are strings or that can be cast to strings. These arguments are passed positionally when

Table 169-31 (Cont.) CREATE_PROGRAM Procedure Parameters

Parameter	Description
	the script is called. If the arguments are named, they are also bound to named variables in the SQL*Plus session. • For a backup script, the action must be either the path to a RMAN script or an inline RMAN script. If the `program_action` is a path to a script, then the script must reside on every computer that the program runs on. The job can only have arguments that are strings or that can be cast to strings. These arguments are passed positionally when the script is called.
`number_of_arguments`	This attribute specifies the number of arguments the program takes. If this parameter is not specified, then the default is 0. A program can have a maximum of 255 arguments. If the `program_type` is `PLSQL_BLOCK`, then this parameter is ignored.
`enabled`	This flag specifies whether the program should be created as enabled or not. If the flag is set to `TRUE`, then validity checks are made and the program is created as `ENABLED` if all the checks be successful. By default, this flag is set to `FALSE`, meaning not created enabled. You can also call the `ENABLE` procedure to enable the program before it can be used.
`comments`	A comment about the program. By default, this attribute is `NULL`.

Usage Notes

To create a program in their own schema, users need the `CREATE JOB` privilege. A user with the `CREATE ANY JOB` privilege can create a program in any schema. A program is created in a disabled state by default (unless the enabled parameter is set to `TRUE`). It cannot be executed by a job until it is enabled.

To use your programs, other users must have `EXECUTE` privileges, therefore once a program has been created, you have to grant `EXECUTE` privileges on it.

> **See Also:**
>
> "DEFINE_PROGRAM_ARGUMENT Procedure"

169.6.20 CREATE_RESOURCE Procedure

This procedure allows users to specify the resources used by jobs or to create a new resource.

Syntax

```
DBMS_SCHEDULER.CREATE_RESOURCE (
  resource_name    IN VARCHAR2,
  units            IN PLS_INTEGER,
  status           IN VARCHAR2 DEFAULT 'ENFORCE_CONSTRAINTS',
  constraint_level IN VARCHAR2 DEFAULT 'JOB_LEVEL',
  comments         IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-32 CREATE_RESOURCE Procedure Parameters

Parameter	Description
resource_name	The name of the resource.
units	The number of units of this resource that the job or program uses.
status	The status of the resource. • 'ENFORCE_CONSTRAINTS'. This is the default value, and when set, will force the scheduler to enforce resource limits. When the maximum number of units of this resource has been reached no additional jobs using this resource will get started. • 'IGNORE_CONSTRAINTS'. When set, the scheduler will ignore any constraints on this resource. • 'BLOCKED_ALL_JOBS'. No jobs having a constraint on this resource will be allowed to run. The resource is considered to be permanently blocking until switched to one of the other two states.
constraint_level	Level of the constraint: JOB_LEVEL or PROGRAM_LEVEL. For incompatibilities, for JOB_LEVEL, the incompatibility members must be jobs; for PROGRAM_LEVEL the incompatibility members must be programs.
comments	Descriptive comment about the resource.

Usage Notes

The following example creates a new resource.

```
BEGIN
   DBMS_SCHEDULER.CREATE_RESOURCE(
      resource_name => 'my_resource',
      units => 3,
      state => 'ENFORCE_CONSTRAINTS',
      comments => 'Resource1'
   )
END;
/
```

See Also:

- Creating or Dropping a Resource in *Oracle Database Administrator's Guide*
- SET_RESOURCE_CONSTRAINT Procedure

169.6.21 CREATE_SCHEDULE Procedure

This procedure creates a schedule.

Syntax

```
DBMS_SCHEDULER.CREATE_SCHEDULE (
   schedule_name          IN VARCHAR2,
```

```
   start_date             IN TIMESTAMP WITH TIMEZONE DEFAULT NULL,
   repeat_interval        IN VARCHAR2,
   end_date               IN TIMESTAMP WITH TIMEZONE DEFAULT NULL,
   comments               IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-33 CREATE_SCHEDULE Procedure Parameters

Parameter	Description
schedule_name	The name to assign to the schedule. The name must be unique in the SQL namespace. For example, a schedule cannot have the same name as a table in a schema. If no name is specified, then an error occurs.
start_date	This attribute specifies the first date and time on which this schedule becomes valid. For a repeating schedule, the value for start_date is a reference date. In this case, the start of the schedule is not the start_date; it depends on the repeat interval specified. start_date is used to determine the first instance of the schedule. If start_date is specified in the past and no value for repeat_interval is specified, the schedule is invalid. For a repeating job or window, start_date can be derived from the repeat_interval if it is not specified. If start_date is null, then the date that the job or window is enabled is used. start_date and repeat_interval cannot both be null.
repeat_interval	This attribute specifies how often the schedule repeats. It is expressed using calendaring syntax. See "Calendaring Syntax" for further information. PL/SQL expressions are not allowed as repeat intervals for named schedules.
end_date	The date and time after which jobs will not run and windows will not open. A non-repeating schedule that has no end_date is valid forever. end_date has to be after the start_date. If this is not the case, then an error is generated when the schedule is created.
comments	This attribute specifies an optional comment about the schedule. By default, this attribute is NULL.

Usage Notes

This procedure requires the CREATE JOB privilege to create a schedule in your own schema or the CREATE ANY JOB privilege to create a schedule in someone else's schema by specifying schema.schedule_name. Once a schedule has been created, it can be used by other users. The schedule is created with access to PUBLIC. Therefore, there is no need to explicitly grant access to the schedule.

169.6.22 CREATE_WINDOW Procedure

This procedure creates a recurring time window and associates it with a resource plan. You can then use the window to schedule jobs that run under the associated resource plan. Windows are created in the SYS schema.

The procedure is overloaded.

Syntax

Creates a window using a named schedule object:

```
DBMS_SCHEDULER.CREATE_WINDOW (
   window_name             IN VARCHAR2,
```

```
   resource_plan           IN VARCHAR2,
   schedule_name           IN VARCHAR2,
   duration                IN INTERVAL DAY TO SECOND,
   window_priority         IN VARCHAR2 DEFAULT 'LOW',
   comments                IN VARCHAR2 DEFAULT NULL);
```

Creates a window using an inlined schedule:

```
DBMS_SCHEDULER.CREATE_WINDOW (
   window_name             IN VARCHAR2,
   resource_plan           IN VARCHAR2,
   start_date              IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   repeat_interval         IN VARCHAR2,
   end_date                IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   duration                IN INTERVAL DAY TO SECOND,
   window_priority         IN VARCHAR2 DEFAULT 'LOW',
   comments                IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-34 CREATE_WINDOW Procedure Parameters

Parameter	Description
window_name	The name to assign to the window. The name must be unique in the SQL namespace. All windows are in the SYS schema, so the preface 'SYS' is optional.
resource_plan	This attribute specifies the resource plan that automatically activates when the window opens. When the window closes, the system switches to the appropriate resource plan, which is usually the plan that was in effect before the window opened, but can also be the plan of a different window. Only one resource plan can be associated with a window. It may be NULL or the empty string (""). When it is NULL, the resource plan in effect when the window opens stays in effect for the duration of the window. When it is the empty string, the resource manager is disabled for the duration of the window. If the window is open and the resource plan is dropped, then the resource allocation for the duration of the window is not affected.
start_date	This attribute specifies the first date and time on which this window is scheduled to open. If the value for start_date specified is in the past or is not specified, the window opens as soon as it is created. For repeating windows that use a calendaring expression to specify the repeat interval, the value for start_date is a reference date. The first time the window opens depends on the repeat interval specified and the value for start_date.
duration	This attribute specifies how long the window stays open. For example, 'interval '5' hour' for five hours. There is no default value for this attribute. Therefore, if no value is specified when the window is created, an error occurs. The duration is of type interval day to seconds and ranges from one minute to 99 days.
schedule_name	This attribute specifies the name of the schedule associated with the window.

Table 169-34 (Cont.) CREATE_WINDOW Procedure Parameters

Parameter	Description
repeat_interval	This attribute specifies how often the window repeats. It is expressed using the Scheduler calendaring syntax. See "Calendaring Syntax" for more information. A PL/SQL expression cannot be used to specify the repeat interval for a window. The expression specified is evaluated to determine the next time the window opens. If no repeat_interval is specified, the window opens only once at the specified start date.
end_date	This attribute specifies the date and time after which the window no longer opens. When the value for end_date is reached, the window is disabled. In the *_SCHEDULER_WINDOWS views, the enabled flag of the window is set to FALSE. A non-repeating window that has no value for end_date opens only once for the duration of the window. For a repeating window, if no end_date is specified, then the window keeps repeating forever. The end_date must be after the start_date. If it is not, then an error is generated when the window is created.
window_priority	This attribute is only relevant when two windows overlap. Because only one window can be in effect at one time, the window priority determines which window opens. The two possible values for this attribute are 'HIGH' and 'LOW'. A high priority window has precedence over a low priority window, therefore, the low priority window does not open if it overlaps a high priority window. By default, windows are created with priority 'LOW'.
comments	This attribute specifies an optional comment about the window. By default, this attribute is NULL.

Usage Notes

Creating a window requires the MANAGE SCHEDULER privilege.

Scheduler windows are the principal mechanism used to automatically switch resource plans according to a schedule. You can also manually activate a resource plan by using the ALTER SYSTEM SET RESOURCE_MANAGER_PLAN statement or the DBMS_RESOURCE_MANAGER.SWITCH_PLAN package procedure. Note that either of these manual methods can also disable resource plan switching by Scheduler windows. For more information, see *Oracle Database Administrator's Guide* and "SWITCH_PLAN Procedure".

169.6.23 DEFINE_ANYDATA_ARGUMENT Procedure

This procedure defines a name or default value for a program argument that is of a complex type and must be encapsulated within an ANYDATA object. A job that references the program can override the default value.

Syntax

```
DBMS_SCHEDULER.DEFINE_ANYDATA_ARGUMENT (
   program_name            IN VARCHAR2,
   argument_position       IN PLS_INTEGER,
   argument_name           IN VARCHAR2 DEFAULT NULL,
   argument_type           IN VARCHAR2,
```

```
   default_value           IN SYS.ANYDATA,
   out_argument            IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-35 DEFINE_ANYDATA_ARGUMENT Procedure Parameters

Parameter	Description
program_name	The name of the program to be altered. A program with this name must exist.
argument_position	The position of the argument as it is passed to the executable. Argument numbers go from one to the number_of_arguments specified for the program. This must be unique, so it can replace any argument already defined at this position.
argument_name	The name to assign to the argument. It is optional, but must be unique for the program if it is specified. If you assign a name, the name can then be used by other package procedures, including the SET_JOB_ANYDATA_VALUE Procedure.
argument_type	The datatype of the argument being defined. This is not verified or used by the Scheduler. It is only used by the user of the program when deciding what value to assign to the argument.
default_value	The default value to be assigned to the argument encapsulated within an AnyData object. This is optional.
out_argument	This parameter is reserved for future use. It must be set to FALSE.

Usage Notes

All program arguments from one to the number_of_arguments value must be defined before a program can be enabled. If a default value for an argument is not defined with this procedure, a value must be defined in the job.

Defining a program argument requires that you be the owner of the program or have ALTER privileges on that program. You can also define a program argument if you have the CREATE ANY JOB privilege.

See Also:

- "DEFINE_PROGRAM_ARGUMENT Procedure"
- "SET_JOB_ANYDATA_VALUE Procedure"

169.6.24 DEFINE_CHAIN_EVENT_STEP Procedure

This procedure adds or replaces a chain step and associates it with an event schedule or an inline event.

Once started in a running chain, this step does not complete until the specified event has occurred. Every step in a chain must be defined before the chain can be enabled and used. Defining a step gives it a name and specifies what happens during the step. If a step already exists with this name, the new step replaces the old one.

Syntax

```
DBMS_SCHEDULER.DEFINE_CHAIN_EVENT_STEP (
   chain_name              IN VARCHAR2,
   step_name               IN VARCHAR2,
   event_schedule_name     IN VARCHAR2,
   timeout                 IN INTERVAL DAY TO SECOND DEFAULT NULL);

DBMS_SCHEDULER.DEFINE_CHAIN_EVENT_STEP (
   chain_name              IN VARCHAR2,
   step_name               IN VARCHAR2,
   event_condition         IN VARCHAR2,
   queue_spec              IN VARCHAR2,
   timeout                 IN INTERVAL DAY TO SECOND DEFAULT NULL);
```

Parameters

Table 169-36 DEFINE_CHAIN_EVENT_STEP Procedure Parameters

Parameter	Description
chain_name	The name of the chain that the step is in
step_name	The name of the step
event_schedule_name	The name of the event schedule that the step waits for
timeout	This parameter is reserved for future use
event_condition	See the CREATE_EVENT_SCHEDULE Procedure
queue_spec	See the CREATE_EVENT_SCHEDULE Procedure

Usage Notes

Defining a chain step requires ALTER privileges on the chain either as the owner of the chain, or as a user with the ALTER object privilege on the chain or the CREATE ANY JOB system privilege.

You can base a chain step on a file watcher as well. To do this, provide the file watcher name directly in the queue_spec parameter, or use a file watcher schedule for the event_schedule_name parameter.

> **See Also:**
>
> "DEFINE_CHAIN_STEP Procedure"

169.6.25 DEFINE_CHAIN_RULE Procedure

This procedure adds a new rule to an existing chain, specified as a condition-action pair. The condition is expressed using either SQL or the Scheduler chain condition syntax and indicates the prerequisites for the action to occur. The action is a result of the condition being met.

An actual rule object is created to store the rule in the schema where the chain resides. If a rule name is given, this name is used for the rule object. If an existing rule name in the schema of the chain is given, the existing rule is altered. (A schema different than the schema of the chain cannot be specified). If no rule name is given, one is generated in the form SCHED_RULE${N}.

Syntax

```
DBMS_SCHEDULER.DEFINE_CHAIN_RULE (
   chain_name              IN VARCHAR2,
   condition               IN VARCHAR2,
   action                  IN VARCHAR2,
   rule_name               IN VARCHAR2 DEFAULT NULL,
   comments                IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-37 DEFINE_CHAIN_RULE Procedure Parameters

Parameter	Description
chain_name	The name of the chain to alter
condition	A boolean expression which must evaluate to TRUE for the action to be performed. Every chain must have a rule that evaluates to TRUE to start the chain. For this purpose, you can use a rule that has 'TRUE' as its condition if you are using Scheduler chain condition syntax, or '1=1' as its condition if you are using SQL syntax. • Scheduler Chain Condition Syntax See "Scheduler Chain Condition Syntax" and *Oracle Database Administrator's Guide* for details • SQL WHERE Clause Syntax Conditions expressed with SQL must use the syntax of a SELECT statement WHERE clause. You can refer to chain step attributes by using the chain step name as a bind variable. The bind variable syntax is :*step_name.attribute*. (*step_name* refers to a typed object.) Possible attributes are: completed, state, start_date, end_date, error_code, and duration. Possible values for the state attribute include: 'NOT_STARTED', 'SCHEDULED', 'RUNNING', 'PAUSED', 'STALLED', 'SUCCEEDED', 'FAILED', and 'STOPPED'. If a step is in the state 'SUCCEEDED', 'FAILED', or 'STOPPED', its completed attribute is set to 'TRUE', otherwise completed is 'FALSE'.
action	The action to be performed when the rule evaluates to TRUE. The action must consist of at least one keyword with an optional value and an optional delay clause. Possible actions include: • [AFTER *delay_interval*] START *step_1*[,*step_2* ...] • STOP *step_1*[,*step_2* ...] • END [{*end_value*|*step_name.error_code*}] At the beginning of the START action, a delay clause can specify a delay interval before performing the action. *delay_interval* is a formatted datetime interval of the form HH:MM:SS. The END action ends the chain with an error code equal to either the supplied end_value or the error code that step_name completes with. The default error code is 0, indicating a successful chain run.
rule_name	The name of the rule being created. If no rule_name is given, one is generated in the form SCHED_RULE$_{N}.

Table 169-37 (Cont.) DEFINE_CHAIN_RULE Procedure Parameters

Parameter	Description
comments	An optional comment describing the rule. This is stored in the rule object created.

Scheduler Chain Condition Syntax

The Scheduler chain condition syntax provides an easy way to construct a condition using the states and error codes of steps in the current chain.

Chain Condition Syntax

The following are the available constructs for Scheduler chain condition syntax, which are all boolean expressions:

```
TRUE
FALSE
stepname [NOT] SUCCEEDED
stepname [NOT] FAILED
stepname [NOT] STOPPED
stepname [NOT] COMPLETED
stepname ERROR_CODE IN (integer, integer, integer ...)
stepname ERROR_CODE NOT IN (integer, integer, integer ...)
stepname ERROR_CODE = integer
stepname ERROR_CODE != integer
stepname ERROR_CODE <> integer
stepname ERROR_CODE > integer
stepname ERROR_CODE >= integer
stepname ERROR_CODE < integer
stepname ERROR_CODE <= integer
```

These boolean operators are available to create more complex conditions:

```
expression AND expression
expression OR expression
NOT (expression)
```

integer can be positive or negative. Parentheses may be used for clarity or to enforce ordering. You must use parentheses with the NOT operator.

PL/SQL code that runs as part of a step can set the value of ERROR_CODE for that step with the RAISE_APPLICATION_ERROR statement.

Usage Notes

Defining a chain rule requires ALTER privileges on the chain (either as the owner, or as a user with ALTER privileges on the chain or the CREATE ANY JOB system privilege).

You must define at least one rule that starts the chain and at least one that ends it. See the section "Adding Rules to a Chain" in *Oracle Database Administrator's Guide* for more information.

Examples

The following are examples of using rule conditions and rule actions.

Rule Conditions Using Scheduler Chain Condition Syntax

```
'step1 completed'
-- satisfied when step step1 has completed. (step1 completed is also TRUE when any
-- of the following are TRUE: step1 succeeded, step1 failed, step1 stopped.)

'step1 succeeded and step2 succeeded'
-- satisfied when steps step1 and step2 have both succeeded

'step1 error_code > 100'
-- satisfied when step step1 has failed with an error_code greater than 100

'step1 error_code IN (1, 3, 5, 7)'
-- satisfied when step step1 has failed with an error_code of 1, 3, 5, or 7
```

Rule Conditions Using SQL Syntax

```
':step1.completed = ''TRUE'' AND :step1.end_date >SYSDATE-1/24'
--satisfied when step step1 completed less than an hour ago

':step1.duration > interval ''5'' minute'
-- satisfied when step step1 has completed and took longer than 5 minutes to complete
```

Rule Actions

```
'AFTER 01:00:00 START step1, step2'
--After an hour start steps step1 and step2

'STOP step1'
--Stop step step1

END step4.error_code'
--End the chain with the error code that step step4 finished with. If step4 has not
completed, the chain will be ended unsuccessfully with error code 27435.

'END' or 'END 0'
--End the chain successfully (with error_code 0)

'END 100'
--End the chain unsuccessfully with error code 100.
```

169.6.26 DEFINE_CHAIN_STEP Procedure

This procedure adds or replaces a chain step and associates it with a program or a nested chain. When the chain step is started, the specified program or chain is run. If a step already exists with the name supplied in the `chain_name` argument, the new step replaces the old one.

The chain owner must have `EXECUTE` privileges on the program or chain associated with the step. Only one program or chain can run during a step.

You cannot set all possible step attributes with this procedure. Use the `ALTER_CHAIN` procedure to set additional chain step attributes, such as `credential_name` and `destination_name`.

Syntax

```
DBMS_SCHEDULER.DEFINE_CHAIN_STEP (
   chain_name              IN VARCHAR2,
   step_name               IN VARCHAR2,
   program_name            IN VARCHAR2);
```

Parameters

Table 169-38 DEFINE_CHAIN_STEP Procedure Parameters

Parameter	Description
chain_name	The name of the chain to alter.
step_name	The name of the step being defined. If a step already exists with this name, the new step replaces the old one.
program_name	The name of a program or chain to run during this step. The chain owner must have EXECUTE privileges on this program or chain.

Usage Notes

Defining a chain step requires ALTER privileges on the chain (either as the owner, or a user with ALTER privileges on the chain or the CREATE ANY JOB system privilege).

> **See Also:**
>
> - "ALTER_CHAIN Procedure"
> - "DEFINE_CHAIN_EVENT_STEP Procedure"

169.6.27 DEFINE_METADATA_ARGUMENT Procedure

This procedure defines a special metadata argument for the program. The Scheduler can pass Scheduler metadata through this argument to your stored procedure or other executable. You cannot set values for jobs using this argument.

Syntax

```
DBMS_SCHEDULER.DEFINE_METADATA_ARGUMENT (
  program_name            IN VARCHAR2,
  metadata_attribute      IN VARCHAR2,
  argument_position       IN PLS_INTEGER,
  argument_name           IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-39 DEFINE_METADATA_ARGUMENT Procedure Parameters

Parameter	Description
program_name	The name of the program to be altered
metadata_attribute	The metadata to be passed. Valid metadata attributes are: 'job_name', 'job_subname', 'job_owner', 'job_start', 'window_start', 'window_end', and 'event_message'. Table 169-40 describes these attributes in detail.
argument_position	The position of the argument as it is passed to the executable. The position cannot be greater than the number_of_arguments specified for the program. It must be unique, so it replaces any argument already defined at this position.

Table 169-39 (Cont.) DEFINE_METADATA_ARGUMENT Procedure Parameters

Parameter	Description
argument_name	The name to assign to the argument. It is optional, but must be unique for the program if it is specified. If you assign a name, the name can then be used by other package procedures.

Table 169-40 Metadata Attributes

Metadata Attribute	Datatype	Description
job_name	VARCHAR2	Name of the currently running job
job_subname	VARCHAR2	Subname of the currently running job. The name + subname form a unique identifier for a job that is running a chain step. NULL if the job is not part of a chain.
job_owner	VARCHAR2	Owner of the currently running job
job_scheduled_start	TIMESTAMP WITH TIME ZONE	When the currently running job was scheduled to start
job_start	TIMESTAMP WITH TIME ZONE	When the currently running job started
window_start	TIMESTAMP WITH TIME ZONE	If the job was started by a window, the time that the window opened
window_end	TIMESTAMP WITH TIME ZONE	If the job was started by a window, the time that the window is scheduled to close
event_message	(See Description)	For an event-based job, the message content of the event that started the job. The datatype of this attribute depends on the queue used for the event. It has the same type as the USER_DATA column of the queue table. In the case of a file arrival event, event_message is of type SYS.SCHEDULER_FILEWATCHER_RESULT. See "SCHEDULER_FILEWATCHER_RESULT Object Type".

Usage Notes

Defining a program argument requires that you be the owner of the program or have ALTER privileges on that program. You can also define a program argument if you have the CREATE ANY JOB privilege.

All metadata attributes except event_message can be used in PL/SQL blocks that you enter into the job_action or program_action attributes of jobs or programs, respectively. You use the attribute name as you use any other PL/SQL identifier, and the Scheduler assigns it a value.

169.6.28 DEFINE_PROGRAM_ARGUMENT Procedure

This procedure defines a name or default value for a program argument. If no default value is defined for a program argument, the job that references the program must supply an argument value. (The job can also override a default value.)

This procedure is overloaded.

Syntax

Defines a program argument without a default value:

```
PROCEDURE define_program_argument(
   program_name            IN VARCHAR2,
   argument_position       IN PLS_INTEGER,
   argument_name           IN VARCHAR2 DEFAULT NULL,
   argument_type           IN VARCHAR2,
   out_argument            IN BOOLEAN DEFAULT FALSE);
```

Defines a program argument with a default value:

```
PROCEDURE define_program_argument(
   program_name            IN VARCHAR2,
   argument_position       IN PLS_INTEGER,
   argument_name           IN VARCHAR2 DEFAULT NULL,
   argument_type           IN VARCHAR2,
   default_value           IN VARCHAR2,
   out_argument            IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-41 DEFINE_PROGRAM_ARGUMENT Procedure Parameters

Parameter	Description
program_name	The name of the program to be altered. A program with this name must exist.
argument_position	The position of the argument as it is passed to the executable. Argument numbers go from one to the number_of_arguments specified for the program. This must be unique so it replaces any argument already defined at this position.
argument_name	The name to assign to the argument. It is optional, but must be unique for the program if specified. If you assign a name, the name can then be used by other package procedures, including the SET_JOB_ARGUMENT_VALUE Procedure.
argument_type	The datatype of the argument being defined. This is not verified or used by the Scheduler. The program user uses argument_type when deciding what value to assign to the argument. Any valid SQL datatype is allowed.
default_value	The default value to be assigned to the argument if none is specified by the job.
out_argument	This parameter is reserved for future use. It must be set to FALSE.

Usage Notes

All program arguments from 1 to the number_of_arguments value must be defined before a program can be enabled. If a default value for an argument is not defined with this procedure, a value must be defined in the job.

Defining a program argument requires that you be the owner of the program or have ALTER privileges on that program. You can also define a program argument if you have the CREATE ANY JOB privilege.

DEFINE_PROGRAM_ARGUMENT only supports arguments of SQL type. Therefore, argument values that are not of SQL type, such as booleans, are not supported as program or job arguments.

See Also:

- "DEFINE_ANYDATA_ARGUMENT Procedure"
- "SET_JOB_ARGUMENT_VALUE Procedure"

169.6.29 DISABLE Procedure

This procedure disables a program, job, chain, window, database destination, external destination, file watcher, or group. When an object is disabled, its `enabled` attribute is set to `FALSE`.

Syntax

```
DBMS_SCHEDULER.DISABLE (
   name              IN VARCHAR2,
   force             IN BOOLEAN DEFAULT FALSE,
   commit_semantics  IN VARCHAR2 DEFAULT 'STOP_ON_FIRST_ERROR');
```

Parameters

Table 169-42 DISABLE Procedure Parameters

Parameter	Description
`name`	The name of the object being disabled. Can be a comma-delimited list. If a job class name is specified, then all the jobs in the job class are disabled. The job class is not disabled. If a group name is specified, then the group is disabled, but the enabled state of the group members is unaffected.
`force`	If `TRUE`, objects are disabled even if other objects depend on them. See the usage notes for more information.
`commit_semantics`	The commit semantics. The following types are supported: • `STOP_ON_FIRST_ERROR`: The procedure returns on the first error and the previous disable operations that were successful are committed to disk. This is the default. • `TRANSACTIONAL`: The procedure returns on the first error and everything that happened before that error is rolled back. This type is only supported when disabling a job or a list of jobs. In addition, this type is not supported when `force` is set to `TRUE`. • `ABSORB_ERRORS`: The procedure tries to absorb any errors and disable the rest of the jobs and commits all the disable operations that were successful. If errors occur, you can query the view `SCHEDULER_BATCH_ERRORS` for details. This type is only supported when disabling a job or a list of jobs.

Usage Notes

Windows must be preceded by `SYS`.

Disabling an object that is already disabled does not generate an error.

The purpose of the `force` option is to point out dependencies. No dependent objects are altered.

To run `DISABLE` for a window or a group of type `WINDOW`, you must have the `MANAGE SCHEDULER` privilege.

You can use `DISABLE` with any schema except the `SYS` schema.

Jobs

Disabling a job means that, although the metadata of the job is there, it should not run and the job coordinator will not pick up these jobs for processing. When a job is disabled, its `state` in the job queue is changed to `disabled`.

If `force` is set to `FALSE` and the job is currently running, an error is returned.

If `force` is set to `TRUE`, the job is disabled, but the currently running instance is allowed to finish.

For jobs with multiple destinations, you cannot disable a child job at a specific destination. Instead, you can disable the destination.

Programs

When a program is disabled, the status is changed to disabled. A disabled program implies that, although the metadata is still there, jobs that point to this program cannot run.

If `force` is set to `FALSE`, the program must not be referenced by any job, otherwise an error will occur.

If `force` is set to `TRUE`, those jobs that point to the program will not be disabled, however, they will fail at runtime because their program will not be valid.

Running jobs that point to the program are not affected by the `DISABLE` call and are allowed to continue

No arguments that pertain to the program are affected when the program is disabled.

File Watchers

If `force` is set to `FALSE`, the file watcher must not be referenced by any job, otherwise an error will occur. If you force disabling a file watcher, jobs that depend on it become disabled.

Windows

This means that the window will not open, however, the metadata of the window is still there, so it can be reenabled.

If `force` is set to `FALSE`, the window must not be open or referenced by any job otherwise an error occurs.

If `force` is set to `TRUE`, disabling a window that is open will succeed but the window will not be closed. It will prevent the window from opening in the future until it is reenabled.

When the window is disabled, those jobs that have the window as their schedule will not be disabled.

Window Groups

When a group of type `WINDOW` is disabled, jobs (other than a running job) that have the window group as their schedule will not run when the member windows open. However, a job that has one of the window group members as its schedule still runs.

The metadata of the window group is still there, so it can be reenabled. Note that the members of the window group will still open.

If `force` is set to `FALSE`, the window group must not have any members that are open or referenced by any job, otherwise an error will occur.

If `force` is set to `TRUE`:

- The window group is disabled and the open window will be not closed or disabled. It will be allowed to continue to its end.

- The window group is disabled but those jobs that have the window group as their schedule will not be disabled.

Job Chains

When a chain is disabled, the metadata for the chain is still there, but jobs that point to it will not be able to be run. This allows changes to the chain to be made safely without the risk of having an incompletely specified chain run.If `force` is set to `FALSE`, the chain must not be referenced by any job, otherwise an error will occur.If `force` is set to `TRUE`, those jobs that point to the chain will not be disabled, however, they will fail at runtime.Running jobs that point to this chain are not affected by the `DISABLE` call and are allowed to complete.

Database Destinations

When you disable a database destination:

- The destination is skipped when a multiple destination job runs.
- If all destinations are disabled for a job, the Scheduler generates an error when it attempts to run the job.
- The `REFS_ENABLED` column in `*_SCHEDULER_JOB_DESTS` is set to `FALSE` for all jobs that reference the database destination.

External Destinations

When you disable an external destination:

- Dependent database destinations remain enabled, but the Scheduler generates an error when it attempts to run a job with a database destination that depends on the external destination.
- The `REFS_ENABLED` column in `*_SCHEDULER_JOB_DESTS` is set to `FALSE` for all external jobs that reference the external destination and for all database jobs with a database destination that depends on the external destination.

Groups

If you disable an external destination group or database destination group, the Scheduler generates an error when it attempts to run a job that names the group as its destination.

169.6.30 DROP_AGENT_DESTINATION Procedure

This procedure drops one or more external destinations, also known as agent destinations. It should be used only when the preferred method of dropping an external destination, using the `schagent` utility to unregister a Scheduler agent with a database, is unavailable due to failures.

This procedure can be called only by the `SYS` user or a user with the `MANAGE SCHEDULER` privilege.

> **Note:**
>
> External destinations are created on a source database only implicitly by registering an agent with the database. There is no user-callable `CREATE_AGENT_DESTINATION` procedure.

Syntax

```
DBMS_SCHEDULER.DROP_AGENT_DESTINATION (
   destination_name        IN VARCHAR2);
```

Parameters

Table 169-43 DROP_AGENT_DESTINATION Procedure Parameters

Parameter	Description
destination_name	A comma-separated list of external destinations to drop. Because user SYS owns all external destinations, do not prefix them with a schema name. The procedure stops processing if it encounters an external destination that does not exist. All external destinations processed before the error are dropped. Cannot be NULL.

Usage Notes

When an external destination is dropped:

- All database destinations that refer to the external destination are disabled and their agent attribute is set to NULL.
- Members of external destination groups that refer to the destination are removed from the group.
- All job instances in the *_SCHEDULER_JOB_DESTS views that refer to the external destination are also dropped.
- Jobs running against the destination are stopped.

169.6.31 DROP_CHAIN Procedure

This procedure drops an existing chain.

Syntax

```
DBMS_SCHEDULER.DROP_CHAIN (
   chain_name              IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-44 DROP_CHAIN Procedure Parameters

Parameter	Description
chain_name	The name of the chain to drop. Can also be a comma-delimited list of chains.
force	If force is set to FALSE, the chain must not be referenced by any job, otherwise an error will occur. If force is set to TRUE, all jobs pointing to the chain are disabled before the chain is dropped.Running jobs that point to this chain are stopped before the chain is dropped.

Usage Notes

Dropping a chain requires alter privileges on the chain (either as the owner, or a user with ALTER privileges on the chain or the CREATE ANY JOB system privilege).

All steps associated with the chain are dropped. If no rule set was specified when the chain was created, then the automatically created rule set and evaluation context associated with the chain are also dropped, so the user must have the privileges required to do this. See the DBMS_RULE_ADM.DROP_RULE_SET and DBMS_RULE_ADM.DROP_EVALUATION_CONTEXT procedures for more information.

If force is FALSE, no jobs may be using this chain. If force is TRUE, any jobs that use this chain are disabled before the chain is dropped (and any of these jobs that are running will be stopped).

169.6.32 DROP_CHAIN_RULE Procedure

This procedure removes a rule from an existing chain. The rule object corresponding to this rule will also be dropped. The chain will not be disabled. If dropping this rule makes the chain invalid, the user should first disable the chain to ensure that it does not run.

Syntax

```
DBMS_SCHEDULER.DROP_CHAIN_RULE (
   chain_name              IN VARCHAR2,
   rule_name               IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-45 DROP_CHAIN_RULE Procedure Parameters

Parameter	Description
chain_name	The name of the chain to alter
rule_name	The name of the rule to drop
force	If force is set to TRUE, the drop operation proceeds even if the chain is currently running. The running chain is not stopped or interrupted. If force is set to FALSE and the chain is running, an error is generated.

Usage Notes

Dropping a chain rule requires alter privileges on the chain (either as the owner or as a user with ALTER privileges on the chain or the CREATE ANY JOB system privilege).

Dropping a chain rule also drops the underlying rule database object so you must have the privileges to drop this rule object. See the DBMS_RULE_ADM.DROP_RULE procedure for more information.

169.6.33 DROP_CHAIN_STEP Procedure

This procedure drops a chain step. If this chain step is still used in the chain rules, the chain will be disabled.

Syntax

```
DBMS_SCHEDULER.DROP_CHAIN_STEP (
   chain_name              IN VARCHAR2,
   step_name               IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-46 DROP_CHAIN_STEP Procedure Parameters

Parameter	Description
`chain_name`	The name of the chain to alter
`step_name`	The name of the step being dropped. Can be a comma-separated list.
`force`	If `force` is set to `TRUE`, this succeeds even if this chain is currently running. The running chain will not be stopped or interrupted.If `force` is set to `FALSE` and this chain is currently running, an error is thrown.

Usage Notes

Dropping a chain step requires `ALTER` privileges on the chain (either as the owner or as a user with `ALTER` privileges on the chain or the `CREATE ANY JOB` system privilege).

169.6.34 DROP_CREDENTIAL Procedure

This deprecated procedure drops a credential.

> **Note:**
>
> This procedure is deprecated with Oracle Database 12*c* Release 1 (12.1). While the procedure remains available in this package, for reasons of backward compatibility, Oracle recommends using the alternative enhanced functionality provided in the DBMS_CREDENTIAL package, specifically the DROP_CREDENTIAL Procedure.

Syntax

```
DBMS_SCHEDULER.DROP_CREDENTIAL (
   credential_name         IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-47 DROP_CREDENTIAL Procedure Parameters

Parameter	Description
credential_name	The name of the credential being dropped. This can optionally be prefixed with a schema name. This cannot be set to NULL.
force	If set to FALSE, the credential must not be referenced by any job, or an error will occur. If set to TRUE, the credential is dropped whether or not there are jobs referencing it. Jobs that reference the credential will continue to point to a nonexistent credential and throw an error at runtime.

Usage Notes

Only the owner of a credential or a user with the CREATE ANY JOB system privilege may drop the credential.

Running jobs that point to the credential are not affected by this procedure and are allowed to continue.

169.6.35 DROP_DATABASE_DESTINATION Procedure

This procedure drops one or more database destinations.

Syntax

```
DBMS_SCHEDULER.DROP_DATABASE_DESTINATION (
   destination_name        IN VARCHAR2);
```

Parameters

Table 169-48 DROP_DATABASE_DESTINATION Procedure Parameters

Parameter	Description
destination_name	The name of the destination to drop. Can be a comma-separated list of database destinations to drop. Each database destination can optionally be prefixed with a schema name. The procedure stops processing if it encounters a database destination that does not exist. All database destinations processed before the error are dropped. Cannot be NULL.

Usage Notes

Only the owner or a user with the CREATE ANY JOB system privilege may drop the database destination.

When a database destination is dropped:

- All job instances that refer to the destination in the *_SCHEDULER_JOB_DESTS views are also dropped.
- Jobs running against the destination are stopped.

- Members of database destination groups that refer to the destination are removed from the group.

> **See Also:**
>
> CREATE_DATABASE_DESTINATION Procedure

169.6.36 DROP_FILE_WATCHER Procedure

This procedure drops one or more file watchers.

Syntax

```
DBMS_SCHEDULER.DROP_FILE_WATCHER (
   file_watcher_name       IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-49 DROP_FILE_WATCHER Procedure Parameters

Parameter	Description
`file_watcher_name`	The file watcher to drop. Can be a comma-separated list of file watchers. Each file watcher name can optionally be prefixed with a schema name. Cannot be `NULL`.
`force`	If set to `FALSE`, the file watcher must not be referenced by any job, or an error occurs. If set to `TRUE`, the file watcher is dropped whether or not there are jobs referencing it. In this case, jobs that reference the dropped file watcher are disabled.

Usage Notes

Only the owner of a file watcher or a user with the `CREATE ANY JOB` system privilege may drop the file watcher.

Running jobs that point to the file watcher are not affected by this procedure and are allowed to continue.

> **See Also:**
>
> "CREATE_FILE_WATCHER Procedure"

169.6.37 DROP_GROUP Procedure

This procedure drops one or more groups.

Syntax

```
DBMS_SCHEDULER.DROP_GROUP (
   group_name       IN VARCHAR2,
   force            IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-50 DROP_GROUP Procedure Parameters

Parameter	Description
`group_name`	A group to drop. Can be a comma-separated list of group names. Each group name can optionally be prefixed with a schema name. The procedure stops processing if it encounters a group that does not exist. All groups processed before the error are dropped. Cannot be `NULL`.
`force`	If `FALSE`, the group must not be referenced by any job, otherwise an error occurs. If `TRUE`, the group is dropped whether or not there are jobs referencing it. In this case, all jobs referencing the group are disabled and all job instances that reference the group are removed from the `*_SCHEDULER_JOB_DESTS` views.

Usage Notes

Only the owner or a user with the `CREATE ANY JOB` system privilege may drop a group. You must have the `MANAGE SCHEDULER` privilege to drop a group of type `WINDOW`.

> **See Also:**
>
> "CREATE_FILE_WATCHER Procedure"

169.6.38 DROP_INCOMPATIBILITY Procedure

This procedure drops an existing incompatibility definition.

Syntax

```
DBMS_SCHEDULER.DROP_INCOMPATIBILITY (
   incompatibility_name    IN VARCHAR2);
```

Parameters

Table 169-51 DROP_INCOMPATIBILITY Procedure Parameters

Parameter	Description
incompatibility_name	The name of the incompatibility definition.

Usage Notes

> **See Also:**
>
> Using Incompatibility Definitions in *Oracle Database Administrator's Guide*

169.6.39 DROP_JOB Procedure

This procedure drops one or more jobs or all jobs in one or more job classes. Dropping a job also drops all argument values set for that job.

Syntax

```
DBMS_SCHEDULER.DROP_JOB (
   job_name                IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE,
   defer                   IN BOOLEAN DEFAULT FALSE,
   commit_semantics        IN VARCHAR2 DEFAULT 'STOP_ON_FIRST_ERROR');
```

Parameters

Table 169-52 DROP_JOB Procedure Parameters

Parameter	Description
job_name	The name of a job or job class. Can be a comma-delimited list. For a job class, the SYS schema should be specified. If the name of a job class is specified, the jobs that belong to that job class are dropped, but the job class itself is not dropped.
force	If force is set to TRUE, the Scheduler first attempts to stop the running job instances (by issuing the STOP_JOB call with the force flag set to false), and then drops the jobs.
defer	If defer is set to TRUE, the Scheduler allows the running jobs to complete and then drops the jobs.

Table 169-52 (Cont.) DROP_JOB Procedure Parameters

Parameter	Description
commit_semantics	The commit semantics. The following types are supported: • STOP_ON_FIRST_ERROR returns on the first error and previous successful drop operations are committed to disk. This is the default. • TRANSACTIONAL returns on the first error. Everything that happened before that error is rolled back. This type is not supported when force is set to TRUE. • ABSORB_ERRORS tries to absorb any errors and drop the rest of the jobs, and commits all the successful drops. If errors occur, you can query the view SCHEDULER_BATCH_ERRORS for details. Only STOP_ON_FIRST_ERROR is permitted when job classes are included in the job_name list.

Usage Notes

If both force and defer are set to FALSE and a job is running at the time of the call, the attempt to drop that job fails. The entire call to DROP_JOB may then fail, depending on the setting of commit_semantics.

Setting both force and defer to TRUE results in an error.

Dropping a job requires ALTER privileges on the job either as the owner of the job or as a user with the ALTER object privilege on the job or the CREATE ANY JOB system privilege.

169.6.40 DROP_JOB_CLASS Procedure

This procedure drops a job class. Dropping a job class means that all the metadata about the job class is removed from the database.

Syntax

```
DBMS_SCHEDULER.DROP_JOB_CLASS (
   job_class_name          IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-53 DROP_JOB_CLASS Procedure Parameters

Parameter	Description
job_class_name	The name of the job class. Can be a comma-delimited list.
force	If force is set to FALSE, a class being dropped must not be referenced by any jobs, otherwise an error occurs. If force is set to TRUE, jobs belonging to the class are disabled and their class is set to the default class. Only if this is successful is the class dropped. Running jobs that belong to the job class are not affected.

Usage Notes

Dropping a job class requires the MANAGE SCHEDULER system privilege.

169.6.41 DROP_PROGRAM Procedure

This procedure drops a program. Any arguments that pertain to the program are also dropped when the program is dropped.

Syntax

```
DBMS_SCHEDULER.DROP_PROGRAM (
   program_name            IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-54 DROP_PROGRAM Procedure Parameters

Parameter	Description
program_name	The name of the program to be dropped. Can be a comma-delimited list.
force	If force is set to FALSE, the program must not be referenced by any job, otherwise an error occurs. If force is set to TRUE, all jobs referencing the program are disabled before the program is dropped. Running jobs that point to the program are not affected by the DROP_PROGRAM call and are allowed to continue.

Usage Notes

Dropping a program requires that you be the owner of the program or have ALTER privileges on that program. You can also drop a program if you have the CREATE ANY JOB privilege.

169.6.42 DROP_PROGRAM_ARGUMENT Procedure

This procedure drops a program argument. An argument can be specified by either name (if one has been given) or position.

The procedure is overloaded.

Syntax

Drops a program argument by position:

```
DBMS_SCHEDULER.DROP_PROGRAM_ARGUMENT (
   program_name            IN VARCHAR2,
   argument_position       IN PLS_INTEGER);
```

Drops a program argument by name:

```
DBMS_SCHEDULER.DROP_PROGRAM_ARGUMENT (
   program_name            IN VARCHAR2,
   argument_name           IN VARCHAR2);
```

Parameters

Table 169-55 DROP_PROGRAM_ARGUMENT Procedure Parameters

Parameter	Description
program_name	The name of the program to be altered. A program with this name must exist.
argument_name	The name of the argument being dropped
argument_position	The position of the argument to be dropped

Usage Notes

Dropping a program argument requires that you be the owner of the program or have ALTER privileges on that program. You can also drop a program argument if you have the CREATE ANY JOB privilege.

169.6.43 DROP_RESOURCE Procedure

This procedure drops a resource.

Syntax

```
DBMS_SCHEDULER.DROP_RESOURCE (
   resource_name  IN VARCHAR2,
   force          IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-56 DROP_RESOURCE Procedure Parameters

Parameter	Description
resource_name	The name of the resource to be dropped. Can be a comma-delimited list.
force	If force is set to FALSE, the resource must not have any existing constraints, otherwise an error occurs. If force is set to TRUE, the resource will be dropped and any constraints defined on this resource will also be dropped.

Usage Notes

Only the owner or a user with the CREATE ANY JOB system privilege may drop the resource.

> **See Also:**
>
> Creating or Dropping a Resource in *Oracle Database Administrator's Guide*

169.6.44 DROP_SCHEDULE Procedure

This procedure drops a schedule.

Syntax

```
DBMS_SCHEDULER.DROP_SCHEDULE (
   schedule_name    IN VARCHAR2,
   force            IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-57 DROP_SCHEDULE Procedure Parameters

Parameter	Description
schedule_name	The name of the schedule. Can be a comma-delimited list.
force	If force is set to FALSE, the schedule must not be referenced by any job or window, otherwise an error will occur. If force is set to TRUE, any jobs or windows that use this schedule are disabled before the schedule is dropped Running jobs and open windows that point to the schedule are not affected.

Usage Notes

You must be the owner of the schedule being dropped or have ALTER privileges for the schedule or the CREATE ANY JOB privilege.

169.6.45 DROP_WINDOW Procedure

This procedure drops a window. All metadata about the window is removed from the database. The window is removed from any groups that reference it.

Syntax

```
DBMS_SCHEDULER.DROP_WINDOW (
   window_name             IN VARCHAR2,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-58 DROP_WINDOW Procedure Parameters

Parameter	Description
window_name	The name of the window. Can be a comma-delimited list.

Table 169-58 (Cont.) DROP_WINDOW Procedure Parameters

Parameter	Description
force	If force is set to FALSE, the window must be not be open or referenced by any job, otherwise an error occurs. If force is set to TRUE, the window is dropped and those jobs that have the window as their schedule are disabled. However, jobs that have a window group, of which the dropped window is a member, as their schedule, are not disabled. If the window is open then, the Scheduler attempts to first close the window and then drop it. When the window is closed, normal close window rules apply. Running jobs that have the window as their schedule is allowed to continue, unless the stop_on_window_close flag is set to TRUE for the job. If this is the case, the job is stopped when the window is dropped.

Usage Notes

Dropping a window requires the MANAGE SCHEDULER privilege.

169.6.46 DUMP_IN_MEMORY_TRACE Procedure

This procedure dumps the scheduler in-memory trace buffer of the specified process state object address into the current trace file of the requester process.

Syntax

```
DBMS_SCHEDULER.DUMP_IN_MEMORY_TRACE (
   PROCESS_ADDRESS IN RAW);
```

Parameters

Table 169-59 DUMP_IN_MEMORY_TRACE Procedure Parameters

Parameter	Description
PROCESS_ADDRESS	State object address of the process being dumped.

Usage Notes

Either connecting as a SYS user or having the DBA role is required to execute dump_in_memory_trace.

169.6.47 ENABLE Procedure

This procedure enables a program, job, chain, window, database destination, external destination, file watcher, or group.

When an object is enabled, its enabled attribute is set to TRUE. By default, jobs, chains, and programs are created disabled and database destinations, external destinations, file watchers, windows, and groups are created enabled.

If a job was disabled and you enable it, the Scheduler begins to automatically run the job according to its schedule. Enabling a disabled job also resets the job RUN_COUNT, FAILURE_COUNT and RETRY_COUNT columns in the *_SCHEDULER_JOBS data dictionary views.

Validity checks are performed before enabling an object. If the check fails, the object is not enabled, and an appropriate error is returned. This procedure does not return an error if the object was already enabled.

Syntax

```
DBMS_SCHEDULER.ENABLE (
   name              IN VARCHAR2,
   commit_semantics  IN VARCHAR2 DEFAULT 'STOP_ON_FIRST_ERROR');
```

Parameters

Table 169-60 ENABLE Procedure Parameters

Parameter	Description
name	The name of the Scheduler object being enabled. Can be a comma-delimited list of names. If a job class name is specified, then all the jobs in the job class are enabled. If a group name is specified, then the group is enabled, but the enabled state of the group members is unaffected.
commit_semantics	The commit semantics. The following types are supported: • STOP_ON_FIRST_ERROR - The procedure returns on the first error and previous successful enable operations are committed to disk. This is the default. • TRANSACTIONAL - The procedure returns on the first error and everything that happened before that error is rolled back. This type is only supported when enabling a job or a list of jobs. • ABSORB_ERRORS - The procedure tries to absorb any errors and enable the rest of the jobs. It commits all the enable operations that were successful. If errors occur, you can query the view SCHEDULER_BATCH_ERRORS for details. This type is only supported when enabling a job or a list of jobs.

Usage Notes

Window names must be preceded by SYS.

To run ENABLE for a window or group of type WINDOW, you must have the MANAGE SCHEDULER privilege. For a job of type EXECUTABLE (or for a job that points to a program of type EXECUTABLE), the job owner must have the CREATE EXTERNAL JOB system privilege before the job can be enabled or run.

To enable a file watcher, the file watcher owner must have the EXECUTE privilege on the designated credential.

You can use ENABLE with any schema except the SYS schema.

169.6.48 END_DETACHED_JOB_RUN Procedure

This procedure ends a detached job run. A detached job points to a detached program, which is a program with the detached attribute set to TRUE.

A detached job run does not end until this procedure or the STOP_JOB Procedure is called.

Syntax

```
DBMS_SCHEDULER.END_DETACHED_JOB_RUN (
   job_name          IN VARCHAR2,
   error_number      IN PLS_INTEGER DEFAULT 0,
   additional_info   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-61 END_DETACHED_JOB_RUN Procedure Parameters

Parameter	Description
job_name	The name of the job to end. Must be a detached job that is running.
error_number	If zero, then the job run is logged as succeeded. If -1013, then the job run is logged as stopped. If any other number, then the job run is logged as failed with that error number.
additional_info	This text is stored in the additional_info column of the *_scheduler_job_run_details views for this job run.

Usage Notes

This procedure requires that you either own the job or have ALTER privileges on it. You can also end any detached job run if you have the CREATE ANY JOB privilege.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for information about detached jobs.

169.6.49 EVALUATE_CALENDAR_STRING Procedure

You can define repeat intervals of jobs, windows or schedules using the Scheduler calendaring syntax. This procedure evaluates the calendar expression and tells you the next execution date and time of a job or window. This is very useful for testing the correct definition of the calendar string without actually scheduling the job or window.

This procedure can also get multiple steps of the repeat interval by passing the next_run_date returned by one invocation as the return_date_after argument of the next invocation.

See the calendaring syntax described in "Operational Notes".

Syntax

```
DBMS_SCHEDULER.EVALUATE_CALENDAR_STRING (
   calendar_string    IN  VARCHAR2,
   start_date         IN  TIMESTAMP WITH TIME ZONE,
   return_date_after  IN  TIMESTAMP WITH TIME ZONE,
   next_run_date      OUT TIMESTAMP WITH TIME ZONE);
```

Parameters

Table 169-62 EVALUATE_CALENDAR_STRING Procedure Parameters

Parameter	Description
calendar_string	The calendar expression to be evaluated. The string must be in the calendaring syntax described in "Operational Notes".
start_date	The date and time after which the repeat interval becomes valid. It can also be used to fill in specific items that are missing from the calendar string. Can optionally be NULL.
return_date_after	The return_date_after argument helps the Scheduler determine which one of all possible matches (all valid execution dates) to return from those determined by the start_date and the calendar string. When a NULL value is passed for this argument, the Scheduler automatically fills in systimestamp as its value.
next_run_date	The first timestamp that matches the calendar string and start date that occur after the value passed in for the return_date_after argument.

Examples

The following code fragment can be used to determine the next five dates a job will run given a specific calendar string.

```
SET SERVEROUTPUT ON;
ALTER SESSION set NLS_DATE_FORMAT = 'DD-MON-YYYY HH24:MI:SS';
Session altered.

DECLARE
start_date        TIMESTAMP;
return_date_after TIMESTAMP;
next_run_date     TIMESTAMP;
BEGIN
start_date :=
  to_timestamp_tz('01-JAN-2003 10:00:00','DD-MON-YYYY HH24:MI:SS');
return_date_after := start_date;
FOR i IN 1..5 LOOP
  DBMS_SCHEDULER.EVALUATE_CALENDAR_STRING(
    'FREQ=DAILY;BYHOUR=9;BYMINUTE=30;BYDAY=MON,TUE,WED,THU,FRI',
    start_date, return_date_after, next_run_date);
DBMS_OUTPUT.PUT_LINE('next_run_date: ' || next_run_date);
return_date_after := next_run_date;
END LOOP;
END;
/

next_run_date: 02-JAN-03 09.30.00.000000 AM
next_run_date: 03-JAN-03 09.30.00.000000 AM
next_run_date: 06-JAN-03 09.30.00.000000 AM
next_run_date: 07-JAN-03 09.30.00.000000 AM
next_run_date: 08-JAN-03 09.30.00.000000 AM

PL/SQL procedure successfully completed.
```

Usage Notes

No specific Scheduler privileges are required.

169.6.50 EVALUATE_RUNNING_CHAIN Procedure

This procedure forces reevaluation of the rules of a running chain to trigger any rules for which the conditions have been satisfied. The job passed as an argument must point to a chain and must be running. If the job is not running, an error is thrown. (`RUN_JOB` can be used to start the job.)

If any of the steps of the chain are themselves running chains, another `EVALUATE_RUNNING_CHAIN` is performed on each of the nested running chains.

Syntax

```
DBMS_SCHEDULER.EVALUATE_RUNNING_CHAIN (
   job_name              IN VARCHAR2);
```

Parameters

Table 169-63 EVALUATE_RUNNING_CHAIN Procedure Parameter

Parameter	Description
`job_name`	The name of the running job (pointing to a chain) to reevaluate the rules for

Usage Notes

Running `EVALUATE_RUNNING_CHAIN` on a job requires alter privileges on the job (either as the owner, or as a user with `ALTER` privileges on the job or the `CREATE ANY JOB` system privilege).

Note:

The Scheduler automatically evaluates a chain:

- At the start of the chain job
- When a chain step completes
- When an event occurs that is associated with an event step in the chain

For most chains, this is sufficient. `EVALUATE_RUNNING_CHAIN` should be used only under the following circumstances:

- After manual intervention of a running chain with the `ALTER_RUNNING_CHAIN` procedure
- When chain rules use SQL syntax and the rule conditions contain elements that are not under the control of the Scheduler.

In these cases, `EVALUATE_RUNNING_CHAIN` may not be needed if you set the `evaluation_interval` attribute when you created the chain.

169.6.51 GENERATE_JOB_NAME Function

This function returns a unique name for a job.

The name will be of the form `{prefix}N` where `N` is a number from a sequence. If no prefix is specified, the generated name will, by default, be `JOB$_1`, `JOB$_2`, `JOB$_3`, and so on. If `'SCOTT'` is specified as the prefix, the name will be `SCOTT1`, `SCOTT2`, and so on.

Syntax

```
DBMS_SCHEDULER.GENERATE_JOB_NAME (
   prefix        IN VARCHAR2 DEFAULT 'JOB$_') RETURN VARCHAR2;
```

Parameters

Table 169-64 GENERATE_JOB_NAME Function Parameter

Parameter	Description
`prefix`	The prefix to use when generating the job name

Usage Notes

If the prefix is explicitly set to `NULL`, the name is just the sequence number. In order to successfully use such numeric names, they must be surrounded by double quotes throughout the `DBMS_SCHEDULER` calls. A prefix cannot be longer than 18 characters and cannot end with a digit.

Note that, even though the `GENERATE_JOB_NAME` function never returns the same job name twice, there is a small chance that the returned name matches an already existing database object.

No specific Scheduler privileges are required to use this function.

169.6.52 GET_AGENT_INFO Function

This function can return job information specific to an agent, such as how many are running and so on, depending on the attribute selected.

Syntax

```
DBMS_SCHEDULER.GET_AGENT_INFO (
   agent_name        IN VARCHAR2,
   attribute         IN VARCHAR2) RETURN VARCHAR2;
```

Parameters

Table 169-65 GET_AGENT_INFO Function Parameter

Parameter	Description
`agent_name`	The name of an external destination where the agent is running

Table 169-65 (Cont.) GET_AGENT_INFO Function Parameter

Parameter	Description
attribute	Possible Attributes values • VERSION:. Returns the agent version number. Requires the CREATE JOB system privilege. • UPTIME: Returns the time the agent has been up and running. Requires the CREATE JOB system privilege. • NUMBER_OF_RUNNING_JOBS: Returns the number of jobs that the agent is currently running. Requires the CREATE JOB system privilege. • TOTAL_JOBS_RUN: Returns the number of jobs run by the agent since it was started. Requires the CREATE JOB system privilege. • RUNNING_JOBS: Returns a comma-separated list of the names of the jobs running currently. Requires the MANAGE SCHEDULER system privilege. • ALL: Returns all the information the previous options return. It requires the MANAGE SCHEDULER system privilege.

Usage Notes

This function returns the same information as the schagent utility status option. See *Oracle Database Administrator's Guide*.

169.6.53 GET_AGENT_VERSION Function

This function returns the version string of a Scheduler agent that is registered with the database and is currently running. GET_AGENT_VERSION throws an error if the agent is not registered with the database or if the agent is not currently running.

Syntax

```
DBMS_SCHEDULER.GET_AGENT_VERSION (
   agent_host        IN VARCHAR2) RETURN VARCHAR2;
```

Parameters

Table 169-66 GET_AGENT_VERSION Function Parameter

Parameter	Description
agent_host	Either the hostname and port on which the agent is running in the form hostname:port or the name of the agent as shown in the destination_name column of the ALL_SCHEDULER_EXTERNAL_DESTS view which lists all Scheduler agents registered with the database.

Usage Notes

This function requires the CREATE EXTERNAL JOB system privilege.

169.6.54 GET_ATTRIBUTE Procedure

This procedure retrieves the value of an attribute of a Scheduler object. It is overloaded to retrieve values of various types.

Syntax

```
DBMS_SCHEDULER.GET_ATTRIBUTE (
   name           IN VARCHAR2,
   attribute      IN VARCHAR2,
   value          OUT {VARCHAR2|PLS_INTEGER|BOOLEAN|DATE|TIMESTAMP|
                       TIMESTAMP WITH TIME ZONE|TIMESTAMP WITH LOCAL TIME ZONE|
                       INTERVAL DAY TO SECOND});

DBMS_SCHEDULER.GET_ATTRIBUTE (
   name           IN VARCHAR2,
   attribute      IN VARCHAR2,
   value          OUT VARCHAR2,
   value2         OUT VARCHAR2);
```

Parameters

Table 169-67 GET_ATTRIBUTE Procedure Parameters

Parameter	Description
`name`	The name of the object
`attribute`	The attribute being retrieved. See the SET_ATTRIBUTE Procedure for tables of attribute values.
`value`	The existing value of the attribute
`value2`	The `value2` argument is for an optional second value. Most attributes have only one value associated with them, but some can have two.

Usage Notes

To run `GET_ATTRIBUTE` for a job class, you must have the `MANAGE SCHEDULER` privilege or have `EXECUTE` privileges on the class. For a schedule, window, or group, no privileges are necessary. Otherwise, you must be the owner of the object or have `ALTER` or `EXECUTE` privileges on that object or have the `CREATE ANY JOB` privilege.

See the SET_ATTRIBUTE Procedure for tables of attribute values that you can retrieve for the various Scheduler object types.

169.6.55 GET_FILE Procedure

This procedure retrieves a file from the operating system file system of a specified host. The file is copied to a destination, or its contents are returned in a procedure output parameter.

You can also use this procedure to retrieve the standard output or error text for a run of an external job that has an associated credential.

This procedures differs from the equivalent `UTL_FILE` procedure in that it uses a credential and can retrieve files from remote hosts that have only a Scheduler agent (and not an Oracle database) installed.

Syntax

```
DBMS_SCHEDULER.GET_FILE (
   source_file                 IN VARCHAR2,
   source_host                 IN VARCHAR2,
   credential_name             IN VARCHAR2,
   file_contents               IN OUT NOCOPY {BLOB|CLOB});

DBMS_SCHEDULER.GET_FILE (
   source_file                 IN VARCHAR2,
   source_host                 IN VARCHAR2,
   credential_name             IN VARCHAR2,
   destination_file_name       IN VARCHAR2,
   destination_directory_object IN VARCHAR2,
   destination_permissions     IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-68 GET_FILE Procedure Parameters

Parameter	Description
source_file	Fully qualified path name of the file to retrieve from the operating system. The file name is case-sensitive and is not converted to uppercase. If the file name starts with a question mark ('?'), the question mark is replaced by the path to the Oracle home if getting a file from the local host, or to the Scheduler agent home if getting a file from a remote host. If the format of this parameter is *external_log_id*_stdout, then the stdout from the designated external job run is returned. If the format of this parameter is *external_log_id*_stderr, the error text from the designated external job run is returned. You obtain the value of *external_log_id* from the ADDITIONAL_INFO column of the *_SCHEDULER_JOB_RUN_DETAILS views. This column contains a set of name/value pairs in an indeterminate order, so you must parse this column for the external_log_id name/value pair, and then append either "_stdout" or "_stderr" to its value. The external job must have an associated credential. The credential_name parameter of GET_FILE must name the same credential that is used by the job, and the source_host parameter must be the same as the destination attribute of the job.
source_host	If the file is to be retrieved from a remote host, then this parameter must be a valid an external destination name. (An external destination is created when you register a remote Scheduler agent with the database. You can view external destination names in the views *_SCHEDULER_EXTERNAL_DESTS.) If source_host is NULL or set to 'localhost', then the file is retrieved from the file system of the local host. To determine the port number of a Scheduler agent, view the schagent.conf file, which is located in the Scheduler agent home directory on the remote host.
credential_name	The name of the credential to use for accessing the file system.
file_contents	The variable into which the file contents is read.
destination_file_name	The file to which the file contents is written.
destination_directory_object	The directory object that specifies the path to the destination file, when destination_file_name is used. The caller must have the necessary privileges on the directory object.

Table 169-68 (Cont.) GET_FILE Procedure Parameters

Parameter	Description
destination_permissions	Reserved for future use

Usage Notes

The caller must have the CREATE EXTERNAL JOB system privilege and have EXECUTE privileges on the credential.

169.6.56 GET_SCHEDULER_ATTRIBUTE Procedure

This procedure retrieves the value of a Scheduler attribute.

Syntax

```
DBMS_SCHEDULER.GET_SCHEDULER_ATTRIBUTE (
   attribute      IN VARCHAR2,
   value          OUT VARCHAR2);
```

Parameters

Table 169-69 GET_SCHEDULER_ATTRIBUTE Procedure Parameters

Parameter	Description
attribute	The name of the attribute
value	The existing value of the attribute

Usage Notes

To run GET_SCHEDULER_ATTRIBUTE, you must have the MANAGE SCHEDULER privilege.

Table 169-70 lists the Scheduler attributes that you can retrieve. For more detail on these attributes, see Table 169-102 and the section "Configuring the Scheduler" in *Oracle Database Administrator's Guide*.

Table 169-70 Scheduler Attributes Retrievable with GET_SCHEDULER_ATTRIBUTE

Scheduler Attribute	Description
current_open_window	Name of the currently open window
default_timezone	Default time zone used by the Scheduler for repeat intervals and windows
email_sender	The default e-mail address of the sender for job state e-mail notifications
email_server	The SMTP server address that the Scheduler uses to send e-mail notifications for job state events. E-mail notifications cannot be sent if this attribute is NULL.
event_expiry_time	Time in seconds before an event generated by the Scheduler and enqueued onto the Scheduler event queue expires. May be NULL.

Table 169-70 (Cont.) Scheduler Attributes Retrievable with GET_SCHEDULER_ATTRIBUTE

Scheduler Attribute	Description
log_history	Retention period in days for job and window logs. The range of valid values is 0 through 1000000.
max_job_slave_processes	This Scheduler attribute is not used.

169.6.57 OPEN_WINDOW Procedure

This procedure manually opens a window, unrelated to its schedule.

The window opens and the resource plan associated with it takes effect immediately, for the duration specified or for the normal duration of the window, if no duration is given. Only an enabled window can be manually opened.

Syntax

```
DBMS_SCHEDULER.OPEN_WINDOW (
   window_name             IN VARCHAR2,
   duration                IN INTERVAL DAY TO SECOND,
   force                   IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 169-71 OPEN_WINDOW Procedure Parameters

Parameter	Description
window_name	The name of the window
duration	The duration of the window. It is of type interval day to second. If it is NULL, then the window opens for the regular duration as specified in the window metadata.
force	If force is set to FALSE, then opening an already open window generates an error. If force is set to TRUE: You can open a window that is already open. The window stays open for the duration specified in the call, from the time the OPEN_WINDOW command was issued. For example: window1 was created with a duration of four hours. It has how been open for two hours. If, at this point, you reopen window1 using the OPEN_WINDOW call and do not specify a duration, then window1 stays open for four hours because it was created with that duration. If you specified a duration of 30 minutes, the window will close in 30 minutes. The Scheduler automatically closes any window that is open at that time, even if it has a higher priority. For the duration of this manually opened window, the Scheduler does not open any other scheduled windows even if they have a higher priority.

Usage Notes

Opening a window manually has no impact on regular scheduled runs of the window. The next open time of the window is not updated and is determined by the regular scheduled opening.

When a window that was manually opened closes, the rules about overlapping windows are applied to determine which other window should be opened at that time if any at all.

If there are jobs running when the window opens, the resources allocated to them might change if there is a switch in resource plan.

If a window fails to switch resource plans because the designated resource plan no longer exists or because resource plan switching by windows is disabled (for example, by using the `ALTER SYSTEM` statement with the `force` option), the failure to switch resource plans is recorded in the window log.

Opening a window requires the `MANAGE SCHEDULER` privilege.

169.6.58 PURGE_LOG Procedure

The `PURGE_LOG` procedure purges rows from the job and window log that were not purged automatically by the scheduler.

By default, the Scheduler automatically purges all rows in the job log and window log that are older than 30 days. The `PURGE_LOG` procedure can be used to purge additional rows from the job and window log.

Rows in the job log table pertaining to the steps of a chain are purged only when the entry for the main chain job is purged (either manually or automatically).

Syntax

```
DBMS_SCHEDULER.PURGE_LOG (
   log_history             IN PLS_INTEGER  DEFAULT 0,
   which_log               IN VARCHAR2     DEFAULT 'JOB_AND_WINDOW_LOG',
   job_name                IN VARCHAR2     DEFAULT NULL);
```

Parameters

Table 169-72 PURGE_LOG Procedure Parameters

Parameter	Description
`log_history`	This specifies how much history (in days) to keep. The valid range is 0 - 1000000. If set to 0, no history is kept.
`which_log`	This specifies the log type. Valid values are: `job_log`, `window_log`, and `job_and_window_log`.
`job_name`	This specifies which job-specific entries must be purged from the jog log. This can be a comma-delimited list of job names and job classes. Whenever `job_name` has a value other than `NULL`, the `which_log` argument implicitly includes the job log.

Usage Notes

This procedure requires the `MANAGE SCHEDULER` privilege.

Examples

The following completely purges all rows from both the job log and the window log:

```
DBMS_SCHEDULER.PURGE_LOG();
```

The following purges all rows from the window log that are older than 5 days:

```
DBMS_SCHEDULER.PURGE_LOG(5, 'window_log');
```

The following purges all rows from the window log that are older than 1 day and all rows from the job log that are related to jobs in `jobclass1` and older than 1 day:

```
DBMS_SCHEDULER.PURGE_LOG(1, 'job_and_window_log', 'sys.jobclass1');
```

169.6.59 PUT_FILE Procedure

This procedure saves a file to the operating system file system of a specified remote host or of the local computer.

It differs from the equivalent `UTL_FILE` procedure in that it uses a credential and can save files to a remote host that has only a Scheduler agent (and not an Oracle Database) installed.

Syntax

```
DBMS_SCHEDULER.PUT_FILE (
   destination_file         IN VARCHAR2,
   destination_host         IN VARCHAR2,
   credential_name          IN VARCHAR2,
   file_contents            IN {BLOB|CLOB},
   destination_permissions  IN VARCHAR2 DEFAULT NULL);

DBMS_SCHEDULER.PUT_FILE (
   destination_file         IN VARCHAR2,
   destination_host         IN VARCHAR2,
   credential_name          IN VARCHAR2,
   source_file_name         IN VARCHAR2,
   source_directory_object  IN VARCHAR2,
   destination_permissions  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-73 PUT_FILE Procedure Parameters

Parameter	Description
`destination_file`	Fully qualified path name of the file to save to the operating system file system. The file name is case-sensitive. If the file name starts with a question mark ('?'), the question mark is replaced by the path to the Oracle home if saving to the local host, or to the Scheduler agent home if saving to a remote host.
`destination_host`	If `NULL` or set to '`localhost`', the file is saved to the file system of the local computer. To save to a remote host, this parameter must be a valid external destination name. (An external destination is created when you register a remote Scheduler agent with the database. You can view external destination names in the views `*_SCHEDULER_EXTERNAL_DESTS`.)
`credential_name`	The name of the credential to use for accessing the destination file system.
`file_contents`	The variable from which the file contents is read.
`source_file_name`	The file from which the file contents is written
`source_directory_object`	The directory object that specifies the path to the source file, when `source_file_name` is used. The caller must have the necessary privileges on the directory object.
`destination_permissions`	Reserved for future use

Usage Notes

The caller must have the `CREATE EXTERNAL JOB` system privilege and have `EXECUTE` privileges on the credential.

169.6.60 REMOVE_EVENT_QUEUE_SUBSCRIBER Procedure

This procedure unsubscribes a user from the Scheduler event queue `SYS.SCHEDULER$_EVENT_QUEUE`.

Syntax

```
DBMS_SCHEDULER.REMOVE_EVENT_QUEUE_SUBSCRIBER (
   subscriber_name          IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-74 REMOVE_EVENT_QUEUE_SUBSCRIBER Procedure Parameters

Parameter	Description
`subscriber_name`	Name of the Oracle Advanced Queuing (AQ) agent to remove the subscription from. If `NULL`, the user name of the calling user is used.

Usage Notes

After the agent is unsubscribed, it is deleted. If the agent does not exist or is not currently subscribed to the Scheduler event queue, an error is raised.

169.6.61 REMOVE_FROM_INCOMPATIBILITY Procedure

This procedure removes jobs or programs from an existing incompatibility definition.

Syntax

```
DBMS_SCHEDULER.REMOVE_FROM_INCOMPATIBILITY (
   incompatibility_name     IN VARCHAR2,
   object_name              IN VARCHAR2);
```

Parameters

Table 169-75 REMOVE_FROM_INCOMPATIBILITY Procedure Parameters

Parameter	Description
`incompatibility_name`	The name of the incompatibility definition.
`object_name`	One or more (comma-separated) programs or jobs

Usage Notes

This procedure does not raise an error if any specified objects do not already exist in the incompatibility definition.

> **See Also:**
>
> Using Incompatibility Definitions in *Oracle Database Administrator's Guide*

169.6.62 REMOVE_GROUP_MEMBER Procedure

This procedure removes one or more members from an existing group.

Syntax

```
DBMS_SCHEDULER.REMOVE_GROUP_MEMBER (
   group_name              IN VARCHAR2,
   member                  IN VARCHAR2);
```

Parameters

Table 169-76 REMOVE_GROUP_MEMBER Procedure Parameters

Parameter	Description
`group_name`	The name of the group.
`member_name`	The name of the member to remove from group. Comma-separated list of members to remove. An error is returned if any of the members is not part of the group. A group of the same type can be named as a member. The Scheduler immediately expands the included group name into its list of members. If the member is a destination, any job instances that run on this destination are removed from the `*_SCHEDULER_JOB_DESTS` views.

Usage Notes

The following users may remove members from a group:

- The group owner
- A user that has been granted the `ALTER` object privilege on the group
- A user with the `CREATE ANY JOB` system privilege

You must have the `MANAGE SCHEDULER` privilege to remove a member from a group of type `WINDOW`.

> **See Also:**
>
> "CREATE_GROUP Procedure"

169.6.63 REMOVE_JOB_EMAIL_NOTIFICATION Procedure

This procedure removes e-mail notifications for a job. You can remove all e-mail notifications or remove notifications only for specified recipients or specified events.

Syntax

```
DBMS_SCHEDULER.REMOVE_JOB_EMAIL_NOTIFICATION (
    job_name             IN VARCHAR2,
    recipients           IN VARCHAR2 DEFAULT NULL,
    events               IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-77 ADD_JOB_EMAIL_NOTIFICATION Procedure Parameters

Parameter	Description
`job_name`	Name of the job to remove e-mail notifications for. Cannot be `NULL`.
`recipients`	E-mail address to remove e-mail notification for. Comma-separated list of e-mail addresses.
`events`	Job state event to remove e-mail notification for. Comma-separate list of job state events.

Usage Notes

When you specify multiple recipients and multiple events, the notification for each specified event is removed for each specified recipient. The procedure ignores any recipients or events that are specified but that were not previously added.

If `recipients` is `NULL`, e-mail notifications for the specified events are removed for all existing recipients. If `events` is `NULL`, notifications for all events are removed for the specified recipients. If both `recipients` and `events` are `NULL`, all e-mail notifications are removed for the job.

For example, if `recipients` is '`jsmith@example.com,rjones@example.com`' and `events` is '`JOB_FAILED,JOB_BROKEN`', then notifications for both the `JOB_FAILED` and `JOB_BROKEN` events are removed for both jsmith and rjones. If `recipients` is `NULL`, then notifications for both the `JOB_FAILED` and `JOB_BROKEN` events are removed for jsmith, rjones, and any other previously defined recipients for these events.

To call this procedure, you must be the job owner or a user with the `CREATE ANY JOB` system privilege or `ALTER` object privilege on the job.

> **See Also:**
>
> "ADD_JOB_EMAIL_NOTIFICATION Procedure"

169.6.64 RESET_JOB_ARGUMENT_VALUE Procedure

This procedure resets (clears) the value previously set to an argument for a job.

`RESET_JOB_ARGUMENT_VALUE` is overloaded.

Syntax

Clears a previously set job argument value by argument position:

```
DBMS_SCHEDULER.RESET_JOB_ARGUMENT_VALUE (
   job_name                IN VARCHAR2,
   argument_position       IN PLS_INTEGER);
```

Clears a previously set job argument value by argument name:

```
DBMS_SCHEDULER.RESET_JOB_ARGUMENT_VALUE (
   job_name                IN VARCHAR2,
   argument_name           IN VARCHAR2);
```

Parameters

Table 169-78 RESET_JOB_ARGUMENT_VALUE Procedure Parameters

Parameter	Description
job_name	The name of the job being altered
argument_position	The position of the program argument being reset
argument_name	The name of the program argument being reset

Usage Notes

If the corresponding program argument has no default value, the job is disabled. Resetting a program argument of a job belonging to another user requires ALTER privileges on that job. Arguments can be specified by position or by name.

RESET_JOB_ARGUMENT_VALUE requires that you be the owner of the job or have ALTER privileges on that job. You can also reset a job argument value if you have the CREATE ANY JOB privilege.

RESET_JOB_ARGUMENT_VALUE only supports arguments of SQL type. Therefore, argument values that are not of SQL type, such as booleans, are not supported as program or job arguments.

169.6.65 RUN_CHAIN Procedure

This procedure immediately runs a chain or part of a chain by creating a run-once job with the job name given.

If no job_name is given, one is generated of the form RUN_CHAIN$_*chainnameN*, where *chainname* is the first 8 characters of the chain name and *N* is an integer.

If a list of start steps is given, only those steps are started when the chain begins running. Steps not in the list that would normally have started are skipped and paused (so that they or the steps after them do not run).

If start_steps is NULL, then the chain starts normally—that is, it performs an initial evaluation to see which steps to start running).

If a list of initial step states is given, the newly created chain job sets every listed step to the state specified for that step before evaluating the chain rules to see which steps to start. (Steps in the list are not started.)

Syntax

Runs a chain, with a list of start steps.

```
DBMS_SCHEDULER.RUN_CHAIN (
   chain_name                IN VARCHAR2,
   start_steps               IN VARCHAR2,
   job_name                  IN VARCHAR2 DEFAULT NULL);
```

Runs a chain, with a list of initial step states.

```
DBMS_SCHEDULER.RUN_CHAIN (
   chain_name               IN VARCHAR2,
   step_state_list          IN SYS.SCHEDULER$_STEP_TYPE_LIST,
   job_name                 IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-79 RUN_CHAIN Procedure Parameters

Parameter	Description
`chain_name`	The name of the chain to run
`job_name`	The name of the job to create to run the chain
`start_steps`	Comma-separated list of the steps to start when the chain starts running
`step_state_list`	List of chain steps with an initial state (`SUCCEEDED` or `FAILED`) to set for each. Set the attributes of `sys.scheduler$_step_type` as follows: `step_name` The name of the step`step_type` 'SUCCEEDED' or 'FAILED *error_number*' where *error_number* is a positive or negative integer.

Usage Notes

Running a chain requires `CREATE JOB` if the job is being created in the user's schema, or `CREATE ANY JOB` otherwise. In addition, the owner of the job being created needs execute privileges on the chain (as the owner of the chain, or as a user with the `EXECUTE` privilege on the chain or the `EXECUTE ANY PROGRAM` system privilege).

Examples

The following example illustrates how to start a chain in the middle by providing the initial state of some chain steps.

```
declare
  initial_step_states sys.scheduler$_step_type_list;
begin
  initial_step_states := sys.scheduler$_step_type_list(
    sys.scheduler$_step_type('step1', 'SUCCEEDED'),
    sys.scheduler$_step_type('step2', 'FAILED 27486'),
    sys.scheduler$_step_type('step3', 'SUCCEEDED'),
    sys.scheduler$_step_type('step5', 'SUCCEEDED'));
  dbms_scheduler.run_chain('my_chain', initial_step_states);
end;
/
```

169.6.66 RUN_JOB Procedure

This procedure runs a job immediately.

If a job is enabled, the Scheduler runs it automatically. It is not necessary to call `RUN_JOB` to run a job according to its schedule. Use `RUN_JOB` to run a job outside of its normal schedule.

Syntax

```
DBMS_SCHEDULER.RUN_JOB (
   job_name                IN VARCHAR2,
   use_current_session     IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 169-80 RUN_JOB Procedure Parameters

Parameter	Description
`job_name`	A job name or a comma-separate list of entries, where each is the name of an existing job, optionally preceded by a schema name and dot separator. If you specify a multiple-destination job, the job runs on all destinations. In this case, the `use_current_session` argument must be `FALSE`.
`use_current_session`	This specifies whether or not the job run should occur in the same session that the procedure was invoked from.The job always runs as the job owner, in the job owner's schema, unless it has credential specified, then the job runs using the user named in the credential. When `use_current_session` is set to `TRUE`: • You can test a job and see any possible errors on the command line. • `state, run_count`, `last_start_date`, `last_run_duration`, and `failure_count` of `*_scheduler_jobs` are not updated. • `RUN_JOB` can be run in parallel with a regularly scheduled job run. When `use_current_session` is set to `FALSE`: • You need to check the job log to find error information. • All relevant fields in `*_scheduler_jobs` are updated. • `RUN_JOB` fails if a regularly scheduled job is running. For jobs that have a specified destination or destination group, or point to chains or programs with the detached attribute set to `TRUE`, `use_current_session` must be `FALSE`.

Usage Notes

Jobs do not have to be enabled. If a job is disabled, the following validity checks are performed before running it:

- The job points to a valid job class.
- The job owner has `EXECUTE` privileges on the job class.
- If a program or chain is referenced, the program/chain exists.
- If a program or chain is referenced, the job owner has privileges to execute the program/chain.
- All argument values have been set (or have defaults).
- The job owner has the `CREATE EXTERNAL JOB` privilege if this is an external job.

A TRUE value for use_current_session is not permitted for the following types of jobs:

- Jobs that specify a destination or destination group in the destination_name attribute
- Jobs that point to chains (chain jobs)
- Jobs that make use of detached programs (detached jobs).

  above bug fix 1261887 6.12.11

When use_current_session is TRUE, the call to RUN_JOB blocks until the job completes. Any errors that occur during the execution of the job are returned as errors to the RUN_JOB procedure.

Using RUN_JOB with use_current_session=TRUE does not update the job state and the job will not appear in *_SCHEDULER_RUNNING_JOBS views.

above bug fix 19185117 9.15.14

When use_current_session is FALSE, RUN_JOB returns immediately, and the job is picked up by the job coordinator process and passed on to a job secondary process for execution. The Scheduler views and logs must be queried for the outcome of the job.

Multiple user sessions can use RUN_JOB in their sessions simultaneously when use_current_session is set to TRUE.

RUN_JOB requires that you own the job or have ALTER privileges on that job. You can also run a job if you have the CREATE ANY JOB privilege.

Example

The following is an example of using RUN_JOB.

```
BEGIN
  DBMS_SCHEDULER.RUN_JOB(
    JOB_NAME            => 'EODJOB, DSS.ETLJOB',
    USE_CURRENT_SESSION => FALSE);
END;
```

169.6.67 SET_AGENT_REGISTRATION_PASS Procedure

This procedure sets the agent registration password for a database.

A Scheduler agent must register with the database before the database can submit jobs to the agent. The agent must provide this password when registering.

Syntax

```
DBMS_SCHEDULER.SET_AGENT_REGISTRATION_PASS (
   registration_password   IN VARCHAR2,
   expiration_date         IN TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   max_uses                IN NUMBER DEFAULT NULL);
```

Parameters

Table 169-81 SET_AGENT_REGISTRATION_PASS Procedure Parameters

Parameter	Description
registration_password	This is the password that remote agents must specify in order to successfully register with the database. If this is NULL, then no agents will be able to register with the database.
expiration_date	If this is set to a non-NULL value, then the registration_password is not valid after this date. After this date, no agents can register with the database. This cannot be set to a date in the past.
max_uses	This is the maximum number of successful registrations that can be performed with this password. After the number of successful registrations has been performed with this password, then no agents can register with the database. This cannot be set to 0 or a negative value. If this is set to NULL, then there will be no limit on the number of successful registrations.

Usage Notes

To prevent abuse, this password can be set to expire after a given date or a maximum number of successful registrations. This procedure will overwrite any password already set. This requires the MANAGE SCHEDULER system privilege.

By default, max_uses is set to NULL, which means that there is no limit to the number of successful registrations.

Oracle recommends that an agent registration password be reset after every agent registration or every known set of agent registrations. Furthermore, Oracle recommends that this password be set to NULL if no new agents are being registered.

169.6.68 SET_ATTRIBUTE Procedure

This procedure modifies an attribute of a Scheduler object. It is overloaded to accept values of various types.

To set an attribute to NULL, use the SET_ATTRIBUTE_NULL procedure. The attributes that can be set depend on the object being altered. All object attributes can be changed, except the object name.

Syntax

```
DBMS_SCHEDULER.SET_ATTRIBUTE (
   name           IN VARCHAR2,
   attribute      IN VARCHAR2,
   value          IN {BOOLEAN|DATE|TIMESTAMP|
                      TIMESTAMP WITH TIME ZONE|TIMESTAMP WITH LOCAL TIME ZONE|
                      INTERVAL DAY TO SECOND});

DBMS_SCHEDULER.SET_ATTRIBUTE (
   name           IN VARCHAR2,
   attribute      IN VARCHAR2,
   value          IN VARCHAR2,
   value2         IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 169-82 SET_ATTRIBUTE Procedure Parameters

Parameter	Description
name	The name of the object.
attribute	See Table 169-84 through Table 169-94.
value	The new value being set for the attribute. This cannot be NULL. To set an attribute value to NULL, use the SET_ATTRIBUTE_NULL procedure.
value2	The value2 argument is for an optional second value. Most attributes have only one value associated with them, but some can have two.

Table 169-83 is a directory of Scheduler object types and tables of attributes for the object types.

These object types can be viewed with Scheduler Data Dictionary Views, listed in *Oracle Database Administrator's Guide*.

Table 169-83 Attribute Tables for Scheduler Object Types

Scheduler Object Type	Table of Attributes
Job	Table 169-84
Program	Table 169-86
Schedule	Table 169-87
File Watcher	Table 169-88
Job Class	Table 169-89
Window	Table 169-90
Chain	Table 169-91
Database Destination	Table 169-92
External Destination	Table 169-93
Group	Table 169-94
Credential	Table 169-95
Resource	Table 169-96

Usage Notes

If an object is altered and it was in the enabled state, the Scheduler first disables it, then makes the change and reenables it. If any errors are encountered during the enable process, the object is not reenabled and an error is generated.

If an object is altered and it was in the disabled state, it remains disabled after it is altered.

To run SET_ATTRIBUTE for a window, a group of type WINDOW, or job class, you must have the MANAGE SCHEDULER privilege. Otherwise, you must be the owner of the object being altered or have ALTER privileges on that object or have the CREATE ANY JOB privilege.

Job

If there is a running instance of the job when the SET_ATTRIBUTE call is made, it is not affected by the call. The change is only affects future runs of the job.

If any of the schedule attributes of a job are altered while the job is running, the time of the next job run is scheduled using the new schedule attributes. Schedule attributes of a job include schedule_name, start_date, end_date, and repeat_interval.

If any of the program attributes of a job are altered while the job is running, the new program attributes take effect the next time the job runs. Program attributes of a job include program_name, job_action, job_type, and number_of_arguments.

If any job argument values are altered while the job is running, the new values take effect the next time the job runs.

Granting the ALTER privilege on a job lets a user alter all attributes of that job except its program attributes (program_name, job_type, job_action, program_action, and number_of_arguments) and does not allow a user to use a PL/SQL expression to specify the schedule for a job.

Oracle recommends that you not alter a job that was automatically created for you by the database. Jobs that were created by the database have the column SYSTEM set to TRUE in job views.

Program

If any currently running jobs use the program that was altered, they continue to run with the program definition prior to the alter. The job runs with the new program definition the next time the job executes.

Schedule

If a schedule is altered, the change does not affect running jobs and open windows that use this schedule. The change only goes into effect the next time the jobs runs or the window opens.

File Watcher

If a file watcher is altered, any currently running event-based jobs started by the file arrival event are not affected. On the local system, the new file watcher attributes take effect the next time that the file watcher checks for the arrival of the file (every ten minutes by default). On remote systems, there may be an additional delay before the new file watcher attributes take effect.

Job Class

With the exception of the default job class, all job classes can be altered. To alter a job class, you must have the MANAGE SCHEDULER privilege.

When a job class is altered, running jobs that belong to the class are not affected. The change only takes effect for jobs that have not started running yet. Job Class names must be preceded by SYS.

Window

When a window is altered, it does not affect an active window. The changes only take effect the next time the window opens.

If there is no current resource plan, when a window with a designated resource plan opens, the Resource Manager activates with that plan. Window names must be preceded by SYS.

Job Attribute Values

Table 169-84 lists attribute values for jobs.

> **Note:**
>
> See the CREATE_JOB procedure and the CREATE_JOBS procedure for more complete descriptions of the attributes in this table.

Table 169-84 Job Attribute Values

Name	Description
allow_runs_in_restricted_mode	If TRUE, the job is permitted to run when the database is in restricted mode, provided that the job owner is permitted to log in during this mode. FALSE by default.
auto_drop	This attribute, if TRUE, causes a job to be automatically dropped after it completes or is automatically disabled. A job is considered completed if: • Its end date (or the end date of the schedule) has passed. • It has run max_runs number of times. max_runs must be set with SET_ATTRIBUTE. • It is not a repeating job and has run once. A job is automatically disabled when it has failed max_failures times. max_failures is also set with SET_ATTRIBUTE. If this attribute is set to FALSE, the jobs are not dropped and their metadata is kept until the job is explicitly dropped with the DROP_JOB procedure. By default, jobs are created with auto_drop set to TRUE.
comments	An optional comment.
connect_credential_name	feature 25230 This attribute may be set to point to a database credential. For a SQL*Plus or backup script job, the credential connects to the database before running the script. For other job types, it is ignored. The job owner must have execute privileges on the credential, otherwise the job fails. Using a connect_credential_name is recommended since it allows the password to be stored securely in a credential in the database rather than in plain view in the job, program action, or script.
credential_name	This attribute specifies the name of the credential object (credential) to use for a remote database job, a remote external job, a local external job, or an event-based job that processes a file arrival event. For local external jobs only, if this attribute is NULL (the default), then a preferred (default) credential is selected. See *Oracle Database Administrator's Guide* for information about preferred credentials for local external jobs.
database_role	This attribute applies when the database participates in an Oracle Data Guard environment. If this attribute is set to 'PRIMARY', the job runs only when the database is in the role of the primary database. If set to 'LOGICAL STANDBY', the job runs only when the database is in the role of a logical standby. The default is 'PRIMARY' when the database is the primary database, and 'LOGICAL STANDBY' when the database is a logical standby. Note: If you want a job to run for all database roles on a particular host, you must create two copies of the job on that host: one with a database_role of 'PRIMARY', and the other with a database_role of 'LOGICAL STANDBY'.

Table 169-84 (Cont.) Job Attribute Values

Name	Description
destination	***** Deprecated in Oracle Database 11*g* Release 2. Use `destination_name` instead.** This attribute specifies a host on which to run a remote external job. It must be set to the host name or IP address of the destination host. It can optionally be followed by a port number, in the following format: *hostname*:*port* This attribute is set to `NULL` by default.
destination_name	The database destination or external destination for the job. Use for remote database jobs and remote external jobs only. For jobs running on the local database or for local external jobs (executables), must be `NULL`. See Table 169-28 for details about this attribute.
end_date	Specifies the date and time after which the job expires and is no longer run. After the `end_date`, if is `TRUE`, the job is dropped. If `auto_drop` is `FALSE`, the job is disabled and the `STATE` of the job is set to `COMPLETED`. If no value for `end_date` is specified, the job repeats forever unless `max_runs` or `max_failures` is set, in which case the job stops when either value is reached. The value for `end_date` must be after the value for `start_date`. If `end_date` is less than `start_date`, then an error will be generated. If `end_date` is the same as `start_date`, then the job will not execute and no error will be generated.
event_spec	This attribute takes two values: the `value` argument specifies the event condition and the `value2` argument specifies the queue specification. For more details, see the descriptions for the `event_condition` and `queue_spec` arguments in the "CREATE_JOB Procedure".
follow_default_timezone	If `TRUE` and if the job `start_date` is null, then when the `default_timezone` scheduler attribute is changed, the Scheduler recomputes the next run date and time for this job so that it is in accordance with the new time zone. For example, if the job was set to run at 02:00 in the previous time zone, it will run at 02:00 in the new time zone. If the job `start_date` is not null, then the time zone for the run date and time for the job is always specified by the time zone of the `start_date`. If `FALSE`, the next start date and time for the job is not recomputed when the `default_timezone` scheduler attribute is changed. In this case, if the old time zone is three hours earlier than the new time zone, then a job scheduled to run at 02:00 in the old time zone runs at 05:00 in the new time zone. Summer and winter transitions do not change the default time zone name.
instance_id	Valid only in an Oracle Real Application Clusters environment. Indicates the instance on which the job is to be run.

Table 169-84 (Cont.) Job Attribute Values

Name	Description
instance_stickiness	This attribute should only be used for a database running in an Oracle Real Application Clusters (Oracle RAC) environment. By default, it is set to TRUE. If you set instance_stickiness to TRUE, jobs start running on the instance with the lightest load and the Scheduler thereafter attempts to run on the instance that it last ran on. If that instance is either down or so overloaded that it does not start new jobs for a significant period of time, another instance runs the job. If the interval between runs is large, instance_stickiness is ignored and the job is handled as if it were a non-sticky job. If instance_stickiness is set to FALSE, each instance of the job runs on the first instance available. For environments other than Oracle RAC, this attribute is not useful because there is only one instance.
job_action	The action that the job performs, depending on the job_type attribute. For example, if job_type is 'STORED_PROCEDURE', job_action contains the name of the stored procedure.
job_class	The class this job is associated with.
job_priority	This attribute specifies the priority of this job relative to other jobs in the same class as this job. If multiple jobs within a class are scheduled to be executed at the same time, the job priority determines the order in which jobs from that class are picked up for execution by the job coordinator. It can be a value from 1 through 5, with 1 being the first to be picked up for job execution. If no job priority is specified when creating a job, the default priority of 3 is assigned to it.
job_type	The type of this job.Valid values are: 'PLSQL_BLOCK', 'STORED_PROCEDURE', 'EXECUTABLE', CHAIN', 'EXTERNAL_SCRIPT', 'SQL_SCRIPT', and 'BACKUP_SCRIPT'. If this is set, program_name must be NULL.
job_weight	*** Deprecated in Oracle Database 11*g*R2. Do not change the value of this attribute from the default, which is 1. Weight of the job for parallel execution.
logging_level	This attribute specifies how much information is logged. The possible options are: DBMS_SCHEDULER.LOGGING_OFF (The default) No logging is performed for this job. However, the logging level of the job class takes precedence and job logging may occur. DBMS_SCHEDULER.LOGGING_FAILED_RUNS The Scheduler logs only jobs that failed, with the reason for failure. If the job class has a higher logging level, then the higher logging level takes precedence. DBMS_SCHEDULER.LOGGING_RUNS The Scheduler writes detailed information to the job log for all runs of each job in this class. If the job class has a higher logging level, then the higher logging level takes precedence. DBMS_SCHEDULER.LOGGING_FULL In addition to recording every run of a job, the Scheduler records all operations performed on the job, including create, enable, disable, alter (with SET_ATTRIBUTE), stop, and so on.

Table 169-84 (Cont.) Job Attribute Values

Name	Description
max_failures	This attribute specifies the number of times a job can fail on consecutive scheduled runs before it is automatically disabled. Once a job is disabled, it is no longer executed and its STATE is set to BROKEN in the *_SCHEDULER_JOB views. max_failures can be an integer between 1 to 1,000,000. By default, it is set to NULL, which indicates that new instances of the job are started regardless of how many previous instances have failed.
max_run_duration	This attribute specifies the maximum amount of time that the job should be allowed to run. Its datatype is INTERVAL DAY TO SECOND. If this attribute is set to a non-zero and non-NULL value, and job duration exceeds this value, the Scheduler raises an event of type JOB_OVER_MAX_DUR. It is then up to your event handler to decide whether or not to allow the job to continue.
max_runs	This attribute specifies the maximum number of consecutive scheduled runs of the job. Once max_runs is reached, the job is disabled and its state is changed to COMPLETED. max_runs can be an integer between 1 and 1,000,000. By default, it is set to NULL, which means that it repeats forever or until end_date or max_failures is reached.
number_of_arguments	The number of arguments if the program is inlined. If this is set, program_name should be NULL.
parallel_instances	This is a boolean attribute that can be set only for event-based jobs. If FALSE (the default), then if an event is raised and the event-based job that processes that event is already running, the new event is ignored. If TRUE, then an instance of the job is started for every instance of the event, and each job instance is a lightweight job so multiple instances of the same event-based job can run in parallel. Each lightweight job takes its attributes (such as action, maximum run duration, and so on) from the definition of the event-based job (its *parent job*). After the lightweight job completes, it is dropped. There is no explicit limit to the number of lightweight jobs that can run simultaneously to process multiple instances of the event. However, limitations may be imposed by available system resources. The lightweight jobs are not visible in any of the *_SCHEDULER_JOBS views. However, they are visible in the *_SCHEDULER_RUNNING_JOBS views. The name of each lightweight job is the same as that of the parent job, and a subname is automatically generated to distinguish each lightweight job from its parent and from its siblings.
program_name	The name of a program object to use with this job. If this is set, job_action, job_type and number_of_arguments should be NULL.

Table 169-84 (Cont.) Job Attribute Values

Name	Description
raise_events	This attribute tells the Scheduler at what stages of the job execution to raise events. It is a bit vector in which zero or more of the following bits can be set. Each bit has a package constant corresponding to it. • job_started CONSTANT PLS_INTEGER := 1 • job_succeeded CONSTANT PLS_INTEGER := 2 • job_failed CONSTANT PLS_INTEGER :=4 • job_broken CONSTANT PLS_INTEGER :=8 • job_completed CONSTANT PLS_INTEGER :=16 • job_stopped CONSTANT PLS_INTEGER :=32 • job_sch_lim_reached CONSTANT PLS_INTEGER :=64 • job_disabled CONSTANT PLS_INTEGER :=128 • job_chain_stalled CONSTANT PLS_INTEGER :=256 • job_all_events CONSTANT PLS_INTEGER := 511 • job_run_completed CONSTANT PLS_INTEGER := job_succeeded + job_failed + job_stopped Table 169-85 describes these event types in detail.
repeat_interval	Either a PL/SQL function returning the next date and time on which to run, or calendaring syntax expression. If this is set, schedule_name should be NULL. See "Calendaring Syntax" for more information.
restartable	This attribute specifies whether or not a job can be restarted in case of failure. By default, jobs are not restartable and this attribute is set to FALSE. Setting this to TRUE means that if a job fails while running, it is restarted from the beginning point of the job. In the case of a chain job, if this attribute is TRUE, the chain is restarted from the beginning after an application failure. If this attribute is FALSE, or if there has been a database failure, the chain is restarted at the last running step. The restart_on_recovery attribute of that step then determines if the step is restarted or marked as stopped. (If marked as stopped, the chain evaluates rules and continues.) Note that setting this attribute to TRUE might lead to data inconsistencies in some situations, for example, if data is committed within a job. Retries on errors are not counted as regular runs. The run count or failure count is not incremented until the job succeeds or has failed all its six retries. The restartable attribute is used by the Scheduler to determine whether to retry the job not only on regular application errors, but after a database malfunction as well. The Scheduler retries the job a maximum of six times. The first time, it waits for one second and multiplies this wait time with a factor of 10 each time thereafter. Both the run count and failure count are incremented by 1 if the job has failed all its six retries. If the job immediately succeeds, or it succeeds on one of its retries, run count is incremented by 1. The Scheduler stops retrying a job when: • One of the retries succeeds. • All of its six retries have failed. • The next retry would occur after the next regularly scheduled run of the job. The Scheduler no longer retries the job if the next scheduled retry is past the next regularly scheduled run for repeating jobs.

Table 169-84 (Cont.) Job Attribute Values

Name	Description
schedule_limit	In heavily loaded systems, jobs are not always started at their scheduled time. This attribute enables you to have the Scheduler not start a job at all if the delay in starting the job is larger than the interval specified. It can be a value of 1 minute to 99 days. For example, if a job was supposed to start at noon and the schedule limit is set to 60 minutes, the job will not be run if it has not started to run by 1:00 p.m. If schedule_limit is not specified, the job is executed at some later date as soon as there are resources available to run it. By default, this attribute is set to null, which indicates that the job can be run at any time after its scheduled time. A scheduled job run that is skipped because of this attribute does not count against the number of runs and failures of the job. An entry in the job log reflects the skipped run.
schedule_name	The name of a schedule, window, or group of type WINDOW to use as the schedule for this job. If this is set, end_date, start_date and repeat_interval should all be NULL.
start_date	The original date and time on which this job started or is scheduled to start. If this is set, schedule_name should be NULL.
stop_on_window_close	This attribute only applies if the schedule of a job is a window or a window group. Setting this attribute to TRUE implies that the job should stop once the associated window is closed. The job is stopped using the stop_job procedure with force set to FALSE. By default, stop_on_window_close is set to FALSE. Therefore, if you do not set this attribute, the job continues after the window closes. Note that, although the job is allowed to continue, its resource allocation will probably change because closing a window generally also implies a change in resource plans.
store_output	This is a boolean attribute. If set to TRUE, then for job runs that are logged, all job output and error messages are stored in the *_JOB_RUN_DETAILS views. If set to FALSE, then the output and messages are not stored. For new jobs, this is set, by default, to TRUE.

The following event types are valid values for the raise_events attribute in Table 169-84.

Table 169-85 Event Types Raised by the Scheduler

Event Type	Description
job_all_events	Not an event, but a constant that provides an easy way for you to enable all events
job_broken	The job has been disabled and has changed to the BROKEN state because it exceeded the number of failures defined by the max_failures job attribute
job_chain_stalled	A job running a chain is in the CHAIN_STALLED state. A running chain becomes stalled if there are no steps running or scheduled to run and the chain evaluation_interval is set to NULL. No progress is made in the chain unless there is manual intervention.
job_completed	The job completed because it reached its max_runs or end_date
job_disabled	The job was disabled by the Scheduler or by a call to SET_ATTRIBUTE
job_failed	The job failed, either due to an error or an unusual termination.

Table 169-85 (Cont.) Event Types Raised by the Scheduler

Event Type	Description
job_over_max_dur	The job exceeded the maximum run duration specified by its max_run_duration attribute. (Note: you do not need to enable this event with the raise_events job attribute; it is always enabled.)
job_run_completed	A job run either failed, succeeded, or was stopped
job_sch_lim_reached	The schedule limit of the job was reached. The job was not started because the delay in starting the job exceeded the value of the schedule_limit job attribute.
job_started	The job started
job_stopped	The job was stopped by a call to STOP_JOB
job_succeeded	The job completed successfully

Program Attribute Values

Table 169-86 lists program attribute values.

> **Note:**
>
> See the "CREATE_PROGRAM Procedure" for more complete descriptions of the attributes in this table.

Table 169-86 Program Attribute Values

Name	Description
comments	An optional comment. This can describe what the program does or give usage details.
detached	If TRUE, the program is a detached program. See *Oracle Database Administrator's Guide* for information about detached jobs and detached programs.
number_of_arguments	The number of arguments required by the stored procedure or other executable that the program invokes
program_action	The action that the program performs, indicated by the program_type attribute. For example, if program_type is 'STORED_PROCEDURE', program_action contains the name of the stored procedure.
program_type	The type of program. This must be one of these supported program types: 'PLSQL_BLOCK', 'STORED_PROCEDURE', and 'EXECUTABLE'.

Schedule Attribute Values

Table 169-87 lists schedule attribute values.

> **Note:**
>
> See "CREATE_SCHEDULE Procedure" for more complete descriptions of the attributes in this table.

Table 169-87 Schedule Attribute Values

Name	Description
comments	An optional comment.
end_date	The cutoff date and time after which the schedule does not specify any dates.
event_spec	This attribute takes two values: the value argument should contain the event condition and the value2 argument should contain the queue specification. For more details, see the descriptions for the event_condition and queue_spec arguments to the "CREATE_JOB Procedure".
repeat_interval	An attribute specifying how often the schedule should repeat, using the calendaring syntax. See "Calendaring Syntax" for more information.
start_date	The start or reference date and time used by the calendaring syntax.

File Watcher Attribute Values

Table 169-88 lists file watcher attribute values.

Table 169-88 File Watcher Attribute Values

Parameter	Description
destination	Remote host name or IP address where the file is expected to arrive. If NULL, destination is the local host.
directory_path	Directory in which the file is expected to arrive. The single wildcard '?' at the beginning of the path denotes the Oracle home path. For example, '?/rdbms/log' denotes the rdbms/log subdirectory of the Oracle home directory.
file_name	Name of the file being looked for. Two wildcards are permitted anywhere in the file name: '?' denotes any single character, and '*' denotes zero or more characters. This attribute cannot be NULL.
credential_name	Name of a valid credential object. The file watcher uses the credential to authenticate itself with the host operating system to access the watched-for file. The file watcher owner must have the EXECUTE privilege on the credential. Cannot be NULL.
min_file_size	Minimum file size in bytes before the file watcher considers the file found. Default is 0.
steady_state_duration	Minimum time interval that the file must remain unchanged before the file watcher considers the file found. If NULL, an internal value is used. The lower limit for this attribute is 10 seconds.
comments	Optional comment.

Job Class Attribute Values

Table 169-89 lists job class attribute values.

> **Note:**
>
> See the "CREATE_JOB_CLASS Procedure" for more complete descriptions of the attributes in this table.

Table 169-89 Job Class Attribute Values

Name	Description
comments	An optional comment about the class.
log_history	This attribute controls the number of days that job log entries for jobs in this class are retained. It helps prevent the job log from growing indiscriminately. The range of valid values is 0 through 1000000. If set to 0, no history is kept. If NULL, retention days are set by the log_history Scheduler attribute (set with SET_SCHEDULER_ATTRIBUTE).
logging_level	This attribute specifies how much information is logged. The valid values are: • DBMS_SCHEDULER.LOGGING_OFF No logging is performed for any jobs in this class. • DBMS_SCHEDULER.LOGGING_FAILED_RUNS The Scheduler logs only jobs in the class that failed, with the reason for failure. • DBMS_SCHEDULER.LOGGING_RUNS The Scheduler writes detailed information to the job log for all runs of each job in this class. This is the default. • DBMS_SCHEDULER.LOGGING_FULL The Scheduler records all operations performed on all jobs in this class, in addition to recording every run of a job. Every time a job is created, enabled, disabled, altered (with SET_ATTRIBUTE), stopped, and so on, an entry is recorded in the log.
resource_consumer_group	The resource consumer group that a class is associated with. All jobs in the class run under this resource consumer group. See *Oracle Database Administrator's Guide* for a description of resource consumer groups and the Database Resource Manager.
service	The database service that the jobs in the job class have affinity to. If both the resource_consumer_group and service attributes are set for a job class, and if the service is mapped to a resource consumer group, the resource_consumer_group attribute takes precedence.

Window Attribute Values

Table 169-90 lists window attribute values.

> **Note:**
>
> See the "CREATE_WINDOW Procedure" for more complete descriptions of the attributes in this table.

Table 169-90 Window Attribute Values

Name	Description
comments	An optional comment about the window.
duration	The duration of the window.
end_date	The date after which the window no longer opens. If this is set, schedule_name must be NULL.
repeat_interval	An attribute specifying how often the schedule should repeat, using the calendaring syntax. PL/SQL date functions are not allowed. If this is set, schedule_name must be NULL. See "Calendaring Syntax" for more information.
resource_plan	The resource plan to be associated with a window. When the window opens, the system switches to this resource plan. When the window closes, the original resource plan is restored. If a resource plan has been made active with the force option, no resource plan switch occurs. Only one resource plan can be associated with a window. It may be NULL or the empty string (""). When it is NULL, the resource plan that is in effect when the window opens stays in effect for the duration of the window. When it is the empty string, the resource manager is disabled for the duration of the window.
schedule_name	The name of a schedule to use with this window. If this is set, start_date, end_date, and repeat_interval must all be NULL.
start_date	The next date and time on which this window is scheduled to open. If this is set, schedule_name must be NULL.
window_priority	The priority of the window. Must be either 'LOW' (default) or 'HIGH'.

Chain Attribute Values

Table 169-91 lists chain attribute values.

> **Note:**
>
> See the "CREATE_CHAIN Procedure" for more complete descriptions of the attributes in this table.

Table 169-91 Chain Attribute Values

Name	Description
comments	An optional comment describing the purpose of the chain.
evaluation_interval	If not NULL, provides an additional evaluation of the chain at this interval, as well as at normal evaluation times (when the job starts, when a step completes, or when an event that is associated with an event step arrives) This attribute should only to be used when chain rules use SQL syntax and the rule conditions contain elements that are not under the control of the Scheduler, because the extra interval is CPU intensive. For most chains, the normal evaluation times are sufficient.

Table 169-91 (Cont.) Chain Attribute Values

Name	Description
rule_set_name	In the normal case, no rule set should be passed in. The Scheduler automatically creates a rule set and associated empty evaluation context. You then use DEFINE_CHAIN_RULE to add rules and DROP_CHAIN_RULE to remove them. Advanced users can create a rule set that describes their chain dependencies and pass it in here. This allows greater flexibility in defining rules. For example, conditions can refer to external variables, and tables can be exposed through the evaluation context. If you pass in a rule set, you must ensure that it is in the format of a chain rule set. (For example, all steps must be listed as variables in the evaluation context). If no rule set is passed in, the rule set created is of the form SCHED_RULESET${N} and the evaluation context created is of the form SCHED_EVCTX${N}

Database Destination Attribute Values

Table 169-92 lists database destination attribute values.

> **Note:**
>
> See the "CREATE_DATABASE_DESTINATION Procedure" for more complete descriptions of the attributes in this table.

Table 169-92 Database Destination Attribute Values

Name	Description
agent	The name of the external destination (also known as agent destination) that is used to connect to the remote database. You can obtain valid external destination names from the view ALL_SCHEDULER_EXTERNAL_DESTS.
connect_info	The TNS connect descriptor that identifies the remote database to connect to, or the net service name (alias) in tnsnames.ora that resolves to the connect descriptor. Note: This corresponds to the tns_name argument of CREATE_DATABASE_DESTINATION.
enabled	If TRUE, the database destination is enabled.
comments	An optional comment about the database destination.

External Destination Attribute Values

Table 169-93 lists external destination attribute values.

> **Note:**
>
> External destinations are created only implicitly by registering a remote Scheduler agent with the local database.

Table 169-93 External Destination Attribute Values

Name	Description
hostname	(GET_ATTRIBUTE only) The fully qualified host name (including domain) or IP address of the computer on which the Scheduler agent resides.
port	(GET_ATTRIBUTE only) The TCP port number on which the agent listens.
ip_address	(GET_ATTRIBUTE only) The IP address of the host on which the agent resides.
enabled	If TRUE, the external destination is enabled.
comments	An optional comment about the external destination.

Group Attribute Values

Table 169-94 lists group attribute values.

> **Note:**
>
> See the "CREATE_GROUP Procedure" for more complete descriptions of the attributes in this table.

Table 169-94 Group Attribute Values

Name	Description
group_type	(GET_ATTRIBUTE only) The group type (either WINDOW, DB_DEST, or EXTERNAL_DEST).
member_name	Comma-separated list of members. Replaces the existing list of members. To add one or more members to the existing list, use ADD_GROUP_MEMBER. Note: this attribute corresponds to the member argument of CREATE_GROUP.
enabled	If TRUE, the group is enabled.
comments	An optional comment about the group.
number_of_members	(GET_ATTRIBUTE only) The number of members in the group.

Credential Attribute Values

Table 169-95 lists credential attribute values.

> **Note:**
>
> Credential attribute values for the SET_ATTRIBUTE and GET_ATTRIBUTE procedures are deprecated with Oracle Database Release 12*c* Release 1 (12.1). While these attribute values remain available in this package, for reasons of backward compatibility, Oracle recommends using the alternative enhanced functionality provided in the DBMS_CREDENTIAL package, specifically the attribute parameter in the UPDATE_CREDENTIAL Procedure.

Table 169-95 Credential Attribute Values

Name	Description
username	The user name for logging into to the host operating system or remote Oracle database. Maximum length is 64.
password	The password for the user name. Maximum length is 128.
comments	A description of the credential. Maximum length is 240.
windows_domain	For a Windows remote executable target, this is the domain that the specified user belongs to. Maximum length is 64.
database_role	The value of the database_role attribute is used as the system privilege for logging into a remote database to run a remote database job. Valid values are: SYSDBA and SYSOPER.

Resource Attribute Values

Table 169-96 lists resource attribute values.

Table 169-96 Resource Attribute Values

Name	Description
resource_name	The name of the resource
units	The number of units of this resource that the job or program uses.
status	The status of the resource. ENFORCE_CONSTRAINTS. This is the default value, and when set, will force the scheduler to enforce resource limits. When the maximum number of units of this resource has been reached, no additional jobs using this resource will get started. IGNORE_CONSTRAINTS. When set, the scheduler will ignore any constraints on this resource. BLOCKED_ALL_JOBS. No jobs having a constraint on this resource will be allowed to run. The resource is considered to be permanently blocking until switched to one of the other two states.
constraint_level	Level of the constraint: JOB_LEVEL or PROGRAM_LEVEL For incompatibilities, for JOB_LEVEL, the incompatibility members must be jobs; for PROGRAM_LEVEL the incompatibility members must be programs.
comments	Descriptive comment about the resource.

169.6.69 SET_ATTRIBUTE_NULL Procedure

This procedure sets an attribute of an object to NULL.

The attributes that can be set depend on the object being altered. If the object is enabled, it is disabled before being altered and reenabled afterward. If the object cannot be reenabled, an error is generated and the object is left in a disabled state.

Syntax

```
DBMS_SCHEDULER.SET_ATTRIBUTE_NULL (
   name              IN VARCHAR2,
   attribute         IN VARCHAR2);
```

Parameters

Table 169-97 SET_ATTRIBUTE_NULL Procedure Parameters

Parameter	Description
name	The name of the object
attribute	The attribute being changed

Usage Notes

To run SET_ATTRIBUTE_NULL for a window, group of type WINDOW, or job class, you must have the MANAGE SCHEDULER privilege. Otherwise, you must be the owner of the object being altered or have ALTER privileges on that object or have the CREATE ANY JOB privilege.

169.6.70 SET_JOB_ANYDATA_VALUE Procedure

This procedure sets the value for an argument of the associated program for a job, encapsulated in an AnyData object.

It overrides any default value set for the program argument. NULL is a valid assignment for a program argument.

The argument can be specified by position or by name. You can specify by name only when:

- The job points to a saved program object
- The argument was assigned a name with the DEFINE_ANYDATA_ARGUMENT Procedure

Scheduler does no type checking of the argument at any time.

SET_JOB_ANYDATA_VALUE is overloaded.

Syntax

Sets a program argument by its position.

```
DBMS_SCHEDULER.SET_JOB_ANYDATA_VALUE (
   job_name                IN VARCHAR2,
   argument_position       IN PLS_INTEGER,
   argument_value          IN SYS.ANYDATA);
```

Sets a program argument by its name.

```
DBMS_SCHEDULER.SET_JOB_ANYDATA_VALUE (
   job_name                IN VARCHAR2,
   argument_name           IN VARCHAR2,
   argument_value          IN SYS.ANYDATA);
```

Parameters

Table 169-98 SET_JOB_ANYDATA_VALUE Procedure Parameters

Parameter	Description
job_name	The name of the job to be altered
argument_name	The name of the program argument being set

Table 169-98 (Cont.) SET_JOB_ANYDATA_VALUE Procedure Parameters

Parameter	Description
`argument_position`	The position of the program argument being set
`argument_value`	The new value to be assigned to the program argument, encapsulated in an `AnyData` object

Usage Notes

`SET_JOB_ANYDATA_VALUE` requires that you own the job or have `ALTER` privileges on that job. You can also set a job argument value if you have the `CREATE ANY JOB` privilege.

`SET_JOB_ANYDATA_VALUE` does not apply to lightweight jobs because lightweight jobs cannot take `AnyData` arguments.

See Also:

- "SET_JOB_ARGUMENT_VALUE Procedure"
- "DEFINE_ANYDATA_ARGUMENT Procedure"

169.6.71 SET_JOB_ARGUMENT_VALUE Procedure

This procedure sets the value of an argument for a job.

It overrides any default value set for the corresponding program or stored procedure argument. The argument can be specified by position or by name. You can specify by name only when:

- The job points to a saved program object
- The argument was assigned a name with the DEFINE_PROGRAM_ARGUMENT Procedure or the DEFINE_METADATA_ARGUMENT Procedure

Scheduler does no type checking of the argument at any time.

`SET_JOB_ARGUMENT_VALUE` is overloaded.

Syntax

Sets an argument value by position:

```
DBMS_SCHEDULER.SET_JOB_ARGUMENT_VALUE (
   job_name                IN VARCHAR2,
   argument_position       IN PLS_INTEGER,
   argument_value          IN VARCHAR2);
```

Sets an argument value by name:

```
DBMS_SCHEDULER.SET_JOB_ARGUMENT_VALUE (
   job_name                IN VARCHAR2,
   argument_name           IN VARCHAR2,
   argument_value          IN VARCHAR2);
```

Parameters

Table 169-99 SET_JOB_ARGUMENT_VALUE Procedure Parameters

Parameter	Description
job_name	The name of the job to be altered
argument_name	The name of the program argument being set
argument_position	The position of the program argument being set
argument_value	The new value to be set for the program argument. To set a non-VARCHAR value, use the SET_JOB_ANYDATA_VALUE procedure.

Usage Notes

SET_JOB_ARGUMENT_VALUE requires that you be the owner of the job or have ALTER privileges on that job. You can also set a job argument value if you have the CREATE ANY JOB privilege.

SET_JOB_ARGUMENT_VALUE only supports arguments of SQL type. Therefore, argument values that are not of SQL type, such as booleans, are not supported as program or job arguments.

SET_JOB_ARGUMENT_VALUE can be used to set arguments of lightweight jobs but only if the argument is of type VARCHAR2.

> **See Also:**
>
> - "SET_JOB_ANYDATA_VALUE Procedure"
> - "DEFINE_PROGRAM_ARGUMENT Procedure"

169.6.72 SET_JOB_ATTRIBUTES Procedure

This procedure changes an attribute of a job.

Syntax

```
DBMS_SCHEDULER.SET_JOB_ATTRIBUTES (
   jobattr_array     IN JOBATTR_ARRAY,
   commit_semantics  IN VARCHAR2 DEFAULT 'STOP_ON_FIRST_ERROR');
```

Parameters

Table 169-100 SET_JOB_ATTRIBUTES Procedure Parameters

Parameter	Description
jobattr_array	The array of job attribute changes.

Table 169-100 (Cont.) SET_JOB_ATTRIBUTES Procedure Parameters

Parameter	Description
commit_semantics	The commit semantics. The following types are supported: • STOP_ON_FIRST_ERROR returns on the first error and commits previous successful attribute changes to disk. This is the default. • TRANSACTIONAL returns on the first error and rolls back everything that happened before that error. • ABSORB_ERRORS tries to absorb any errors and complete the rest of the job attribute changes on the list. It commits all the successful changes. If errors occur, you can query the view SCHEDULER_BATCH_ERRORS for details.

Usage Notes

Calling SET_ATTRIBUTE on an enabled job disables the job, changes the attribute value, and reenables the job. SET_JOB_ATTRIBUTES changes the attribute values in the context of a single transaction.

169.6.73 SET_RESOURCE_CONSTRAINT Procedure

This procedure allows users to specify the resources used by jobs.

Syntax

```
DBMS_SCHEDULER.SET_RESOURCE_CONSTRAINT (
   object_name       IN VARCHAR2,
   resource_name     IN VARCHAR2,
   units             IN NUMBER DEFAULT 1);
```

Parameters

Table 169-101 SET_RESOURCE_CONSTRAINT Procedure Parameters

Parameter	Description
object_name	The name of a program or a job, or a comma separated list of these objects.
resource_name	The name of the resource.
units	The number of units of this resource that the job or program uses.

Usages Notes

object_name can be the name or comma-separated list of names of either programs or jobs. This creates a constraint on the named resource for these programs or jobs.

units specifies the number of units of the resource that the program or job can use. If units is set to 0, then the program or job does not use this resource anymore, and the resulting constraint is deleted. Setting units to 0 on a resource with no previous constraint results in an error.

When multiple constraints are defined on the same resource, the object types must match. When one or more existing constraints for a resource are based on jobs and a new constraint

is added for the same resource that is based on a program (or vice versa) an error will be raised.

169.6.74 SET_SCHEDULER_ATTRIBUTE Procedure

This procedure sets the value of a Scheduler attribute. This takes effect immediately but the resulting changes may not be seen immediately, depending on the attribute affected.

Table 169-102 provides short attribute descriptions for the `SET_SCHEDULER_ATTRIBUTE` procedure. For complete descriptions, see section "Setting Scheduler Preferences" in *Oracle Database Administrator's Guide*.

Syntax

```
DBMS_SCHEDULER.SET_SCHEDULER_ATTRIBUTE (
   attribute      IN VARCHAR2,
   value          IN VARCHAR2);
```

Parameters

Table 169-102 SET_SCHEDULER_ATTRIBUTE Procedure Parameters

Parameter	Description
`attribute`	The name of the Scheduler attribute. Possible values are: • `'default_timezone'`: Repeating jobs and windows that use the calendaring syntax retrieve the time zone from this attribute when `start_date` is not specified. See "Calendaring Syntax" for more information. • `'email_server'`: The SMTP server address that the Scheduler uses to send e-mail notifications for job state events. E-mail notifications cannot be sent if this attribute is `NULL`. • `'email_sender'`: The default e-mail address of the sender of job state e-mail notifications. • `'email_server_credential'`: The schema and name of an existing credential object that `SYS` has execute object privileges on. Default is `NULL`. The username and password stored in this credential are used to authenticate with the e-mail server when sending e-mail notifications. • `'email_server_encryption'`: This attribute indicates whether or not encryption is enabled for this email server connection, and if so, at what point encryption starts, and with which protocol. Values are: – `NONE`: the default, indicating no encryption used – `SSL_TLS`: indicating that either `SSL` or `TLS` are used, from the beginning of the connection – `STARTTLS`: indicating that the connection starts unencrypted, but the command `STARTTLS` is sent to the e-mail server and starts encryption • `'event_expiry_time'`: The time, in seconds, before a job state event generated by the Scheduler expires from the Scheduler event queue. If `NULL`, job state events expire after 24 hours. • `'log_history'`: The number of days that log entries for both the job log and the window log are retained. Default is 30 and the range of valid values is 0 through 1000000. • `'max_job_slave_processes'`: This Scheduler attribute is not used.
`value`	The new value of the attribute

Usage Notes

To run `SET_SCHEDULER_ATTRIBUTE`, you must have the `MANAGE SCHEDULER` privilege.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more detailed descriptions of Scheduler attributes

169.6.75 STOP_JOB Procedure

This procedure stops currently running jobs or all jobs in a job class.

After stopping the job, the state of a one-time job is set to `STOPPED`, whereas the state of a repeating job is set to `SCHEDULED` or `COMPLETED`, depending on whether the next run of the job is scheduled.

If a job pointing to a chain is stopped, all running steps of the running chain are stopped.

If a job has multiple destinations, the database attempts to stop the job at all destinations.

For external jobs, `STOP_JOB` stops only the external process that was directly started by the job action. It does not stop child processes of external jobs.

For in-memory full jobs in an Oracle Real Application Clusters environment, `STOP_JOB` uses the `instance_id` attribute of the job definition to determine in which instance (or all of them if the attribute is left null) to stop the in-memory full job. (In-memory full jobs are kept cached in memory, and as such are limited to the instance currently caching them. Because of this, the same `job_name` can in some conditions be used for different jobs on different instances.)

Syntax

```
DBMS_SCHEDULER.STOP_JOB (
   job_name         IN VARCHAR2
   force            IN BOOLEAN DEFAULT FALSE
   commit_semantics IN VARCHAR2 DEFAULT 'STOP_ON_FIRST_ERROR');
```

Parameters

Table 169-103 STOP_JOB Procedure Parameters

Parameter	Description
job_name	Name of a job to stop. Can be a comma-separate list of jobs, where each entry can be one of the following: • Job name: the name of an existing job, optionally preceded by a schema name and dot separator. • Job destination ID: a number, obtained from the JOB_DEST_ID column of the *_SCHEDULER_JOB_DESTS views, that represents the unique combination of a job, a credential, and a destination. • Job class: the name of a job class. Must be preceded by the SYS schema name and a dot separator. If you specify a job class, all jobs that belong to that job class are stopped. If you specify a job that was created with a destination group as its destination_name attribute, all job instances on all destinations are stopped.
force	If force is set to FALSE, the Scheduler tries to gracefully stop the job using an interrupt mechanism. This method gives control back to the secondary process, which can update the status of the job in the job queue to stopped. If this fails, an error is returned. If force is set to TRUE, the Scheduler immediately terminates the secondary process. Oracle recommends that STOP_JOB with force set to TRUE be used only after a STOP_JOB with force set to FALSE has failed. Use of the force option requires the MANAGE SCHEDULER system privilege.
commit_semantics	The commit semantics. The following two types are supported: • STOP_ON_FIRST_ERROR: The procedure returns on the first error and commits previous successful stop operations to disk. This is the default. • ABSORB_ERRORS: The procedure tries to absorb any errors, stops the rest of the jobs, and commits all the successful stop operations. This type is available only if no job classes are specified in the job_name list. If errors occur, you can query the view SCHEDULER_BATCH_ERRORS for details.

Usage Notes

STOP_JOB without the force option requires that you be the owner of the job or have ALTER privileges on that job. You can also stop a job if you have the CREATE ANY JOB or MANAGE SCHEDULER privilege.

STOP_JOB with the force option requires that you have the MANAGE SCHEDULER privilege.

Example

The following is an example of using STOP_JOB.

```
BEGIN
  DBMS_SCHEDULER.STOP_JOB('DSS.ETLJOB, 984, 1223, SYS.ETL_JOBCLASS');
END;
```

170
DBMS_SEARCH

The `DBMS_SEARCH` package enables the indexing of multiple schema objects in a single index.

This Oracle Text package provides procedures and functions to create, maintain, and query a ubiquitous search index. Ubiquitous database search is a full-text and range-based search across multiple tables and views. To add data sources from different schemas to the index, the index owner must have `SELECT` and `DML` access to the sources.

A complete description of this package is available in *Oracle Text Reference*.

Related Topics

- *Oracle Text Reference*

171
DBMS_SERVER_ALERT

The DBMS_SERVER_ALERT package enables you to configure the Oracle Database server to issue an alert when a threshold for a specified server metric has been violated. You can configure both warning and critical thresholds for a large number of predefined metrics.

If a warning threshold is reached, the server generates a severity level 5 alert. If a critical threshold is reached, the server generates a severity level 1 alert.

The chapter contains the following topics:

- Security Model
- Object Types
- Relational Operators
- Supported Metrics
- Summary of DBMS_SERVER_ALERT Subprograms

171.1 DBMS_SERVER_ALERT Security Model

The user needs DBA or IMP_FULL_DATABASE roles to use the DBMS_SERVER_ALERT package.

171.2 DBMS_SERVER_ALERT Object Types

You qualify the metric by an individual object for the listed object types.

Table 171-1 Object Types Defined as Constants

Constant	Description
OBJECT_TYPE_SYSTEM	Metrics collected on the system level for each instance.
OBJECT_TYPE_FILE	Metrics collected on the file level. These are used for AVG_FILE_READ_TIME and AVG_FILE_WRITE_TIME metrics.
OBJECT_TYPE_SERVICE	Metrics collected on the service level. Currently ELAPSED_TIME_PER_CALL and CPU_TIME_PER_CALL are collected.
OBJECT_TYPE_TABLESPACE	Metrics collected on the tablespace level. **Note**: Dictionary managed tablespaces are not supported.
OBJECT_TYPE_EVENT_CLASS	Metrics collected on wait event class level. Currently supported metrics are AVG_USERS_WAITING and DB_TIME_WAITING.
OBJECT_TYPE_SESSION	Metrics collected on the session level. Currently only BLOCKED_USERS is collected. The threshold can only be set at the instance level, which means that no object name should be specified when setting the threshold for this type of metric.
OBJECT_TYPE_WRCLIENT	Refers to a group of metrics (WCR_. . .) used during replay to monitor the replay clients' performance

171.3 DBMS_SERVER_ALERT Relational Operators

You can specify a relational comparison operator to determine whether or not a given metric's value violates the threshold setting. The server supports the following operators.

Table 171-2 Relational Operators Defined as Constants

Constant	Description
OPERATOR_CONTAINS	A metric value matching an entry in a list of threshold values is considered a violation.
OPERATOR_DO_NOT_CHECK	The metric value is not compared to the threshold value, and no alerts are generated. Use this operator to disable alerts for a metric.
OPERATOR_EQ	A metric value equal to the threshold value is considered a violation.
OPERATOR_GE	A metric value greater than or equal to the threshold value is considered a violation.
OPERATOR_GT	A metric value greater than the threshold value is considered a violation.
OPERATOR_LE	A metric value less than or equal to the threshold value is considered a violation.
OPERATOR_LT	A metric value less than the threshold value is considered a violation.
OPERATOR_NE	A metric value not equal to the threshold value is considered a violation.

171.4 DBMS_SERVER_ALERT Supported Metrics

These metrics are supported. All internal metric names are supplied as package constants.

Table 171-3 List of Supported Metrics

Metric Name (Internal)	Metric Name (External)	Units
AVG_FILE_READ_TIME	Average File Read Time	Microseconds
AVG_FILE_WRITE_TIME	Average File Write Time	Microseconds
AVG_USERS_WAITING	Average Number of Users Waiting on a Class of Wait Events	Count of sessions
BLOCKED_USERS	Number of Users blocked by some Session	Number of Users
BRANCH_NODE_SPLITS_SEC	Branch Node Splits (for each second)	Splits for each Second
BRANCH_NODE_SPLITS_TXN	Branch Node Splits (for each transaction)	Splits for each Transaction
BUFFER_CACHE_HIT	Buffer Cache Hit (%)	% of cache accesses
CONSISTENT_CHANGES_SEC	Consistent Changes (for each second)	Changes for each Second
CONSISTENT_CHANGES_TXN	Consistent Changes (for each transaction)	Changes for each Transaction
CONSISTENT_GETS_SEC	Consistent Gets (for each second)	Gets for each Second

Table 171-3 (Cont.) List of Supported Metrics

Metric Name (Internal)	Metric Name (External)	Units
CONSISTENT_GETS_TXN	Consistent Gets (for each transaction)	Gets for each Transaction
CR_BLOCKS_CREATED_SEC	CR Blocks Created (for each second)	Blocks for each Second
CR_BLOCKS_CREATED_TXN	CR Blocks Created (for each transaction)	Blocks for each Transaction
CR_RECORDS_APPLIED_SEC	CR Undo Records Applied (for each second)	Records for each Second
CR_RECORDS_APPLIED_TXN	CR Undo Records Applied (for each transaction)	Records for each Transaction
CURSOR_CACHE_HIT	Cursor Cache Hit (%)	% of soft parses
DATABASE_WAIT_TIME	Database Wait Time (%)	% of all database time
DATABASE_CPU_TIME	Database CPU Time (%)	% of all database time
DB_BLKGETS_SEC	DB Block Gets (for each second)	Gets for each Second
DB_BLKGETS_TXN	DB Block Gets (for each transaction)	Gets for each Transaction
DB_TIME_WAITING	Percent of Database Time Spent Waiting on a Class of Wait Events	% of Database Time
DBWR_CKPT_SEC	DBWR Checkpoints (for each second)	Checkpoints for each Second
DISK_SORT_SEC	Sorts to Disk (for each second)	Sorts for each Second
DISK_SORT_TXN	Sorts to Disk (for each transaction)	Sorts for each Transaction
ELAPSED_TIME_PER_CALL	Elapsed time for each user call for each service	Microseconds for each call
ENQUEUE_DEADLOCKS_SEC	Enqueue Deadlocks (for each second)	Deadlocks for each Second
ENQUEUE_DEADLOCKS_TXN	Enqueue Deadlocks (for each transaction)	Deadlocks for each Transaction
ENQUEUE_REQUESTS_SEC	Enqueue Requests (for each second)	Requests for each Second
ENQUEUE_REQUESTS_TXN	Enqueue Requests (for each transaction)	Requests for each Transaction
ENQUEUE_TIMEOUTS_SEC	Enqueue Timeouts (for each second)	Timeouts for each Second
ENQUEUE_TIMEOUTS_TXN	Enqueue Timeouts (for each transaction)	Timeouts for each Transaction
ENQUEUE_WAITS_SEC	Enqueue Waits (for each second)	Waits for each Second
ENQUEUE_WAITS_TXN	Enqueue Waits (for each transaction)	Waits for each Transaction
EXECUTE_WITHOUT_PARSE	Executes Performed Without Parsing	% of all executes
FULL_INDEX_SCANS_SEC	Fast Full Index Scans (for each second)	Scans for each Second
FULL_INDEX_SCANS_TXN	Fast Full Index Scans (for each transaction)	Scans for each Transaction
GC_AVG_CR_GET_TIME	Global Cache CR Request	Milliseconds
GC_AVG_CUR_GET_TIME	Global Cache Current Request	Milliseconds
GC_BLOCKS_CORRUPT	Global Cache Blocks Corrupt	Blocks
GC_BLOCKS_LOST	Global Cache Blocks Lost	Blocks
HARD_PARSES_SEC	Hard Parses (for each second)	Parses for each Second

Table 171-3 (Cont.) List of Supported Metrics

Metric Name (Internal)	Metric Name (External)	Units
HARD_PARSES_TXN	Hard Parses (for each transaction)	Parses for each Transaction
LEAF_NODE_SPLITS_SEC	Leaf Node Splits (for each second)	Splits for each Second
LEAF_NODE_SPLITS_TXN	Leaf Node Splits (for each transaction)	Splits for each Transaction
LIBRARY_CACHE_HIT	Library Cache Hit (%)	% of cache accesses
LIBRARY_CACHE_MISS	Library Cache Miss (%)	% of cache accesses
LOGONS_CURRENT	Current Number of Logons	Number of Logons
LOGONS_SEC	Cumulative Logons (for each second)	Logons for each Second
LOGONS_TXN	Cumulative Logons (for each transaction)	Logons for each Transaction
LONG_TABLE_SCANS_SEC	Scans on Long Tables (for each second)	Scans for each Second
LONG_TABLE_SCANS_TXN	Scans on Long Tables (for each transaction)	Scans for each Transaction
OPEN_CURSORS_SEC	Cumulative Open Cursors (for each second)	Cursors for each Second
MEMORY_SORTS_PCT	Sorts in Memory (%)	% of sorts
NETWORK_BYTES_SEC	Network Bytes, for each second	Bytes for each Second
OPEN_CURSORS_CURRENT	Current Number of Cursors	Number of Cursors
OPEN_CURSORS_TXN	Cumulative Open Cursors (for each transaction)	Cursors for each Transaction
OS_SCHED_CPU_WAIT_TIME	Operating System Scheduler CPU Wait (by time)	Microseconds
PARSE_FAILURES_SEC	Parse Failures (for each second)	Parses for each Second
PARSE_FAILURES_TXN	Parse Failures (for each transaction)	Parses for each Transaction
PGA_CACHE_HIT	PGA Cache Hit (%)	% bytes processed in PGA
PHYS_DESGN_WAIT_SCT	Physical Design Wait (by session count)	Count of sessions
PHYSICAL_READS_SEC	Physical Reads (for each second)	Reads for each Second
PHYSICAL_READS_TXN	Physical Reads (for each transaction)	Reads for each Transaction
PHYSICAL_WRITES_SEC	Physical Writes (for each second)	Writes for each Second
PHYSICAL_WRITES_TXN	Physical Writes (for each transaction)	Writes for each Transaction
PHYSICAL_READS_DIR_SEC	Direct Physical Reads (for each second)	Reads for each Second
PHYSICAL_READS_DIR_TXN	Direct Physical Reads (for each transaction)	Reads for each Transaction
PHYSICAL_WRITES_DIR_SEC	Direct Physical Writes (for each second)	Writes for each Second
PHYSICAL_WRITES_DIR_TXN	Direct Physical Writes (for each transaction)	Writes for each Transaction
PHYSICAL_READS_LOB_SEC	Direct LOB Physical Reads (for each second)	Reads for each Second
PHYSICAL_READS_LOB_TXN	Direct LOB Physical Reads (for each transaction)	Reads for each Transaction
PHYSICAL_WRITES_LOB_SEC	Direct LOB Physical Writes (for each second)	Writes for each Second

Table 171-3 (Cont.) List of Supported Metrics

Metric Name (Internal)	Metric Name (External)	Units
PHYSICAL_WRITES_LOB_TXN	Direct LOB Physical Writes (for each transaction)	Writes for each Transaction
PROCESS_LIMIT_PCT	Process Limit Usage (%)	% of maximum value
PX_DOWNGRADED_SEC	Downgraded Parallel Operations (for each second)	Operations for each Second
PX_DOWNGRADED_25_SEC	Downgraded to 25% and more (for each second)	Operations for each Second
PX_DOWNGRADED_50_SEC	Downgraded to 50% and more (for each second)	Operations for each Second
PX_DOWNGRADED_75_SEC	Downgraded to 75% and more (for each second)	Operations for each Second
PX_DOWNGRADED_SER_SEC	Downgraded to serial (for each second)	Operations for each Second
RB_RECORDS_APPLIED_SEC	Rollback Undo Records Applied (for each second)	Records for each Second
RB_RECORDS_APPLIED_TXN	Rollback Undo Records Applied (for each transaction)	Records for each Transaction
REDO_ALLOCATION_HIT	Redo Log Allocation Hit	% of redo allocations
REDO_GENERATED_SEC	Redo Generated (for each second)	Redo Bytes for each Second
REDO_GENERATED_TXN	Redo Generated (for each transaction)	Redo Bytes for each Transaction
REDO_WRITES_SEC	Redo Writes (for each second)	Writes for each Second
REDO_WRITES_TXN	Redo Writes (for each transaction)	Writes for each Transaction
RECURSIVE_CALLS_SEC	Recursive Calls (for each second)	Calls for each Second
RECURSIVE_CALLS_TXN	Recursive Calls (for each transaction)	Calls for each Transaction
RESPONSE_TXN	Response (for each transaction)	Seconds for each Transaction
ROWS_PER_SORT	Rows Processed for each Sort	Rows for each Sort
SESS_LOGICAL_READS_SEC	Session Logical Reads (for each second)	Reads for each Second
SESS_LOGICAL_READS_TXN	Session Logical Reads (for each transaction)	Reads for each Transaction
SESSION_CPU_SEC	Database CPU (for each second)	Microseconds for each Second
SESSION_CPU_TXN	Database CPU (for each transaction)	Microseconds for each Transaction
SESSION_LIMIT_PCT	Session Limit Usage (%)	% of maximum value
SHARED_POOL_FREE_PCT	Shared Pool Free(%)	% of shared pool
SOFT_PARSE_PCT	Soft Parse (%)	% of all parses
SQL_SRV_RESPONSE_TIME	Service Response (for each execution)	Seconds
TABLESPACE_PCT_FULL	Tablespace space usage	% full
TABLESPACE_BYT_FREE	Tablespace bytes space usage	Kilobytes free
TOTAL_TABLE_SCANS_SEC	Total Table Scans (for each second)	Scans for each Second
TOTAL_TABLE_SCANS_TXN	Total Table Scans (for each transaction)	Scans for each Transaction
TOTAL_INDEX_SCANS_SEC	Total Index Scans (for each second)	Scans for each Second

Table 171-3 (Cont.) List of Supported Metrics

Metric Name (Internal)	Metric Name (External)	Units
TOTAL_INDEX_SCANS_TXN	Total Index Scans (for each transaction)	Scans for each Transaction
TOTAL_PARSES_SEC	Total Parses (for each second)	Parses for each Second
TOTAL_PARSES_TXN	Total Parses (for each transaction)	Parses for each Transaction
USER_COMMITS_SEC	User Commits (for each second)	Commits for each Second
USER_COMMITS_TXN	User Commits (for each transaction)	Commits for each Transaction
USER_ROLLBACKS_SEC	User Rollbacks (for each second)	Rollbacks for each Second
USER_ROLLBACKS_TXN	User Rollbacks (for each transaction)	Rollbacks for each Transaction
USER_CALLS_SEC	User Calls (for each second)	Calls for each Second
USER_CALLS_TXN	User Calls (for each transaction)	Calls for each Transaction
USER_CALLS_PCT	User Calls (%)	% of all calls
USER_LIMIT_PCT	User Limit Usage (%)	% of maximum value
WCR_AVG_IO_LAT	Average IO response time (for a WRC client)	Milliseconds
WCR_PCPU	Percentage of replay threads on CPU (for a WRC client)	% of total replay threads
WCR_PIO	Percentage of replay threads doing IOs (for a WRC client)	% of total replay threads

171.5 Summary of DBMS_SERVER_ALERT Subprograms

This table lists the DBMS_SERVER_ALERT subprograms and briefly describes them.

Table 171-4 DBMS_SERVER_ALERT Package Subprograms

Subprogram	Description
EXPAND_MESSAGE Function	Expands alert messages
GET_THRESHOLD Procedure	Gets the current threshold settings for a specified metric
SET_THRESHOLD Procedure	Sets the warning and critical thresholds for a specified metric

171.5.1 EXPAND_MESSAGE Function

This function expands alert messages.

Syntax

```
DBMS_SERVER_ALERT.EXPAND_MESSAGE(
   user_language            IN   VARCHAR2,
   message_id               IN   NUMBER,
   argument_1               IN   VARCHAR2,
   argument_2               IN   VARCHAR2,
   argument_3               IN   VARCHAR2,
   argument_4               IN   VARCHAR2,
   argument_5               IN   VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 171-5 *EXPAND_MESSAGE Function Parameters*

Parameter	Description
user_language	The language of the current session.
message_id	Id of the alert message
argument_1	The first argument in the alert message.
argument_2	The second argument in the alert message.
argument_3	The third argument in the alert message.
argument_4	The fourth argument in the alert message.
argument_5	The fifth argument in the alert message.

171.5.2 GET_THRESHOLD Procedure

This procedure gets the current threshold settings for the specified metric.

Syntax

```
DBMS_SERVER_ALERT.GET_THRESHOLD(
   metrics_id               IN   BINARY_INTEGER,
   warning_operator         OUT  BINARY_INTEGER,
   warning_value            OUT  VARCHAR2,
   critical_operator        OUT  BINARY_INTEGER,
   critical_value           OUT  VARCHAR2,
   observation_period       OUT  BINARY_INTEGER,
   consecutive_occurrences  OUT  BINARY_INTEGER,
   instance_name            IN   VARCHAR2,
   object_type              IN   BINARY_INTEGER,
   object_name              IN   VARCHAR2);
```

Parameters

Table 171-6 GET_THRESHOLD Procedure Parameters

Parameter	Description
metrics_id	The internal name of the metric. See "Supported Metrics".
warning_operator	The operator for the compa3ring the actual value with the warning threshold.
warning_value	The warning threshold value.
critical_operator	The operator for the comparing the actual value with the critical threshold.
critical_value	The critical threshold value.
observation_period	The period at which the metric values are computed and verified against the threshold setting.
consecutive_occurrences	The number of observation periods the metric value should violate the threshold value before the alert is issued.

Table 171-6 (Cont.) GET_THRESHOLD Procedure Parameters

Parameter	Description
instance_name	The name of the instance for which the threshold is set. This is NULL for database-wide alerts. In cases in which this parameter is not NULL, this should be set to one of the INSTANCE_NAME values found in the GV$INSTANCE View.
object_type	Either OBJECT_TYPE_SYSTEM or OBJECT_TYPE_SERVICE.
object_name	The name of the object.

Usage Notes

Note that this subprogram does not check if the value of the instance_name parameter is meaningful or valid.

171.5.3 SET_THRESHOLD Procedure

This procedure sets the warning and critical thresholds for a specified metric.

Syntax

```
DBMS_SERVER_ALERT.SET_THRESHOLD(
   metrics_id               IN   BINARY_INTEGER,
   warning_operator         IN   BINARY_INTEGER,
   warning_value            IN   VARCHAR2,
   critical_operator        IN   BINARY_INTEGER,
   critical_value           IN   VARCHAR2,
   observation_period       IN   BINARY_INTEGER,
   consecutive_occurrences  IN   BINARY_INTEGER,
   instance_name            IN   VARCHAR2,
   object_type              IN   BINARY_INTEGER,
   object_name              IN   VARCHAR2);
```

Parameters

Table 171-7 SET_THRESHOLD Procedure Parameters

Parameter	Description
metrics_id	The internal name of the metric. See "Supported Metrics".
warning_operator	The operator for the comparing the actual value with the warning threshold (such as OPERATOR_GE). See "Relational Operators".
warning_value	The warning threshold value. This is NULL if no warning threshold is set. A list of values may be specified for OPERATOR_CONTAINS.
critical_operator	The operator for the comparing the actual value with the critical threshold. See "Relational Operators".
critical_value	The critical threshold value. This is NULL if not set. A list of values may be specified for OPERATOR_CONTAINS.
observation_period	The period at which the metric values are computed and verified against the threshold setting. The valid range is 1 to 60 minutes.
consecutive_occurrences	The number of observation periods the metric value should violate the threshold value before the alert is issued.

Table 171-7 (Cont.) SET_THRESHOLD Procedure Parameters

Parameter	Description
instance_name	The name of the instance for which the threshold is set. This is NULL for database-wide alerts.
object_type	See "Object Types".
object_name	The name of the object. This is NULL for SYSTEM.

Usage Notes

Note that this subprogram does not check if the value of the instance_name parameter is meaningful or valid. Passing a name that does not identify a valid instance will result in a threshold that is not used by any by any instance although the threshold setting will be visible in the DBA_THRESHOLDS view. The exception is the lower-case string 'database_wide' which is semantically equivalent to passing NULL for the instance name, the latter being the preferred usage.

172
DBMS_SERVICE

The `DBMS_SERVICE` package lets you create, delete, activate, and deactivate services for a single instance.

The chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operating Procedures
- Exceptions
- Summary of DBMS_SERVICE Subprograms

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide* for administering services in Oracle Real Application Clusters.

172.1 DBMS_SERVICE Overview

`DBMS_SERVICE` supports the workload management of high availability, quality of service, job scheduling, and other planned operations in the RDBMS for the purposes of workload measurement, management, prioritization, and XA and distributed transaction management.

Oracle Real Application Clusters (RAC) can manage service names across instances as administered through `SRVCTL`. The `DBMS_SERVICE` package allows the creation, deletion, starting, and stopping of services in a single instance. Additionally, it provides the ability to disconnect all sessions that connect to an instance with a service name.

> **See Also:**
>
> For more information about Oracle Real Application Clusters, *Oracle Real Application Clusters Administration and Deployment Guide*.

172.2 DBMS_SERVICE Security Model

The DBMS_SERVICE package has certain security requirements.

Privileges

The client using this package must have the ALTER SYSTEM execution privilege and the V$SESSION table read privilege.

Schemas

This package must be installed under SYS schema.

Roles

The EXECUTE privilege of the package is granted to the DBA role only.

172.3 DBMS_SERVICE Constants

The DBMS_SERVICE package provides constants that can be used for specifying parameter values.

- Constants used in calling arguments are described in Table 172-1
- Constants used in connection balancing goal arguments are described in Table 172-2
- Constants used in TAF failover attribute arguments are described in Table 172-3

Table 172-1 Constants Used in Calling Arguments

Name	Type	Value	Description
GOAL_NONE	NUMBER	0	Disables Load Balancing Advisory
GOAL_SERVICE_TIME	NUMBER	1	Load Balancing Advisory is based on elapsed time for work done in the service plus available bandwidth to the service
GOAL_THROUGHPUT	NUMBER	2	Load Balancing Advisory is based on the rate that work is completed in the service plus available bandwidth to the service

Table 172-2 Constants Used in Connection Balancing Goal Arguments

Name	Type	Value	Description
CLB_GOAL_SHORT	NUMBER	1	Connection load balancing uses Load Balancing Advisory, when Load Balancing Advisory is enabled (either goal_service_time or goal_throughput). When GOAL=NONE (no load balancing advisory), connection load balancing uses an abridged advice based on CPU utilization.

Table 172-2 (Cont.) Constants Used in Connection Balancing Goal Arguments

Name	Type	Value	Description
CLB_GOAL_LONG	NUMBER	2	Balances the number of connections for each instance using session count for each service. This setting is recommended for applications with long connections such as forms. This setting can be used with Load Balancing Advisory when the connection pool is sized to accommodate gravitation within the pool itself (without adding or removing connections). The latter is the most efficient design.

Table 172-3 Constants Used in High Availability Attribute Arguments for FAN, Application Continuity, Transaction Guard and TAF

Name	Type	Value	Description
FAILOVER_METHOD_NONE	VARCHAR2	0	Server side TAF is not enabled for this service
FAILOVER_METHOD_BASIC	VARCHAR2	1	Server side TAF method is BASIC. BASIC is the only value currently supported. This means that a new connection is established at failure time.
FAILOVER_TYPE_NONE	VARCHAR		Server side TAF type is NONE
FAILOVER_TYPE_SESSION	VARCHAR		Server side TAF failover type is SESSION. At failure time, if the failover type is SESSION, TAF reconnects to a surviving node and re-establish a vanilla database session. Customizations (for example, ALTER SESSION) must be re-executed in a failover callback.
FAILOVER_TYPE_SELECT	VARCHAR		Server side TAF failover type is SELECT
FAILOVER_RETRIES	NUMBER		Number of connection attempts when failover occurs. Specifies the number of times for Application Continuity and TAF to attempt the reconnect and re-authenticate pair. The value must be an integer greater than 0. The default in Oracle Database 12c Release 1 (12.1) for Application Continuity is 30.
FAILOVER_RESTORE_NONE	CONSTANT VARCHAR2	NONE	The initial state is not restored before replaying for Application Continuity and TAF. This is recommended for OCI applications that use Application Continuity and build their own state in the request. For example, SQLPLUS.
FAILOVER_RESTORE_BASIC	VARCHAR2	LEVEL1	This is the recommended value for Java and ODP.NET applications using Application Continuity. The initial states that the user knows are restored automatically before replaying. If the user needs additional states, a callback must be registered.

Table 172-3 (Cont.) Constants Used in High Availability Attribute Arguments for FAN, Application Continuity, Transaction Guard and TAF

Name	Type	Value	Description
FAILOVER_DELAY	NUMBER		Number of seconds delay between each connection attempt. This is the delay that Application Continuity and TAF waits if a reconnect and re-authentication fails. The value must be an integer greater than 0. The default in Oracle Database 12c Release 1 (12.1)is 10s when using Application Continuity. Using FAILOVER_DELAY the failover can be delayed until the service is next available. This can work well in conjunction with a planned outage that may make a service temporarily unavailable (such as for several minutes).
STOP_OPTION_NONE	VARCHAR		Sessions are not disconnected.
STOP_OPTION_IMMEDIATE	VARCHAR		Sessions are disconnected immediately after the drain_timeout expires.
STOP_OPTION_TRANSACTIONAL			Sessions are disconnected after the transactions during the drain_timeout. The sessions disconnect immediately when drain_timeout expires.
DYNAMIC	NUMBER		For Application Continuity, this parameter specifies whether the session state that is not transactional is changed by the application during request execution. A value of DYNAMIC is recommended for all applications. If you are in any doubt, or the application can be customized, you must use DYNAMIC.

Usage Notes

- If a TAF callback has been registered, then the failover retries and failover delay are ignored. If an error occurs, TAF continues to re-attempt the connect and authentication as long as the callback returns a value of OCI_FO_RETRY. Any delay must be coded into the callback logic
- Server side TAF settings override client-side counterparts that might be configured in TNS connect descriptors. If TAF is not configured on the client side, then at a minimum, the failover type must be set to enable TAF. If the failover type is set on the server side, then the failover method defaults to BASIC. Delay and retries are optional and may be specified independently.

172.4 DBMS_SERVICE Operating Procedures

You cannot use the following procedures with Oracle Real Applications Clusterware, Oracle Restart, and Oracle Global Data Services.

- – CREATE_SERVICE Procedure
  - – DELETE_SERVICE Procedure
  - – MODIFY_SERVICE Procedure
  - – START_SERVICE Procedure

  - STOP_SERVICE Procedure

- With Oracle Database 12*c* release 1, you are advised to use the parameter interface in all service -related subprorgrams.
- If you wish to use `DBMS_SERVICE` on a pluggable database (PDB) in a single instance, you must connect to that PDB first.

172.5 DBMS_SERVICE Exceptions

This table lists the exceptions raised by the `DBMS_SERVICE` package.

Table 172-4 DBMS_SERVICE Exceptions

Exception	Error Code	Description
NULL_SERVICE_NAME	44301	Service name argument was found to be NULL
NULL_NETWORK_NAME	44302	Network name argument was found to be NULL
SERVICE_EXISTS	44303	Service name already exists
SERVICE_DOES_NOT_EXIST	44304	Specified service does not exist
SERVICE_IN_USE	44305	Specified service was running
SERVICE_NAME_TOO_LONG	44306	Service name was too long
NETWORK_PREFIX_TOO_LONG	44307	Network name, excluding the domain, was too long
NOT_INITIALIZED	44308	Services layer was not yet initialized
GENERAL_FAILURE	44309	An unknown failure
MAX_SERVICES_EXCEEDED	44310	Maximum number of services has been reached
SERVICE_NOT_RUNNING	44311	Specified service was not running
DATABASE_CLOSED	44312	Database was closed
INVALID_INSTANCE	44313	Instance name argument was not valid
NETWORK_EXISTS	44314	Network name already exists
NULL_ATTRIBUTES	44315	All attributes specified were NULL
INVALID_ARGUMENT	44316	Invalid argument supplied
DATABASE_READONLY	44317	Database is open read-only
MAX_SN_LENGTH	44318	Total length of all running service network names exceeded the maximum allowable length
ERR_AQ_SERVICE	44319	Cannot delete AQ service
ERR_GLB_SERVICE	44320	Cannot delete global service
ERR_INVALID_PDB_NAME	44771	Invalid name for a pluggable database
ERR_CRS_API	44772	Cluster ready services (CRS) operation failed
ERR_PDB_CLOSED	44773	Cannot perform requested service operation

Table 172-4 (Cont.) DBMS_SERVICE Exceptions

Exception	Error Code	Description
ERR_PDB_INVALID	44774	Pluggable database attribute cannot be changed
ERR_PDB_NAME	44775	Pluggable database service cannot be created
ERR_PDB_EXP	44776	Pluggable database service cannot be deleted
ERR_PDB_FAIL	44777	Pluggable database service cannot be started

172.6 Summary of DBMS_SERVICE Subprograms

This table lists the DBMS_SERVICE subprograms and briefly describes them.

Table 172-5 DBMS_SERVICE Package Subprograms

Subprogram	Description
CREATE_SERVICE Procedure	Creates service
DELETE_SERVICE Procedure	Deletes service
DISCONNECT_SESSION Procedure	Disconnects sessions running under this service
MODIFY_SERVICE Procedure	Modifies service
START_SERVICE Procedure	Activates service
STOP_SERVICE Procedure	Stops service

172.6.1 CREATE_SERVICE Procedure

This procedure creates a service name in the data dictionary. Services are also created in the data dictionary implicitly when you set the service in the `service_name` parameter or by means of the `ALTER SYSTEM SET SERVICE_NAMES` command.

> **Note:**
>
> The service attribute values `FAILOVER_TYPE = TRANSACTION` with `SESSION_STATE_CONSISTENCY = STATIC` are no longer a supported service attribute combination.
> In previous releases, you could use the service parameter `SESSION_STATE_CONSISTENCY` to manage session state automatically using Application Continuity by setting `SESSION_STATE_CONSISTENCY` to `DYNAMIC` or `STATIC`. However, starting with Oracle Database 23ai, you can no longer use the `STATIC` option. Instead, use one of the following failover options:
>
> - `FAILOVER_TYPE = AUTO` with `SESSION_STATE_CONSISTENCY = AUTO`
> - `FAILOVER_TYPE = TRANSACTION` with `SESSION_STATE_CONSISTENCY = DYNAMIC`
>
> These configurations enforce session state tracking in Oracle Database, ensuring that session state is preserved at session migration and session failover.

Syntax

```
DBMS_SERVICE.CREATE_SERVICE(
   service_name               IN VARCHAR2,
   network_name               IN VARCHAR2,
   parameter_array            IN TABLE OF VARCHAR2(100));
```

This overload is maintained for backward compatibility:

```
DBMS_SERVICE.CREATE_SERVICE(
   service_name               IN VARCHAR2,
   network_name               IN VARCHAR2,
   goal                       IN NUMBER DEFAULT NULL,
   dtp                        IN BOOLEAN DEFAULT NULL,
   true_cache_service         IN VARCHAR2,
   aq_ha_notifications        IN BOOLEAN DEFAULT NULL,
   failover_method            IN VARCHAR2 DEFAULT NULL,
   failover_type              IN VARCHAR2 DEFAULT NULL,
   failover_retries           IN NUMBER DEFAULT NULL,
   failover_delay             IN NUMBER DEFAULT NULL,
   clb_goal                   IN NUMBER DEFAULT NULL,
   edition                    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 172-6 CREATE_SERVICE Procedure Parameters

Parameter	Description
`service_name`	Name of the service, limited to 64 characters in the Data Dictionary

Table 172-6 (Cont.) CREATE_SERVICE Procedure Parameters

Parameter	Description
network_name	Network name of the service as used in SQLNet connect descriptors for client connections. This is limited to the NET service_names character set (see *Oracle Database Net Services Reference*).
parameter_array	Associative array with name-value pairs of the service attributes. Supported names: • goal • dtp • true_cache_service • aq_ha_notifications • failover_method • failover_type • failover_retries • failover_restore • failover_delay • clb_goal • edition • commit_outcome • reset_state • retention_timeout • replay_initiation_timeout • session_state_consistency • sql_translation_profile • drain_timeout • stop_option • placement_policy
goal	Workload management goal directive for the service. Valid values: • DBMS_SERVICE.GOAL_SERVICE_TIME • DBMS_SERVICE.GOAL_THROUGHPUT • DBMS_SERVICE.GOAL_NONE
dtp	Declares the service to be for X/Open Distributed Transaction Processing (DTP) or any distributed transaction (especially XA)
true_cache_service	Name of the True Cache service being registered with the primary service.
aq_ha_notifications	Determines whether Fast Application Notification (FAN) is enabled for OCI/OCCI/ODP. In Oracle Database12*c*, FAN uses Oracle Notification Services (ONS). This parameter is still used to enable FAN. FAN is recommended for all High Availability systems, and is on by default for Application Continuity
failover_method	Failover TYPE for the service for Application Continuity and TAF. If the failover_type is set to TRANSACTION on the service, this automatically sets COMMIT_OUTCOME to TRUE. JDBC Replay Driver uses the FAILOVER_TYPE service attribute setting of TRANSACTION for TRANSACTION failover. OCI uses the older settings of SELECT and SESSION. The server only accepts FAILOVER_METHOD = BASIC with the TRANSACTION setting.
failover_type	TRANSACTION for Application Continuity. SELECT or BASIC for TAF.

Table 172-6 (Cont.) CREATE_SERVICE Procedure Parameters

Parameter	Description
failover_retries	Number of connection retries for Application Continuity and TAF. Using the failover_retries and failover_delay parameters, the failover can be delayed until the service is next available. This parameter is for connecting. It does not control the number of failovers, which is 3 for each incident for Application Continuity.
failover_delay	Delay in seconds between connection retries for Application Continuity and TAF. The default is 10 seconds for Application Continuity. Do not use a 0-second delay if the service needs time to failover and register. Long delays are good for planned outages and to failover to Data Guard. Short delays work well with RAC when the service is already available.
edition	If this argument has a non-NULL value, this provides the initial session edition for subsequent database connections using this service that do not specify an edition. If no value is specified, this argument has no effect. During service creation or modification, no validation is performed on this parameter. At connection time, if the connecting user does not have USE privilege on the edition, or the edition does not exist, this raises the error ORA-38802 (edition does not exist).
drain_timeout	If this parameter is defined, all sessions connected to that service are drained by the client drivers and pools using Fast Connection Failover (FCF). The drain_timeout can be set on the service, so stopping and relocating drains for this time by default.
stop_option	Provides options to terminate a session. The supported values are: • TRANSACTIONAL drain_timeout applies to the transactions. After the drain_timeout expire, the sessions are immediately killed. • IMMEDIATE The sessions are killed after drain_timeout expires. • NONE Sessions are not terminated. However, these parameter values can be overridden using the command line.
commit_outcome	Determines whether transaction COMMIT outcome is accessible after the COMMIT has executed. While the database guarantees that COMMIT is durable, this ensures that the outcome of the COMMIT is durable. Applications use the feature to probe the status of the commit last executed after an outage, and is available to applications to determine an outcome. Note: • Invoking the GET_LTXID_OUTCOME Procedure of the DBMS_APP_CONT package requires that the commit_outcome attribute be set. • commit_outcome has no effect on active Data Guard and read-only databases. • commit_outcome is only allowed on the database service and on user-defined database services
reset_state	Clears the session state usage between requests, so that each new request starts clean (usage web and stateless applications).

Table 172-6 (Cont.) CREATE_SERVICE Procedure Parameters

Parameter	Description
retention_timeout	Used in conjunction with commit_outcome, it determines the amount of time (in seconds) that the COMMIT_OUTCOME is retained. Default is 24 hours (86400). Maximum value is 30 days (2592000).
replay_initiation_timeout	For Application Continuity, replay_initiation_timeout is the difference between the time of original execution of first operation of a request, and the time that the replay is ready to start after a successful reconnect. Replay initiation time is measured from the time that the request was originally submitted until the time that replay has connected and is ready to replay. When replay is expected, keep this value high. Default is 900 seconds.
session_state_consistency	Describes how nontransactional is changed during a request (DYNAMIC). This parameter is considered only if failover_type is set to TRANSACTION for Application Continuity. Examples of session state are NLS settings, optimizer preferences, event settings, PL/SQL global variables, temporary tables, advanced queues, LOBs, and result cache. If these values change after the request starts, set to DYNAMIC (default). Almost all applications should use DYNAMIC mode. If you are unsure, use DYNAMIC mode.
sql_translation_profile	Name of SQL translation profile.
clb_goal	Method used for Connection Load Balancing (see Table 172-2)
placement_policy	NUMBER Placement policy for the service. Possible values: • 0: PDB-NONE • 1: PDB-SINGLETON • 2: PDB-UNIFORM

> **Note:**
>
> Values other than 0 are applicable only in the ATP-Dedicated Cloud in an Oracle RAC environment.

Examples

```
DBMS_SERVICE.CREATE_SERVICE('ernie.example.com','ernie.example.com');

DECLARE
   params dbms_service.svc_parameter_array;
   BEGIN
      params('TRUE_CACHE_SERVICE')       :='C_SALES';
      params('FAILOVER_TYPE')            :='TRANSACTION';
      params('REPLAY_INITIATION_TIMEOUT'):=1800;
      params('RETENTION_TIMEOUT')        :=86400;
      params('FAILOVER_DELAY')           :=10;
      params('FAILOVER_RETRIES')         :=30;
      params('DRAIN_TIMEOUT')            :=60;
      params('STOP_OPTION')              :='DBMS_SERVICE.STOP_OPTION_IMMEDIATE';
      params('FAILOVER_RESTORE')         :='DBMS_SERVICE.FAILOVER_RESTORE_BASIC';
      params('commit_outcome')           :='true';
```

```
   params('aq_ha_notifications')      :='true';
   DBMS_SERVICE.MODIFY_SERVICE('GOLD',params);
END;
```

172.6.2 DELETE_SERVICE Procedure

This procedure deletes a service from the data dictionary.

> **Note:**
>
> Starting with Oracle Database 19c, customer use of the `SERVICE_NAME` parameter is deprecated. It can be desupported in a future release. It must not be used for high availability (HA) deployments. It is not supported to use service name parameter for any HA operations. This restriction includes FAN, load balancing, `FAILOVER_TYPE`, `FAILOVER_RESTORE`, `SESSION_STATE_CONSISTENCY`, and any other uses.
>
> You cannot use this subprogram if your services are managed by Oracle Clusterware, Oracle Restart, or Oracle Global Data Services.

Syntax

```
DBMS_SERVICE.DELETE_SERVICE(
   service_name    IN VARCHAR2);
```

Parameters

Table 172-7 DELETE_SERVICE Procedure Parameters

Parameter	Description
`service_name`	Name of the service, limited to 64 characters in the Data Dictionary

Examples

```
DBMS_SERVICE.DELETE_SERVICE('ernie.example.com');
```

172.6.3 DISCONNECT_SESSION Procedure

This procedure disconnects sessions with the named service at the current instance.

Syntax

```
DBMS_SERVICE.DISCONNECT_SESSION(
   service_name         IN VARCHAR2,
   disconnect_option    IN NUMBER DEFAULT POST_TRANSACTION;
```

Parameters

Table 172-8 DISCONNECT_SESSION Procedure Parameters

Parameter	Description
`service_name`	Name of the service, limited to 64 characters in the Data Dictionary

Table 172-8 (Cont.) DISCONNECT_SESSION Procedure Parameters

Parameter	Description
disconnect_option	The options, package constants, are expressed as NUMBER: • POST_TRANSACTION = 0: session disconnects after the current transaction commits or rolls back • IMMEDIATE = 1: session disconnects immediately • NOREPLAY = 2: session disconnects immediately and be flagged to not be replayed by application continuity, that is IMMEDIATE and NOREPLAY together **Note**: IMMEDIATE or POST_TRANSACTION and NOREPLAY is automatically translated as 1 or 0 or 2 respectively. However, passing a string literal (quoted using either the ' or " characters, such as "IMMEDIATE" or 'POST_TRANSACTION' or 'NOREPLAY') raises an error.

Usage Notes

- This procedure can be used in the context of a single instance as well as with Oracle Real Application Clusters.
- This subprogram does not return until all corresponding sessions are disconnected. Therefore, use the DBMS_JOB package or put the SQL session in background if the caller does not want to wait for all corresponding sessions to be disconnected.

Examples

This disconnects sessions with service_name 'ernie.example.com'.

```
DBMS_SERVICE.DISCONNECT_SESSION('ernie.example.com');
```

If a service is using application continuity, and you do not want the sessions replayed but simply terminated, use the following:

```
EXECUTE DBMS_SERVICE.DISCONNECT_SESSION('service name', DBMS_SERVICE.NOREPLAY);
```

172.6.4 MODIFY_SERVICE Procedure

This procedure modifies an existing service.

> **Note:**
>
> You cannot use the second version of subprogram if your services are managed by Oracle Clusterware, Oracle Restart, or Oracle Global Data Services. The version with the parameter array interface applies to databases that are not managed by Oracle Clusterware, Oracle Restart or Oracle Global Data Services. New attributes are only available using the parameter interface.

Syntax

```
DBMS_SERVICE.MODIFY_SERVICE(
   service_name               IN VARCHAR2,
   parameter_array            IN svc_parameter_array);
```

This overload is maintained for backward compatibility:

```
DBMS_SERVICE.MODIFY_SERVICE(
   service_name               IN VARCHAR2,
   goal                       IN NUMBER DEFAULT NULL,
   dtp                        IN BOOLEAN DEFAULT NULL,
   true_cache_service         IN VARCHAR2,
   aq_ha_notifications        IN BOOLEAN DEFAULT NULL,
   failover_method            IN VARCHAR2 DEFAULT NULL,
   failover_type              IN VARCHAR2 DEFAULT NULL,
   failover_retries           IN NUMBER DEFAULT NULL,
   failover_delay             IN NUMBER DEFAULT NULL,
   clb_goal                   IN NUMBER DEFAULT NULL,
   edition                    IN VARCHAR2 DEFAULT NULL,
   modify_edition             IN BOOLEAN DEFAULT FALSE;
```

Parameters

Table 172-9 MODIFY_SERVICE Procedure Parameters

Parameter	Description
service_name	Name of the service, limited to 64 characters in the Data Dictionary
parameter_array	Associative array with name/value pairs of the service attributes. Supported names: • goal • dtp • true_cache_service • aq_ha_notifications • failover_method • failover_type • failover_restore • failover_retries • failover_delay • drain_timeout • stop_option • edition • commit_outcome • retention_timeout • replay_initiation_timeout • session_state_consistency • sql_translation_profile • placement_policy
goal	Workload management goal directive for the service. Valid values: • DBMS_SERVICE.GOAL_SERVICE_TIME • DBMS_SERVICE.GOAL_THROUGHPUT • DBMS_SERVICE.GOAL_NONE
dtp	Declares the service to be for X/Open Distributed Transaction Processing (DTP) or any distributed transaction (especially XA)
true_cache_service	Name of the True Cache service being registered with the primary service.

Table 172-9 (Cont.) MODIFY_SERVICE Procedure Parameters

Parameter	Description
aq_ha_notifications	Determines whether Fast Application Notification (FAN) is enabled for OCI/OCCI/ODP. In Oracle Database12*c*, FAN uses Oracle Notification Services (ONS). This parameter is still used to enable FAN. FAN is recommended for all High Availability systems, and is on by default for Application Continuity
failover_method	Failover TYPE for the service for Application Continuity and TAF. If the failover_type is set to TRANSACTION on the service, this automatically sets COMMIT_OUTCOME to TRUE. JDBC Replay Driver uses the FAILOVER_TYPE service attribute setting of TRANSACTION for TRANSACTION failover. OCI uses the older settings of SELECT and SESSION. The server only accepts FAILOVER_METHOD = BASIC with the TRANSACTION setting.
failover_type	Failover TYPE for the service for Application Continuity and TAF.
failover_restore	For Application Continuity, when the failover_restore parameter is set, the session states are restored before replaying for ODP.NET and Java. Use LEVEL1 for ODP.NET and Java with Application Continuity to restore the initial state. For AC OCI, use NONE for applications that are not STATIC.
failover_retries	Number of connection retries for Application Continuity and TAF. Using the failover_retries and failover_delay parameters, the failover can be delayed until the service is next available. This parameter is for connecting. It does not control the number of failovers, which is 3 for each incident for Application Continuity.
failover_delay	Delay in seconds between connection retries for Application Continuity and TAF. The default is 10 seconds for Application Continuity. Do not use a 0-second delay if the service needs time to failover and register. Long delays are good for planned outages and to failover to Data Guard. Short delays work well with Oracle RAC when the service is already available.
drain_timeout	When this parameter is set, all sessions connected to that service are drained by the client drivers and pools using Fast Connection Failover (FCF). The drain_timeout can be set on the service, to stop and relocate drains for this time by default.
edition	If this argument has a non-NULL value, this provides the initial session edition for subsequent database connections using this service that do not specify an edition. If no value is specified, this argument has no effect. During service creation or modification, no validation is performed on this parameter. At connection time, if the connecting user does not have USE privilege on the edition, or the edition does not exist, this raises the error ORA-38802 (edition does not exist).

Table 172-9 (Cont.) MODIFY_SERVICE Procedure Parameters

Parameter	Description
commit_outcome	Determines whether transaction COMMIT outcome is accessible after the COMMIT has executed. While the database guarantees that COMMIT is durable, this ensures that the outcome of the COMMIT is durable.Applications use the feature to probe the status of the commit last executed after an outage, and is available to applications to determine an outcome. Note: • Invoking the GET_LTXID_OUTCOME Procedure of the DBMS_APP_CONT package requires that the commit_outcome attribute be set. • commit_outcome has no effect on active Data Guard and read-only databases. • commit_outcome is allowed only on user-defined database services
retention_timeout	Used in conjunction with commit_outcome, it determines the amount of time (in seconds) that the COMMIT_OUTCOME is retained. Default is 24 hours (86400). Maximum value is 30 days (2592000).
replay_initiation_timeout	For Application Continuity, replay_initiation_timeout is the difference between the time of original execution of first operation of a request, and the time that the replay is ready to start after a successful reconnect. Replay initiation time is measured from the time that the request was originally submitted until the time that replay has connected and is ready to replay. When replay is expected, keep this value high. Default is 900 seconds.
session_state_consistency	Describes how nontransactional is changed during a request. This parameter is considered only if failover_type is set to TRANSACTION for Application Continuity. Examples of session state are NLS settings, optimizer preferences, event settings, PL/SQL global variables, temporary tables, advanced queues, LOBs, and result cache. If these values change after the request starts, set to DYNAMIC (default). Almost all applications should use DYNAMIC mode. If you are unsure, use DYNAMIC mode.
sql_translation_profile	Name of SQL translation profile.
modify_edition	If TRUE, the edition service attribute is updated to use the edition argument value. If FALSE or NULL, the edition attribute is not updated.
clb_goal	Method used for Connection Load Balancing (see Table 172-2)
placement_policy	NUMBER Placement policy for the service. Possible values: • 0: PDB-NONE • 1: PDB-SINGLETON • 2: PDB-UNIFORM

> **Note:**
>
> Values other than 0 are applicable only in the ATP-Dedicated Cloud in an Oracle RAC environment.

Usage Notes

- If you are using Clustered Managed Services with Oracle Clusterware, or using Oracle Restart with your single instance database, you must modify services using the `srvctl` command rather than `DBMS_SERVICE`. When the service is started by Oracle Clusterware or Oracle Restart, the service is modified in the database to match the resource defined to either Oracle Clusterware or Oracle Restart. Any changes made with `DBMS_SERVICE` are lost unless they are also made with the corresponding `srvctl` command. Starting with 11.2.0.2, service attribute modifications take effect immediately when the service is started or modified by `srvctl`.
- Although users can modify the edition attribute while the service is up and running, it may not be safe to do so. Users must proceed with caution because this causes new connections to be connected at the new edition, while the existing connection is not affected. This can cause mid-tier operations to connect to the wrong edition.

172.6.5 START_SERVICE Procedure

This procedure starts a service. In Oracle RAC, implementing this option acts on the instance specified.

> **Note:**
>
> You cannot use this subprogram if your services are managed by Oracle Clusterware, Oracle Restart or Oracle Global Data Services.

Syntax

```
DBMS_SERVICE.START_SERVICE(
   service_name  IN VARCHAR2,
   instance_name IN VARCHAR2);
```

Parameters

Table 172-10 *START_SERVICE Procedure Parameters*

Parameter	Description
`service_name`	Name of the service limited to 64 characters in the Data Dictionary
`instance_name`	Name of the instance where the service must be activated (optional). `NULL` results in starting of the service on the local instance. In single instance, this can only be the current instance or `NULL`. Specify `DBMS_SERVICE.ALL_INSTANCES` to start the service on all configured instances.

Examples

```
DBMS_SERVICE.START_SERVICE('ernie.example.com');
```

172.6.6 STOP_SERVICE Procedure

This procedure stops a service.

> **Note:**
>
> You cannot use this subprogram if your services are managed by Oracle Clusterware, Oracle Restart or Oracle Global Data Services.

Syntax

```
DBMS_SERVICE.STOP_SERVICE(
   service_name   IN VARCHAR2,
   instance_name  IN VARCHAR2  DEFAULT NULL,
   stop_option    IN VARCHAR2  DEFAULT NULL,
   drain_timeout  IN NUMBER    DEFAULT NULL,
   replay         IN BOOLEAN   DEFAULT TRUE);
```

Parameters

Table 172-11 STOP_SERVICE Procedure Parameters

Parameter	Description
`service_name`	Name of the service limited to 64 characters in the Data Dictionary
`instance_name`	Name of the instance where the service must be stopped (optional). `NULL` results in stopping of the service locally. In single instance, this can only be the current instance or `NULL`. The default in Oracle RAC and exclusive case is `NULL`. Specify `DBMS_SERVICE.ALL_INSTANCES` to stop the service on all configured instances.
`stop_option`	To specify how sessions are stopped with draining. The possible values are as follows: • `IMMEDIATE`: sessions are aborted immediately after the time specified in `drain_timeout`. • `TRANASCTIONAL`: applies for transactions. After the transaction expires, the sessions are immediately terminated. • `NONE`: sessions are not terminated. These values can be overridden on the command line using `SRVCTL`.
`drain_timeout`	The time in seconds for the session to drain.
`replay`	Enable application continuity replay.

Examples

```
DBMS_SERVICE.STOP_SERVICE('ernie.example.com');
```

173
DBMS_SESSION

This package provides access to SQL `ALTER SESSION` and `SET ROLE` statements, and other session information, from PL/SQL. You can use `DBMS_SESSION` to set preferences and security levels.

This chapter contains the following topics:

- Security Model
- Operational Notes
- Data Structures
- Summary of DBMS_SESSION Subprograms

173.1 DBMS_SESSION Security Model

This package runs with the privileges of the calling user, rather than the package owner `SYS`.

173.2 DBMS_SESSION Operational Notes

You should not attempt to turn `close_cached_open_cursors` on or off.

173.3 DBMS_SESSION Data Structures

The `DBMS_SESSION` package defines `TABLE` types.

Table Types

- INTEGER_ARRAY Table Type
- LNAME_ARRAY Table Type

173.3.1 DBMS_SERVICE INTEGER_ARRAY Table Type

`INTEGER_ARRAY` is a table type of `BINARY_INTEGER`.

Syntax

```
TYPE integer_array IS TABLE OF BINARY_INTEGER INDEX BY BINARY_INTEGER;
```

173.3.2 DBMS_SERVICE LNAME_ARRAY Table Type

`LNAME_ARRAY` is a table type of `VARCHAR2`.

Syntax

```
TYPE lname_array IS TABLE OF VARCHAR2(4000) INDEX BY BINARY_INTEGER;
```

173.4 Summary of DBMS_SESSION Subprograms

This table lists the DBMS_SESSION subprograms in alphabetical order and briefly describes them.

Table 173-1 ***DBMS_SESSION Package Subprograms***

Subprogram	Description
CLEAR_ALL_CONTEXT Procedure	Clears all context information.
CLEAR_ALL_LOCAL_CONTEXTS Procedure	Clears all session based application context values that are set in the current session.
CLEAR_CONTEXT Procedure	Clears the context.
CLEAR_IDENTIFIER Procedure	Clears the identifier.
CLOSE_DATABASE_LINK Procedure	Closes database link.
CURRENT_IS_ROLE_ENABLED Function	Determines if the named role is currently enabled.
FREE_UNUSED_USER_MEMORY Procedure	Lets you reclaim unused memory after performing operations requiring large amounts of memory.
GET_PACKAGE_MEMORY_UTILIZATION Procedure	Describes static package memory usage.
IS_ROLE_ENABLED Function	Determines if the named role is enabled for the session.
IS_SESSION_ALIVE Function	Determines if the specified session is active.
LIST_CONTEXT Procedures	Returns a list of active namespace and context for the current session.
MODIFY_PACKAGE_STATE Procedure	Used to perform various actions (as specified by the `action_flags` parameter) on the session state of all PL/SQL program units active in the session.
RESET_PACKAGE Procedure	De-instantiates all packages in the session.
SESSION_IS_ROLE_ENABLED Function	Determines if the named role is enabled at the session level.
SESSION_TRACE_DISABLE Procedure	Resets the session-level SQL trace for the session from which it was called.
SESSION _TRACE_ENABLE Procedure	Enables session-level SQL trace for the invoking session.
SET_CONTEXT Procedure	Sets or resets the value of a context attribute.
SET_EDITION_DEFERRED Procedure	Requests a switch to the specified edition.
SET_IDENTIFIER Procedure	Sets the identifier.
SET_NLS Procedure	Sets Globalization Support (NLS).
SET_ROLE Procedure	Sets role.
SET_SQL_TRACE Procedure	Turns tracing on or off.
SLEEP Procedure	Suspends the session for a specified period of time.
SWITCH_CURRENT_CONSUMER_GROUP Procedure	Facilitates changing the current resource consumer group of a user's current session.
UNIQUE_SESSION_ID Function	Returns an identifier that is unique for all sessions currently connected to this database.

173.4.1 CLEAR_ALL_CONTEXT Procedure

This procedure clears application context information in the specified namespace.

Syntax

```
DBMS_SESSION.CLEAR_ALL_CONTEXT
   namespace         VARCHAR2);
```

Parameters

Table 173-2 CLEAR_ALL_CONTEXT Procedure Parameters

Parameter	Description
namespace	The namespace where the application context information is to be cleared. Required.

Usage Notes

- This procedure must be invoked directly or indirectly by the trusted package.
- Any changes in context value are reflected immediately and subsequent calls to access the value through `SYS_CONTEXT` return the most recent value.

173.4.2 CLEAR_ALL_LOCAL_CONTEXTS Procedure

This procedure clears all session-based application context values that are set in the current session and are initialized locally, externally, and globally, but cannot be accessed globally.

Syntax

```
DBMS_SESSION.CLEAR_ALL_LOCAL_CONTEXTS;
```

Usage Notes

Any public user with the `CLEAR ALL LOCAL CONTEXTS` privilege can execute this procedure. If a user without the `CLEAR ALL LOCAL CONTEXTS` privilege executes this procedure, it fails with the following error:

```
ORA-01056: User $USER_NAME does not have CLEAR ALL LOCAL CONTEXTS privilege
```

173.4.3 CLEAR_CONTEXT Procedure

This procedure clears application context in the specified namespace.

Syntax

```
DBMS_SESSION.CLEAR_CONTEXT
   namespace         VARCHAR2,
   client_identifier VARCHAR2
   attribute         VARCHAR2);
```

Parameters

Table 173-3 CLEAR_CONTEXT Procedure Parameters

Parameter	Description
namespace	Namespace in which the application context is to be cleared. Required. For a session-local context, namespace must be specified. If namespace is defined as Session Local Context, then client_identifier is optional since it is only associated with a globally accessed context. For a globally accessed context, namespace must be specified. NULL is a valid value for client_identifier because a session with no identifier set can see a context that looks like the (namespace, attribute, value, username, null) set using SET_CONTEXT.
client_identifier	Applies to a global context and is optional for other types of contexts; 64-byte maximum
attribute	Specific attribute in the namespace to be cleared. Optional. the default is NULL. If you specify attribute as NULL, then (namespace, attribute, value) for that namespace are cleared from the session. If attribute is not specified, then all context information that has the namespace and client_identifier arguments is cleared.

Usage Notes

- This procedure must be invoked directly or indirectly by the trusted package.
- Any changes in context value are reflected immediately and subsequent calls to access the value through SYS_CONTEXT return the most recent value.

173.4.4 CLEAR_IDENTIFIER Procedure

This procedure removes the set_client_id in the session.

Syntax

```
DBMS_SESSION.CLEAR_IDENTIFIER;
```

Usage Notes

This procedure is executable by public.

173.4.5 CLOSE_DATABASE_LINK Procedure

This procedure closes an open database link. It is equivalent to the following SQL statement:
ALTER SESSION CLOSE DATABASE LINK <name>

Syntax

```
DBMS_SESSION.CLOSE_DATABASE_LINK (
   dblink VARCHAR2);
```

Parameters

Table 173-4 CLOSE_DATABASE_LINK Procedure Parameters

Parameter	Description
dblink	Name of the database link to close

173.4.6 CURRENT_IS_ROLE_ENABLED Function

This function determines if the named role is currently enabled.

Syntax

```
DBMS_SESSION.CURRENT_IS_ROLE_ENABLED (
   rolename    VARCHAR2)
  RETURN BOOLEAN;
```

Parameters

Table 173-5 CURRENT_IS_ROLE_ENABLED Function Parameters

Parameter	Description
rolename	Name of the role.

Return Values

- TRUE-if the role is enabled.
- FALSE-if the role is not enabled.

173.4.7 FREE_UNUSED_USER_MEMORY Procedure

This procedure reclaims unused memory after performing operations requiring large amounts of memory (more than 100K).

Examples of operations that use large amounts of memory include:

- Large sorting where entire sort_area_size is used and sort_area_size is hundreds of KB.
- Compiling large PL/SQL packages, procedures, or functions.
- Storing hundreds of KB of data within PL/SQL indexed tables.

You can monitor user memory by tracking the statistics "session UGA memory" and "session PGA memory" in the v$sesstat or v$statname fixed views. Monitoring these statistics also shows how much memory this procedure has freed.

> **Note:**
>
> This procedure should only be used in cases where memory is at a premium. It should be used infrequently and judiciously.

Syntax

```
DBMS_SESSION.FREE_UNUSED_USER_MEMORY;
```

Return Values

The behavior of this procedure depends upon the configuration of the server operating on behalf of the client:

- **Dedicated server**: This returns unused PGA memory and session memory to the operating system. Session memory is allocated from the PGA in this configuration.
- **Shared server**: This returns unused session memory to the `shared_pool`. Session memory is allocated from the `shared_pool` in this configuration.

Usage Notes

In order to free memory using this procedure, the memory must not be in use.

After an operation allocates memory, only the same type of operation can reuse the allocated memory. For example, after memory is allocated for sort, even if the sort is complete and the memory is no longer in use, only another sort can reuse the sort-allocated memory. For both sort and compilation, after the operation is complete, the memory is no longer in use, and the user can call this procedure to free the unused memory.

An indexed table implicitly allocates memory to store values assigned to the indexed table's elements. Thus, the more elements in an indexed table, the more memory the RDBMS allocates to the indexed table. As long as there are elements within the indexed table, the memory associated with an indexed table is in use.

The scope of indexed tables determines how long their memory is in use. Indexed tables declared globally are indexed tables declared in packages or package bodies. They allocate memory from session memory. For an indexed table declared globally, the memory remains in use for the lifetime of a user's login (lifetime of a user's session), and is freed after the user disconnects from ORACLE.

Indexed tables declared locally are indexed tables declared within functions, procedures, or anonymous blocks. These indexed tables allocate memory from PGA memory. For an indexed table declared locally, the memory remains in use for as long as the user is still running the procedure, function, or anonymous block in which the indexed table is declared.After the procedure, function, or anonymous block is finished running, the memory is then available for other locally declared indexed tables to use (in other words, the memory is no longer in use).

Assigning an uninitialized, "empty" indexed table to an existing index table is a method to explicitly re-initialize the indexed table and the memory associated with the indexed table. After this operation, the memory associated with the indexed table is no longer in use, making it available to be freed by calling this procedure. This method is particularly useful on indexed tables declared globally which can grow during the lifetime of a user's session, as long as the user no longer needs the contents of the indexed table.

The memory rules associated with an indexed table's scope still apply; this method and this procedure, however, allow users to intervene and to explicitly free the memory associated with an indexed table.

Examples

The following PL/SQL illustrates the method and the use of procedure `FREE_UNUSED_USER_MEMORY`.

```
CREATE PACKAGE foobar
   type number_idx_tbl is table of number indexed by binary_integer;

   store1_table  number_idx_tbl;     --  PL/SQL indexed table
   store2_table  number_idx_tbl;     --  PL/SQL indexed table
   store3_table  number_idx_tbl;     --  PL/SQL indexed table
   ...
END;            --  end of foobar

DECLARE
   ...
   empty_table   number_idx_tbl;     --  uninitialized ("empty") version
BEGIN
   FOR i in 1..1000000 loop
     store1_table(i) := i;           --  load data
   END LOOP;
   ...
   store1_table := empty_table;      --  "truncate" the indexed table
   ...
   -
   dbms_session.free_unused_user_memory;  -- give memory back to system

   store1_table(1) := 100;           --  index tables still declared;
   store2_table(2) := 200;           --  but truncated.
   ...
END;
```

173.4.8 GET_PACKAGE_MEMORY_UTILIZATION Procedure

This procedure describes static package memory usage.

The output collections describe memory usage in each instantiated package. Each package is described by its owner name, package name, used memory amount, and unused allocated memory amount. The amount of unused memory is greater than zero because of memory fragmentation and also because once used free memory chunks initially go to a free list owned by the package memory heap. They are released back to the parent heap only when the FREE_UNUSED_USER_MEMORY Procedure is invoked.

Syntax

```
DBMS_SESSION.GET_PACKAGE_MEMORY_UTILIZATION (
   owner_names     OUT NOCOPY LNAME_ARRAY,
   unit_names      OUT NOCOPY LNAME_ARRAY,
   unit_types      OUT NOCOPY INTEGER_ARRAY,
   used_amounts    OUT NOCOPY INTEGER_ARRAY,
   free_amounts    OUT NOCOPY INTEGER_ARRAY);
```

Parameters

Table 173-6 GET_PACKAGE_MEMORY_UTILIZATION Function Parameters

Parameter	Description
owner_name	Owner of package
unit_name	Name of package
unit_types	Value of the type# columns of the dictionary table obj$
used_amounts	Amount of allocated memory specified in bytes
free_amounts	Amount of available memory specified in bytes

173.4.9 IS_ROLE_ENABLED Function

This function determines if the named role is enabled for this session.

> **Note:**
>
> This function is deprecated starting in Oracle Database 19c. Use `DBMS_SESSION.CURRENT_IS_ROLE_ENABLED` or `DBMS_SESSION.SESSION_IS_ROLE_ENABLED` instead.

Syntax

```
DBMS_SESSION.IS_ROLE_ENABLED (
   rolename    VARCHAR2)
  RETURN BOOLEAN;
```

Parameters

Table 173-7 IS_ROLE_ENABLED Function Parameters

Parameter	Description
`rolename`	Name of the role.\

Return Values

Table 173-8 IS_ROLE_ENABLED Function Return Values

Return	Description
`is_role_enabled`	`TRUE` or `FALSE`, depending on whether the role is enabled

173.4.10 IS_SESSION_ALIVE Function

This function determines if the specified session is active.

Syntax

```
DBMS_SESSION.IS_SESSION_ALIVE (
   uniqueid VARCHAR2)
  RETURN BOOLEAN;
```

Parameters

Table 173-9 IS_SESSION_ALIVE Function Parameters

Parameter	Description
`uniqueid`	Unique ID of the session: This is the same one as returned by `UNIQUE_SESSION_ID`.

Return Values

Table 173-10 IS_SESSION_ALIVE Function Return Values

Return	Description
`is_session_alive`	`TRUE` or `FALSE`, depending on whether the session is active

173.4.11 LIST_CONTEXT Procedures

This procedure returns a list of active namespaces and contexts for the current session.

Syntax

```
TYPE AppCtxRecTyp IS RECORD (
   namespace VARCHAR2(30),
   attribute VARCHAR2(30),
   value     VARCHAR2(256));

TYPE AppCtxTabTyp IS TABLE OF AppCtxRecTyp INDEX BY BINARY_INTEGER;

DBMS_SESSION.LIST_CONTEXT (
   list OUT AppCtxTabTyp,
   size OUT NUMBER);
```

Parameters

Table 173-11 LIST_CONTEXT Procedure Parameters

Parameter	Description
`list`	Buffer to store a list of application context set in the current session

Return Values

Table 173-12 LIST_CONTEXT Procedure Return Values

Return	Description
`list`	A list of (namespace, attribute, values) set in current session
`size`	Returns the number of entries in the buffer returned

Usage Notes

The context information in the list appears as a series of <`namespace`> <`attribute`> <`value`>. Because `list` is a table type variable, its size is dynamically adjusted to the size of returned list.

173.4.12 MODIFY_PACKAGE_STATE Procedure

This procedure is used to perform various actions (as specified by the `action_flags` parameter) on the session state of all PL/SQL program units active in the session. This takes effect after the PL/SQL call that made the current invocation finishes running.

The procedure uses the `DBMS_SESSION` constants listed in Table 173-14.

Syntax

```
DBMS_SESSION.MODIFY_PACKAGE_STATE(
   action_flags IN PLS_INTEGER);
```

Parameters

Table 173-13 MODIFY_PACKAGE_STATE Procedure Parameters

Parameter	Description
action_flags	Bit flags that determine the action taken on PL/SQL program units: DBMS_SESSION.FREE_ALL_RESOURCES (or 1)—frees all memory associated with each of the previously run PL/SQL programs from the session. Clears the current values of any package globals and closes cached cursors. On subsequent use, the PL/SQL program units are reinstantiated and package globals are reinitialized. Invoking MODIFY_PACKAGE_STATE with the DBMS_SESSION.FREE_ALL_RESOURCES parameter provides functionality identical to the DBMS_SESSION.RESET_PACKAGE() interface. DBMS_SESSION.REINITIALIZE (or 2)—reinitializes packages without actually being freed and recreated from scratch. Instead the package memory is reused. In terms of program semantics, the DBMS_SESSION.REINITIALIZE flag is similar to the DBMS_SESSION.FREE_ALL_RESOURCES flag in that both have the effect of reinitializing all packages. However, DBMS_SESSION.REINITIALIZE should exhibit better performance than the DBMS_SESSION.FREE_ALL_RESOURCES option because: • Packages are reinitialized without actually being freed and recreated from scratch. Instead the package memory gets reused. • Any open cursors are closed, semantically speaking. However, the cursor resource is not actually freed. It is simply returned to the PL/SQL cursor cache. The cursor cache is not flushed. Hence, cursors corresponding to frequently accessed static SQL in PL/SQL remains cached in the PL/SQL cursor cache and the application does not incur the overhead of opening, parsing, and closing a new cursor for those statements on subsequent use. • The session memory for PL/SQL modules without global state (such as types, stored-procedures) are not freed and recreated.

Usage Notes

See the parameter descriptions in Table 173-16 for the differences between the flags and why DBMS_SESSION.REINITIALIZE exhibits better performance than DBMS_SESSION.FREE_ALL_RESOURCES.

Table 173-14 Action_flags Constants for MODIFY_PACKAGE_STATE

Constant	Description
FREE_ALL_RESOURCES	PLS_INTEGER:= 1
REINITIALIZE	PLS_INTEGER:= 2

- Reinitialization refers to the process of resetting all package variables to their initial values and running the initialization block (if any) in the package bodies. Consider the package:

```
package P is
  n number;
  m number := P2.foo;
  d date := SYSDATE;
  cursor c is select * from emp;
  procedure bar;
end P;
/
package body P is
  v   varchar2(20) := 'hello';
  procedure bar is
  begin
    ...
  end;
  procedure init_pkg is
  begin
    ....
  end;
begin
  -- initialization block
  init_pkg;
  ...
  ...
end P;
/
```

For the package P, reinitialization involves:

- Setting `P.n` to `NULL`
- Invoking function `P2.foo` and setting `P.m` to the value returned from `P2.foo`
- Setting `P.d` to the return value of `SYSDATE` built-in
- Closing cursor `P.c` if it was previously opened
- Setting `P.v` to 'hello'
- Running the initialization block in the package body
- The reinitialization for a package is done only if the package is actually referenced subsequently. Furthermore, the packages are reinitialized in the order in which they are referenced subsequently.
- When using `FREE_ALL_RESOURCES` or `REINITIALIZE`, make sure that resetting package variable values does not affect the application.
- Because `DBMS_SESSION.REINITIALIZE` does not actually cause all the package state to be freed, in some situations, the application could use significantly more session memory than if the `FREE_ALL_RESOURCES` flag or the `RESET_PACKAGE` procedure had been used. For instance, after performing `DBMS_SESSION.MODIFY_PACKAGE_STATE(DBMS_SESSION.REINITIALIZE)`, if the application does not refer to many of the packages that were previously referenced, then the session memory for those packages remains until the end of the session (or until `DBMS_SESSION.RESET_PACKAGE` is called).
- Because the client-side PL/SQL code cannot reference remote package variables or constants, you must explicitly use the values of the constants. For example, `DBMS_SESSION.MODIFY_PACKAGE_STATE(DBMS_SESSION.REINITIALIZE)` does not compile on the client because it uses the constant `DBMS_SESSION.REINITIALIZE`.

  Instead, use `DBMS_SESSION.MODIFY_PACKAGE_STATE(2)` on the client, because the argument is explicitly provided.

Examples

This example illustrates the use of `DBMS_SESSION.MODIFY_PACKAGE_STATE`. Consider a package `P` with some global state (a cursor `c` and a number `cnt`). When the package is first initialized, the package variable `cnt` is `0` and the cursor `c` is `CLOSED`. Then, in the session, change the value of `cnt` to `111` and also execute an `OPEN` operation on the cursor. If you call `print_status` to display the state of the package, you see that `cnt` is `111` and that the cursor is `OPEN`. Next, call `DBMS_SESSION.MODIFY_PACKAGE_STATE`. If you print the status of the package `P` again using `print_status`, you see that `cnt` is `0` again and the cursor is `CLOSED`. If the call to `DBMS_SESSION.MODIFY_PACKAGE_STATE` had not been made, then the second `print_status` would have printed `111` and `OPEN`.

```
create or replace package P is
  cnt    number := 0;
  cursor c is select * from emp;
  procedure print_status;
end P;
/
show errors;

create or replace package body P is
  procedure print_status is
  begin
    dbms_output.put_line('P.cnt = ' || cnt);
    if c%ISOPEN then
      dbms_output.put_line('P.c is OPEN');
    else
      dbms_output.put_line('P.c is CLOSED');
    end if;
  end;
end P;
/
show errors;

SQL> set serveroutput on;
SQL> begin
  2   P.cnt := 111;
  3   open p.c;
  4   P.print_status;
  5  end;
  6  /
P.cnt = 111
P.c is OPEN

PL/SQL procedure successfully completed.

SQL> begin
  2   dbms_session.modify_package_state(dbms_session.reinitialize);
  3  end;
  4  /

PL/SQL procedure successfully completed.

SQL> set serveroutput on;
SQL>
SQL> begin
  2   P.print_status;
  3  end;
  4  /
P.cnt = 0
```

```
P.c is CLOSED

PL/SQL procedure successfully completed.
```

173.4.13 RESET_PACKAGE Procedure

This procedure de-instantiates all packages in this session. It frees the package state.

> **Note:**
>
> See "SESSION _TRACE_ENABLE Procedure" . The MODIFY_PACKAGE_STATE interface, introduced in Oracle9i, provides an equivalent of the RESET_PACKAGE capability. It is an efficient, lighter-weight variant for reinitializing the state of all PL/SQL packages in the session.

Memory used for caching the execution state is associated with all PL/SQL functions, procedures, and packages that were run in a session.

For packages, this collection of memory holds the current values of package variables and controls the cache of cursors opened by the respective PL/SQL programs. A call to `RESET_PACKAGE` frees the memory associated with each of the previously run PL/SQL programs from the session, and, consequently, clears the current values of any package globals and closes any cached cursors.

`RESET_PACKAGE` can also be used to reliably restart a failed program in a session. If a program containing package variables fails, then it is hard to determine which variables need to be reinitialized. `RESET_PACKAGE` guarantees that all package variables are reset to their initial values.

Syntax

```
DBMS_SESSION.RESET_PACKAGE;
```

Usage Notes

Because the amount of memory consumed by all executed PL/SQL can become large, you might use `RESET_PACKAGE` to trim down the session memory footprint at certain points in your database application. However, make sure that resetting package variable values does not affect the application. Also, remember that later execution of programs that have lost their cached memory and cursors will perform slower, because they need to re-create the freed memory and cursors.

`RESET_PACKAGE` does not free the memory, cursors, and package variables immediately when called.

> **Note:**
>
> `RESET_PACKAGE` only frees the memory, cursors, and package variables after the PL/SQL call that made the invocation finishes running.

For example, PL/SQL procedure `P1` calls PL/SQL procedure `P2`, and `P2` calls `RESET_PACKAGE`. The `RESET_PACKAGE` effects do not occur until procedure `P1` finishes execution (the PL/SQL call ends).

Examples

This SQL*Plus script runs a large program with many PL/SQL program units that may or may not use global variables, but it doesn't need them beyond this execution:

```
EXCECUTE large_plsql_program1;
```

To free up PL/SQL cached session memory:

```
EXECUTE DBMS_SESSION.RESET_PACKAGE;
```

To run another large program:

```
EXECUTE large_plsql_program2;
```

173.4.14 SESSION_IS_ROLE_ENABLED Function

This function determines if the named role is enabled for the login user session.

Syntax

```
DBMS_SESSION.SESSION_IS_ROLE_ENABLED (
   rolename    VARCHAR2)
  RETURN BOOLEAN;
```

Parameters

Table 173-15 SESSION_IS_ROLE_ENABLED Function Parameters

Parameter	Description
`rolename`	Name of the role.

Return Values

- `TRUE`-if the role is enabled.
- `FALSE`-if the role is not enabled.

173.4.15 SESSION_TRACE_DISABLE Procedure

This procedure resets the session-level SQL trace for the session from which it was called. Client ID and service/module/action traces are not affected.

Syntax

```
DBMS_SESSION.SESSION_TRACE_DISABLE;
```

173.4.16 SESSION _TRACE_ENABLE Procedure

This procedure enables session-level SQL trace for the invoking session. Invoking this procedure results in SQL tracing of every SQL statement issued by the session.

Syntax

```
DBMS_SESSION.SESSION_TRACE_ENABLE(
   waits     IN   BOOLEAN DEFAULT TRUE,
   binds     IN   BOOLEAN DEFAULT FALSE,
   plan_stat IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 173-16 SESSION_TRACE_ENABLE Procedure Parameters

Parameter	Description
`waits`	Specifies if wait information is to be traced
`binds`	Specifies if bind information is to be traced
`plan_stat`	Frequency at which we dump row source statistics. Value should be 'NEVER', 'FIRST_EXECUTION' (equivalent to NULL) or 'ALL_EXECUTIONS'.

173.4.17 SET_CONTEXT Procedure

This procedure sets the context, of which there are four types: session local, globally initialized, externally initialized, and globally accessed.

Of the five parameters, only the first three are required; the final two parameters are optional, used only in globally accessed contexts. Further parameter information appears in the parameter table and the usage notes.

Syntax

```
DBMS_SESSION.SET_CONTEXT (
   namespace VARCHAR2,
   attribute VARCHAR2,
   value     VARCHAR2,
   username  VARCHAR2,
   client_id VARCHAR2 );
```

Parameters

Table 173-17 SET_CONTEXT Procedure Parameters

Parameter	Description
`namespace`	The namespace of the application context to be set, limited to 128 bytes. Exceeding the maximum permissible length will result in an error during the execution of the procedure.
`attribute`	The attribute of the application context to be set, limited to 128 bytes. Exceeding the maximum permissible length will result in an error during the execution of the procedure.
`value`	Value of the application context to be set, limited to 4 kilobytes.

Table 173-17 (Cont.) SET_CONTEXT Procedure Parameters

Parameter	Description
username	Database username attribute of the application context, limited to 128 bytes. Exceeding the maximum permissible length will result in an error during the execution of the procedure. The default value is NULL.
client_id	Application-specific client_id attribute of the application context, limited to 64 bytes. The default value is NULL.

Usage Notes

- The first three parameters are required for all types of context.
- The username parameter must be a valid SQL identifier.
- The client_id parameter must be a string of at most 64 bytes. It is case-sensitive and must match the argument provided for set_identifier.
- If the namespace parameter is a global context namespace, then the username parameter is matched against the current database user name in the session, and the client_id parameter is matched against the current client_id in the session. If these parameters are not set, NULL is assumed, enabling any user to see the context values.
- This procedure must be invoked directly or indirectly by the trusted package.
- The caller of SET_CONTEXT must be in the calling stack of a procedure that has been associated to the context namespace through a CREATE CONTEXT statement. The checking of the calling stack does not cross a DBMS boundary.
- No limit applies to the number of attributes that can be set in a namespace. An attribute retains its value during the user's session unless it is reset by the user.
- If the value of the parameter in the namespace has been set, SET_CONTEXT overwrites this value.
- Any changes in context value are reflected immediately and subsequent calls to access the value through SYS_CONTEXT return the most recent value.

See Also:

Oracle Database Security Guide for more information about

- "Setting the username and client ID"
- "Example: Creating a Global Application Context that Uses a Client Session ID"

173.4.18 SET_EDITION_DEFERRED Procedure

This procedure requests a switch to the specified edition. The switch takes effect at the end of the current client call.

Syntax

```
DBMS_SESSION.SET_EDITION_DEFERRED (
   edition    IN    VARCHAR2);
```

Parameters

Table 173-18 SET_EDITION_DEFERRED Procedure Parameters

Parameter	Description
`edition`	Name of the edition to which to switch. The contents of the string are processed as a SQL identifier; double quotation marks must surround the remainder of the string if special characters or lower case characters are present in the edition's actual name and, if double quotation marks are not used, the contents are set in uppercase. The caller must have `USE` privilege on the named edition.

173.4.19 SET_IDENTIFIER Procedure

This procedure sets the client ID in the session.

Syntax

```
DBMS_SESSION.SET_IDENTIFIER (
   client_id VARCHAR2);
```

Parameters

Table 173-19 SET_IDENTIFIER Procedure Parameters

Parameter	Description
`client_id`	Case-sensitive application-specific identifier of the current database session. The maximum number of bytes for this parameter is 64 bytes. If the input exceeds 64 bytes, then `ORA-28264` is raised.

Usage Notes

- `SET_IDENTIFIER` sets the session's client id to the given value. This value can be used to identify sessions in `v$session` by means of `v$session.client_identifier`. It can also be used to identify sessions by means of `sys_context('USERENV','CLIENT_IDENTIFIER')`.
- This procedure is executable by `PUBLIC`.

173.4.20 SET_NLS Procedure

This procedure sets up your Globalization Support (NLS). It is equivalent to the SQL statement: `ALTER SESSION SET <nls_parameter> = <value>`.

Syntax

```
DBMS_SESSION.SET_NLS (
   param VARCHAR2,
   value VARCHAR2);
```

Parameters

Table 173-20 SET_NLS Procedure Parameters

Parameter	Description
`param`	Globalization Support parameter. The parameter name must begin with 'NLS'.
`value`	Parameter value. If the parameter is a text literal, then it needs embedded single-quotes. For example, "`set_nls ('nls_date_format','"DD-MON-YY"')`".

173.4.21 SET_ROLE Procedure

This procedure enables and disables roles. It is equivalent to the `SET ROLE` SQL statement.

Syntax

```
DBMS_SESSION.SET_ROLE (
   role_cmd VARCHAR2);
```

Parameters

Table 173-21 SET_ROLE Procedure Parameters

Parameter	Description
`role_cmd`	Text is appended to "set role" and then run as SQL

Usage Notes

Note that the procedure creates a new transaction if it is not invoked from within an existing transaction.

173.4.22 SET_SQL_TRACE Procedure

This procedure turns tracing on or off. It is equivalent to the SQL statement `ALTER SESSION SET SQL_TRACE ...`.

Syntax

```
DBMS_SESSION.SET_SQL_TRACE (
   sql_trace boolean);
```

Parameters

Table 173-22 SET_SQL_TRACE Procedure Parameters

Parameter	Description
sql_trace	TRUE turns tracing on, FALSE turns tracing off

173.4.23 SLEEP Procedure

This procedure suspends the session for a specified period of time.

Syntax

```
DBMS_SESSION.SLEEP (
   seconds  IN NUMBER);
```

Parameters

Table 173-23 SLEEP Procedure Parameters

Parameter	Description
seconds	Amount of time, in seconds, to suspend the session. The smallest increment can be entered in hundredths of a second; for example, 1.95 is a legal time value.

173.4.24 SWITCH_CURRENT_CONSUMER_GROUP Procedure

This procedure changes the current resource consumer group of a user's current session.

This lets you switch to a consumer group if you have the switch privilege for that particular group. If the caller is another procedure, then this enables the user to switch to a consumer group for which the owner of that procedure has switch privilege.

Syntax

```
DBMS_SESSION.switch_current_consumer_group (
   new_consumer_group     IN  VARCHAR2,
   old_consumer_group     OUT VARCHAR2,
   initial_group_on_error IN  BOOLEAN);
```

Parameters

Table 173-24 SWITCH_CURRENT_CONSUMER_GROUP Procedure Parameters

Parameter	Description
new_consumer_group	Name of consumer group to which you want to switch
old_consumer_group	Name of the consumer group from which you just switched out
initial_group_on_error	If TRUE, then sets the current consumer group of the caller to his/her initial consumer group in the event of an error

Return Values

This procedure outputs the old consumer group of the user in the parameter `old_consumer_group`.

> **Note:**
>
> - The `old_consumer_group` parameter returns the name of old consumer group only if it were set explicitly. That is, you might get NULL if the old consumer group was set by some mapping rules.
> - You can switch back to the old consumer group later using the value returned in `old_consumer_group`.

Exceptions

Table 173-25 SWITCH_CURRENT_CONSUMER_GROUP Procedure Exceptions

Exception	Description
29368	Non-existent consumer group
1031	Insufficient privileges
29396	Cannot switch to `OTHER_GROUPS` consumer group

Usage Notes

The owner of a procedure must have privileges on the group from which a user was switched (`old_consumer_group`) in order to switch them back. There is one exception: The procedure can always switch the user back to his/her initial consumer group (skipping the privilege check).

By setting `initial_group_on_error` to `TRUE`, `SWITCH_CURRENT_CONSUMER_GROUP` puts the current session into the default group, if it can't put it into the group designated by `new_consumer_group`. The error associated with the attempt to move a session into `new_consumer_group` is raised, even though the current consumer group has been changed to the initial consumer group.

Examples

```
CREATE OR REPLACE PROCEDURE high_priority_task is
   old_group varchar2(30);
   prev_group varchar2(30);
   curr_user varchar2(30);
BEGIN
  -- switch invoker to privileged consumer group. If we fail to do so, an
  -- error is thrown, but the consumer group does not change
  -- because 'initial_group_on_error' is set to FALSE

  dbms_session.switch_current_consumer_group('tkrogrp1', old_group, FALSE);
  -- set up exception handler (in the event of an error, we do not want to
  -- return to caller while leaving the session still in the privileged
  -- group)

  BEGIN
```

```
    -- perform some operations while under privileged group

  EXCEPTION
    WHEN OTHERS THEN
     -- It is possible that the procedure owner does not have privileges
     -- on old_group. 'initial_group_on_error' is set to TRUE to make sure
     -- that the user is moved out of the privileged group in such a
     -- situation

     dbms_session.switch_current_consumer_group(old_group,prev_group,TRUE);
     RAISE;
    END;

  -- we've succeeded. Now switch to old_group, or if cannot do so, switch
  -- to caller's initial consumer group

  dbms_session.switch_current_consumer_group(old_group,prev_group,TRUE);
END high_priority_task;
/
```

173.4.25 UNIQUE_SESSION_ID Function

This function returns an identifier that is unique for all sessions currently connected to this database. Multiple calls to this function during the same session always return the same result.

Syntax

```
DBMS_SESSION.UNIQUE_SESSION_ID
  RETURN VARCHAR2;
```

Pragmas

```
pragma restrict_references(unique_session_id,WNDS,RNDS,WNPS);
```

Return Values

Table 173-26 UNIQUE_SESSION_ID Function Return Values

Return	Description
unique_session_id	Returns up to 24 bytes

173.4.26 USE_DEFAULT_EDITION_ALWAYS procedure

This procedure turns a mode 'on' or 'off' that disassociates the session from an edition at the end of each and every client call.

Following each call and before the next top level call, no edition will be in use by the session. The session will use the reigning database default edition on its next operation.

Syntax

```
DBMS_SESSION.USE_DEFAULT_EDITION_ALWAYS(
                      mode_on IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 173-27 USE_DEFAULT_EDITION_ALWAYS Procedure Parameters

Parameter	Description
mode_on	TRUE to turn on this mode, FALSE to turn off this mode.

Usage Notes

- A choice of session edition that might otherwise occur due to use of ALTER SESSION SET EDITION, SET_EDITION_DEFERRED, USE_DEFAULT_EDITION_DEFERRED, or the edition-choosing aspect of ALTER SESSION SET CONTAINER will be overridden by this mode.
- Turn this mode off before invoking the above mechanisms to avoid the overriding effect.

173.4.27 USE_DEFAULT_EDITION_DEFERRED procedure

This procedure disassociates the session from an edition.

The reset takes effect at the end of the current client call. Following this call, and before the next top level call, no edition will be in use by the session. The session will use the reigning database default edition on its next operation.

Syntax

```
DBMS_SESSION.USE_DEFAULT_EDITION_DEFERRED
```

174
DBMS_SFW_ACL_ADMIN

The DBMS_SFW_ACL_ADMIN package provides interfaces for administering and managing access control policies for the "database service firewall" feature. Each policy is represented by an access control list (ACL) containing hosts that are allowed access to a specific database service. Local listeners and server processes validate all inbound client connections against the ACL.

This chapter contains the following topics:

- DBMS_SFW_ACL_ADMIN Security Model
- DBMS_SFW_ACL_ADMIN Operational Notes
- DBMS_SFW_ACL_ADMIN Examples
- Summary of DBMS_SFW_ACL_ADMIN Subprograms

174.1 DBMS_SFW_ACL_ADMIN Security Model

This package is owned by the DBSFWUSER schema. The procedures in this package can be run only by the DBSFWUSER user.

174.2 DBMS_SFW_ACL_ADMIN Operational Notes

These operation notes apply to `DBMS_SFW_ACL_ADMIN`.

- An ACL contains entries, which are called "ace", for "access control entries".
- You add entries to an ACL using the IP_ADD_ACE Procedure and IP_ADD_PDB_ACE Procedure. After calling these procedures, you call the COMMIT_ACL Procedure to send the updated ACL to the listeners. Similarly, if you remove entries from an ACL using the IP_REMOVE_ACE Procedure, IP_REMOVE_ACL Procedure, IP_REMOVE_PDB_ACE Procedure, or IP_REMOVE_PDB_ACL Procedure, you need to call theCOMMIT_ACL Procedure to update the ACL on the listeners.
- Access control must be enabled using the new `FIREWALL` endpoint attribute and the new `LOCAL_REGISTRATION_ADDRESS_`*listenerName* parameter. The configuration can be done manually in `listener.ora` or through the Server Control Utility (SRVCTL). Refer to the *Oracle Database Net Services Reference* and the *Oracle Real Application Clusters Administration and Deployment Guide* for configuration instructions.
- You can query the `IP_ACL` table to show the committed ACLs. But to see the ACLs that have been sent to the local listeners, you have to query the `V$IP_ACL` or `GV$IP_ACL` view. The `IP_ACL` table can contain ACLs that are not in `[G]V$IP_ACL` because the database services for those ACLs were not running at the time of the commit. When the services are running, you can call the COMMIT_ACL procedure again to send the committed ACLs in the IP_ACL table to the local listeners.

  In an Oracle RAC environment, `GV$IP_ACL` can be used to query ACLs across the database cluster, and `V$IP_ACL` to query ACLs in the connected instance.

174.3 DBMS_SFW_ACL_ADMIN Examples

These three examples show how `DBMS_SFW_ACL_ADMIN` can be used to administer and manage access control policies.

The following example adds three access control entries to the ACL and commits them.

```
## Connect to DBSFWUSER
SQL> connect dbsfwuser/<password>
Connected.

## Create an ACL for database service SVC1
SQL> exec dbms_sfw_acl_admin.ip_add_ace('svc1','192.168.12.1');
PL/SQL procedure successfully completed.

SQL> exec dbms_sfw_acl_admin.ip_add_ace('svc1','192.168.12.2');
PL/SQL procedure successfully completed.

SQL> exec dbms_sfw_acl_admin.ip_add_ace('svc1','test02.example.com');
PL/SQL procedure successfully completed.

## Commit the ACLs to the DB ACL table.
## This sends the ACLs for running services to ALL local Listeners
SQL> exec dbms_sfw_acl_admin.commit_acl;
PL/SQL procedure successfully completed.
```

The following example retrieves the ACLs committed from the previous example.

```
SQL> select * from ip_acl;
SERVICE_NAME              HOST
------------------------  -------------------
"SVC1.EXAMPLE.COM"        192.168.12.1
"SVC1.EXAMPLE.COM"        192.168.12.2
"SVC1.EXAMPLE.COM"        TEST02.EXAMPLE.COM

## View ACLs sent to the local Listeners
## NOTE: ACLs are sent ONLY to running services
SQL> select * from v$ip_acl;
SERVICE_NAME              HOST                  CON_ID
------------------------  ------------------    --------
SVC1.EXAMPLE.COM          192.168.12.1          1
SVC1.EXAMPLE.COM          192.168.12.2          1
SVC1.EXAMPLE.COM          TEST02.EXAMPLE.COM    1
```

The following example adds access control entries for pluggable database "PDB1" using various host formats.

```
SQL> exec dbms_sfw_acl_admin.ip_add_pdb_ace('pdb1','192.168.12.3');
PL/SQL procedure successfully completed.

SQL> exec dbms_sfw_acl_admin.ip_add_pdb_ace('pdb1','192.168.12.0/23');
PL/SQL procedure successfully completed.

SQL> exec dbms_sfw_acl_admin.ip_add_pdb_ace('pdb1','192.168.12.*');
PL/SQL procedure successfully completed.

SQL> exec dbms_sfw_acl_admin.commit_acl;
PL/SQL procedure successfully completed.
```

174.4 Summary of DBMS_SFW_ACL_ADMIN Subprograms

This table lists the `DBMS_SFW_ACL_ADMIN` subprograms and briefly describes them.

Table 174-1 DBMS_SFW_ACL_ADMIN Package Subprograms

Subprogram	Description
COMMIT_ACL Procedure	Commits changes to the ACL tables, and propagates the changes to the local listeners for database instances.
IP_ADD_ACE Procedure	Adds an access control entry to the ACL for a database service.
IP_ADD_PDB_ACE Procedure	Adds an access control entry to the ACL for all the database services in a pluggable database (PDB).
IP_REMOVE_ACE Procedure	Removes an entry from the ACL for a database service.
IP_REMOVE_ACL Procedure	Removes all entries from the ACL for a database service.
IP_REMOVE_PDB_ACE Procedure	Removes an access control entry from the ACL for all the database services in a pluggable database (PDB).
IP_REMOVE_PDB_ACL Procedure	Removes all entries from the ACL for all the database services in a pluggable database (PDB).

174.4.1 COMMIT_ACL Procedure

This procedure commits changes to the ACL tables. It also propagates the changes to the local listeners for database instances.

If you have changed access entries for database services, but the database services were not running at the time when you called the COMMIT_ACL procedure, then those changes will be committed to the ACL tables, but they will not be sent to the local listener. To send the entries to the listener, start up the database services, and call the COMMIT_ACL procedure again.

This procedure returns when the operation has completed successfully.

Syntax

```
DBMS_SFW_ACL_ADMIN.COMMIT_ACL;
```

Parameters

None

174.4.2 IP_ADD_ACE Procedure

This procedure adds an access control entry to the ACL for a database service.

Syntax

```
DBMS_SFW_ACL_ADMIN.IP_ADD_ACE (
  p_service_name  IN  VARCHAR2,
  p_host          IN  VARCHAR2);
```

Parameters

Table 174-2 IP_ADD_ACE Procedure Parameters

Parameter	Description
p_service_name	The name of the database service for which you want to add an access control entry.
p_host	The host of the client that is allowed access to the service. This value can be a host name, an IPv4 address, or an IPv6 address. Wildcard "*" for IPv4 and CIDR format are also allowed.

174.4.3 IP_ADD_PDB_ACE Procedure

This procedure adds an access control entry to the ACL for all the database services in a pluggable database (PDB).

Syntax

```
DBMS_SFW_ACL_ADMIN.IP_ADD_PDB_ACE (
  p_pdb_name  IN  VARCHAR2,
  p_host      IN  VARCHAR2);
```

Parameters

Table 174-3 IP_ADD_PDB_ACE Procedure Parameters

Parameter	Description
p_pdb_name	The name of the PDB.
p_host	The host of the client that is allowed access to the database services in the PDB. This value can be a host name, an IPv4 address, or an IPv6 address. Wildcard "*" for IPv4 and CIDR format are also allowed.

174.4.4 IP_REMOVE_ACE Procedure

This procedure removes an entry from the ACL for a database service.

Syntax

```
DBMS_SFW_ACL_ADMIN.IP_REMOVE_ACE (
  p_service_name  IN  VARCHAR2,
  p_host          IN  VARCHAR2);
```

Parameters

Table 174-4 IP_REMOVE_ACE Procedure Parameters

Parameter	Description
p_service_name	The name of the database service from which you want to remove an access control entry.

Table 174-4 (Cont.) IP_REMOVE_ACE Procedure Parameters

Parameter	Description
p_host	The host that you want to remove from the ACL. This value can be a host name, an IPv4 address, or an IPv6 address. Wildcard "*" for IPv4 and CIDR format are also allowed. This has to match the existing value exactly. You can query the `IP_ACL` table to get the list of entries for a database service.

174.4.5 IP_REMOVE_ACL Procedure

This procedure removes all entries from the ACL for a database service.

Syntax

```
DBMS_SFW_ACL_ADMIN.IP_REMOVE_ACL (
  p_service_name  IN  VARCHAR2);
```

Parameters

Table 174-5 IP_REMOVE_ACL Procedure Parameters

Parameter	Description
p_service_name	The name of the database service whose ACL you want to clear.

174.4.6 IP_REMOVE_PDB_ACE Procedure

This procedure removes an access control entry from the ACL for all the database services in the specified pluggable database (PDB).

Syntax

```
DBMS_SFW_ACL_ADMIN.IP_REMOVE_PDB_ACE (
  p_pdb_name  IN  VARCHAR2,
  p_host      IN  VARCHAR2);
```

Parameters

Table 174-6 IP_REMOVE_PDB_ACE Procedure Parameters

Parameter	Description
p_pdb_name	The name of the PDB.
p_host	The host that you want to remove from the ACL. This value can be a host name, an IPv4 address, or an IPv6 address. Wildcard "*" for IPv4 and CIDR format are also allowed. This has to match the existing value exactly. You can query the `IP_ACL` table to get the list of entries for a database service.

174.4.7 IP_REMOVE_PDB_ACL Procedure

This procedure removes all entries from the ACL for all the database services in the specified pluggable database (PDB).

Syntax

```
DBMS_SFW_ACL_ADMIN.IP_REMOVE_PDB_ACL (
  p_pdb_name  IN  VARCHAR2);
```

Parameters

Table 174-7 IP_REMOVE_PDB_ACL Procedure Parameters

Parameter	Description
p_pdb_name	The name of the PDB.

175
DBMS_SHARED_POOL

The DBMS_SHARED_POOL package provides access to the shared pool, which is the shared memory area where cursors and PL/SQL objects are stored. DBMS_SHARED_POOL enables you to display the sizes of objects in the shared pool, and mark them for keeping or not-keeping in order to reduce memory fragmentation.

This chapter contains the following topics:

- Overview
- Operational Notes
- Summary of DBMS_SHARED_POOL Subprograms

175.1 DBMS_SHARED_POOL Overview

The procedures provided here may be useful when loading large PL/SQL objects. When large PL/SQL objects are loaded, users response time is affected because of the large number of smaller objects that need to be aged out from the shared pool to make room (due to memory fragmentation). In some cases, there may be insufficient memory to load the large objects.

DBMS_SHARED_POOL is also useful for frequently executed triggers. You may want to keep compiled triggers on frequently used tables in the shared pool.

Additionally, DBMS_SHARED_POOL supports sequences. Sequence numbers are lost when a sequence is aged out of the shared pool. DBMS_SHARED_POOL is useful for keeping sequences in the shared pool and thus preventing the loss of sequence numbers.

175.2 DBMS_SHARED_POOL Operational Notes

To create DBMS_SHARED_POOL, run the DBMSPOOL.SQL script. The PRVTPOOL.PLB script is automatically executed after DBMSPOOL.SQL runs. These scripts are *not* run by as part of standard database creation.

175.3 Summary of DBMS_SHARED_POOL Subprograms

This table lists the DBMS_SHARED_POOL subprograms and briefly describes them.

Table 175-1 ***DBMS_SHARED_POOL Package Subprograms***

Subprogram	Description
ABORTED_REQUEST_THRESHOLD Procedure	Sets the aborted request threshold for the shared pool
KEEP Procedure	Keeps an object in the shared pool
MARKHOT Procedure	Marks a library cache object as a hot object
PURGE Procedure	Purges the named object or specified heap(s) of the object

Table 175-1 (Cont.) *DBMS_SHARED_POOL Package Subprograms*

Subprogram	Description
SIZES Procedure	Shows objects in the shared pool that are larger than the specified size
UNKEEP Procedure	Unkeeps the named object
UNMARKHOT Procedure	Unmarks a library cache object as a hot object

175.3.1 ABORTED_REQUEST_THRESHOLD Procedure

This procedure sets the aborted request threshold for the shared pool.

Syntax

```
DBMS_SHARED_POOL.ABORTED_REQUEST_THRESHOLD (
   threshold_size NUMBER);
```

Parameters

Table 175-2 ABORTED_REQUEST_THRESHOLD Procedure Parameters

Parameter	Description
`threshold_size`	Size, in bytes, of a request which does not try to free unpinned (not "unkeep-ed") memory within the shared pool. The range of `threshold_size` is 5000 to ~2 GB inclusive.

Exceptions

An exception is raised if the threshold is not in the valid range.

Usage Notes

Usually, if a request cannot be satisfied on the free list, then the RDBMS tries to reclaim memory by freeing objects from the LRU list and checking periodically to see if the request can be fulfilled. After finishing this step, the RDBMS has performed a near equivalent of an '`ALTER SYSTEM FLUSH SHARED_POOL`'.

Because this impacts all users on the system, this procedure "localizes" the impact to the process failing to find a piece of shared pool memory of size greater than `thresh_hold` size. This user gets the 'out of memory' error without attempting to search the LRU list.

175.3.2 KEEP Procedure

This procedure keeps an object in the shared pool. Once an object has been kept in the shared pool, it is not subject to aging out of the pool. This may be useful for frequently used large objects. When large objects are brought into the shared pool, several objects may need to be aged out to create a contiguous area large enough.

Syntax

```
DBMS_SHARED_POOL.KEEP (
   name         VARCHAR2,
   flag         CHAR DEFAULT 'P');
```

```
DBMS_SHARED_POOL.KEEP (
   schema          VARCHAR2,
   objname         VARCHAR2,
   namespace       NUMBER,
   heaps           NUMBER,
   edition_name    VARCHAR2 DEFAULT NULL);

DBMS_SHARED_POOL.KEEP (
   hash            VARCHAR2,
   namespace       NUMBER,
   heaps           NUMBER);
```

Parameters

Table 175-3 KEEP Procedure Parameters

Parameter	Description
name	The name of the object to keep.
flag	A character string indicating the kind of object to keep in the shared pool. The string is not case sensitive. This parameter is optional. If the parameter is not specified, the package assumes that the first parameter is the name of a package/procedure/function and will resolve the name. • Set to 'P' or 'p' to fully specify that the input is the name of a package/procedure/function. • Set to 'Q' or 'q' to specify that the input is the name of a sequence. • Set to 'R' or 'r' to specify that the input is the name of a trigger. • Set to 'T' or 't' to specify that the input is the name of a type. • Set to 'JS' or 'js' to specify that the input is the name of a java source. • Set to 'JC' or 'jc' to specify that the input is the name of a java class. • Set to 'JD' or 'jd' to specify that the input is the name of a java shared data. • Set to 'JR' or 'jr' to specify that the input is the name of a java resource. • Set to 'C' or 'c' to specify that the input is the name of a cursor.
schema	The user name or the schema to which the object belongs to.
objname	The name of the object to keep.
namespace	A number indicating the library cache namespace in which the object has to be searched for.
heaps	The heaps to keep. For example, if heap 0 and heap 6 are to be kept.
edition_name	The name of the edition that the target object resides in. This parameter is optional.
hash	A 16-byte hash value for the object.

Exceptions

An exception is raised if the named object is not found.

Usage Notes

There are two kinds of objects:

- PL/SQL objects, triggers, sequences, types, and Java objects, which are specified by name.

- SQL cursor objects which are specified by a two-part number (indicating a location in the shared pool).

For example:

```
DBMS_SHARED_POOL.KEEP('scott.hispackage')
```

This keeps package `HISPACKAGE`, owned by `SCOTT`. The names for PL/SQL objects follow SQL rules for naming objects (for example, delimited identifiers and multibyte names are allowed). A cursor can be kept by `DBMS_SHARED_POOL.KEEP('0034CDFF, 20348871','C')`, `0034CDFF` being the `ADDRESS` and `20348871` the `HASH_VALUE`. Note that the complete hexadecimal address must be in the first 8 characters.

175.3.3 MARKHOT Procedure

This procedure marks a library cache object as a hot object.

Syntax

```
DBMS_SHARED_POOL.MARKHOT (
   schema         VARCHAR2,
   objname        VARCHAR2,
   namespace      NUMBER DEFAULT 1,
   global         BOOLEAN DEFAULT TRUE,
   edition_name   VARCHAR2 DEFAULT NULL);

DBMS_SHARED_POOL.MARKHOT (
   hash          VARCHAR2,
   namespace     NUMBER DEFAULT 1,
   global        BOOLEAN DEFAULT TRUE);
```

Parameters

Table 175-4 MARKHOT Procedure Parameters

Parameter	Description
`schema`	User name or the schema to which the object belongs
`objname`	Name of the object
`namespace`	Number indicating the library cache namespace in which the object is to be searched. Views, such as `USER_OBJECTS` and `DBA_OBJECTS`, reflect the namespace as a number column, as do most dictionary tables such as `OBJ$`.
`global`	If `TRUE` (default), mark the object hot on all Oracle RAC instances
`hash`	16-byte hash value for the object
`edition_name`	Denotes the name of the edition that the target object resides in. This parameter is optional.

Exceptions

`ORA-06502`: An exception is raised if the named object cannot be found due to incorrect input

`ORA-04043`: An exception is raised if the named object cannot be found (bad namespace, or hash input)

Usage Notes

If a package or type's specification is marked hot or unhot, then the corresponding package or type body will be implicitly marked as hot or unhot.

Users can examine column, `V$DB_OBJECT_CACHE.PROPERTY`, to see whether or not the object has been marked hot. The values for `PROPERTY` are:

- `HOTCOPYnnn`- An object that is a hot copy with integer identifier of 'nnn'. For example, `HOTCOPY5`, `HOTCOPY94`, and `HOTCOPY125`.
- `HOTCOPY`-As above but the identifier is unknown
- `HOT`-The "root" kgl object that has been marked as hot
- `NULL`-A normal object

175.3.4 PURGE Procedure

This procedure purges the named object or specified heaps of the object.

Syntax

```
DBMS_SHARED_POOL.PURGE (
   name         VARCHAR2,
   flag         CHAR DEFAULT 'P',
   heaps        NUMBER DEFAULT 1);

DBMS_SHARED_POOL.PURGE (
   schema       VARCHAR2,
   objname      VARCHAR2,
   namespace    NUMBER,
   heaps        NUMBER,
   edition_name VARCHAR2 DEFAULT NULL);

DBMS_SHARED_POOL.PURGE (
   hash         VARCHAR2,
   namespace    NUMBER,
   heaps        NUMBER);
```

Parameters

Table 175-5 PURGE Procedure Parameters

Parameter	Description
`name`	Name of the object to purge. The value for this identifier is the concatenation of the address and `hash_value` columns from the `v$sqlarea` view. This is displayed by the `SIZES` procedure. Currently, `TABLE` and `VIEW` objects may not be purged.

Table 175-5 (Cont.) PURGE Procedure Parameters

Parameter	Description
flag	A character string indicating the kind of object to purge from the shared pool. The string is not case sensitive. This parameter is optional. If the parameter is not specified, the package assumes that the first parameter is the name of a package/procedure/function and will resolve the name. • Set to 'P' or 'p' to fully specify that the input is the name of a package/procedure/function. • Set to 'Q' or 'q' to specify that the input is the name of a sequence. • Set to 'R' or 'r' to specify that the input is the name of a trigger. • Set to 'T' or 't' to specify that the input is the name of a type. • Set to 'JS' or 'js' to specify that the input is the name of a java source. • Set to 'JC' or 'jc' to specify that the input is the name of a java class. • Set to 'JD' or 'jd' to specify that the input is the name of a java shared data. • Set to 'JR' or 'jr' to specify that the input is the name of a java resource. • Set to 'C' or 'c' to specify that the input is the name of a cursor.
heaps	Heaps to be purged. For example, if heap 0 and heap 6 are to be purged: 1<<0 | 1<<6 => hex 0x41 => decimal 65, so specify heaps =>65. Default value is 1, that is, heap 0 which means the whole object would be purged
schema	User name or the schema to which the object belongs
objname	Name of the object to purge
namespace	Parameter is a number indicating the library cache namespace in which the object is to be searched
hash	16-byte hash value for the object
edition_name	The name of the edition that the target object resides in. This parameter is optional.

Exceptions

ORA-6570: An exception is raised if the named object cannot be found

ORA-6570: An object cannot be purged it marked as permanently kept

Usage Notes

All objects supported by the KEEP Procedure are supported for PURGE.

175.3.5 PURGE_ROW_CACHE Procedure

This procedure purges the row cache as identified by cache id, or a specific object within the selected cache.

Syntax

```
DBMS_SHARED_POOL.PURGE_ROW_CACHE (
   cacheid number, hash varchar2 DEFAULT null);
```

Parameters

Table 175-6 PURGE_ROW_CACHE Procedure Parameters

Parameter	Description
cacheid	Indicates which cache to purge.
hash	The hash value of a targeted object to purge within the named cache. The default value (NULL) purges the entire cache.

175.3.6 SIZES Procedure

This procedure shows objects in the shared_pool that are larger than the specified size. The name of the object is also given, which can be used as an argument to either the KEEP or UNKEEP calls.

Syntax

```
DBMS_SHARED_POOL.SIZES (
   minsize NUMBER);
```

Parameters

Table 175-7 SIZES Procedure Parameters

Parameter	Description
minsize	Size, in kilobytes, over which an object must be occupying in the shared pool, in order for it to be displayed.

Usage Notes

Issue the SQLDBA or SQLPLUS 'SET SERVEROUTPUT ON SIZE XXXXX' command prior to using this procedure so that the results are displayed.

175.3.7 UNKEEP Procedure

This procedure unkeeps the named object.

Syntax

```
DBMS_SHARED_POOL.UNKEEP (
   name         VARCHAR2,
   flag         CHAR DEFAULT 'P');

DBMS_SHARED_POOL.UNKEEP (
   schema       VARCHAR2,
   objname      VARCHAR2,
   namespace    NUMBER,
   edition_name VARCHAR2 DEFAULT NULL);

DBMS_SHARED_POOL.UNKEEP (
   hash         VARCHAR2,
   namespace    NUMBER);
```

> **WARNING:**
>
> This procedure may not be supported in the future if automatic mechanisms are implemented to make this unnecessary.

Parameters

Table 175-8 UNKEEP Procedure Parameters

Parameter	Description
name	The name of the object to unkeep.
flag	A character string indicating what kind of object to keep the name identifies. The string is not case sensitive. This parameter is optional. If the parameter is not specified, the package assumes that the first parameter is the name of a package/procedure/function and will resolve the name.
schema	The user name or the schema to which the object belongs to.
objname	The name of the object to unkeep.
namespace	A number that indicates the library cache namespace in which the object has to be searched for.
edition_name	The name of the edition that the target object resides in. This parameter is optional.
hash	A 16-byte hash value for the object.

Exceptions

ORA-06502: An exception is raised if the named object cannot be found

175.3.8 UNMARKHOT Procedure

This procedure unmarks a library cache object as a hot object.

Syntax

```
DBMS_SHARED_POOL.UNMARKHOT (
   schema         VARCHAR2,
   objname        VARCHAR2,
   namespace      NUMBER  DEFAULT 1,
   global         BOOLEAN DEFAULT TRUE,
   edition_name   VARCHAR2 DEFAULT NULL);

DBMS_SHARED_POOL.UNMARKHOT (
   hash          VARCHAR2,
   namespace     NUMBER DEFAULT 1,
   global        BOOLEAN DEFAULT TRUE);
```

Parameters

Table 175-9 UNMARKHOT Procedure Parameters

Parameter	Description
schema	User name or the schema to which the object belongs
objname	Name of the object
namespace	Number indicating the library cache namespace in which the object is to be searched
global	If TRUE, unmark the object hot on all Oracle RAC instances. The default value of this parameter is TRUE.
hash	A 16-byte hash value for the object
edition_name	Denotes the name of the edition that the target object resides in. This parameter is optional.

Exceptions

ORA-06502: An exception is raised if the named object cannot be found due to incorrect input

ORA-04043: An exception is raised if the named object cannot be found (bad namespace, or hash input, or non-existent object)

Usage Notes

If a package or type's specification is marked hot or unhot, then the corresponding package or type body will be implicitly marked as hot or unhot.

176
DBMS_SHARDING_DIRECTORY

This package provides procedures to manage an Oracle Globally Distributed Database sharded database created with directory-based data distribution.

This chapter contains the following topics:

- DBMS_SHARDING_DIRECTORY Overview
- DBMS_SHARDING_DIRECTORY Security Model
- Summary of DBMS_SHARDING_DIRECTORY Subprograms

176.1 DBMS_SHARDING_DIRECTORY Overview

PL/SQL package `DBMS_SHARDING_DIRECTORY` provides support operations on the directory table for sharded databases using the directory-based data distribution method.

This package includes procedures to:

- Add a key to a partition entry
- Remove a key mapping entry
- Flag a key for partition split

The APIs in this package can be invoked on the shard catalog.

The directory table is automatically created during root table creation. The definition of the directory table is `shard user schema.root_table$SDIR`.

When adding and removing keys there are APIs that include commit and those that do not. Unless the commit versions of the APIs are used, the directory content is not propagated to the shards until commit is issued explicitly.

> **See Also:**
>
> Directory-Based Sharding in *Oracle Globally Distributed Database Guide* for a detailed description of sharded database directory-based data distribution.

176.2 DBMS_SHARDING_DIRECTORY Security Model

All `DBMS_SHARDING_DIRECTORY` subprograms require the user to have `EXECUTE` privilege over the `DBMS_SHARDING_DIRECTORY` package.

Only the schema owner of the root table is allowed to execute the procedures in this package, and they can only execute it on the shard catalog.

176.3 DBMS_SHARDING_DIRECTORY Public Constants

There are public constants defined for use with the `DBMS_SHARDING_DIRECTORY.setAssignmentRule` procedure for key-to-partition assignment rules.

All such constants are defined as part of the `DBMS_SHARDING_DIRECTORY` package. Any references to these constants must be prefixed by `DBMS_SHARDING_DIRECTORY.` and followed by the symbols in the following lists

- `NONE constant number :=0;` -- turn off rule-based assignment
- `LAST_PARTITION constant number := 1;` -- rule for assigning key only to the last added partition
- `ROUNT_ROBIN constant number :=2;` -- rule for assigning key to partition by round robin
- `RANDOM constant number :=3;` -- rule for assigning key to partition randomly
- `CUSTOM constant number :=4;` -- TBD

176.4 Summary of DBMS_SHARDING_DIRECTORY Subprograms

This table lists and describes the subprograms of the `DBMS_SHARDING_DIRECTORY` package

Table 176-1 DBMS_SHARDING_DIRECTORY Package Subprograms

Subprogram	Description
addKeyToPartition Procedure	This procedure allows you to add a new key to the directory with the specified partition name.
addKeyToPartitionCommit Procedure	This procedure allows you to add a new key to the directory with the specified partition name, and performs a commit at the end.
flagKeyForSplit Procedure	This procedure allows you to mark a key in the directory for split, to be performed later.
removeKey Procedure	This procedure allows you to remove a key from the directory.
removeKeyCommit Procedure	This procedure allows you to remove a key from the directory, and performs a commit at the end.
setAssignmentRule	This procedure allows you to indicate an automatic key-to-partition assignment rule for subsequent new key inserts into the root table.

176.4.1 addKeyToPartition Procedure

This procedure allows you to add a new key to the directory with the specified partition name.

Syntax

```
DBMS_SHARDING_DIRECTORY.addKeyToPartition(
   (schema_name     IN varchar2,
    root_table      IN varchar2,
```

```
    partition_name IN  varchar2,
    key … );
```

Parameters

Table 176-2 addKeyToPartition Procedure Parameters

Parameter	Description
schema_name	Root table schema name.
partition_name	Name of the partition.
root_table	Root table name.
key	Shard key column values.

Usage Notes

Note that the `key` column value needs to be in the same order as specified in the `CREATE TABLE` statement with the correct types. The procedure can only succeed if the provided key does not yet exist in the directory.

176.4.2 addKeyToPartitionCommit Procedure

This procedure allows you to add a new key to the directory with the specified partition name.

Syntax

```
DBMS_SHARDING_DIRECTORY.addKeyToPartitionCommit(
   (schema_name    IN varchar2,
    root_table     IN varchar2,
    partition_name IN  varchar2,
    key … );
```

Parameters

Table 176-3 addKeyToPartitionCommit Procedure Parameters

Parameter	Description
schema_name	Root table schema name.
partition_name	Name of the partition.
root_table	Root table name.
key	Shard key column values.

Usage Notes

The `addKeyToPartitionCommit` procedure is exactly the same as the `addKeyToPartition` procedure with the same parameters, except that it performs a commit automatically at the end.

Note that the `key` column value needs to be in the same order as specified in the `CREATE TABLE` statement with the correct types. The procedure can only succeed if the provided key does not yet exist in the directory.

176.4.3 flagKeyForSplit Procedure

This procedure allows you to mark a key in the directory for split, to be performed later.

Syntax

```
DBMS_SHARDING_DIRECTORY.flagKeyForSplit(
   (schema_name     IN varchar2,
    root_table      IN varchar2,
    key … );
```

Parameters

Table 176-4 flagKeyForSplit Procedure Parameters

Parameter	Description
schema_name	Root table schema name.
root_table	Root table name.
key	Shard key column values.

Usage Notes

Note that the `key` column values need to be in the same order as specified in the `CREATE TABLE` statement with the correct types. The procedure can only succeed if the provided key exists in the directory.

A subsequent `ALTER TABLE SPLIT PARTITION` operation will go through all the keys that have been marked for split in the directory and split the corresponding data out into the new partition.

176.4.4 removeKey Procedure

This procedure allows you to remove a key from the directory.

Syntax

```
DBMS_SHARDING_DIRECTORY.removeKey(
   (schema_name     IN varchar2,
    root_table      IN varchar2,
    key … );
```

Parameters

Table 176-5 removeKey Procedure Parameters

Parameter	Description
schema_name	Root table schema name.
root_table	Root table name.
key	Shard key column values.

Usage Notes

Note that the `key` column values need to be in the same order as specified in the `CREATE TABLE` statement with the correct types. The procedure can only succeed if the provided key exists in the directory, and there are no tables (either root table or child tables) with rows still referencing the key.

176.4.5 removeKeyCommit Procedure

This procedure allows you to remove a key from the directory.

Syntax

```
DBMS_SHARDING_DIRECTORY.removeKeyCommit(
   (schema_name     IN varchar2,
    root_table      IN varchar2,
    key … );
```

Parameters

Table 176-6 removeKeyCommit Procedure Parameters

Parameter	Description
`schema_name`	Root table schema name.
`root_table`	Root table name.
`key`	Shard key column values.

Usage Notes

The `removeKeyCommit` procedure is exactly the same as the `removeKey` procedure with the same parameters, except that it performs a commit automatically at the end.

Note that the `key` column values need to be in the same order as specified in the `CREATE TABLE` statement with the correct types. The procedure can only succeed if the provided key exists in the directory, and there are no tables (either root table or child tables) with rows still referencing the key.

176.4.6 setAssignmentRule

This procedure allows you to indicate an automatic key-to-partition assignment rule for subsequent new key inserts into the root table.

It will be in effect across different sessions and regardless of system restart until another call to this procedure is made with a different `rule_id` value, or with `NONE`, meaning automatic assignment should be turned off.

Rule ID values are defined in DBMS_SHARDING_DIRECTORY Public Constants.

Syntax

```
DBMS_SHARDING_DIRECTORY.setAssignmentRule(
   (schema_name     IN varchar2,
    root_table      IN varchar2,
    rule_id         IN number);
```

Parameters

Table 176-7 setAssignmentRule Procedure Parameters

Parameter	Description
`schema_name`	Root table schema name.
`root_table`	Root table name.
`rule_id`	Rule ID value, as defined in DBMS_SHARDING_DIRECTORY Public Constants.

177
DBMS_SODA

The DBMS_SODA package is a PL/SQL package implementing Simple Oracle Document Access (SODA). SODA allows you to use the Oracle Database as a NoSQL document store. The core abstraction provided by SODA is that of document collections. The DBMS_SODA package allows you to create, list, and delete document collections from PL/SQL, and to perform CRUD (create, replace, update, delete) operations on documents. All DDL functions are encapsulated within this package.

This chapter contains the following topics:

- DBMS_SODA Security Model
- Summary of DBMS_SODA Subprograms
- Summary of SODA Online Redefinition Subprograms

> **See Also:**
>
> - *Oracle Database SODA for PL/SQL Developer's Guide*
> - SODA Types

177.1 DBMS_SODA Security Model

This package is available to users with the SODA_APP role.

All SODA types (packages and types) are SYS types. PUBLIC is granted EXECUTE privilege on the DBMS_SODA described in this chapter.

177.2 Summary of DBMS_SODA Subprograms

This table lists the DBMS_SODA subprograms in alphabetical order and briefly describes them.

Table 177-1 DBMS_SODA Package Subprograms

Subprogram	Purpose
CREATE_COLLECTION Function	Creates a collection using the collection name and metadata.
DROP_COLLECTION Function	Drops an existing collection from the user's schema. This also removes all the documents in the collection.
GET_AS_OF_SCN Function	This function returns the current database SCN number.
GET_AS_OF_TIMESTAMP Function	This function returns the current database timestamp value.

Table 177-1 (Cont.) DBMS_SODA Package Subprograms

Subprogram	Purpose
GET_DEFAULT_METADATA_AS_CLOB Function	This function returns the default metadata.
GET_DEFAULT_METADATA_AS_VARCHAR2 Function	This function returns the default metadata.
LIST_COLLECTION_NAMES Function	Lists the collection names in the user's schema as a table of `NVARCHAR2`.
OPEN_COLLECTION Function	Opens an existing collection.

177.2.1 CREATE_COLLECTION Function

Creates a collection using the collection name and metadata. Uses the settings specified in the metadata and auto-assigns the ones that are not, and returns the collection object. If the metadata argument is omitted or set to `NULL`, a collection is created with default metadata. The returned collection is open for `read` and/or `write` operations. If a collection already exists, the function just opens and returns the collection object.

Syntax

```
DBMS_SODA.CREATE_COLLECTION (
     collection_Name      IN NVARCHAR2,
     metadata             IN VARCHAR2 DEFAULT NULL,
     create_Mode          IN PLS_INTEGER DEFAULT CREATE_MODE_DDL)
 RETURN SODA_Collection_T;
```

Parameters

Table 177-2 CREATE_COLLECTION Parameters

Parameter	Description
`collection_Name`	The name of the collection. The value of `collection_Name` is case-sensitive.
`metadata`	The metadata of the collection in `VARCHAR2` format.
create_Mode	Valid values are: • `DBMS_SODA.CREATE_MODE_DDL` (default). First attempts to create a new table for the collection. If table exists already, tries to map that existing table to the collection. Minimal checking is performed to ensure the table shape matches supplied collection metadata (if not, then an error is returned). • `DBMS_SODA.CREATE_MODE_MAP`. Tries to map an existing table to the collection. Minimal checking is performed to ensure the table shape matches supplied collection metadata (if not, then an error is returned).

Return Values

The function returns a `Soda_Collection_T` object representing the collection.

Exceptions

- `Descriptor Error`—if the input descriptor is invalid
- `Error`—if an error occurs while creating the collection

> **See Also:**
>
> *Oracle Database SODA for PL/SQL Developer's Guide*

177.2.2 DROP_COLLECTION Function

Drops an existing collection from the user's schema. This also removes all the documents in the collection.

Syntax

```
DBMS_SODA.DROP_COLLECTION (
     collection_Name      IN NVARCHAR2,
     purge                IN BOOLEAN DEFAULT FALSE,
     drop_Mapped_Table    IN BOOLEAN DEFAULT FALSE)
 RETURN NUMBER;
```

Parameters

Table 177-3 DROP_COLLECTION Parameters

Parameter	Description
`collection_Name`	The name of the collection. The value of `collection_Name` is case-sensitive.
`purge`	The default value is `FALSE`.
`drop_Mapped_Table`	The default value is `FALSE`.

Return values

This function returns the following values:

- `1`—if the collection was dropped successfully
- `0`—if the collection does not exist

Exceptions

If an error occurs while dropping the collection, for example, due to uncommitted writes to the collection or privilege issues.

> **See Also:**
>
> *Oracle Database SODA for PL/SQL Developer's Guide*

177.2.3 GET_AS_OF_SCN Function

This function returns the current database SCN number.

Syntax

```
DBMS_SODA.GET_AS_OF_SCN ( )
 RETURN NUMBER;
```

Return values

This function returns the current database SCN number.

177.2.4 GET_AS_OF_TIMESTAMP Function

This function returns the current database timestamp value.

Syntax

```
DBMS_SODA.GET_AS_OF_TIMESTAMP ( )
 RETURN VARCHAR2;
```

Return values

This function returns the current database timestamp value.

177.2.5 GET_DEFAULT_METADATA_AS_CLOB Function

This function returns the default metadata.

Syntax

```
DBMS_SODA.GET_DEFAULT_METADATA_AS_CLOB ( )
 RETURN CLOB;
```

Return Values

This function returns the default metadata using the `CLOB` datatype.

177.2.6 GET_DEFAULT_METADATA_AS_VARCHAR2 Function

This function returns the default metadata.

Syntax

```
DBMS_SODA.GET_DEFAULT_METADATA_AS_VARCHAR2 ( )
 RETURN VARCHAR2;
```

Return Values

This function returns the default metadata using the `VARCHAR2` datatype.

177.2.7 LIST_COLLECTION_NAMES Function

This function returns a list of collection names in the user's schema as a table of `NVARCHAR2`.

Syntax

```
DBMS_SODA.LIST_COLLECTION_NAMES ()
 RETURN SODA_CollName_List_T;
```

Return Values

This function returns a list of collection names as a table of `NVARCHAR2(255)`. The collection list is empty if there are no collections in the schema.

Exceptions

`Error`—if an error occurs while listing the collection names.

> **See Also:**
>
> *Oracle Database SODA for PL/SQL Developer's Guide*

177.2.8 OPEN_COLLECTION Function

Opens an existing collection for `read` and/or `write` operations.

Syntax

```
DBMS_SODA.OPEN_COLLECTION (
     collection_Name      IN NVARCHAR2)
 RETURN SODA_Collection_T;
```

Parameters

Table 177-4 OPEN_COLLECTION Parameters

Parameter	Description
`collection_Name`	The name of the collection. The value of `collection_Name` is case-sensitive.

Return Values

This function returns the following values:

- a collection object which is open
- `NULL`, if the collection does not exist

Exceptions

`Error`—if an error occurs while creating the collection

See Also:

- Opening an Existing Document Collection with SODA for PL/SQL
- Checking Whether a Given Collection Exists with SODA for PL/SQL

177.3 Summary of SODA Online Redefinition Subprograms

This table lists the SODA Online Redefinition subprograms in alphabetical order and briefly describes them.

Table 177-5 SODA Online Redefinition Subprograms

Subprogram	Purpose
ABORT_REDEF_COLLECTION Procedure	This procedure reverts the changes made to a collection if there are any errors in `CREATE_INTERIM_COLLECTION`, `START_REDEF_COLLECTION`, `COPY_COLLECTION_DEPENDENTS`, `SYNC_INTERIM_COLLETION`, or due to any other reasons.
CAN_REDEF_COLLECTION Procedure	This procedure checks if the data table can be redefined. If the data table cannot be redefined, errors are raised.
COPY_COLLECTION_DEPENDENTS Procedure	This procedure copies all the dependents that are defined on the original table to the interim table. However, if the interim table is JSON type, constraints like `IS JSON NOT NULL` will not be copied because those are extraneous.
CREATE_INTERIM_COLLECTION Procedure	This procedure creates an interim `SODA` collection. The interim_metadata specifies the changes the user needs to make. For example, a delta applied as a patch to the original matadata.
FINISH_REDEF_COLLECTION Procedure	This procedure performs the following tasks in one atomic transaction. It introduces a downtime. This can only be called after every dependent required for a `SODA` data table is present on the interim collection, otherwise an error is raised.
START_REDEF_COLLECTION Procedure	This procedure starts the redefinition process. Copies the existing data from the original table to the interim table and performs the required transformations on the redefined columns.
SYNC_INTERIM_COLLETION Procedure	This procedure syncronizes the changes to the interim table if the original data table was changed. This procedure can only be called after every dependent required for the `SODA` data table is present on the interim collection, otherwise, an error is raised.
TO_UUID Function	This function is used to remap the version column of the collection table during Data Pump Import to UUID values. This function is only for use with the `REMAP_DATA` feature of Data Pump.

177.3.1 ABORT_REDEF_COLLECTION Procedure

This procedure reverts the changes made to a collection if there are any errors in `CREATE_INTERIM_COLLECTION`, `START_REDEF_COLLECTION`, `COPY_COLLECTION_DEPENDENTS`, `SYNC_INTERIM_COLLETION`, or due to any other reasons.

Syntax

```
DBMS_SODA.ABORT_REDEF_COLLECTION (
   collection_name         IN  NVARCHAR2,
   interim_collection_name IN  NVARCHAR2);
```

Parameters

Table 177-6 ABORT_REDEF_COLLECTION Procedure Parameters

Parameter	Description
`collection_name`	The name of the collection to be redefined.
`interim_collection_name`	The name used for the interim collection.

177.3.2 CAN_REDEF_COLLECTION Procedure

This procedure checks if the data table can be redefined. If the data table cannot be redefined, errors are raised.

Syntax

```
DBMS_SODA.CAN_REDEF_COLLECTION (
   collection_name   IN  NVARCHAR2);
```

Parameters

Table 177-7 CAN_REDEF_COLLECTION Procedure Parameters

Parameter	Description
`collection_name`	The name of the collection to be redefined.

177.3.3 COPY_COLLECTION_DEPENDENTS Procedure

This procedure copies all the dependents that are defined on the original table to the interim table.

Syntax

```
DBMS_SODA.COPY_COLLECTION_DEPENDENTS (
   collection_name         IN  NVARCHAR,
   interim_collection_name IN  NVARCHAR2,
   ignore_error            IN  BOOLEAN DEFAULT NULL,
   num_errors              OUT PLS_INTEGER);
```

Parameters

Table 177-8 COPY_COLLECTION_DEPENDENTS Procedure Parameters

Parameter	Description
collection_name	The name of the collection to be redefined.
interim_collection_name	The name used for the interim collection.
ignore_error	Ignore the errors encountered in the process and proceed.
num_errors	The number of errors encountered in the process.

177.3.4 CREATE_INTERIM_COLLECTION Procedure

This procedure creates an interim SODA collection. The interim_metadata specifies the changes the user needs to make. For example, a delta applied as a patch to the original matadata.

Syntax

```
DBMS_SODA.CREATE_INTERIM_COLLECTION (
   collection_name            IN  NVARCHAR2,
   interim_collection_name    IN  NVARCHAR2,
   interim_metadata           IN  VARCHAR2);
```

Parameters

Table 177-9 CREATE_INTERIM_COLLECTION Procedure Parameters

Parameter	Description
collection_name	The name of the collection to be redefined.
interim_collection_name	The name used for the interim collection.
interim_metadata	The metadata snippet specifying the columns that need to be redefined and how they need to be redefined.

Usage Notes

The interim_metadata parameter can contain the following fields. An error is raised if other fields are present and if the fields are not the same as those in the metadata for the original collection.

- Set contentColumn.sqlType to JSON in order to redefine the CONTENT column to JSON type, automatically generating the VERSION column if needed.
- Set versionColumn.method to UUID to redefine the VERSION column to use UUID.
- Set tableName to specify a name for the data table of this interim collection. If this field is not present, a default data table name will be generated according to the rules of table name defaulting outlined in Default Naming of a Collection Table. However, if the table already exists, the collection is created using the MAP mode. In this scenario, make sure that the mapped table does not have any constraints defined.
- The user can combine the above fields in the supplied interim_metadata snippet. For example, the most common use case is to both redefine the content column to JSON type and redefine the version column to UUID, which can be achieved by setting the

`interim_metadata` parameter to `{"contentColumn" : {"sqlType": "JSON"}, "versionColumn" : {"method": "UUID"}}`.

177.3.5 FINISH_REDEF_COLLECTION Procedure

This procedure performs the following tasks in one atomic transaction. It introduces a downtime. This can only be called after every dependent required for a `SODA` data table is present on the interim collection, otherwise an error is raised.

Syntax

```
DBMS_SODA.FINISH_REDEF_COLLECTION (
   collection_name         IN  NVARCHAR,
   interim_collection_name IN  NVARCHAR2,
   dml_lock_timeout        IN  PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 177-10 FINISH_REDEF_COLLECTION Procedure Parameters

Parameter	Description
`collection_name`	The name of the collection to be redefined.
`interim_collection_name`	The name used for the interim collection.
`dml_lock_timeout`	The wait period in seconds. If the lock is not acquired within this wait period, an error occurs.

Usage Notes

This action cannot be undone. After this procedure is executed, you cannot perform `ROLLBACK` procedure provided by the `DBMS_REDEFINITION` package.

> **Caution:**
>
> Before calling `FINISH_REDEF_COLLECTION Procedure`, the interim collection should be tried out to make sure it can be accessed and is working as expected. In case of any issues, use ABORT_REDEF_COLLECTION Procedure.
>
> It is very important to do this before calling `FINISH_REDEF_COLLECTION` because the latter cannot be reversed.

177.3.6 START_REDEF_COLLECTION Procedure

This procedure starts the redefinition process. Copies the existing data from the original table to the interim table and performs the required transformations on the redefined columns.

Syntax

```
DBMS_SODA.START_REDEF_COLLECTION (
   collection_name         IN  NVARCHAR,
   interim_collection_name IN  NVARCHAR2,
   copy_vpd_opt            IN  DEFAULT NULL,
   refresh_dep_mviews      IN  DEFAULT NULL);
```

Parameters

Table 177-11 START_REDEF_COLLECTION Procedure Parameters

Parameter	Description
collection_name	The name of the collection to be redefined.
interim_collection_name	The name used for the interim collection.
copy_vpd_opt	Can be either DBMS_REDEFINITION.CONS_VPD_MANUAL or DBMS_REDEFINITION.CONS_VPD_NONE. DBMS_REDEFINITION.CONS_VPD_MANUAL is used to indicate to copy VPD policies manually. DBMS_REDEFINITION.CONS_VPD_NONE is used to indicate that there are no VPD policies on the original table. If NULL, DBMS_REDEFINITION.CONS_VPD_NONE is used.
refresh_dep_mviews	Can be N or Y. When set to N, fast refresh of dependent materialized views is performed when the START_REDEF_TABLE procedure is run, each time the SYNC_INTERIM_TABLE procedure is run, and when the FINISH_REDEF_TABLE procedure is run. If NULL, N is used.

177.3.7 SYNC_INTERIM_COLLETION Procedure

This procedure synchronizes the changes to the interim table if the original data table was changed. This procedure can only be called after every dependent required for the SODA data table is present on the interim collection, otherwise, an error is raised.

Syntax

```
DBMS_SODA.SYNC_INTERIM_COLLETION (
   collection_name         IN  NVARCHAR,
   interim_collection_name IN  NVARCHAR2);
```

Parameters

Table 177-12 SYNC_INTERIM_COLLETION Procedure Parameters

Parameter	Description
collection_name	The name of the collection to be redefined.
interim_collection_name	The name used for the interim collection.

177.3.8 TO_UUID Function

This function is used to remap the version column of the collection table during Data Pump Import to UUID values. This function is only for use with the REMAP_DATA feature of Data Pump.

Syntax

```
DBMS_SODA.TO_UUID (
     placeholder      IN   NVARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 177-13 TO_UUID Parameters

Parameter	Description
`placeholder`	This parameter is not used and can be ignored. It is used due to the syntactic constraints for the Data Pump `REMAP_DATA` function.

Return values

This function returns a hexadecimal string that can be used as a `UUID` value.

> **See Also:**
>
> REMAP_DATA in the *Oracle® Database Utilities Guide*.

178
DBMS_SPACE

The DBMS_SPACE package enables you to analyze segment growth and space requirements.

This chapter contains the following topics:

- Security Model
- Data Structures
- Summary of DBMS_SPACE Subprograms

178.1 DBMS_SPACE Security Model

This package runs with SYS privileges. The execution privilege is granted to PUBLIC. Subprograms in this package run under the caller security. The user must have ANALYZE privilege on the object.

178.2 DBMS_SPACE Data Structures

The DBMS_SPACE package defines an OBJECT type, a RECORD type, and a TABLE type.

OBJECT Types

CREATE_TABLE_COST_COLINFO Object Type

RECORD Types

ASA_RECO_ROW Record Type

TABLE Types

ASA_RECO_ROW_TB Table Type

178.2.1 DBMS_SPACE CREATE_TABLE_COST_COLINFO Object Type

This type describes the datatype and size of a column in the table.

Syntax

```
TYPE create_table_cost_colinfo IS OBJECT(
    col_type    VARCHAR(200),
    col_size    NUMBER)
```

Attributes

Table 178-1 CREATE_TABLE_COST_COLINFO Object Type

Attribute	Description
col_type	Column type

Table 178-1 (Cont.) CREATE_TABLE_COST_COLINFO Object Type

Attribute	Description
col_size	Column size

178.2.2 DBMS_SPACE ASA_RECO_ROW Record Type

This type contains the column type of individual columns returned by the ASA_RECOMMENDATIONS Function.

Syntax

```
TYPE asa_reco_row IS RECORD (
   tablespace_name       VARCHAR2(30),
   segment_owner         VARCHAR2(30),
   segment_name          VARCHAR2(30),   segment_type          VARCHAR2(18),
   partition_name        VARCHAR2(30),
   allocated_space       NUMBER,
   used_space            NUMBER,
   reclaimable_space     NUMBER,
   chain_rowexcess       NUMBER,
   recommendations       VARCHAR2(1000),
   c1                    VARCHAR2(1000),
   c2                    VARCHAR2(1000),
   c3                    VARCHAR2(1000),
   task_id               NUMBER,
   mesg_id               NUMBER);
```

Attributes

Table 178-2 ASA_RECO_ROW Attributes

Field	Description
tablespace_name	Name of the tablespace containing the object
segment_owner	Name of the schema
segment_name	Name of the object
segment_type	Type of the segment 'TABLE','INDEX' and so on
partition_name	Name of the partition
allocated_space	Space allocated to the segment
used_space	Space actually used by the segment
reclaimable_space	Reclaimable free space in the segment
chain_rowexcess	Percentage of excess chain row pieces that can be eliminated
recommendations	Recommendation or finding for this segment
c1	Command associated with the recommendation
c2	Command associated with the recommendation
c3	Command associated with the recommendation
task_id	Advisor Task that processed this segment
mesg_id	Message ID corresponding to the recommendation

Related Topics

- DBMS_SPACE ASA_RECOMMENDATIONS Function
  This function returns recommendations using the stored results of the auto segment advisor. This function returns results from the latest run on any given object.

178.2.3 DBMS_SPACE ASA_RECO_ROW_TB Table Type

The type `asa_reco_row_tb` is a table of `asa_reco_row`.

Syntax

```
TYPE asa_reco_row_tb IS TABLE OF asa_reco_row;
```

178.3 Summary of DBMS_SPACE Subprograms

This table lists the `DBMS_SPACE` subprograms and briefly describes them.

Table 178-3 *DBMS_SPACE Package Subprograms*

Subprogram	Description
ASA_RECOMMENDATIONS Function	Returns recommendations/findings of segment advisor run automatically by the system or manually invoked by the user
CREATE_INDEX_COST Procedure	Determines the cost of creating an index on an existing table
CREATE_TABLE_COST Procedures	Determines the size of the table given various attributes
FREE_BLOCKS Procedure	Returns information about free blocks in an object (table, index, or cluster)
ISDATAFILEDROPPABLE_NAME Procedure	Checks whether a datafile is droppable
OBJECT_DEPENDENT_SEGMENTS Function	Returns the list of segments that are associated with the object
OBJECT_GROWTH_TREND Function	A table function where each row describes the space usage of the object at a specific point in time
SHRINK_TABLESPACE Procedure	Analyzes a bigfile tablespace before resizing or resizes a bigfile tablespace and optionally returns information about the resize operation
SPACE_USAGE Procedures	Returns information about free blocks in an auto segment space managed segment
UNUSED_SPACE Procedure	Returns information about unused space in an object (table, index, or cluster)

178.3.1 DBMS_SPACE ASA_RECOMMENDATIONS Function

This function returns recommendations using the stored results of the auto segment advisor. This function returns results from the latest run on any given object.

Syntax

```
DBMS_SPACE.ASA_RECOMMENDATIONS (
   all_runs        IN    VARCHAR2 DEFAULT := TRUE,
   show_manual     IN    VARCHAR2 DEFAULT := TRUE,
```

```
  show_findings   IN    VARCHAR2 DEFAULT := FALSE)
RETURN ASA_RECO_ROW_TB PIPELINED;
```

Parameters

Table 178-4 ASA_RECOMMENDATIONS Procedure Parameters

Parameter	Description
all_runs	Returns the results of all the auto advisor runs or only the results of the latest run. The valid values are TRUE and FALSE. The default value is TRUE. If TRUE, returns recommendations/findings for all runs of auto segment advisor. If FALSE, returns the results of the LATEST run only. LATEST does not make sense for manual invocation of segment advisor. This is applicable only for auto advisor.
show_manual	This parameter is used to indicate if the results of manual jobs should be included. If TRUE, results of manual tasks are shown. If FALSE, results of manual tasks are not shown. Specifying manual=true does not negate the specification of auto advisor tasks. However, the all_runs settings may override manual. If all_runs is FALSE, implying we only want to see the latest of auto advisor job, then manual may not be specified as TRUE. The valid values are TRUE and FALSE. The default value is TRUE.
show_findings	Shows only the findings instead of the recommendations. The valid values are TRUE and FALSE. The default value is FALSE.

Table 178-5 Parameter Usage

all_runs	show_manual	show_ findings	Outcome
TRUE	TRUE	TRUE	All findings from auto advisor and manual tasks.
TRUE	TRUE	FALSE	All recommendations from auto advisor and manual tasks.
TRUE	FALSE	TRUE	All findings from auto advisor tasks.
TRUE	FALSE	FALSE	All recommendations from all auto advisor tasks.
FALSE	TRUE	TRUE	N/A
FALSE	TRUE	FALSE	N/A
FALSE	FALSE	TRUE	Findings for the latest auto advisor task.
FALSE	FALSE	FALSE	Recommendations from the latest auto advisor task.

178.3.2 CREATE_INDEX_COST Procedure

This procedure determines the cost of creating an index on an existing table. The input is the DDL statement that will be used to create the index. The procedure will output the storage required to create the index.

Syntax

```
DBMS_SPACE.CREATE_INDEX_COST (
   ddl             IN    VARCHAR2,
   used_bytes      OUT   NUMBER,
   alloc_bytes     OUT   NUMBER,
   plan_table      IN    VARCHAR2 DEFAULT NULL);
```

Pragmas

```
pragma restrict_references(create_index_cost,WNDS);
```

Parameters

Table 178-6 CREATE_INDEX_COST Procedure Parameters

Parameter	Description
ddl	The create index DDL statement
used_bytes	The number of bytes representing the actual index data
alloc_bytes	Size of the index when created in the tablespace
plan_table	Which plan table to use, default NULL

Usage Notes

- The table on which the index is created must already exist.
- The computation of the index size depends on statistics gathered on the segment.
- It is imperative that the table must have been analyzed recently.
- In the absence of correct statistics, the results may be inaccurate, although the procedure will not raise any errors.

178.3.3 CREATE_TABLE_COST Procedures

This procedure is used in capacity planning to determine the size of the table given various attributes. The size of the object can vary widely based on the tablespace storage attributes, tablespace block size, and so on. There are two overloads of this procedure.

- The first version takes the column information of the table as argument and outputs the table size.
- The second version takes the average row size of the table as argument and outputs the table size.

This procedure can be used on tablespace of dictionary managed and locally managed extent management as well as manual and auto segment space management.

Syntax

```
DBMS_SPACE.CREATE_TABLE_COST (
   tablespace_name     IN VARCHAR2,
   avg_row_size        IN NUMBER,
   row_count           IN NUMBER,
   pct_free            IN NUMBER,
   used_bytes          OUT NUMBER,
   alloc_bytes         OUT NUMBER);

DBMS_SPACE.CREATE_TABLE_COST (
   tablespace_name     IN VARCHAR2,
   colinfos            IN CREATE_TABLE_COST_COLUMNS,
   row_count           IN NUMBER,
   pct_free            IN NUMBER,
   used_bytes          OUT NUMBER,
   alloc_bytes         OUT NUMBER);

CREATE TYPE create_table_cost_colinfo IS OBJECT (
   COL_TYPE   VARCHAR(200),
   COL_SIZE   NUMBER);
```

Parameters

Table 178-7 CREATE_TABLE_COST Procedure Parameters

Parameter	Description
tablespace_name	The tablespace in which the object will be created. The default is SYSTEM tablespace.
avg_row_size	The anticipated average row size in the table
colinfos	The description of the columns
row_count	The anticipated number of rows in the table
pct_free	The percentage of free space in each block for future expansion of existing rows due to updates
used_bytes	The space used by user data
alloc_bytes	The size of the object taking into account the tablespace extent characteristics

Usage Notes

- The used_bytes represent the actual bytes used by the data. This includes the overhead due to the block metadata, pctfree etc.
- The alloc_bytes represent the size of the table when it is created in the tablespace. This takes into account, the size of the extents in the tablespace and tablespace extent management properties.

Examples

```
-- review the parameters
SELECT argument_name, data_type, type_owner, type_name
FROM all_arguments
WHERE object_name = 'CREATE_TABLE_COST'
AND overload = 2

-- examine the input parameter type
```

```
SELECT text
FROM dba_source
WHERE name = 'CREATE_TABLE_COST_COLUMNS';

-- drill down further into the input parameter type
SELECT text
FROM dba_source
WHERE name = 'create_table_cost_colinfo';

set serveroutput on

DECLARE
 ub NUMBER;
 ab NUMBER;
 cl sys.create_table_cost_columns;
BEGIN
  cl := sys.create_table_cost_columns( sys.create_table_cost_colinfo('NUMBER',10),
        sys.create_table_cost_colinfo('VARCHAR2',30),
        sys.create_table_cost_colinfo('VARCHAR2',30),
        sys.create_table_cost_colinfo('DATE',NULL));

  DBMS_SPACE.CREATE_TABLE_COST('SYSTEM',cl,100000,0,ub,ab);

  DBMS_OUTPUT.PUT_LINE('Used Bytes: ' || TO_CHAR(ub));
  DBMS_OUTPUT.PUT_LINE('Alloc Bytes: ' || TO_CHAR(ab));
END;
/
```

178.3.4 FREE_BLOCKS Procedure

This procedure returns information about free blocks in an object (table, index, or cluster).

See SPACE_USAGE Procedures for returning free block information in an auto segment space managed segment.

Syntax

```
DBMS_SPACE.FREE_BLOCKS (
   segment_owner     IN  VARCHAR2,
   segment_name      IN  VARCHAR2,
   segment_type      IN  VARCHAR2,
   freelist_group_id IN  NUMBER,
   free_blks         OUT NUMBER,
   scan_limit        IN  NUMBER DEFAULT NULL,
   partition_name    IN  VARCHAR2 DEFAULT NULL);
```

Pragmas

```
pragma restrict_references(free_blocks,WNDS);
```

Parameters

Table 178-8 FREE_BLOCKS Procedure Parameters

Parameter	Description
segment_owner	Schema name of the segment to be analyzed
segment_name	Segment name of the segment to be analyzed

Table 178-8 (Cont.) FREE_BLOCKS Procedure Parameters

Parameter	Description
segment_type	Type of the segment to be analyzed (TABLE, INDEX, or CLUSTER): • TABLE • TABLE PARTITION • TABLE SUBPARTITION • INDEX • INDEX PARTITION • INDEX SUBPARTITION • CLUSTER • LOB • LOB PARTITION • LOB SUBPARTITION
freelist_group_id	Freelist group (instance) whose free list size is to be computed
free_blks	Returns count of free blocks for the specified group
scan_limit	Maximum number of free list blocks to read (optional). Use a scan limit of X you are interested only in the question, "Do I have X blocks on the free list?"
partition_name	Partition name of the segment to be analyzed. This is only used for partitioned tables. The name of subpartition should be used when partitioning is composite.

Examples

The following uses the CLUS cluster in SCOTT schema with 4 freelist groups. It returns the number of blocks in freelist group 3 in CLUS.

```
DBMS_SPACE.FREE_BLOCKS('SCOTT', 'CLUS', 'CLUSTER', 3, :free_blocks);
```

> **Note:**
>
> An error is raised if scan_limit is not a positive number.

178.3.5 ISDATAFILEDROPPABLE_NAME Procedure

This procedure checks whether a datafile is droppable. This procedure may be called before actually dropping the file.

Syntax

```
DBMS_SPACE.ISDATAFILEDROPPABLE_NAME (
   filename     IN     VARCHAR2,
   value        OUT    NUMBER);
```

Pragmas

```
pragma restrict_references(free_blocks,WNDS);
```

Parameters

Table 178-9 ISDATAFILEDROPPABLE_NAME Procedure Parameters

Parameter	Description
filename	Name of the file
value	Values: 0 if the file is not droppable, 1 if the file is droppable.

Examples

```
DECLARE   fname    VARCHAR2(100);   retval   NUMBER;BEGIN   SELECT file_name   INTO
fname   FROM dba_data_files   WHERE file_name like
'%empty%';DBMS_SPACE.ISDATAFILEDROPPABLE_NAME(fname,
retval);DBMS_OUTPUT.PUT_LINE(retval);END;/
```

178.3.6 OBJECT_DEPENDENT_SEGMENTS Function

This table function, given an object, returns the list of segments that are associated with the object.

Syntax

```
DBMS_SPACE.OBJECT_DEPENDENT_SEGMENTS(
   objowner     IN     VARCHAR2,
   objname      IN     VARCHAR2,
   partname     IN     VARCHAR2,
   objtype      IN     NUMBER)
  RETURN dependent_segments_table PIPELINED;
```

Parameters

Table 178-10 OBJECT_DEPENDENT_SEGMENTS Function Parameters

Parameter	Description
objowner	The schema containing the object
objname	The name of the object
partname	The name of the partition
objtype	Type of the object: • OBJECT_TYPE_TABLE constant positive := 1; • OBJECT_TYPE_NESTED_TABLE constant positive := 2; • OBJECT_TYPE_INDEX constant positive := 3; • OBJECT_TYPE_CLUSTER constant positive := 4; • OBJECT_TYPE_TABLE_PARTITION constant positive := 7; • OBJECT_TYPE_INDEX_PARTITION constant positive := 8; • OBJECT_TYPE_TABLE_SUBPARTITION constant positive := 9; • OBJECT_TYPE_INDEX_SUBPARTITION constant positive := 10; • OBJECT_TYPE_MV constant positive := 13; • OBJECT_TYPE_MVLOG constant positive := 14;

Return Values

The content of one row of a `dependent_segments_table`:

```
TYPE object_dependent_segment IS RECORD (
   segment_owner      VARCHAR2(100),
   segment_name       VARCHAR2(100),
   segment_type       VARCHAR2(100),
   tablespace_name    VARCHAR2(100),
   partition_name     VARCHAR2(100),
   lob_column_name    VARCHAR2(100));
```

Table 178-11 OBJECT_DEPENDENT_SEGMENT Type Parameters

Parameter	Description
`segment_owner`	The schema containing the segment
`segment_name`	The name of the segment
`segment_type`	The type of the segment, such as table, index or LOB
`tablespace_name`	The name of the tablespace
`partition_name`	The name of the partition, if any
`lob_column_name`	The name of the `LOB` column, if any

178.3.7 OBJECT_GROWTH_TREND Function

This is a table function. The output is one or more rows where each row describes the space usage of the object at a specific point in time.

Either the space usage totals will be retrieved from Automatic Workload Repository Facilities (AWRF), or the current space usage will be computed and combined with space usage deltas retrieved from AWRF.

Syntax

```
DBMS_SPACE.OBJECT_GROWTH_TREND (
   object_owner           IN    VARCHAR2,
   object_name            IN    VARCHAR2,
   object_type            IN    VARCHAR2,
   partition_name         IN    VARCHAR2 DEFAULT NULL,
   start_time             IN    TIMESTAMP DEFAULT NULL,
   end_time               IN    TIMESTAMP DEFAULT NULL,
   interval               IN    DSINTERVAL_UNCONSTRAINED DEFAULT NULL,
   skip_interpolated      IN    VARCHAR2 DEFAULT 'FALSE',
   timeout_seconds        IN    NUMBER DEFAULT NULL,
   single_datapoint_flag  IN    VARCHAR2 DEFAULT 'TRUE')
 RETURN object_growth_trend_table PIPELINED;
```

Parameters

Table 178-12 OBJECT_GROWTH_TREND Function Parameters

Parameter	Description
`object_owner`	The schema containing the object
`object_name`	The name of the object

Table 178-12 (Cont.) OBJECT_GROWTH_TREND Function Parameters

Parameter	Description
object_type	The type of the object
partition_name	The name of the partition
start_time	Statistics generated after this time will be used in generating the growth trend
end_time	Statistics generated until this time will be used in generating the growth trend
interval	The interval at which to sample
skip_interpolated	Whether interpolation of missing values should be skipped
timeout_seconds	The time-out value for the function in seconds
single_data_point_flag	Whether in the absence of statistics the segment should be sampled

Return Values

The object_growth_trend_row and object_growth_trend_table are used by the OBJECT_GROWTH_TREND table function to describe its output.

```
TYPE object_growth_trend_row IS RECORD(
   timepoint      TIMESTAMP,
   space_usage    NUMBER,
   space_alloc    NUMBER,
   quality        VARCHAR(20));
```

Table 178-13 OBJECT_GROWTH_TREND_ROW Type Parameters

Parameter	Description
timepoint	The time at which the statistic was recorded
space_usage	The space used by data
space_alloc	The size of the segment including overhead and unused space
quality	The quality of result: "GOOD", "INTERPOLATED", "PROJECTION"

```
TYPE object_growth_trend_table IS TABLE OF object_growth_trend_row;
```

178.3.8 SHRINK_TABLESPACE Procedure

This procedure can resize a bigfile tablespace or analyze a bigfile tablespace before resizing.

Syntax

```
DBMS_SPACE.SHRINK_TABLESPACE (
   ts_name             IN  VARCHAR2,
   shrink_mode         IN  NUMBER,
   target_size         IN  NUMBER,
   shrink_result      OUT CLOB);


DBMS_SPACE.SHRINK_TABLESPACE (
   ts_name             IN  VARCHAR2,
```

```
   shrink_mode         IN  NUMBER,
   target_size         IN  NUMBER);
```

Parameters

Table 178-14 SHRINK_TABLESPACE Procedure Parameters

Parameter	Description
ts_name	The name of the tablespace to be analyzed or resized
shrink_mode	The shrink mode to execute. The values are: • TS_MODE_ANALYZE • TS_MODE_SHRINK • TS_MODE_SHRINK_FORCE The default mode is TS_MODE_SHRINK which moves objects online by default, except for index-organized tables. TS_MODE_SHRINK_FORCE will move objects online by default, but if the online move fails, it will attempt to move them offline.
target_size	The desired tablespace size specified in **bytes**. The default value is TS_TARGET_MAX_SHRINK.
shrink_result	Output result of the procedure returned as a CLOB. The output results for TS_MODE_SHRINK include: • total number and size of moved objects • original and new datafile size • process time The output results for TS_MODE_ANALYZE include: • list of movable objects • total number and size of movable objects • suggested target size • process time

Errors

Table 178-15 SHRINK_TABLESPACE Procedure Errors

Error	Description
ORA-00054	**Message:** Failed to acquire a lock (Type: "Type", Name: "Name", Description: "Description") because it is currently held by another session. The resource being locked can be identified by ID1_value ("ID1_description") and ID2_value ("ID2_description") **Reasons:** • Procedure exited because it can't move an object. • Tablespace is currently help by another session.

Examples

This example analyzes bigfile tablespace TBS_1.

```
set serveroutput on
execute dbms_space.shrink_tablespace('TBS_1', shrink_mode =>
DBMS_SPACE.TS_MODE_ANALYZE);

------------------------------ANALYZE
```

```
RESULT--------------------------------------
1. { BG_TEST.SYS_IL0000081422C00004$$ | type: INDEX | blocks: 256 |
tablespace_name: TBS_1 }
2. { BG_TEST.SYS_IL0000081422C00005$$ | type: INDEX | blocks: 512 |
tablespace_name: TBS_1 }
3. { BG_TEST.T2 | type: TABLE | blocks: 512 | tablespace_name: TBS_1 }
4. { BG_TEST.T2_LOB1 | type: LOBSEGMENT | blocks: 45824 | tablespace_name:
TBS_1}
5. { BG_TEST.T2_LOB2 | type: LOBSEGMENT | blocks: 41216 | tablespace_name:
TBS_1}
Total Movable Objects: 5
Total Movable Size(GB): .67
Orginal Datafile Size(GB): 10
Suggested Target Size(GB): 2.09
Process Time: +00 00:00:03.94897
```

This example shrinks the bigfile tablespace TBS_1 to its current minimum possible size.

```
set serveroutput on
execute dbms_space.shrink_tablespace('TBS_1');

------------------SHRINK RESULT------------------
Total Moved Objects: 5
Total Moved Size(GB): 1.35
Orginal Datafile Size(GB): 10
New Datafile Size(GB): 1.81
Process Time: +00 00:00:50.94897
```

178.3.9 SPACE_USAGE Procedures

This procedure has two variations to show space usage.

The first form of the procedure shows the space usage of data blocks under the segment High Water Mark. You can calculate usage for `LOB`s, `LOB PARTITIONS` and `LOB SUBPARTITIONS`. This procedure can only be used on tablespaces that are created with auto segment space management. The bitmap blocks, segment header, and extent map blocks are not accounted for by this procedure. Note that this overload cannot be used on `SECUREFILE LOB`s.

> **Note:**
>
> For LOB segments, the number of blocks that is returned from `full_blocks` and `unformatted_blocks` is actually the number of chunks for the LOB segment.

The second form of the procedure returns information about `SECUREFILE LOB` space usage. It will return the amount of space in blocks being used by all the `SECUREFILE LOB`s in the `LOB` segment. The procedure displays the space actively used by the LOB column, freed space that has retention expired, and freed space that has retention unexpired. Note that this overload can be used only on `SECUREFILE LOB`s.

Syntax

```
DBMS_SPACE.SPACE_USAGE(
   segment_owner           IN  VARCHAR2,
   segment_name            IN  VARCHAR2,
   segment_type            IN  VARCHAR2,
   unformatted_blocks      OUT NUMBER,
   unformatted_bytes       OUT NUMBER,
   fs1_blocks              OUT NUMBER,
   fs1_bytes               OUT NUMBER,
   fs2_blocks              OUT NUMBER,
   fs2_bytes               OUT NUMBER,
   fs3_blocks              OUT NUMBER,
   fs3_bytes               OUT NUMBER,
   fs4_blocks              OUT NUMBER,
   fs4_bytes               OUT NUMBER,
   full_blocks             OUT NUMBER,
   full_bytes              OUT NUMBER,
   partition_name          IN  VARCHAR2 DEFAULT NULL);

DBMS_SPACE.SPACE_USAGE(
   segment_owner           IN    VARCHAR2,
   segment_name            IN    VARCHAR2,
   segment_type            IN    VARCHAR2,
   segment_size_blocks     OUT   NUMBER,
   segment_size_bytes      OUT   NUMBER,
   used_blocks             OUT   NUMBER,
   used_bytes              OUT   NUMBER,
   expired_blocks          OUT   NUMBER,
   expired_bytes           OUT   NUMBER,
   unexpired_blocks        OUT   NUMBER,
   unexpired_bytes         OUT   NUMBER,
   partition_name          IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 178-16 SPACE_USAGE Procedure Parameters

Parameter	Description
segment_owner	Schema name of the segment to be analyzed
segment_name	Name of the segment to be analyzed
partition_name	Partition name of the segment to be analyzed
segment_type	Type of the segment to be analyzed (TABLE, INDEX, or CLUSTER): • TABLE • TABLE PARTITION • TABLE SUBPARTITION • INDEX • INDEX PARTITION • INDEX SUBPARTITION • CLUSTER • LOB • LOB PARTITION • LOB SUBPARTITION

Table 178-16 (Cont.) SPACE_USAGE Procedure Parameters

Parameter	Description
unformatted_blocks	For LOB segments, the number of blocks that is returned from unformatted_blocks is actually the number of chunks for the LOB segment.
unformatted bytes	Total number of bytes unformatted
fs1_blocks	Number of blocks having at least 0 to 25% free space
fs1_bytes	Number of bytes having at least 0 to 25% free space
fs2_blocks	Number of blocks having at least 25 to 50% free space
fs2_bytes	Number of bytes having at least 25 to 50% free space
fs3_blocks	Number of blocks having at least 50 to 75% free space
fs3_bytes	Number of bytes having at least 50 to 75% free space
fs4_blocks	Number of blocks having at least 75 to 100% free space
fs4_bytes	Number of bytes having at least 75 to 100% free space
ful1_blocks	The number of blocks that is returned from full_blocks is actually the number of chunks for the LOB segment
full_bytes	Total number of bytes full in the segment
segment_size_blocks	Number of blocks allocated to the segment
segment_size_bytes	Number of bytes allocated to the segment
used_blocks	Number blocks allocated to the LOB that contains active data
used_bytes	Number bytes allocated to the LOB that contains active data
expired_blocks	Number of expired blocks used by the LOB to keep version data
expired_bytes	Number of expired bytes used by the LOB to keep version data
unexpired_blocks	Number of unexpired blocks used by the LOB to keep version data
unexpired_bytes	Number of unexpired bytes used by the LOB to keep version data
partition_name	Name of the partition (NULL if not a partition)

Examples

```
variable unf number;
variable unfb number;
variable fs1 number;
variable fs1b number;
variable fs2 number;
variable fs2b number;
variable fs3 number;
variable fs3b number;
variable fs4 number;
variable fs4b number;
variable full number;
variable fullb number;

begin
dbms_space.space_usage('U1','T',
                       'TABLE',
                       :unf, :unfb,
                       :fs1, :fs1b,
                       :fs2, :fs2b,
```

```
                        :fs3, :fs3b,
                        :fs4, :fs4b,
                        :full, :fullb);
end;
/
print unf ;
print unfb ;
print fs4 ;
print fs4b;
print fs3 ;
print fs3b;
print fs2 ;
print fs2b;
print fs1 ;
print fs1b;
print full;
print fullb;
```

178.3.10 UNUSED_SPACE Procedure

This procedure returns information about unused space in an object (table, index, or cluster).

Syntax

```
DBMS_SPACE.UNUSED_SPACE (
   segment_owner              IN  VARCHAR2,
   segment_name               IN  VARCHAR2,
   segment_type               IN  VARCHAR2,
   total_blocks               OUT NUMBER,
   total_bytes                OUT NUMBER,
   unused_blocks              OUT NUMBER,
   unused_bytes               OUT NUMBER,
   last_used_extent_file_id   OUT NUMBER,
   last_used_extent_block_id  OUT NUMBER,
   last_used_block            OUT NUMBER,
   partition_name             IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 178-17 UNUSED_SPACE Procedure Parameters

Parameter	Description
segment_owner	Schema name of the segment to be analyzed
segment_name	Segment name of the segment to be analyzed
segment_type	Type of the segment to be analyzed (TABLE, INDEX, or CLUSTER): • TABLE • TABLE PARTITION • TABLE SUBPARTITION • INDEX • INDEX PARTITION • INDEX SUBPARTITION • CLUSTER • LOB • LOB PARTITION • LOB SUBPARTITION

Table 178-17 (Cont.) UNUSED_SPACE Procedure Parameters

Parameter	Description
total_blocks	Returns total number of blocks in the segment
total_bytes	Returns total number of blocks in the segment, in bytes
unused_blocks	Returns number of blocks which are not used
unused_bytes	Returns, in bytes, number of blocks which are not used
last_used_extent_file_id	Returns the file ID of the last extent which contains data
last_used_extent_block_id	Returns the starting block ID of the last extent which contains data
last_used_block	Returns the last block within this extent which contains data
partition_name	Partition name of the segment to be analyzed. This is only used for partitioned tables; the name of subpartition should be used when partitioning is compose.

Examples

The following declares the necessary bind variables and executes.

```
DBMS_SPACE.UNUSED_SPACE('SCOTT', 'EMP', 'TABLE', :total_blocks,
   :total_bytes,:unused_blocks, :unused_bytes, :lastextf,
   :last_extb, :lastusedblock);
```

179
DBMS_SPACE_ADMIN

The DBMS_SPACE_ADMIN package provides functionality for locally managed tablespaces.

This chapter contains the following topics:

- Security Model
- Constants
- Operational Notes
- Summary of DBMS_SPACE_ADMIN Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide* for an example and description of using DBMS_SPACE_ADMIN.

179.1 DBMS_SPACE_ADMIN Security Model

This package runs with SYS privileges; therefore, any user who has privilege to execute the package can manipulate the bitmaps.

179.2 DBMS_SPACE_ADMIN Constants

The DBMS_SPACE_ADMIN package provides constants that can be used for specifying parameter values.

Table 179-1 DBMS_SPACE_ADMIN Constants

Constant	Type	Value	Description
SEGMENT_VERIFY_EXTENTS	POSITIVE	1	Verifies that the space owned by segment is appropriately reflected in the bitmap as used
SEGMENT_VERIFY_EXTENTS_GLOBAL	POSITIVE	2	Verifies that the space owned by segment is appropriately reflected in the bitmap as used and that no other segment claims any of this space to be used by it
SEGMENT_MARK_CORRUPT	POSITIVE	3	Marks a temporary segment as corrupt whereby facilitating its elimination from the dictionary (without space reclamation)

Table 179-1 (Cont.) DBMS_SPACE_ADMIN Constants

Constant	Type	Value	Description
SEGMENT_MARK_VALID	POSITIVE	4	Marks a corrupt temporary segment as valid. It is useful when the corruption in the segment extent map or elsewhere has been resolved and the segment can be dropped normally.
SEGMENT_DUMP_EXTENT_MAP	POSITIVE	5	Dumps the extent map for a given segment
TABLESPACE_VERIFY_BITMAP	POSITIVE	6	Verifies the bitmap of the tablespace with extent maps of the segments in that tablespace to make sure everything is consistent
TABLESPACE_EXTENT_MAKE_FREE	POSITIVE	7	Marks the block range (extent) as free in the bitmaps
TABLESPACE_EXTENT_MAKE_USED	POSITIVE	8	Marks the block range (extent) as used in the bitmaps
SEGMENT_VERIFY_BASIC	POSITIVE	9	Performs the basic metadata checks
SEGMENT_VERIFY_DEEP	POSITIVE	10	Performs deep verification
SEGMENT_VERIFY_SPECIFIC	POSITIVE	11	Performs a specific check for the segment
HWM_CHECK	POSITIVE	12	Checks high water mark (HWM)
BMB_CHECK	POSITIVE	13	Checks integrity among L1, L2 and L3 BMBs (Bit Map Blocks)
SEG_DICT_CHECK	POSITIVE	14	Checks consistency of segment header with corresponding SEG entry
EXTENT_TS_BITMAP_CHECK	POSITIVE	15	Checks whether the tablespace bitmaps corresponding to the extent map are marked used
DB_BACKPOINTER_CHECK	POSITIVE	16	Checks whether the L1 BMBs, L2 BMBs, L3 BMBs and data blocks point to the same parent segment
EXTENT_SEGMENT_BITMAP_CHECK	POSITIVE	17	Checks whether the bitmap blocks are consistent with the extent map
BITMAPS_CHECK	POSITIVE	18	Checks from the datablocks that the bitmap states representing the blocks are consistent
TS_VERIFY_BITMAPS	POSITIVE	19	Checks whether the tablespace bitmaps are consistent with the extents belonging to that tablespace
TS_VERIFY_DEEP	POSITIVE	20	Performs TS_VERIFY_BITMAPS and TS_VERIFY_SEGMENTS with DEEP option
TS_VERIFY_SEGMENTS	POSITIVE	21	Performs ASSM_SEGMENT_VERIFY on all segments in the tablespace, taking either the BASIC or the DEEP option
SEGMENTS_DUMP_BITMAP_SUMMARY	POSITIVE	27	Dumps only bitmap block summaries

179.3 DBMS_SPACE_ADMIN Operational Notes

Before migrating the `SYSTEM` tablespace, certain conditions must be met. These conditions are enforced by the `TABLESPACE_MIGRATE_TO_LOCAL` procedure, except for the cold backup.

- The database must have a default temporary tablespace that is not `SYSTEM`.
- Dictionary-managed tablespaces cannot have any rollback segments.
- A locally managed tablespace must have at least one online rollback segment. If you are using automatic undo management, then an undo tablespace must be online.
- All tablespaces—except the tablespace containing the rollback segment or the undo tablespace—must be read-only.
- You must have a cold backup of the database.
- The system must be in restricted mode.

179.4 Summary of DBMS_SPACE_ADMIN Subprograms

This table lists the `DBMS_SPACE_ADMIN` subprograms and briefly describes them.

Table 179-2 DBMS_SPACE_ADMIN Package Subprograms

Subprogram	Description
ASSM_SEGMENT_VERIFY Procedure	Verifies segments created in ASSM (Automatic Segment-Space Management) tablespaces
ASSM_TABLESPACE_VERIFY Procedure	Verifies ASSM tablespaces
DROP_EMPTY_SEGMENTS Procedure	Drops segments from empty tables or table fragments and dependent objects
GET_SEGADV_ATTRIB Procedure	returns the values of attributes of `DBSM_SPACE_ADMIN` package
MATERIALIZE_DEFERRED_SEGMENTS Procedure	Materializes segments for tables and table fragments with deferred segment creation and their dependent objects
SEGMENT_CORRUPT Procedure	Marks the segment corrupt or valid so that appropriate error recovery can be done
SEGMENT_DROP_CORRUPT Procedure	Drops a segment currently marked corrupt (without reclaiming space)
SEGMENT_DUMP Procedure	Dumps the segment header and extent maps of a given segment
SEGMENT_VERIFY Procedure	Verifies the consistency of the extent map of the segment
SET_SEGADV_ATTRIB Procedure	Sets/changes the values of attributes of `DBSM_SPACE_ADMIN` package
TABLESPACE_FIX_BITMAPS Procedure	Marks the appropriate block range (extent) as free or used in bitmap
TABLESPACE_FIX_SEGMENT_STATES Procedure	Fixes the state of the segments in a tablespace in which migration was aborted
TABLESPACE_MIGRATE_FROM_LOCAL Procedure	Migrates a locally managed tablespace to dictionary-managed tablespace
TABLESPACE_MIGRATE_TO_LOCAL Procedure	Migrates a tablespace from dictionary-managed format to locally managed format

Table 179-2 (Cont.) DBMS_SPACE_ADMIN Package Subprograms

Subprogram	Description
TABLESPACE_REBUILD_BITMAPS Procedure	Rebuilds the appropriate bitmaps
TABLESPACE_REBUILD_QUOTAS Procedure	Rebuilds quotas for given tablespace
TABLESPACE_RELOCATE_BITMAPS Procedure	Relocates the bitmaps to the destination specified
TABLESPACE_VERIFY Procedure	Verifies that the bitmaps and extent maps for the segments in the tablespace are synchronized

179.4.1 ASSM_SEGMENT_VERIFY Procedure

Given a segment definition, the procedure verifies the basic consistency of the space metadata blocks as well as consistency between space metadata and segment data blocks. This procedure verifies segments created in Automatic Segment Space Management (ASSM) tablespaces.

There is however a difference between basic verification and deep verification:

- Basic verification involves consistency checks of space metadata, such as integrity among level 1, level 2, level 3 bitmap blocks, consistency of segment extent map and level 1 bitmap ranges.
- Deep verification involves consistency checks between datablocks and space metadata blocks such as whether the datablocks point correctly to the parent level 1 bitmap blocks, and whether the freeness states in the datablocks are consistent with the freeness states of bits in level 1 bitmap blocks corresponding to the datablocks.

Syntax

```
DBMS_SPACE_ADMIN.ASSM_SEGMENT_VERIFY (
   segment_owner    IN VARCHAR2,
   segment_name     IN VARCHAR2,
   segment_type     IN VARCHAR2,
   partition_name   IN VARCHAR2,
   verify_option    IN POSITIVE  DEFAULT SEGMENT_VERIFY_BASIC,
   attrib           IN POSITIVE  DEFAULT NULL);
```

Parameters

Table 179-3 ASSM_SEGMENT_VERIFY Procedure Parameters

Parameter	Description
segment_owner	Schema that owns the segment
segment_name	Name of the segment to be verified
segment_type	Segment namespace is one of TABLE, TABLE PARTITION, TABLE SUBPARTITION, INDEX, INDEX PARTITION, INDEX SUBPARTITION, LOB, LOB PARTITION, LOB SUBPARTITION, CLUSTER
partition_name	Name of the partition or subpartition

Table 179-3 (Cont.) ASSM_SEGMENT_VERIFY Procedure Parameters

Parameter	Description
verify_option	One of the following options: • SEGMENT_VERIFY_BASIC := 9. Performs the basic metadata checks (Default) • SEGMENT_VERIFY_DEEP := 10. Performs deep verification • SEGMENT_VERIFY_SPECIFIC := 11. Performs a specific check for the segment
attrib	When option SEGMENT_VERIFY_SPECIFIC is specified as option, attrib can be one of the following: • HWM_CHECK := 12. Checks whether high water mark information is accurate • BMB_CHECK := 13. Checks whether space bitmap blocks have correct backpointers to the segment header • SEG_DICT_CHECK := 14. Checks whether dictionary information for segment is accurate • EXTENT_TS_BITMAP_CHECK := 15. Checks whether extent maps are consistent with file level bitmaps • DB_BACKPOINTER_CHECK := 16. Checks whether datablocks have correct backpointers to the space metadata blocks • EXTENT_SEGMENT_BITMAP_CHECK := 17. Checks whether extent map in the segment matches the bitmaps in the segment • BITMAPS_CHECK := 18. Checks whether space bitmap blocks are accurate

Usage Notes

- Using this procedure requires SYSDBA privileges.
- You can determine the relative file # and header block # (header_relative_file and header_block parameters) by querying DBA_SEGMENTS.
- This procedure outputs a dump file named sid_ora_*process_ID*.trc to the location specified in the USER_DUMP_DEST initialization parameter.

179.4.2 ASSM_TABLESPACE_VERIFY Procedure

This procedures verifies all the segments created in an ASSM tablespace. The verification for each segment performs basic consistency checks of the space metadata blocks as well as consistency checks between space metadata and segment data blocks.

Syntax

```
DBMS_SPACE_ADMIN.ASSM_TABLESPACE_VERIFY (
   tablespace_name   IN VARCHAR2,
   ts_option         IN POSITIVE,
   segment_option    IN POSITIVE DEFAULT NULL);
```

Parameters

Table 179-4 ASSM_TABLESPACE_VERIFY Procedure Parameters

Parameter	Description
tablespace_name	Name of the tablespace to verify. The tablespace must be an ASSM tablespace.
ts_option	• 19: SecureFiles supports DBMS_SPACE_ADMIN.ASSM_TABLESPACE_VERIFY only when you set the value of TS_OPTION to 19. • TS_VERIFY_BITMAPS := 19. The bitmaps are verified against the extents. This detects bits that are marked used or free wrongly and detects multiple allocation of extents. The file metadata is validated against file$ and control file. • TS_VERIFY_DEEP := 20. This option is used to verify the file bitmaps as well perform checks on all the segments. • TS_VERIFY_SEGMENTS := 21. This option is used to invoke SEGMENT_VERIFY on all the segments in the tablespace. Optionally you can write a script that queries all the segments in the tablespace and invoke SEGMENT_VERIFY.
segment_option	When TS_VERIFY_SEGMENTS is specified, segment_option can be one of the following: • SEGMENT_VERIFY_BASIC := 9 • SEGMENT_VERIFY_DEEP := 10 The value of segment_option is NULL when TS_VERIFY_DEEP or TS_VERIFY_BITMAPS is specified.

Usage Notes

- Using this procedure requires SYSDBA privileges.
- This procedure outputs a dump file named sid_ora_*process_ID*.trc to the location specified in the USER_DUMP_DEST initialization parameter.

179.4.3 DROP_EMPTY_SEGMENTS Procedure

This procedures drops segments from empty tables or table fragments and dependent objects.

Syntax

```
DBMS_SPACE_ADMIN.DROP_EMPTY_SEGMENTS (
   schema_name       IN     VARCHAR2   DEFAULT NULL,
   table_name        IN     VARCHAR2   DEFAULT NULL,
   partition_name    IN     VARCHAR2   DEFAULT NULL);
```

Parameters

Table 179-5 DROP_EMPTY_SEGMENTS Procedure Parameters

Parameter	Description
schema_name	Name of schema
table_name	Name of table

Table 179-5 (Cont.) DROP_EMPTY_SEGMENTS Procedure Parameters

Parameter	Description
partition_name	Name of partition

Usage Notes

Given a schema name, this procedure scans all tables in the schema. For each table, if the table or any of its fragments are found to be empty, and the table satisfies certain criteria (restrictions being the same as those described in "Restrictions on Deferred Segment Creation"), then the empty table fragment and associated index segments are dropped along with the corresponding LOB data and index segments. A subsequent insert creates segments with the same properties.

Optionally:

- No schema_name is specified, in which case tables belonging to all schemas are scanned
- Both schema_name and table_name are specified to perform the operation on a specified table
- All three arguments are supplied, restricting the operation to the partition and its dependent objects

179.4.4 GET_SEGADV_ATTRIB Procedure

This procedure returns the values of attributes of DBMS_SPACE_ADMIN package.

Syntax

```
DBMS_SPACE_ADMIN.GET_SEGADV_ATTRIB(
   attribute  IN  NUMBER,
   value      OUT NUMBER);
```

Parameters

Table 179-6 GET_SEGADV_ATTRIB Procedure Parameters

Parameter	Description
attribute	Supported attributes: • COMP_ADVISOR — Provides an option to enable or disable Compression Advisor for Automatic Segment Advisor. By default Compression Advisor is enabled for Automatic Segment Advisor. • COMP_LOB — Provides an option to enable or disable Compression Advisor for the tables with LOB columns while Automatic Segment Advisor is running. By default Compression Advisor is enabled for tables with LOB columns.
value	Supported values: • ATTR_ENABLE : 1 • ATTR_DISABLE : 0

179.4.5 MATERIALIZE_DEFERRED_SEGMENTS Procedure

This procedure materializes segments for tables and table fragments with deferred segment creation and their dependent objects.

Syntax

```
DBMS_SPACE_ADMIN.MATERIALIZE_DEFERRED_SEGMENTS (
   schema_name       IN     VARCHAR2   DEFAULT NULL,
   table_name        IN     VARCHAR2   DEFAULT NULL,
   partition_name    IN     VARCHAR2   DEFAULT NULL);
```

Parameters

Table 179-7 MATERIALIZE_DEFERRED_SEGMENTS Procedure Parameters

Parameter	Description
schema_name	Name of schema
table_name	Name of table
partition_name	Name of partition

Usage Notes

Given a schema name, this procedure scans all tables in the schema. For each table, if the deferred or delayed segment property is set for the table or any of its fragments, then a new segment is created for those fragments and their dependent objects.

Optionally:

- No `schema_name` is specified, in which case tables belonging to all schemas are scanned
- Both `schema_name` and `table_name` are specified to perform the operation on a specified table
- All three arguments are supplied, restricting the operation to the partition and its dependent objects

179.4.6 SEGMENT_CORRUPT Procedure

This procedure marks the segment corrupt or valid so that appropriate error recovery can be performed.

It cannot be used on the `SYSTEM` tablespace.

Syntax

```
DBMS_SPACE_ADMIN.SEGMENT_CORRUPT (
   tablespace_name         IN    VARCHAR2,
   header_relative_file    IN    POSITIVE,
   header_block            IN    NUMBER,
   corrupt_option          IN    POSITIVE  DEFAULT SEGMENT_MARK_CORRUPT);
```

Parameters

Table 179-8 SEGMENT_CORRUPT Procedure Parameters

Parameter	Description
tablespace_name	Name of tablespace in which segment resides
header_relative_file	Relative file number of segment header
header_block	Block number of segment header
corrupt_option	SEGMENT_MARK_CORRUPT (default) or SEGMENT_MARK_VALID

Usage Noes

You can determine the relative file number and block number (header_relative_file and header_block parameter) of the segment header block by querying DBA_SEGMENTS.

Examples

The following example marks the segment as corrupt:

```
EXECUTE DBMS_SPACE_ADMIN.SEGMENT_CORRUPT('USERS', 4, 33,
DBMS_SPACE_ADMIN.SEGMENT_MARK_CORRUPT);
```

Alternately, the next example marks a corrupt segment valid:

```
EXECUTE DBMS_SPACE_ADMIN.SEGMENT_CORRUPT('USERS', 4, 33,
DBMS_SPACE_ADMIN.SEGMENT_MARK_VALID);
```

179.4.7 SEGMENT_DROP_CORRUPT Procedure

This procedure drops a segment currently marked corrupt (without reclaiming space).

For this to work, the segment must be marked *temporary*. To mark a corrupt segment as temporary, issue a DROP command on the segment.

Syntax

```
DBMS_SPACE_ADMIN.SEGMENT_DROP_CORRUPT (
   tablespace_name         IN    VARCHAR2,
   header_relative_file    IN    POSITIVE,
   header_block            IN    NUMBER);
```

Parameters

Table 179-9 SEGMENT_DROP_CORRUPT Procedure Parameters

Parameter	Description
tablespace_name	Name of tablespace in which segment resides
header_relative_file	Relative file number of segment header
header_block	Block number of segment header

Usage Notes

- The space for the segment is not released, and it must be fixed by using the TABLESPACE_FIX_BITMAPS Procedure or the TABLESPACE_REBUILD_BITMAPS Procedure.
- The procedure cannot be used on the `SYSTEM` tablespace.
- You can determine the relative file number and block number (`header_relative_file` and `header_block` parameter) of the segment header block by querying `DBA_SEGMENTS`.

Examples

```
EXECUTE DBMS_SPACE_ADMIN.SEGMENT_DROP_CORRUPT('USERS', 4, 33);
```

179.4.8 SEGMENT_DUMP Procedure

This procedure dumps the segment header and bitmap blocks of a specific segment to the location specified in the `USER_DUMP_DEST` initialization parameter.

Syntax

```
DBMS_SPACE_ADMIN.SEGMENT_DUMP (
   tablespace_name         IN    VARCHAR2,
   header_relative_file    IN    POSITIVE,
   header_block            IN    NUMBER,
   dump_option             IN    POSITIVE  DEFAULT SEGMENT_DUMP_EXTENT_MAP);
```

Parameters

Table 179-10 SEGMENT_DUMP Procedure Parameters

Parameter	Description
`tablespace_name`	Name of tablespace in which segment resides
`header_relative_file`	Relative file number of segment header
`header_block`	Block number of segment header
`dump_option`	One of the following options: • `SEGMENT_DUMP_EXTENT_MAP` • `SEGMENT_DUMP_BITMAP_SUMMARY`

Usage Notes

- You can produce a slightly abbreviated dump, which includes the segment header and bitmap block summaries, without percent-free states of each block if you pass `SEGMENT_DUMP_BITMAP_SUMMARY` as the `dump_option` parameter.
- You can determine the relative file number and block number (`header_relative_file` and `header_block` parameter) of the segment header block by querying `DBA_SEGMENTS.HEADER_FILE`. If `HEADER_FILE` is greater than 1023 then use `DBA_DATA_FILES.RELATIVE_FNO`.

Examples

```
EXECUTE DBMS_SPACE_ADMIN.SEGMENT_DUMP('USERS', 4, 33);
```

179.4.9 SEGMENT_VERIFY Procedure

This procedure checks the consistency of the segment extent map with the tablespace file bitmaps.

Syntax

```
DBMS_SPACE_ADMIN.SEGMENT_VERIFY (
   tablespace_name         IN    VARCHAR2,
   header_relative_file    IN    POSITIVE,
   header_block            IN    NUMBER,
   verify_option           IN    POSITIVE  DEFAULT SEGMENT_VERIFY_EXTENTS);
```

Parameters

Table 179-11 SEGMENT_VERIFY Procedure Parameters

Parameters	Description
tablespace_name	Name of tablespace in which segment resides
header_relative_file	Relative file number of segment header
header_block	Block number of segment header
verify_option	What kind of check to do: SEGMENT_VERIFY_EXTENTS or SEGMENT_VERIFY_EXTENTS_GLOBAL

Usage Notes

- Anomalies are output as block range, bitmap-block, bitmap-block-range, anomaly-information, in the trace file for all block ranges found to have incorrect space representation. The kinds of problems which would be reported are free space not considered free, used space considered free, and the same space considered used by multiple segments.
- You can determine the relative file number and block number (header_relative_file and header_block parameter) of the segment header block by querying DBA_SEGMENTS.

Examples

The following example verifies that the segment with segment header at relative file number 4, block number 33, has its extent maps and bitmaps synchronized.

```
EXECUTE DBMS_SPACE_ADMIN.SEGMENT_VERIFY('USERS', 4, 33,
DBMS_SPACE_ADMIN.SEGMENT_VERIFY_EXTENTS);
```

179.4.10 SET_SEGADV_ATTRIB Procedure

This procedure sets the values of attributes of DBMS_SPACE_ADMIN package.

Syntax

```
DBMS_SPACE_ADMIN.SET_SEGADV_ATTRIB(
   attribute  IN  NUMBER,
   value      IN  NUMBER);
```

Parameters

Table 179-12 SET_SEGADV_ATTRIB Procedure Parameters

Parameter	Description
attribute	Supported attributes: • COMP_ADVISOR — Provides an option to enable or disable Compression Advisor for Automatic Segment Advisor. By default Compression Advisor is enabled for Automatic Segment Advisor. • COMP_LOB — Provides an option to enable or disable Compression Advisor for the tables with LOB columns while Automatic Segment Advisor is running. By default Compression Advisor is enabled for tables with LOB columns.
value	Supported values: • ATTR_ENABLE : 1 • ATTR_DISABLE : 0

179.4.11 TABLESPACE_FIX_BITMAPS Procedure

This procedure marks the appropriate block range (extent) as free or used in bitmap. It cannot be used on the SYSTEM tablespace.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_FIX_BITMAPS (
   tablespace_name         IN    VARCHAR2,
   dbarange_relative_file  IN    POSITIVE,
   dbarange_begin_block    IN    POSITIVE,
   dbarange_end_block      IN    POSITIVE,
   fix_option              IN    POSITIVE);
```

Parameters

Table 179-13 TABLESPACE_FIX_BITMAPS Procedure Parameters

Parameter	Description
tablespace_name	Name of tablespace
dbarange_relative_file	Relative file number of block range (extent)
dbarange_begin_block	Block number of beginning of extent
dbarange_end_block	Block number (inclusive) of end of extent
fix_option	One of the following options: • TABLESPACE_EXTENT_MAKE_FREE • TABLESPACE_EXTENT_MAKE_USED

Examples

The following example marks bits for 51 blocks for relative file number 4, beginning at block number 33 and ending at 83, as USED in bitmaps.

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_FIX_BITMAPS('USERS', 4, 33, 83,
DBMS_SPACE_ADMIN.EXTENT_MAKE_USED);
```

Alternatively, specifying an option of TABLESPACE_EXTENT_MAKE_FREE marks the bits free in bitmaps. The BEGIN and END blocks must be in extent boundary and be extent multiple; otherwise, an error is raised.

179.4.12 TABLESPACE_FIX_SEGMENT_STATES Procedure

This procedure fixes the state of the segments in a tablespace in which migration was aborted.

During tablespace migration to or from local, the segments are put in a transient state. If migration is aborted, then the segment states are corrected by SMON when event 10906 is set. A database with segments in such a transient state cannot be downgraded. The procedure can be used to fix the state of such segments.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_FIX_SEGMENT_STATES (
   tablespace_name     IN    VARCHAR);
```

Parameters

Table 179-14 TABLESPACE_FIX_SEGMENT_STATES Procedure Parameters

Parameter Name	Description
tablespace_name	Name of the tablespace whose segments must be fixed

Usage Notes

The tablespace must be kept online and read/write when this procedure is called.

Examples

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_FIX_SEGMENT_STATES('TS1')
```

179.4.13 TABLESPACE_MIGRATE_FROM_LOCAL Procedure

This procedure migrates a locally managed tablespace to a dictionary-managed tablespace.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_FROM_LOCAL (
   tablespace_name         IN    VARCHAR2);
```

Parameter

Table 179-15 TABLESPACE_MIGRATE_FROM_LOCAL Procedure Parameter

Parameter	Description
tablespace_name	Name of tablespace

Usage Notes

The tablespace must be kept online and read/write during migration. Migration of temporary tablespaces and migration of SYSTEM tablespaces are not supported.

Examples

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_FROM_LOCAL('USERS');
```

179.4.14 TABLESPACE_MIGRATE_TO_LOCAL Procedure

This procedure migrates the tablespace from a dictionary-managed format to a locally managed format. Tablespaces migrated to locally managed format are user managed.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_TO_LOCAL (
   tablespace_name     IN     VARCHAR2,
   unit_size           IN     POSITIVE DEFAULT NULL,
   rfno                IN     POSITIVE DEFAULT NULL);
```

Parameters

Table 179-16 TABLESPACE_MIGRATE_TO_LOCAL Procedure Parameters

Parameter Name	Description
`tablespace_name`	Name of the tablespace to be migrated
`unit_size`	Bitmap unit size (which is the size of the smallest possible chunk of space that can be allocated) in the tablespace specified in number of blocks
`rfno`	Relative File Number of the file where the bitmap blocks are placed

Usage Notes

- Before you migrate the `SYSTEM` tablespace, migrate any dictionary-managed tablespaces that you want to use in read/write mode to locally managed. After the `SYSTEM` tablespace is migrated, you cannot change dictionary-managed tablespaces to read/write.

  > **See Also:**
  >
  > *Oracle Database Administrator's Guide*

- The tablespace must be kept online and read/write during migration. Note that temporary tablespaces cannot be migrated.
- Allocation Unit may be specified optionally. The default is calculated by the system based on the highest common divisor of all extents (used or free) for the tablespace. This number is further trimmed based on the `MINIMUM EXTENT` for the tablespace (5 if `MINIMUM EXTENT` is not specified). Thus, the calculated value will not be larger than the `MINIMUM EXTENT` for the tablespace. The last free extent in every file is ignored for GCD calculation. If you specify the unit size, then it must be a factor of the `unit_size` calculated by the system; otherwise an error message is returned.
- The Relative File Number parameter is used to place the bitmaps in a desired file. If space is not found in the file, then an error is issued. The data file specified must be part of the tablespace being migrated. If the dataflow is not specified, then the system chooses a dataflow in which to place the initial bitmap blocks. If space is not found for the initial bitmaps, then an error is raised.

Examples

To migrate a tablespace 'TS1' in 2KB blocksize with minimum extent size 1MB:

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_TO_LOCAL('TS1', 512, 2);
```

The bitmaps are placed in file with relative file number 2.

179.4.15 TABLESPACE_REBUILD_BITMAPS Procedure

This procedure rebuilds the appropriate bitmaps. If no bitmap block is specified, then it rebuilds all bitmaps for the given tablespace.

The procedure cannot be used on the `SYSTEM` tablespace.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_REBUILD_BITMAPS (
   tablespace_name         IN    VARCHAR2,
   bitmap_relative_file    IN    POSITIVE   DEFAULT NULL,
   bitmap_block            IN    POSITIVE   DEFAULT NULL);
```

Parameters

Table 179-17 TABLESPACE_REBUILD_BITMAPS Procedure Parameters

Parameter	Description
`tablespace_name`	Name of tablespace
`bitmap_relative_file`	Relative file number of bitmap block to rebuild
`bitmap_block`	Block number of bitmap block to rebuild

Usage Notes

Only full rebuild is supported.

Examples

The following example rebuilds bitmaps for all the files in the `USERS` tablespace.

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_REBUILD_BITMAPS('USERS');
```

179.4.16 TABLESPACE_REBUILD_QUOTAS Procedure

This procedure rebuilds quotas for the given tablespace.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_REBUILD_QUOTAS (
   tablespace_name         IN    VARCHAR2);
```

Parameters

Table 179-18 TABLESPACE_REBUILD_QUOTAS Procedure Parameters

Parameter	Description
tablespace_name	Name of tablespace

Examples

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_REBUILD_QUOTAS('USERS');
```

179.4.17 TABLESPACE_RELOCATE_BITMAPS Procedure

This procedure relocates the bitmaps to the destination specified.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_RELOCATE_BITMAPS (
   tablespace_name    IN      VARCHAR2,
   filno              IN      POSITIVE,
   blkno              IN      POSITIVE);
```

Parameters

Table 179-19 TABLESPACE_RELOCATE_BITMAPS Procedure Parameters

Parameter Name	Description
tablespace_name	Name of tablespace
filno	Relative File Number of the destination file
blkno	Block Number of the destination range

Usage Notes

- Migration of a tablespace from dictionary-managed to locally managed format could result in the creation of SPACE HEADER segment that contains the bitmap blocks. The SPACE HEADER segment is treated as user data. If you explicitly resize a file at or below the space header segment, then an error is issued. Use the TABLESPACE_RELOCATE_BITMAPS command to move the control information to a different destination and then resize the file.
- This procedure cannot be used on the SYSTEM tablespace.
- The tablespace must be kept online and read/write during relocation of bitmaps. This can be done only on migrated locally managed tablespaces.

Examples

```
EXECUTE  DBMS_SPACE_ADMIN.TABLESPACE_RELOCATE_BITMAPS('TS1', 3, 4);
```

Moves the bitmaps to file 3, block 4.

> **Note:**
>
> The source and the destination addresses must not overlap. The destination block number is rounded down to the unit boundary. If there is user data in that location, then an error is raised.

179.4.18 TABLESPACE_VERIFY Procedure

This procedure verifies that the bitmaps and extent maps for the segments in the tablespace are synchronized.

Syntax

```
DBMS_SPACE_ADMIN.TABLESPACE_VERIFY (
   tablespace_name         IN    VARCHAR2,
   verify_option           IN    POSITIVE DEFAULT TABLESPACE_VERIFY_BITMAP);
```

Parameters

Table 179-20 TABLESPACE_VERIFY Procedure Parameters

Parameter	Description
`tablespace_name`	Name of tablespace
`verify_option`	One option is supported: `TABLESPACE_VERIFY_BITMAP`

Examples

```
EXECUTE DBMS_SPACE_ADMIN.TABLESPACE_VERIFY('USERS');
```

180
DBMS_SPD

The `DBMS_SPD` package provides subprograms for managing SQL plan directives (SPD).

This chapter contains the following topics:

- Overview
- Security Model
- Views
- Summary of DBMS_SPD Subprograms

> **See Also:**
>
> - DBMS_STATS
> - *Oracle Database SQL Tuning Guide* regarding SQL plan directives

180.1 DBMS_SPD Overview

This package provides subprograms for managing SQL plan directives (SPD).

SPD are objects generated automatically by Oracle. For example, if Oracle detects that the single table cardinality estimated made by the optimizer is different from the actual number of rows returned when accessing the table, it will automatically create a directive to perform dynamic statistics for the table. When any SQL statement referencing the table is compiled, the optimizer will perform dynamic statistics for the table to get a more accurate estimate.

180.2 DBMS_SPD Security Model

`DBMS_SPD` is an invoker-rights package. The invoker requires `ADMINISTER SQL MANAGEMENT OBJECT` privilege for executing most of the subprograms in this package. Also, the subprograms commit the current transaction (if any), perform the operation, and then commit it again.

180.3 DBMS_SPD Views

The DBA view `DBA_SQL_PLAN_DIRECTIVES` shows all the directives created in the system and the view `DBA_SQL_PLAN_DIR_OBJECTS` displays the objects that are included in the directives.

180.4 Summary of DBMS_SPD Subprograms

This table lists and briefly describes the `DBMS_SPD` package subprograms.

Table 180-1 DBMS_SPD Package Subprograms

Subprogram	Description
ALTER_SQL_PLAN_DIRECTIVE Procedure	Changes different attributes of a SQL plan directive
CREATE_STGTAB_DIRECTIVE Procedure	Creates a staging table into which to pack (export) SQL plan directives.
DROP_SQL_PLAN_DIRECTIVE Procedure	Drops a SQL plan directive
FLUSH_SQL_PLAN_DIRECTIVE Procedure	Allows for manual flushing of the SQL plan directives that are automatically recorded in SGA memory while executing SQL statements.
GET_PREFS Function	Gets the values for preferences for SQL plan directives
PACK_STGTAB_DIRECTIVE Function	Packs (exports) SQL plan directives into a staging table.
SET_PREFS Procedure	Allows the setting of different preferences for SQL plan directives
UNPACK_STGTAB_DIRECTIVE Function	Unpacks (imports) SQL plan directives from a staging table.

180.4.1 ALTER_SQL_PLAN_DIRECTIVE Procedure

This procedure changes different attributes of a SQL plan directive.

Syntax

```
DBMS_SPD.ALTER_SQL_PLAN_DIRECTIVE (
   directive_id        IN     NUMBER,
   attribute_name      IN     VARCHAR2,
   attribute_value     IN     VARCHAR2);
```

Parameters

Table 180-2 *ALTER_SQL_PLAN_DIRECTIVE Procedure Parameters*

Parameter	Description
directive_id	SQL plan directive ID
attribute_name	• ENABLED • AUTO_DROP
attribute_value	Possible values: • ENABLED: - If YES directive is enabled and may be used - If NO directive is not enabled and will not be used • AUTO_DROP: - If YES directive will be dropped automatically if not used for SPD_RETENTION_WEEKS. This is the default behavior. - If NO directive will not be dropped automatically

Exceptions

- ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.

- ORA-28104 INVALID_INPUT: The input value is not valid.
- ORA-13158 OBJECT_DOES_NOT_EXIST: The specified object does not exist.

Usage Notes

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.

Examples

```
BEGIN
  DBMS_SPD.ALTER_SQL_PLAN_DIRECTIVE (12345, 'STATE', 'PERMANENT');
END;
```

180.4.2 CREATE_STGTAB_DIRECTIVE Procedure

This procedure creates a staging table into which to pack (export) SQL plan directives.

Syntax

```
DBMS_SPD.CREATE_STGTAB_DIRECTIVE (
   table_name         IN VARCHAR2,
   table_owner        IN VARCHAR2  := USER,
   tablespace_name    IN VARCHAR2  := NULL);
```

Parameters

Table 180-3 *CREATE_STGTAB_DIRECTIVE Procedure Parameters*

Parameter	Description
table_name	Name of staging table
table_owner	Name of schema owner of staging table. Default is current schema.
tablespace_name	Name of tablespace. Default NULL means create staging table in the default tablespace:

Exceptions

- ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.
- ORA-28104 INVALID_INPUT: The input value is not valid.
- ORA-44001 INVALID_SCHEMA: The input schema does not exist.
- ORA-13159 TABLE_ALREADY_EXISTS: The specified table already exists.
- ORA-29304 TABLESPACE_MISSING: The specified tablespace does not exist.

Usage Notes

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.

180.4.3 DROP_SQL_PLAN_DIRECTIVE Procedure

This procedure drops a SQL plan directive.

Syntax

```
DBMS_SPD.DROP_SQL_PLAN_DIRECTIVE (
   directive_id        IN     NUMBER);
```

Parameters

Table 180-4 *DROP_SQL_PLAN_DIRECTIVE Procedure Parameters*

Parameter	Description
directive_id	SQL plan directive ID

Exceptions

- ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.
- ORA-28104 INVALID_INPUT: The input value is not valid.
- ORA-13158 OBJECT_DOES_NOT_EXIST: The specified object does not exist.

Usage Notes

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.

Examples

```
BEGIN
  DBMS_SPD.DROP_SQL_PLAN_DIRECTIVE (12345);
END;
```

180.4.4 FLUSH_SQL_PLAN_DIRECTIVE Procedure

This procedure allows for manual flushing of the SQL plan directives that are automatically recorded in SGA memory while executing SQL statements.

The information recorded in the SGA is periodically flushed by an Oracle background process. This procedure provides a way to flush the information manually.

Syntax

```
DBMS_SPD.FLUSH_SQL_PLAN_DIRECTIVE;
```

Exceptions

ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.

Usage Notes

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.

Examples

```
BEGIN
  DBMS_SPD.FLUSH_SQL_PLAN_DIRECTIVE;
END;
```

180.4.5 GET_PREFS Function

This function returns the value for the specified preferences for SQL plan directives.

Syntax

```
DBMS_SPD.GET_PREFS (
   pname     IN   VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 180-5 *GET_PREFS Function Parameters*

Parameter	Description
pname	Preference name. The procedure supports the preference SPD_RETENTION_WEEKS.

Return Values

Preference value

Exceptions

- ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.
- ORA-28104 INVALID_INPUT: The input value is not valid.

Usage Notes

- The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.
- SPD_RETENTION_WEEKS - SQL plan directives are purged if not used for more than the value set for this preference.

Examples

```
SELECT DBMS_SPD.GET_PREFS('SPD_RETENTION_WEEKS') FROM DUAL;
```

180.4.6 PACK_STGTAB_DIRECTIVE Function

This function packs (exports) SQL plan directives into a staging table.

Syntax

```
DBMS_SPD.PACK_STGTAB_DIRECTIVE (
   table_name         IN VARCHAR2,
   table_owner        IN VARCHAR2  := USER
   directive_id       IN NUMBER    := NULL,
```

```
  obj_list           IN OBJECTTAB := NULL)
 RETURN NUMBER
```

Parameters

Table 180-6 *PACK_STGTAB_DIRECTIVE Function Parameters*

Parameter	Description
table_name	Name of staging table
table_owner	Name of schema owner of staging table. Default is current schema.
directive_id	SQL plan directive ID. Default NULL means all directives in the system.
obj_list	Used to filter the directives to be packed based on the objects used in directives. If obj_list is not NULL, a directive is packed only if all the objects in the directive exist in obj_list.

Return Values

Number of SQL plan directives packed.

Exceptions

- ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.
- ORA-28104 INVALID_INPUT: The input value is not valid.
- ORA-44001 INVALID_SCHEMA: The input schema does not exist.
- ORA-29304 INVALID_STGTAB: The specified staging table is invalid or does not exist.
- ORA-13158 OBJECT_DOES_NOT_EXIST: The specified object does not exist.

Usage Notes

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.

Examples

```
-- Pack all directives in the system
SELECT DBMS_SPD.PACK_STGTAB_DIRECTIVE('mydirtab') FROM DUAL;

SET SERVEROUTPUT ON;
-- Pack directives relevant to objects in SH schema
DECLARE
  my_list  DBMS_SPD.OBJECTTAB := DBMS_SPD.ObjectTab();
  dir_cnt  NUMBER;
BEGIN
  my_list.extend(1);
  my_list(1).owner := 'SH';           -- schema name
  my_list(1).object_name := NULL;     -- all tables in SH
  my_list(1).object_type := 'TABLE';  -- type of object

  dir_cnt :=
    DBMS_SPD.PACK_STGTAB_DIRECTIVE('mydirtab', obj_list => my_list);
  DBMS_OUTPUT.PUT_LINE('dir_cnt = ' || dir_cnt);
END;

-- Pack directives relevant to tables SALES and CUSTOMERS in SH schema
DECLARE
```

```
  my_list  DBMS_SPD.OBJECTTAB := DBMS_SPD.ObjectTab();
  dir_cnt  NUMBER;
BEGIN
  my_list.extend(2);

  -- SALES table
  my_list(1).owner := 'SH';
  my_list(1).object_name := 'SALES';
  my_list(1).object_type := 'TABLE';

  -- CUSTOMERS table
  my_list(2).owner := 'SH';
  my_list(2).object_name := 'CUSTOMERS';
  my_list(2).object_type := 'TABLE';

  dir_cnt :=
    DBMS_SPD.PACK_STGTAB_DIRECTIVE('mydirtab', obj_list => my_list);
  DBMS_OUTPUT.PUT_LINE('dir_cnt = ' || dir_cnt);
END;
```

180.4.7 SET_PREFS Procedure

This procedure allows the setting of different preferences for SQL plan directives.

Syntax

```
DBMS_SPD.SET_PREFS (
   pname     IN   VARCHAR2,
   pvalue    IN   VARCHAR2);
```

Parameters

Table 180-7 *SET_PREFS Procedure Parameters*

Parameter	Description
pname	Preference name. The procedure supports the preference SPD_RETENTION_WEEKS.
pvalue	Preference value. • SPD_RETENTION_WEEKS: SQL plan directives are purged if not used for more than the value set for this preference. Default is 53 (SPD_RETENTION_WEEKS_DEFAULT) weeks, which means a directive is purged if it has been left unused for little over a year. It can be set to any value greater than or equal to 0. Also value NULL can be passed to set the preference to default.

Exceptions

- ORA-38171 INSUFFICIENT_PRIVILEGE: The user does not have proper privilege to perform the operation.
- ORA-28104 INVALID_INPUT: The input value is not valid.

Usage Notes

- The ADMINISTER SQL MANAGEMENT OBJECT privilege is required to execute this procedure.
- SPD_RETENTION_WEEKS - SQL plan directives are purged if not used for more than the value set for this preference.

Examples

```
BEGIN
  DBMS_SPD.SET_PREFS('SPD_RETENTION_WEEKS', '4');
END;
```

180.4.8 UNPACK_STGTAB_DIRECTIVE Function

This procedure unpacks (imports) SQL plan directives from a staging table.

Syntax

```
DBMS_SPD.UNPACK_STGTAB_DIRECTIVE (
   table_name         IN VARCHAR2,
   table_owner        IN VARCHAR2  := USER,
   directive_id       IN NUMBER    := NULL,
   obj_list           IN OBJECTTAB := NULL)
 RETURN NUMBER
```

Parameters

Table 180-8 *UNPACK_STGTAB_DIRECTIVE Function Parameters*

Parameter	Description
`table_name`	Name of staging table
`table_owner`	Name of schema owner of staging table. Default is current schema.
`directive_id`	SQL plan directive ID. Default `NULL` means all directives in the system.
`obj_list`	Used to filter the directives to be unpacked based on the objects used in directives. If `obj_list` is not `NULL`, a directive is unpacked only if all the objects in the directive exist in `obj_list`.

Return Values

Number of SQL plan directives unpacked.

Exceptions

- `ORA-38171 INSUFFICIENT_PRIVILEGE`: The user does not have proper privilege to perform the operation.
- `ORA-28104 INVALID_INPUT`: The input value is not valid.
- `ORA-44001 INVALID_SCHEMA`: The input schema does not exist.
- `ORA-29304 INVALID_STGTAB`: The specified staging table is invalid or does not exist.
- `ORA-13158 OBJECT_DOES_NOT_EXIST`: The specified object does not exist.

Usage Notes

The `ADMINISTER SQL MANAGEMENT OBJECT` privilege is required to execute this procedure.

Examples

```
-- Unack all directives in the staging table
SELECT DBMS_SPD.UNPACK_STGTAB_DIRECTIVE('mydirtab') FROM DUAL;
```

```
SET SERVEROUTPUT ON;
-- Unpack directives relevant to objects in SH schema
DECLARE
  my_list  DBMS_SPD.OBJECTTAB := DBMS_SPD.ObjectTab();
  dir_cnt  number;
BEGIN
  my_list.extend(1);
  my_list(1).owner := 'SH';           -- schema name
  my_list(1).object_name := null;     -- all tables in SH
  my_list(1).object_type := 'TABLE';  -- type of object

  dir_cnt :=
    DBMS_SPD.UNPACK_STGTAB_DIRECTIVE('mydirtab', obj_list => my_list);
  DBMS_OUTPUT.PUT_LINE('dir_cnt = ' || dir_cnt);
END;

-- Unpack directives relevant to tables SALES and CUSTOMERS in SH schema
DECLARE
  my_list  DBMS_SPD.OBJECTTAB := DBMS_SPD.ObjectTab();
  dir_cnt  NUMBER;
begin
  my_list.extend(2);

  -- SALES table
  my_list(1).owner := 'SH';
  my_list(1).object_name := 'SALES';
  my_list(1).object_type := 'TABLE';

  -- CUSTOMERS table
  my_list(2).owner := 'SH';
  my_list(2).object_name := 'CUSTOMERS';
  my_list(2).object_type := 'TABLE';

  dir_cnt :=
    DBMS_SPD.UNPACK_STGTAB_DIRECTIVE('mydirtab', obj_list => my_list);
  DBMS_OUTPUT.PUT_LINE('dir_cnt = ' || dir_cnt);
END;
```

181
DBMS_SPM

The `DBMS_SPM` package supports the SQL plan management feature by providing an interface for the DBA or other user to perform controlled manipulation of plan history and SQL plan baselines maintained for various SQL statements.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Examples
- Data Structures
- Summary of DBMS_SPM Subprograms

> **See Also:**
>
> For more information about "Using SQL Plan Management" in the *Oracle Database SQL Tuning Guide*

181.1 DBMS_SPM Overview

The `DBMS_SPM` package allows the user to manage SQL execution plans using SQL plan management.

SQL plan management prevents performance regressions resulting from sudden changes to the execution plan of a SQL statement by recording and evaluating the execution plans of SQL statements over time, and builds SQL plan baselines composed of a set of existing plans known to be efficient. The SQL plan baselines are then used to preserve performance of corresponding SQL statements, regardless of changes occurring in the system. Common usage scenarios where SQL plan management can improve or preserve SQL performance include:

- A database upgrade that installs a new optimizer version usually results in plan changes for a small percentage of SQL statements, with most of the plan changes resulting in either no performance change or improvement. However, certain plan changes may cause performance regressions. The use of SQL plan baselines significantly minimizes potential performance regressions resulting from a database upgrade.
- Ongoing system and data changes can impact plans for some SQL statements, potentially causing performance regressions. The use of SQL plan baselines helps to minimize performance regressions and stabilize SQL performance.
- Deployment of new application modules means introducing new SQL statements into the system. The application software may use appropriate SQL execution plans developed under a standard test configuration for the new SQL statements. If the system production

configuration differs significantly from the test configuration, SQL plan baselines can be evolved over time to produce better performance.

181.2 DBMS_SPM Security Model

The package is owned by `SYS`. The `EXECUTE` package privilege is required to execute its procedures. Any user granted the `ADMINISTER SQL MANAGEMENT OBJECT` privilege is able to execute the `DBMS_SPM` package.

181.3 DBMS_SPM Constants

The `DBMS_SPM` package provides constants that can be used for specifying parameter values.

These are shown in the following table. These constants are defined as standard input for the `time_limit` parameter of the EVOLVE_SQL_PLAN_BASELINE Function.

Table 181-1 DBMS_SPM Constants

Constant	Type	Value	Description
AUTO_LIMIT	INTEGER	2147483647	Oracle determines the appropriate time spent by the EVOLVE_SQL_PLAN_BASELINE Function.
NO_LIMIT	INTEGER	2147483647 -1	There is no limit to the time spent by the EVOLVE_SQL_PLAN_BASELINE Function.

181.4 DBMS_SPM Examples

These examples will help you understand use of `DBMS_SPM`.

Detailed examples are located under the following topics:

- Migrating Stored Outlines to SQL Plan Baselines
- Migrating Outlines to Utilize SQL Plan Management Features
- Migrating Outlines to Preserve Stored Outline Behavior
- Performing Follow-Up Tasks After Stored Outline Migration

181.5 DBMS_SPM Data Structures

The `DBMS_SPM` package defines a `TABLE` type.

Table Types

- DBMS_SPM NAMELIST Table Type

181.5.1 DBMS_SPM NAMELIST Table Type

This type allows for a list of names as an input parameter.

Syntax

```
TYPE  name_list  IS TABLE OF VARCHAR2(30);
```

181.6 Summary of DBMS_SPM Subprograms

This table lists and briefly describes the `DBMS_SPM` package subprograms.

Table 181-2 DBMS_SPM Package Subprograms

Subprogram	Description
ACCEPT_SQL_PLAN_BASELINE Procedure	Accepts a plan based on the recommendation of an evolve task
ADD_VERIFIED_SQL_PLAN_BASELINE Function	Finds the plans available for the given SQL_ID in different sources such Cursor Cache, Auto SQL Tuning Set, and Automatic Workload Repository.
ALTER_SQL_PLAN_BASELINE Function	Changes an attribute of a single plan or all plans associated with a SQL statement using the attribute name/value format
CANCEL_EVOLVE_TASK Procedure	Cancels a currently executing evolve task
CONFIGURE Procedure	Sets configuration options for SQL management base, in parameter/value format
CREATE_EVOLVE_TASK Function	Creates an advisor task and sets its parameters
CREATE_STGTAB_BASELINE Procedure	Creates a staging table that used for transporting SQL plan baselines from one system to another
DROP_EVOLVE_TASK Procedure	Drops an evolved task
DROP_SQL_PLAN_BASELINE Function	Drops a single plan, or all plans associated with a SQL statement
EVOLVE_SQL_PLAN_BASELINE Function	Evolves SQL plan baselines associated with one or more SQL statements
EXECUTE_EVOLVE_TASK Function	Executes a previously created evolve task
IMPLEMENT_EVOLVE_TASK Function	Implements a plan based on the recommendation of an evolve task
INTERRUPT_EVOLVE_TASK Procedure	Interrupts a currently executing evolve task
LOAD_PLANS_FROM_CURSOR_CACHE Functions	Loads one or more plans present in the cursor cache for a SQL statement
LOAD_PLANS_FROM_AWR Function	Loads the SQL Management Base (SMB) with SQL plan baselines for a set of SQL statements using the plans from the AWR, and returns the number of plans loaded
LOAD_PLANS_FROM_SQLSET Function	Loads plans stored in a SQL tuning set (STS) into SQL plan baselines

Table 181-2 (Cont.) DBMS_SPM Package Subprograms

Subprogram	Description
MIGRATE_STORED_OUTLINE Functions	Migrates existing stored outlines to SQL plan baselines
PACK_STGTAB_BASELINE Function	Packs (exports) SQL plan baselines from SQL management base into a staging table
RESET_EVOLVE_TASK Procedure	Resets an evolve task to its initial state
RESUME_EVOLVE_TASK Procedure	Resumes a previously interrupted task
REPORT_AUTO_EVOLVE_TASK Function	Displays the results of an execution of an automatic evolve task.
REPORT_EVOLVE_TASK Function	Displays the results of an evolved task
SET_EVOLVE_TASK_PARAMETER Procedure	Sets a parameter of an evolve task
UNPACK_STGTAB_BASELINE Function	Unpacks (imports) SQL plan baselines from a staging table into SQL management base

181.6.1 ACCEPT_SQL_PLAN_BASELINE Procedure

The procedure accepts a plan based on the recommendation of an evolve task.

Syntax

```
DBMS_SPM.ACCEPT_SQL_PLAN_BASELINE  (
   task_name        IN  VARCHAR2,
   object_id        IN  NUMBER    := NULL,
   task_owner       IN  VARCHAR2  := NULL,
   force            IN  BOOLEAN   := FALSE);
```

Parameters

Table 181-3 ACCEPT_SQL_PLAN_BASELINE Procedure Parameters

Parameter	Description
task_name	Identifier of task to implement
object_id	Identifier of the advisor framework object that represents a single plan. If NULL, the report is generated for all objects.
task_owner	Owner of the evolve task. Defaults to the current schema owner.
force	Accept the plan even if the advisor did not recommend such an action. The default is FALSE requiring acceptance of the plan only if the plan is verified and shows sufficient improvement in benefit.

181.6.2 ADD_VERIFIED_SQL_PLAN_BASELINE Function

This function finds the plans available for the given SQL_ID in different sources such Cursor Cache, Auto SQL Tuning Set, and Automatic Workload Repository.

Syntax

```
DBMS_SPM.ADD_VERIFIED_SQL_PLAN_BASELINE (
    sql_id IN VARCHAR2
);
```

Parameters

Table 181-4 ADD_VERIFIED_SQL_PLAN_BASELINE Function Parameters

Parameter	Description
sql_id	The SQL statement identifier, which is used to identify the plans in different sources, such as the cursor cache.

Return Value

A CLOB containing a report of SQL plan baselines verified, reproduced and accepted.

Usage Notes

SQL plan baselines are created to ensure that the best-known plans are used for the selected SQL statement. These are the steps that the function executes:

1. Loads plans from Cursor Cache, Automatic Workload Repository, and Auto SQL Tuning Set (SYS_AUTO_STS) into SQL Plan Management SQL plan history in a non-accepted state.
2. Uses the SQL Plan Management Evolve Advisor internally to identify the best-performing execution plans. The best-performing plans are accepted.

When the function has completed, it will have accepted the best plans among the plans available from the Cursor Cache, AWR, and the Automatic SQL Tuning Set.

> **Note:**
>
> "Administer SQL Management Object" privilege. is required.

Examples

Example 1: 36k32wnz0v0fd923079310

```
select sql_id, plan_hash_value, sql_text
from   v$sql
where  sql_text like 'select /* SPM_TEST_QUERY%'
/

SQL_ID          PLAN_HASH_VALUE  SQL_TEXT
--------------- ---------------
--------------------------------------------------------------------
```

```
36k32wnz0v0fd   923079310        select /* SPM_TEST_QUERY */ num from
example_spm_table where id = 100
```

Example 2:9230793102448381833

```
select sql_text, sqlset_name, plan_hash_value
from   dba_sqlset_statements
where  sql_text like 'select /* SPM_TEST_QUERY%';

SQL_TEXT
SQLSET_NAME   PLAN_HASH_VALUE
--------------------------------------------------------------------------------
------------ ---------------
select /* SPM_TEST_QUERY */ num from example_spm_table where id = 100
SYS_AUTO_STS   923079310
select /* SPM_TEST_QUERY */ num from example_spm_table where id = 100
SYS_AUTO_STS   2448381833
```

Example 3:

```
set tab off
set serveroutput on
set pagesize 100
set linesize 250
set  long 100000
column report format a200

var rep clob;

BEGIN
   :rep := DBMS_SPM. ADD_VERIFIED_SQL_PLAN_BASELINE('36k32wnz0v0fd');
END;
/

select :rep report from dual;

SQL Plan Baselines verified for SQL ID: 36k32wnz0v0fd
--------------------------------------------------------------------------------
------------

Plan Hash Value   Plan Name                       Reproduced   Accepted
Source
---------------   -----------------------------   ----------   --------
--------------
923079310         SQL_PLAN_163tr5qgzwmgt05ce4c2e   YES          YES
CURSOR CACHE
2448381833        SQL_PLAN_163tr5qgzwmgt1f191f3e   YES          NO
SQL TUNING SET
--------------------------------------------------------------------------------
-----------

SQL Handle    : SQL_130f372d9ffe4df9
SQL Text      : select /* SPM_TEST_QUERY */ num from example_spm_table where
id = 100
```

```
--------------------------------------------------------------------------------
-----------
```

Example 4:

```
select sql_text, accepted, enabled, sql_handle, plan_name
from   dba_sql_plan_baselines
where  sql_text like 'select /* SPM_TEST_QUERY%';

SQL_TEXT                                                                   ACC
ENA SQL_HANDLE           PLAN_NAME
-------------------------------------------------------------------------- ---
--- -------------------- ------------------------------
select /* SPM_TEST_QUERY */ num from example_spm_table where id = 100  YES
YES SQL_130f372d9ffe4df9 SQL_PLAN_163tr5qgzwmgt05ce4c2e
select /* SPM_TEST_QUERY */ num from example_spm_table where id = 100  NO
YES SQL_130f372d9ffe4df9 SQL_PLAN_163tr5qgzwmgt1f191f3e
```

Example 5

```
select * from
table(dbms_xplan.display_sql_plan_baseline('SQL_130f372d9ffe4df9'));

--------------------------------------------------------------------------------
--
SQL handle: SQL_130f372d9ffe4df9
SQL text: select /* SPM_TEST_QUERY */ num from example_spm_table where id =
100
--------------------------------------------------------------------------------
--

--------------------------------------------------------------------------------
--
Plan name: SQL_PLAN_163tr5qgzwmgt05ce4c2e         Plan id: 923079310
Enabled: YES     Fixed: NO      Accepted: YES     Origin: EVOLVE-LOAD-FROM-
CURSOR-CACHE
Plan rows: From Auto SQL Tuning Set
--------------------------------------------------------------------------------
--

Plan hash value: 923079310

--------------------------------------------------------------------------------
------------------
| Id  | Operation                   | Name              | Rows  | Bytes |
Cost (%CPU)| Time     |
--------------------------------------------------------------------------------
------------------
|   0 | SELECT STATEMENT            |                   |       |       |
2 (100)|          |
|   1 |  TABLE ACCESS BY INDEX ROWID| EXAMPLE_SPM_TABLE |     1 |     8 |
2   (0)| 00:00:01 |
|*  2 |   INDEX UNIQUE SCAN         | SPM_TAB_PK        |     1 |       |
1   (0)| 00:00:01 |
```

```
-------------------------------------------------------------------------------
------------------

Predicate Information (identified by operation id):
---------------------------------------------------

   2 - access("ID"=100)

-------------------------------------------------------------------------------
--
Plan name: SQL_PLAN_163tr5qgzwmgt1f191f3e         Plan id: 2448381833
Enabled: YES     Fixed: NO      Accepted: NO      Origin: EVOLVE-LOAD-FROM-STS
Plan rows: From Auto SQL Tuning Set
-------------------------------------------------------------------------------
--

Plan hash value: 2448381833

-------------------------------------------------------------------------------
---------
| Id  | Operation         | Name              | Rows  | Bytes | Cost (%CPU)|
Time     |
-------------------------------------------------------------------------------
---------
|   0 | SELECT STATEMENT  |                   |       |       |     3
(100)|          |
|*  1 |  TABLE ACCESS FULL| EXAMPLE_SPM_TABLE |     1 |     8 |     3   (0)|
00:00:01 |
-------------------------------------------------------------------------------
---------

Predicate Information (identified by operation id):
---------------------------------------------------

   1 - filter("ID"=100)
```

181.6.3 ALTER_SQL_PLAN_BASELINE Function

This function changes an attribute of a single plan or all plans associated with a SQL statement using the attribute name/value format.

Syntax

```
DBMS_SPM.ALTER_SQL_PLAN_BASELINE (
   sql_handle        IN VARCHAR2 := NULL,
   plan_name         IN VARCHAR2 := NULL,
   attribute_name    IN VARCHAR2,
   attribute_value   IN VARCHAR2)
 RETURN PLS_INTEGER;
```

Parameters

Table 181-5 ALTER_SQL_PLAN_BASELINE Function Parameters

Parameter	Description
sql_handle	SQL statement handle. It identifies plans associated with a SQL statement for an attribute change. If NULL then plan_name must be specified.
plan_name	Name of a specific plan. Default NULL means set the attribute for all plans associated with a SQL statement identified by sql_handle. If NULL then sql_handle must be specified.
attribute_name	Name of the plan attribute to set (see table below).
attribute_value	Value of the plan attribute to use (see table below)

Table 181-6 Names & Values for ALTER_SQL_PLAN_BASELINE Function Parameters

Name	Description	Possible Values
enabled	'YES' means the plan is available for use by the optimizer. It may or may not be used, depending on accepted status.	'YES' or 'NO'
fixed	'YES' means the SQL plan baseline is not evolved over time. A fixed plan takes precedence over a non-fixed plan.	'YES' or 'NO'
autopurge	'YES' means the plan is purged if it is not used for a time period. 'NO' means it is never purged.	'YES' or 'NO'
plan_name	Name of the plan	String of up to 30 characters
description	Plan description.	String of up to 500 bytes

Return Values

The number of plans altered.

Usage Notes

When a single plan is specified, one of the various statuses, or plan name, or description can be altered. When all plans for a SQL statement are specified, one of various statuses, or description can be altered. This function can be called numerous times, each time setting a different plan attribute of same plan(s) or different plan(s).

181.6.4 CANCEL_EVOLVE_TASK Procedure

The procedure cancels a currently executing evolve task. All intermediate results are removed from the task.

Syntax

```
DBMS_SPM.CANCEL_EVOLVE_TASK  (
   task_name        IN  VARCHAR2);
```

Parameters

Table 181-7 CANCEL_EVOLVE_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of task to cancel

181.6.5 CONFIGURE Procedure

This procedure sets configuration options for the SQL management base and for the maintenance of SQL plan baselines. You can call this function multiple times, setting a different configuration option each time.

Syntax

```
DBMS_SPM.CONFIGURE (
   parameter_name    IN VARCHAR2,
   parameter_value   IN VARCHAR2 := NULL,
   allow             IN BOOLEAN := TRUE);
```

Parameters

Table 181-8 CONFIGURE Procedure Parameters

Parameter	Description
parameter_name	Name of parameter to set (see table below).
parameter_value	Value of parameter to use (see table below). The maximum length of parameter_value is 1000 characters.
allow	Whether to include (true) or exclude (false) matching SQL statements and plans for theauto_capture_* parameters. If null, then the procedure ignores the specified parameter.

Table 181-9 Names and Values for CONFIGURE Procedure Parameters

Parameter Name	Description	Possible Values
auto_capture_action	Action to include (=) or exclude (<>) for SQL plan management automatic capture, depending on whether allow is TRUE or FALSE. A null value removes the filter for parameter_name entirely.	Action name, for example, R%

Table 181-9 (Cont.) Names and Values for CONFIGURE Procedure Parameters

Parameter Name	Description	Possible Values
auto_capture_module	Module to include (=) or exclude (<>) for SQL plan management auto capture, depending on whether allow is TRUE or FALSE. A null value removes the filter for parameter_name entirely. The database only uses this filter when OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES is TRUE.	Module name, for example, LOGGER
auto_capture_parsing_schema_name	Parsing schema to include (=) or exclude (<>) for SQL plan management auto capture, depending on whether allow is TRUE or FALSE. A null value removes the filter for parameter_name entirely. The database only uses this filter when OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES is TRUE.	Schema name, for example, HR
auto_capture_sql_text	Search pattern to apply to SQL text of LIKE or NOT LIKE, depending on whether allow is TRUE or FALSE. A null value removes the filter for parameter_name entirely. The database only uses this filter when OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES is TRUE.	Text of a SQL statement, for example, SELECT a%
auto_spm_evolve_task	Setting to enable or disable the high-frequency SPM Evolve Advisor task. OFF -- disables background-verified SQL plan management and real-time SQL plan management. ON -- enables background-verified SQL plan management. AUTO -- enables real-time SQL plan management. **Note:** See About Automatic SQL Plan Management in the *Oracle Database SQL Tuning Guide*, which discusses the high frequency SPM Evolve task and explains the difference between real-time and background-verified Automatic SQL Plan Management.	ON, OFF, AUTO
plan_retention_weeks	Number of weeks to retain unused plans before the database purges them. A null value resets to the default value of 53 weeks, or 1 year plus 1 week. (This retains plans for annually executing queries.) The value of allow is ignored.	5–523 (default is 53)
space_budget_percent	Maximum percent of SYSAUX space that can be used for SQL management base. The database issues alerts when this amount is exceeded. A null value resets the percentage to the default value of 10%. The value of allow is ignored.	1–50 (default is 10)

Exceptions

Table 181-10 CONFIGURE Exceptions

Error Number	Description
ORA-38133	Invalid parameter name
ORA-38134	Invalid parameter value
ORA-38150	Not enough space for new filter
ORA-38151	Module name too long
ORA-38152	Action name too long
ORA-38304	Missing or invalid user name

Usage Notes

- When `parameter_name` is `auto_capture_sql_text`, the `parameter_value` is an automatic search filter. The filter uses the search pattern of `LIKE` *`parameter_name`* when `allow=>true`. The filter uses the pattern `NOT LIKE` *`parameter_name`* when `allow=>false`.

  For all other non-null `parameter_name` values, the search pattern depends on the `allow` setting. The parameter uses an equal sign (=) when `allow=>true`. The parameter uses a not-equal sign (<>) when `allow=>false`.
- You can configure multiple automatic capture parameters of different types. You cannot specify multiple values for the same parameter. Instead, the values specified for a particular parameter are combined. For example, specifying `auto_capture_sql_text` to be `'%TABLE1%'`, `TRUE`, and `'%TABLE2%'`, `FALSE` will result in matching SQL text `LIKE '%TABLE1%'` and `NOT LIKE '%TABLE2%'`. The database uses these configuration settings only when the initialization parameter `OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES` is set to `TRUE`.
- A null value for `parameter_value` removes the filter for `parameter_name` entirely. By using `paramter_value=>'%'` in combination with `allow=FALSE`, you can filter out all values for a parameter, and then create a separate filter to include only specified values. The `DBA_SQL_MANAGEMENT_CONFIG` view shows the current filters.
- The default space budget for SQL management base is no more than ten percent of the size of `SYSAUX` tablespace. The space budget can be set to a maximum of 50%. The default unused plan retention period is one year and one week, which means a plan will be automatically purged if it has not been used for more than a year. The retention period can be set to a maximum of 523 weeks (i.e. a little over 10 years).
- When the space occupied by SQL management base exceeds the defined space budget limit, a weekly database alert is generated.

Examples

The following example creates a filter for SQL text that is like `SELECT a%`:

```
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_SQL_TEXT', 'select a%', 'TRUE');
```

The following example filters out the `HR` parsing schema:

```
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_PARSING_SCHEMA_NAME', 'HR', 'FALSE');
```

The following example removes any existing filters for SQL text:

```
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_SQL_TEXT', NULL, NULL);
```

The following example removes any `LIKE` or `NOT LIKE` filters for the SQL text `select a%`:

```
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_SQL_TEXT', 'select a%', NULL);
```

The following example creates a filter with the predicate `(action LIKE 'R%') OR (action LIKE '%E_')`:

```
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_ACTION', 'R%', 'TRUE');
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_ACTION', '%E_', 'TRUE');
```

The following example creates a filter with the predicate `NOT(module LIKE 'LOGGER') AND NOT(module LIKE 'UTIL__')`:

```
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_MODULE', 'LOGGER', 'FALSE');
EXEC DBMS_SPM.CONFIGURE('AUTO_CAPTURE_MODULE', 'UTIL__', 'FALSE');
```

181.6.6 CREATE_EVOLVE_TASK Function

The function has two overloads, both of which create an advisor task and sets its parameters. This version which takes a SQL handle creates an evolve task in order to evolve one or more plans for a given SQL statement.

Syntax

```
DBMS_SPM.CREATE_EVOLVE_TASK (
   sql_handle     IN  VARCHAR2  := NULL,
   plan_name      IN  VARCHAR2  := NULL,
   time_limit     IN  NUMBER    := DBMS_SPM.AUTO_LIMIT,
   task_name      IN  VARCHAR2  := NULL,
   description    IN  VARCHAR2  := NULL)
  RETURN VARCHAR2;

DBMS_SPM.CREATE_EVOLVE_TASK (
   plan_list      IN  DBMS_SPM.NAME_LIST,
   time_limit     IN  NUMBER    := DBMS_SPM.AUTO_LIMIT,
   task_name      IN  VARCHAR2  := NULL,
   description    IN  VARCHAR2  := NULL)
  RETURN VARCHAR2;
```

Parameters

Table 181-11 CREATE_EVOLVE_TASK Function Parameters

Parameter	Description
`sql_handle`	Handle of a SQL statement. The default `NULL` considers all SQL statements with non-accepted plans.
`plan_list`	List of plan names. The plans may belong to different SQL statements.

Table 181-11 (Cont.) CREATE_EVOLVE_TASK Function Parameters

Parameter	Description
plan_name	Plan identifier. The default NULL considers all non-accepted plans of the specified SQL handle or all SQL statements if the SQL handle is NULL.
time_limit	Time limit in number of minutes. The time limit is global and it is used in the following manner. The time limit for first non-accepted plan is equal to the input value. The time limit for the second non-accepted plan is equal to (input value - time spent in first plan verification) and so on. The default DBMS_SPM.AUTO_LIMIT means let the system choose an appropriate time limit based on the number of plan verifications required to be done. The value DBMS_SPM.NO_LIMIT means no time limit.
task_name	Evolve task name
description	Description of the task (maximum 256 characters)

Return Values

SQL evolve task unique name

181.6.7 CREATE_STGTAB_BASELINE Procedure

This procedure creates a staging table used for transporting SQL plan baselines from one system to another.

Syntax

```
DBMS_SPM.CREATE_STGTAB_BASELINE (
   table_name        IN VARCHAR2,
   table_owner       IN VARCHAR2 := NULL,
   tablespace_name   IN VARCHAR2 := NULL);
```

Parameters

Table 181-12 CREATE_STGTAB_BASELINE Procedure Parameters

Parameter	Description
table_name	Name of staging table to create for the purpose of packing and unpacking SQL plan baselines
table_owner	Name of owner of the staging table. The default, NULL, means that the current schema is the table owner.
tablespace_name	Name of the tablespace. The default, NULL, results in creating the staging table in the default tablespace.

Usage Notes

The creation of staging table is the first step. To migrate SQL plan baselines from one system to another, the user or DBA has to perform a series of steps as follows:

1. Create a staging table in the source system

2. Select SQL plan baselines in the source system and pack them into the staging table
3. Copy the staging table from the source system to the target system. For example, you can use Oracle Data Pump to export and import the staging table.
4. Select SQL plan baselines from the staging table and unpack them into the target system

181.6.8 DROP_EVOLVE_TASK Procedure

The procedure drops an evolved task.

Syntax

```
DBMS_SPM.DROP_EVOLVE_TASK  (
   task_name        IN  VARCHAR2);
```

Parameters

Table 181-13 DROP_EVOLVE_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of the task that you want to drop

181.6.9 DROP_SQL_PLAN_BASELINE Function

This function drops a single plan, or all plans associated with a SQL statement.

Syntax

```
DBMS_SPM.DROP_SQL_PLAN_BASELINE (
   sql_handle     IN VARCHAR2 := NULL,
   plan_name      IN VARCHAR2 := NULL)
RETURN PLS_INTEGER;
```

Parameters

Table 181-14 DROP_SQL_PLAN_BASELINE Function Parameters

Parameter	Description
sql_handle	SQL statement handle. It identifies plans associated with a SQL statement that are to be dropped. If NULL then plan_name must be specified.
plan_name	Plan name. It identifies a specific plan. Default NULL means to drop all plans associated with the SQL statement identified by sql_handle.

Return Values

The number of plans dropped

181.6.10 EVOLVE_SQL_PLAN_BASELINE Function

This function evolves SQL plan baselines associated with one or more SQL statements. A SQL plan baseline is evolved when one or more of its non-accepted plans is changed to an accepted plan or plans.

If interrogated by the user (parameter `verify` = '`YES`'), the execution performance of each non-accepted plan is compared against the performance of a plan chosen from the associated SQL plan baseline. If the non-accepted plan performance is found to be better than SQL plan baseline performance, the non-accepted plan is changed to an accepted plan provided such action is permitted by the user (parameter `commit` = 'YES').

The second form of the function employs a plan list format.

Syntax

```
DBMS_SPM.EVOLVE_SQL_PLAN_BASELINE (
   sql_handle   IN VARCHAR2 := NULL,
   plan_name    IN VARCHAR2 := NULL,
   time_limit   IN INTEGER  := DBMS_SPM.AUTO_LIMIT,
   verify       IN VARCHAR2 := 'YES',
   commit       IN VARCHAR2 := 'YES')
  RETURN CLOB;

DBMS_SPM.EVOLVE_SQL_PLAN_BASELINE (
   plan_list    IN DBMS_SPM.NAME_LIST,
   time_limit   IN INTEGER  := DBMS_SPM.AUTO_LIMIT,
   verify       IN VARCHAR2 := 'YES',
   commit       IN VARCHAR2 := 'YES')
  RETURN CLOB;
```

Parameters

Table 181-15 EVOLVE_SQL_PLAN_BASELINE Function Parameters

Parameter	Description
`sql_handle`	SQL statement identifier. Unless `plan_name` is specified, `NULL` means to consider all statements with non-accepted plans in their SQL plan baselines.
`plan_name`	Plan identifier. Default `NULL` means to consider all non- accepted plans in the SQL plan baseline of either the identified SQL statement or all SQL statements if `sql_handle` is `NULL`.
`plan_list`	A list of plan names. Each plan in the list can belong to same or different SQL statement.
`time_limit`	Time limit in number of minutes. This applies only if `verify` = '`YES`'. The time limit is global and it is used as follows: The time limit for first non-accepted plan verification is set equal to the input value; the time limit for second non-accepted plan verification is set equal to (input value - time spent in first plan verification); and so on. • `DBMS_SPM.AUTO_LIMIT` (Default) lets the system choose an appropriate time limit based on the number of plan verifications required to be done. • `DBMS_SPM.NO_LIMIT` means there is no time limit. • A positive integer value represents a user specified time limit.

Table 181-15 (Cont.) EVOLVE_SQL_PLAN_BASELINE Function Parameters

Parameter	Description
verify	Specifies whether to execute the plans and compare the performance before changing non-accepted plans into accepted plans. A performance verification involves executing a non-accepted plan and a plan chosen from corresponding SQL plan baseline and comparing their performance statistics. If non-accepted plan shows performance improvement, it is changed to an accepted plan. • 'YES' (Default) - verifies that a non-accepted plan gives better performance before changing it to an accepted plan • 'NO' - directs not to execute plans but only to change non-accepted plans into accepted plans
commit	Specifies whether to update the ACCEPTED status of non-accepted plans from 'NO' to 'YES'. • 'YES' (Default) - perform updates of qualifying non-accepted plans and generate a report that shows the updates and the result of performance verification when verify = 'YES'. • 'NO' - generate a report without any updates. Note that commit = 'NO' together with verify = 'NO' represents a no-op.

Return Values

A CLOB containing a formatted text report showing non-accepted plans in sequence, each with a possible change of its ACCEPTED status, and if verify = 'YES' the result of their performance verification.

Usage Notes

Invoking this subprogram requires the ADMINISTER SQL MANAGEMENT OBJECT privilege.

181.6.11 EXECUTE_EVOLVE_TASK Function

The function executes a previously created evolve task.

Syntax

```
DBMS_SPM.EXECUTE_EVOLVE_TASK  (
   task_name        IN  VARCHAR2,
   execution_name   IN  VARCHAR2  := NULL,
   execution_desc   IN  VARCHAR2  := NULL);
 RETURN VARCHAR2;
```

Parameters

Table 181-16 EXECUTE_EVOLVE_TASK Function Parameters

Parameter	Description
task_name	Evolve task name
execution_name	Name to qualify and identify an execution. If not specified, it is generated by the advisor and returned by the function.
execution_desc	Description of the execution (maximum 256 characters)

Return Values

Name of the new execution

181.6.12 IMPLEMENT_EVOLVE_TASK Function

The function implements all the actions recommended by an evolve task.

Syntax

```
DBMS_SPM.IMPLEMENT_EVOLVE_TASK  (
   task_name       IN  VARCHAR2,
   task_owner      IN  VARCHAR2  := NULL,
   execution_name  IN  VARCHAR2  := NULL,
   force           IN BOOLEAN    := FALSE)
 RETURN NUMBER;
```

Parameters

Table 181-17 IMPLEMENT_EVOLVE_TASK Function Parameters

Parameter	Description
`task_name`	Identifier of task to report
`task_owner`	Owner of the evolve task. Defaults to the current schema owner.
`execution_name`	Name to qualify and identify an execution. If `NULL`, the action will be taken for the last task execution.
`force`	Accept all plans even if the advisor did not recommend such an action. The default is `FALSE` requiring acceptance of the plan only if the plan is verified and shows sufficient improvement in benefit.

Return Values

The number of plans accepted

181.6.13 INTERRUPT_EVOLVE_TASK Procedure

The procedure interrupts a currently executing evolve task. The task ends its operations as at a normal exit and the user can access the intermediate results. The task can be resumed later.

Syntax

```
DBMS_SPM.INTERRUPT_EVOLVE_TASK  (
   task_name        IN  VARCHAR2);
```

Parameters

Table 181-18 INTERRUPT_EVOLVE_TASK Procedure Parameters

Parameter	Description
`task_name`	Identifier of task to interrupt

181.6.14 LOAD_PLANS_FROM_AWR Function

This function loads the SQL Management Base (SMB) with SQL plan baselines for a set of SQL statements using the plans from the AWR, and returns the number of plans loaded.

Syntax

```
DBMS_SPM.LOAD_PLANS_FROM_AWR
   begin_snap      IN  NUMBER,
   end_snap        IN  NUMBER,
   basic_filter    IN  VARCHAR2 := NULL,
   fixed           IN  VARCHAR2 := 'NO',
   enabled         IN  VARCHAR2 := 'YES',
   commit_rows     IN  NUMBER := 1000)
 RETURN PLS_INTEGER;
```

Parameters

Table 181-19 LOAD_PLANS_FROM_AWR Function Parameters

Parameter	Description
`begin_snap`	Begin snapshot
`end_snap`	End snapshot
`basic_filter`	SQL predicate to filter the SQL from AWR. `NULL` means all plans in AWR are selected. Specifies the SQL predicate that filters the SQL from the shared SQL area defined on attributes of the `SQLSET_ROW`.
`fixed`	Default `'NO'` means the loaded plans will not change the current 'fixed' property of the SQL plan baseline into which they are loaded.
`enabled`	Default `'YES'` means the loaded plans will be considered by the optimizer
`commit_rows`	Number of SQL plans to load before doing a periodic commit.
`dbid`	The DBID that is used for imported or PDB-level AWR data.

Usage Notes

Requires the `Administer SQL Management Object` privilege

> **See Also:**
>
> For information on the `SQLSET_ROW` objects, see SQLSET_ROW Object Type.

181.6.15 LOAD_PLANS_FROM_CURSOR_CACHE Functions

This function loads one or more plans present in the cursor cache for a SQL statement, or a set of SQL statements. It has four overloads: using SQL statement text, using SQL handle, using SQL ID, or using `attribute_name` and `attribute_value` pair.

Syntax

```
DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE (
   sql_id            IN  VARCHAR2,
   plan_hash_value   IN  NUMBER   := NULL,
   sql_text          IN  CLOB,
   fixed             IN  VARCHAR2 := 'NO',
   enabled           IN  VARCHAR2 := 'YES')
 RETURN PLS_INTEGER;

DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE (
   sql_id            IN  VARCHAR2,
   plan_hash_value   IN  NUMBER   := NULL,
   sql_handle        IN  VARCHAR2,
   fixed             IN  VARCHAR2 := 'NO',
   enabled           IN  VARCHAR2 := 'YES')
 RETURN PLS_INTEGER;

DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE (
   sql_id            IN  VARCHAR2,
   plan_hash_value   IN  NUMBER   := NULL,
   fixed             IN  VARCHAR2 := 'NO',
   enabled           IN  VARCHAR2 := 'YES')
 RETURN PLS_INTEGER;

DBMS_SPM.LOAD_PLANS_FROM_CURSOR_CACHE (
   attribute_name   IN VARCHAR2,
   attribute_value  IN VARCHAR2,
   fixed            IN VARCHAR2 := 'NO',
   enabled          IN VARCHAR2 := 'YES')
  RETURN PLS_INTEGER;
```

Parameters

Table 181-20 LOAD_PLANS_FROM_CURSOR_CACHE Function Parameters

Parameter	Description
`sql_id`	SQL statement identifier. Identifies a SQL statement in the cursor cache. Note: In the third overload the text of identified SQL statement is extracted from cursor cache and is used to identify the SQL plan baseline into which the plan(s) are loaded. If the SQL plan baseline doesn't exist it is created.
`plan_hash_value`	Plan identifier. Default `NULL` means capture all plans present in the cursor cache for the SQL statement identified by `SQL_ID`.
`sql_text`	SQL text to use in identifying the SQL plan baseline into which the plans are loaded. If the SQL plan baseline does not exist, it is created. The use of text is crucial when the user tunes a SQL statement by adding hints to its text and then wants to load the resulting plan(s) into the SQL plan baseline of the original SQL statement.

Table 181-20 (Cont.) LOAD_PLANS_FROM_CURSOR_CACHE Function Parameters

Parameter	Description
sql_handle	SQL handle to use in identifying the SQL plan baseline into which the plans are loaded. The sql_handle must denote an existing SQL plan baseline. The use of handle is crucial when the user tunes a SQL statement by adding hints to its text and then wants to load the resulting plan(s) into the SQL plan baseline of the original SQL statement.
fixed	Default 'NO' means the loaded plans are used as non-fixed plans. Value 'YES' means the loaded plans are used as fixed plans and the SQL plan baseline will not be evolved over time.
attribute_name	One of possible attribute names: • SQL_TEXT" • 'PARSING_SCHEMA_NAME' • 'MODULE' • 'ACTION'
attribute_value	Attribute value is used as a search pattern of LIKE predicate if attribute name is 'SQL_TEXT'. Otherwise, it is used as an equality search value. (for example, for specifying attribute_name => 'SQL_TEXT', and attribute_value => '% HR-123 %' means applying SQL_TEXT LIKE '% HR-123 %' as a selection filter. Similarly, specifying attribute_name => 'MODULE', and attribute_value => 'HR' means applying 'MODULE = 'HR' as a plan selection filter). The attribute value is upper-cased except when it is enclosed in double quotes or attribute name is 'SQL_TEXT'.
enabled	Default 'YES' means the loaded plans are enabled for use by the optimizer

Return Values

Number of plans loaded

Usage Notes

Invoking this subprogram requires the ADMINISTER SQL MANAGEMENT OBJECT privilege.

181.6.16 LOAD_PLANS_FROM_SQLSET Function

This function loads plans stored in a SQL tuning set (STS) into SQL plan baselines. The plans loaded from STS are not verified for performance but added as accepted plans to existing or new SQL plan baselines. This function can be used to seed SQL management base with new SQL plan baselines.

Syntax

```
DBMS_SPM.LOAD_PLANS_FROM_SQLSET (
   sqlset_name      IN  VARCHAR2,
   sqlset_owner     IN  VARCHAR2 := NULL,
   basic_filter     IN  VARCHAR2 := NULL,
   fixed            IN  VARCHAR2 := 'NO',
   enabled          IN  VARCHAR2 := 'YES'
   commit_rows      IN  NUMBER   := 1000)
RETURN PLS_INTEGER;
```

Parameters

Table 181-21 LOAD_PLANS_FROM_SQLSET Function Parameters

Parameter	Description
sqlset_name	Name of the STS from where the plans are loaded into SQL plan baselines
sqlset_owner	Owner of STS. NULL means current schema is the owner.
basic_filter	A filter applied to the STS to select only qualifying plans to be loaded. The filter can take the form of any WHERE clause predicate that can specified against the view DBA_SQLSET_STATEMENTS. For example basic_filter => 'sql_text like ''select /*LOAD_STS*/ %''' or basic_filter => 'sql_id=''b62q7nc33gzwx'''.
fixed	Default 'NO' means the loaded plans are used as non-fixed plans. Value 'YES' means the loaded plans are used as fixed plans and the SQL plan baseline will not be evolved over time.
enabled	Default 'YES' means the loaded plans are enabled for use by the optimizer
commit_rows	Number of SQL plans to load before doing a periodic commit. This helps to shorten the undo log.

Return Values

The number of plans loaded

Usage Notes

- To load plans from a remote system, first load the plans into an STS on the remote system, export/import the STS from remote to local system, and then use this function.
- To load plans from Automatic Workload Repository (AWR), first load the plans stored in AWR snapshots into an STS, and then use this procedure.
- The user can also capture plans resident in the cursor cache for one or more SQL statements into an STS, and then use this procedure.

181.6.17 MIGRATE_STORED_OUTLINE Functions

This function migrates stored outlines for one or more SQL statements to plan baselines in the SQL management base (SMB). Users can specify which stored outline(s) to be migrated based on outline name, SQL text, or outline category, or migrate all stored outlines in the system to SQL plan baselines.

This second overload of the function migrates stored outlines for one or more SQL statements to plan baselines in the SQL management base (SMB) given one or more outline names.

Syntax

```
DBMS_SPM.MIGRATE_STORED_OUTLINE (
   attribute_name     IN  VARCHAR2,
   attribute_value    IN  CLOB,
   fixed              IN  VARCHAR2 := 'NO')
 RETURN CLOB;

DBMS_SPM.MIGRATE_STORED_OUTLINE (
```

```
  outln_list          IN  DBMS_SPM.NAME_LIST,
  fixed               IN  VARCHAR2 := 'NO')
RETURN CLOB;
```

Parameters

Table 181-22 MIGRATE_STORED_OUTLINE Function Parameters

Parameter	Description
attribute_name	Specifies the type of parameter used in attribute_value to identify the migrated stored outlines. It is case insensitive. Possible values: • outline_name • sql_text • category • all
attribute_value	Based on attribute_name, this can be: • Name of stored outline to be migrated • SQL text of stored outlines to be migrated • Category of stored outlines to be migrated • NULL if attribute_name is all
fixed	NO (default) or YES. Specifies the "fixed" status of the plans generated during migration. By default, plans are generated as "non-fixed" plans.
outln_list	List of outline names to be migrated

Return Values

A CLOB containing a formatted report to describe the statistics during the migration, including:

- Number of stored outlines successfully migrated
- Number of stored outlines (and also the corresponding outline names) failed to be migrated and the reasons for the failure

Usage Note

- When the user specifies an outline name, the function migrates stored outlines to plan baseline based on given outline name, which uniquely identifies a single stored outline to be migrated.
- When the user specifies SQL text, the function migrates all stored outlines created for a given SQL statement. A single SQL statement can have multiple stored outlines created for it under different category names. One plan baseline plan is created for each stored outline. The new plan baselines have category names set to DEFAULT. The module name of a plan baseline is set to be the same as the stored outline.
- When the user specifies a category name, the function migrates all stored outlines with the given category name. Only one stored outline exists per category per SQL statement. One plan baseline is created for each stored outline.
- When user specifies to migrate all, the function migrates all stored outlines in the system to plan baselines. One plan baseline is created for each stored outline.

181.6.18 PACK_STGTAB_BASELINE Function

This function packs (exports) SQL plan baselines from SQL management base into a staging table.

Syntax

```
DBMS_SPM.PACK_STGTAB_BASELINE (
   table_name       IN VARCHAR2,
   table_owner      IN VARCHAR2 := NULL,
   sql_handle       IN VARCHAR2 := NULL,
   plan_name        IN VARCHAR2 := NULL,
   sql_text         IN CLOB     := NULL,
   creator          IN VARCHAR2 := NULL,   origin          IN VARCHAR2 := NULL,
   enabled          IN VARCHAR2 := NULL,
   accepted         IN VARCHAR2 := NULL,
   fixed            IN VARCHAR2 := NULL,
   module           IN VARCHAR2 := NULL,
   action           IN VARCHAR2 := NULL)
RETURN NUMBER;
```

Parameters

Table 181-23 PACK_STGTAB_BASELINE Function Parameters

Parameter	Description
table_name	Name of staging table into which SQL plan baselines are packed (case insensitive unless double quoted)
table_owner	Name of staging table owner.Default NULL means current schema is the table owner
sql_handle	SQL handle (case sensitive)
plan_name	Plan name (case sensitive, % wildcards accepted)
sql_text	SQL text string (case sensitive, % wildcards accepted)
creator	Creator of SQL plan baseline (case insensitive unless double quoted)
origin	Origin of SQL plan baseline, should be 'MANUAL-LOAD', 'AUTO-CAPTURE', 'MANUAL_SQLTUNE' or 'AUTO-SQLTUNE' (case insensitive)
enabled	Must be 'YES' or 'NO' (case insensitive)
accepted	Must be 'YES' or 'NO' (case insensitive)
fixed	Must be 'YES' or 'NO' (case insensitive)
module	Module (case sensitive)
action	Action (case sensitive)

Return Values

Number of SQL plan baselines packed

181.6.19 RESET_EVOLVE_TASK Procedure

This procedure resets an evolve task to its initial state.

All intermediate results will be removed from the task. Call this procedure on a task that is not currently executing.

Syntax

```
DBMS_SPM.RESET_EVOLVE_TASK  (
   task_name        IN  VARCHAR2);
```

Parameters

Table 181-24 RESET_EVOLVE_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of task to reset

181.6.20 RESUME_EVOLVE_TASK Procedure

The procedure resumes a previously interrupted task.

Syntax

```
DBMS_SPM.RESUME_EVOLVE_TASK  (
   task_name        IN  VARCHAR2);
```

Parameters

Table 181-25 RESUME_EVOLVE_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of task to resume

181.6.21 REPORT_AUTO_EVOLVE_TASK Function

The procedure displays the results of an execution of an automatic evolve task.

Syntax

```
DBMS_SPM.REPORT_AUTO_EVOLVE_TASK  (
   type            IN  VARCHAR2  := TYPE_TEXT,
   level           IN  VARCHAR2  := LEVEL_TYPICAL,
   section         IN  VARCHAR2  := SECTION_ALL,
   object_id       IN  NUMBER    := NULL,
   execution_name  IN  VARCHAR2  := NULL)
 RETURN CLOB;
```

Parameters

Table 181-26 REPORT_AUTO_EVOLVE_TASK Function Parameters

Parameter	Description
type	Type of the report. Possible values are TEXT, HTML, XML
level	Format of the report. Possible values are BASIC, TYPICAL, ALL.
section	Particular section in the report. Possible values are: SUMMARY, FINDINGS, PLANS, INFORMATION, ERRORS, ALL.
object_id	Identifier of the advisor framework object that represents a single plan. If NULL, the report is generated for all objects.
execution_name	Name to qualify and identify an execution. If NULL, the report is generated for the last task execution.

Return Values

The report

181.6.22 REPORT_EVOLVE_TASK Function

The procedure displays the results of an evolved task.

Syntax

```
DBMS_SPM.REPORT_EVOLVE_TASK  (
   task_name       IN  VARCHAR2,
   type            IN  VARCHAR2  := TYPE_TEXT,
   level           IN  VARCHAR2  := LEVEL_TYPICAL,
   section         IN  VARCHAR2  := SECTION_ALL,
   object_id       IN  NUMBER    := NULL,
   task_owner      IN  VARCHAR2  := NULL,
   execution_name  IN  VARCHAR2  := NULL)
 RETURN CLOB;
```

Parameters

Table 181-27 REPORT_EVOLVE_TASK Function Parameters

Parameter	Description
task_name	Identifier of task to report
type	Type of the report. Possible values are TEXT, HTML, XML
level	Format of the report. Possible values are BASIC, TYPICAL, ALL.
section	Particular section in the report. Possible values are: SUMMARY, FINDINGS, PLANS, INFORMATION, ERRORS, ALL.
object_id	Identifier of the advisor framework object that represents a single plan. If NULL, the report is generated for all objects.
task_owner	Owner of the evolve task. Defaults to the current schema owner.
execution_name	Name to qualify and identify an execution. If NULL, the report is generated for the last task execution.

Return Values

The report

181.6.23 SET_EVOLVE_TASK_PARAMETER Procedure

The procedure sets a parameter of an evolve task, either a `VARCHAR2` or a `NUMBER`.

Syntax

```
DBMS_SPM.SET_EVOLVE_TASK_PARAMETER  (
   task_name     IN  VARCHAR2,
   parameter     IN  VARCHAR2,
   value         IN  NUMBER);

DBMS_SPM.SET_EVOLVE_TASK_PARAMETER  (
   task_name     IN  VARCHAR2  := NULL,
   parameter     IN  VARCHAR2,
   value         IN  VARCHAR2);
```

Parameters

Table 181-28 SET_EVOLVE_TASK_PARAMETER Procedure Parameters

Parameter	Description
`task_name`	Evolve task name
`parameter`	Name of the parameter to set (see following table)
`value`	Value of the parameter (see following table)

The following table describes parameters for the `SET_EVOLVE_TASK_PARAMETER` procedure.

Table 181-29 DBMS_SPM.SET_EVOLVE_TASK_PARAMETER Parameters

Parameter	Description	Default
`alternate_plan_source`	Determines which sources to search for additional plans: • `AUTO` (the database selects the source automatically) • `AUTOMATIC_WORKLOAD_REPOSITORY` • `CURSOR_CACHE` • `SQL_TUNING_SET` You can combine multiple values with the plus sign (+).	The default depends on whether the SPM Evolve Advisor task is automated or manual: • If automated, the default is `AUTO`. • If manual, the default is `CURSOR_CACHE+AUTOMATIC_WORKLOAD_REPOSITORY`.

Table 181-29 (Cont.) DBMS_SPM.SET_EVOLVE_TASK_PARAMETER Parameters

Parameter	Description	Default
`alternate_plan_baseline`	Determines which alternative plans should be loaded: • `AUTO` lets Autonomous Database choose whether to load plans for statements with or without baselines. • `EXISTING` loads alternate plans with for statements with existing baselines. • `NEW` loads alternative plans for statements without a baseline, in which case a new baseline is created. You can combine multiple values with the plus sign (+), as in `EXISTING+NEW`.	`EXISTING`
`alternate_plan_limit`	Specifies the maximum number of plans to load in total (that is, not the limit for each SQL statement).	The default depends on whether the SPM Evolve Advisor task is automated or manual: • If automated, the default is `UNLIMITED`. • If manual, the default is `10`.
`accept_plans`	Specifies whether to accept recommended plans automatically. When `ACCEPT_PLANS` is `true`, SQL plan management automatically accepts all plans recommended by the task. When `ACCEPT_PLANS` is `false`, the task verifies the plans and generates a report of its findings, but does not evolve the plans automatically. You can use a report to identify new SQL plan baselines and accept them manually.	`true` (regardless of whether the advisor is run automatically or manually)
`time_limit`	Global time limit in seconds. This is the total time allowed for the task.	The default depends on whether the SPM Evolve Advisor task is automated or manual: • If automated, the default is `3600`. • If manual, the default is `2147483646`.

See Also:

Oracle Database Licensing Information User Manual for details on which features are supported for different editions and services

181.6.24 UNPACK_STGTAB_BASELINE Function

This function unpacks (imports) SQL plan baselines from a staging table into SQL management base.

Syntax

```
DBMS_SPM.UNPACK_STGTAB_BASELINE (
   table_name       IN VARCHAR2,
   table_owner      IN VARCHAR2 := NULL,
   sql_handle       IN VARCHAR2 := NULL,
   plan_name        IN VARCHAR2 := NULL,
   sql_text         IN CLOB     := NULL,
   creator          IN VARCHAR2 := NULL,   origin          IN VARCHAR2 := NULL,
   enabled          IN VARCHAR2 := NULL,
   accepted         IN VARCHAR2 := NULL,
   fixed            IN VARCHAR2 := NULL,
   module           IN VARCHAR2 := NULL,
   action           IN VARCHAR2 := NULL)
RETURN NUMBER;
```

Parameters

Table 181-30 UNPACK_STGTAB_BASELINE Function Parameters

Parameter	Description
table_name	Name of staging table from which SQL plan baselines are unpacked (case insensitive unless double quoted)
table_owner	Name of staging table owner.Default NULL means current schema is the table owner
sql_handle	SQL handle (case sensitive)
plan_name	Plan name (case sensitive,% wildcards accepted)
sql_text	SQL text string (case sensitive, % wildcards accepted)
creator	Creator of SQL plan baseline (case insensitive unless double quoted)
origin	Origin of SQL plan baseline, should be 'MANUAL-LOAD', 'AUTO-CAPTURE','MANUAL_SQLTUNE' or 'AUTO-SQLTUNE' (case insensitive)
enabled	Must be 'YES' or 'NO' (case insensitive)
accepted	Must be 'YES' or 'NO' (case insensitive)
fixed	Must be 'YES' or 'NO' (case insensitive)
module	Module (case sensitive)
action	Action (case sensitive)

Return Values

Number of plans unpacked

182
DBMS_SQL

The `DBMS_SQL` package provides an interface to use dynamic SQL to parse any data manipulation language (DML) or data definition language (DDL) statement using PL/SQL.

For example, you can enter a `DROP TABLE` statement from within a stored procedure by using the PARSE Procedures supplied with the `DBMS_SQL` package.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Exceptions
- Operational Notes
- Examples
- Data Structures
- Summary of DBMS_SQL Subprograms

> **See Also:**
>
> For more information on native dynamic SQL, see *Oracle Database PL/SQL Language Reference*.

182.1 DBMS_SQL Overview

Oracle lets you write stored procedures and anonymous PL/SQL blocks that use dynamic SQL. Dynamic SQL statements are not embedded in your source program; rather, they are stored in character strings that are input to, or built by, the program at runtime. This enables you to create more general-purpose procedures. For example, dynamic SQL lets you create a procedure that operates on a table whose name is not known until runtime.

Native Dynamic SQL is an alternative to `DBMS_SQL` that lets you place dynamic SQL statements directly into PL/SQL blocks. In most situations, Native Dynamic SQL is easier to use and performs better than `DBMS_SQL`. However, Native Dynamic SQL itself has certain limitations:

- There is no support for so-called Method 4 (for dynamic SQL statements with an unknown number of inputs or outputs)
- There are some tasks that can only be performed using `DBMS_SQL`. For tasks that require `DBMS_SQL`, see *Oracle Database PL/SQL Language Reference*.

The ability to use dynamic SQL from within stored procedures generally follows the model of the Oracle Call Interface (OCI).

> **See Also:**
>
> *Oracle Call Interface Programmer's Guide*

PL/SQL differs somewhat from other common programming languages, such as C. For example, addresses (also called pointers) are not user-visible in PL/SQL. As a result, there are some differences between the Oracle Call Interface and the `DBMS_SQL` package. These differences include the following:

- The OCI binds by address and the `DBMS_SQL` package binds by value.
- With `DBMS_SQL` you must call `VARIABLE_VALUE` to retrieve the value of an `OUT` parameter for an anonymous block, and you must call `COLUMN_VALUE` after fetching rows to retrieve the values of the columns in the rows into your program.
- The current release of the `DBMS_SQL` package does not provide `CANCEL` cursor procedures.
- Indicator variables are not required, because `NULLs` are fully supported as values of a PL/SQL variable.

182.2 DBMS_SQL Security Model

`DBMS_SQL` is a `SYS`-owned package compiled with `AUTHID CURRENT_USER`. Any `DBMS_SQL` subprogram called from an anonymous PL/SQL block runs with the privileges of the current user.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information about using Invoker Rights or Definer Rights

Preventing Malicious or Accidental Access of Open Cursor Numbers

An error, `ORA-29471`, is raised when any `DBMS_SQL` subprogram is called with a cursor number that does not denote an open cursor. When the error is raised, an alert is issued to the alert log and DBMS_SQL becomes inoperable for the life of the session.

If the actual value for the cursor number in a call to the IS_OPEN Function denotes a cursor currently open in the session, the return value is `TRUE`. If the actual value is `NULL`, then the return value is `FALSE`. Otherwise, this raises an `ORA-29471` error.

Preventing Inappropriate Use of a Cursor

Cursors are protected from security breaches that subvert known existing cursors.

Checks are made when binding and executing. Optionally, checks may be performed for every single DBMS_SQL subprogram call. The check is:

- The `current_user` is the same on calling the subprogram as it was on calling the most recent parse.
- The enabled roles on calling the subprogram must be identical to the enabled roles on calling the most recent parse.

- The container is the same on calling the subprogram as it was on calling the most recent parse.

Consistent with the use of definer's rights subprograms, roles do not apply.

If either check fails, then an `ORA-29470` error is raised.

The mechanism for defining when checks are performed is a new overload for the `OPEN_CURSOR` subprogram, which takes a formal parameter, `security_level`, with allowed values `NULL`, `1` and `2`.

- When `security_level` = `1` (or is `NULL`), the checks are made only when binding and executing.
- When `security_level` = `2`, the checks are always made.

Upgrade Considerations

This security regime is stricter than those in the previous releases. As a consequence, users of `DBMS_SQL` may encounter runtime errors on upgrade.

182.3 DBMS_SQL Constants

The `DBMS_SQL Constants` package provides constants that are used with the `language_flag` parameter of the `PARSE Procedures`.

These constants are described in the following table.

Table 182-1 DBMS_SQL Constants

Name	Type	Value	Description
V6	INTEGER	0	Specifies Oracle database version 6 behavior
NATIVE	INTEGER	1	Specifies normal behavior for the database to which the program is connected
V7	INTEGER	2	Specifies Oracle database version 7 behavior
FOREIGN_SYNTAX	INTEGER	4294967295	Specifies a non-Oracle database syntax and behavior. The SQL statement to be parsed needs to be translated first using the SQL translation profile set in the database session. The SQL translation profile is a database schema object that directs how SQL statements are translated to Oracle. An error is raised if a profile is not set.

Related Topics

- PARSE Procedures
  This procedure parses the given statement in the given cursor. All statements are parsed immediately. In addition, DDL statements are run immediately when parsed.

182.4 DBMS_SQL Operational Notes

These operational notes describe processing queries, processing updates, inserts, and deletes, and locating errors.

Processing Queries

If you are using dynamic SQL to process a query, then you must perform the following steps:

1. Specify the variables that are to receive the values returned by the `SELECT` statement by calling the DEFINE_COLUMN Procedures, the DEFINE_COLUMN_LONG Procedure, or the DEFINE_ARRAY Procedure.
2. Run your `SELECT` statement by calling the EXECUTE Function.
3. Call the FETCH_ROWS Function (or `EXECUTE_AND_FETCH`) to retrieve the rows that satisfied your query.
4. Call COLUMN_VALUE Procedure or COLUMN_VALUE_LONG Procedure to determine the value of a column retrieved by the FETCH_ROWS Function for your query. If you used anonymous blocks containing calls to PL/SQL procedures, then you must call the VARIABLE_VALUE Procedures to retrieve the values assigned to the output variables of these procedures.

Processing Updates, Inserts, and Deletes

If you are using dynamic SQL to process an `INSERT`, `UPDATE`, or `DELETE`, then you must perform the following steps:

1. Run your `INSERT`, `UPDATE`, or `DELETE` statement by calling the EXECUTE Function.
2. If statements have the `returning` clause, then you must call the VARIABLE_VALUE Procedures to retrieve the values assigned to the output variables.

Locating Errors

The `DBMS_SQL` package has additional functions for obtaining information about the last referenced cursor in the session. The values returned by these functions are meaningful only immediately after a SQL statement is run. In addition, some error-locating functions are meaningful only after certain `DBMS_SQL` calls. For example, you call the LAST_ERROR_POSITION Function immediately after calling one of the PARSE Procedures.

182.5 DBMS_SQL Execution Flow

These functions comprise the DBMS_SQL execution flow.

1. OPEN_CURSOR
2. PARSE
3. BIND_VARIABLE, BIND_VARIABLE_PKG or BIND_ARRAY
4. DEFINE_COLUMN_ DEFINE_COLUMN_LONG_ or DEFINE_ARRAY
5. EXECUTE
6. FETCH_ROWS or EXECUTE_AND_FETCH
7. VARIABLE_VALUE, VARIABLE_PKG, COLUMN_VALUE or COLUMN_VALUE_LONG
8. CLOSE_CURSOR

182.5.1 OPEN_CURSOR

To process a SQL statement, you must have an open cursor. When you call the OPEN_CURSOR Functions, you receive a cursor `ID` number for the data structure representing a valid cursor maintained by Oracle.

These cursors are distinct from cursors defined at the precompiler, OCI, or PL/SQL level, and are used only by the `DBMS_SQL` package.

Related Topics

- OPEN_CURSOR Functions
  This function opens a new cursor.

182.5.2 PARSE

Every SQL statement must be parsed by calling the PARSE procedures. Parsing the statement checks the statement's syntax and associates it with the cursor in your program.

You can parse any DML or DDL statement. DDL statements are run on the parse, which performs the implied commit.

The execution flow of `DBMS_SQL` is shown in Figure 182-1.

Figure 182-1 DBMS_SQL Execution Flow

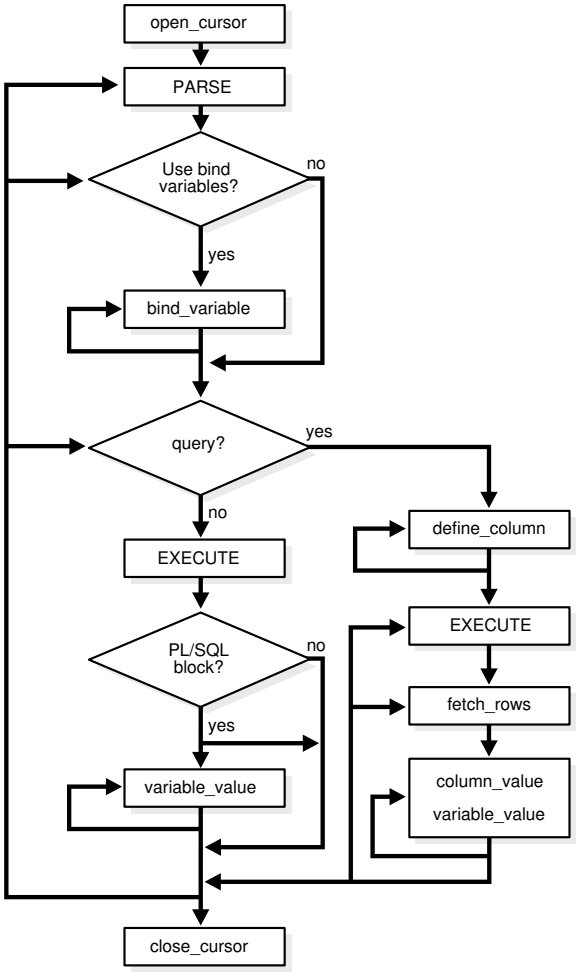

illustration: arpls008
release: 9
caption: DBMS_SQL Execution Flow
date: 7/31/01
platform: pc
ref: ADG8013

Related Topics

- PARSE Procedures
  This procedure parses the given statement in the given cursor. All statements are parsed immediately. In addition, DDL statements are run immediately when parsed.

182.5.3 BIND_VARIABLE, BIND_VARIABLE_PKG or BIND_ARRAY

Many DML statements require that data in your program be input to Oracle. When you define a SQL statement that contains input data to be supplied at runtime, you must use placeholders in the SQL statement to mark where data must be supplied.

For each placeholder in the SQL statement, you must call one of the BIND_ARRAY Procedures, or BIND_VARIABLE Procedures, or the BIND_VARIABLE_PKG Procedure to supply the value of a variable in your program (or the values of an array) to the placeholder. When the SQL statement is subsequently run, Oracle uses the data that your program has placed in the output and input, or bind variables.

`DBMS_SQL` can run a DML statement multiple times — each time with a different bind variable. The `BIND_ARRAY` procedure lets you bind a collection of scalars, each value of which is used as an input variable once for each `EXECUTE`. This is similar to the array interface supported by the OCI.

Note that the datatype of the values bound to placeholders cannot be PL/SQL-only datatypes.

182.5.4 DEFINE_COLUMN, DEFINE_COLUMN_LONG, or DEFINE_ARRAY

The `DEFINE_COLUMN`, `DEFINE_COLUMN_LONG`, and `DEFINE_ARRAY` procedures specify the variables that receive `SELECT` values on a query.

The columns of the row being selected in a `SELECT` statement are identified by their relative positions as they appear in the select list, from left to right. For a query, you must call one of the define procedures (DEFINE_COLUMN Procedures, DEFINE_COLUMN_LONG Procedure, or DEFINE_ARRAY Procedure) to specify the variables that are to receive the `SELECT` values, much the way an `INTO` clause does for a static query.

Use the `DEFINE_COLUMN_LONG` procedure to define `LONG` columns, in the same way that `DEFINE_COLUMN` is used to define non-`LONG` columns. You must call `DEFINE_COLUMN_LONG` before using the COLUMN_VALUE_LONG Procedure to fetch from the `LONG` column.

Use the `DEFINE_ARRAY` procedure to define a PL/SQL collection into which you want to fetch rows in a single `SELECT` statement. `DEFINE_ARRAY` provides an interface to fetch multiple rows at one fetch. You must call `DEFINE_ARRAY` before using the `COLUMN_VALUE` procedure to fetch the rows.

182.5.5 EXECUTE

Call the `EXECUTE` Function to run your SQL statement.

Related Topics

- EXECUTE Function
  This function executes a given cursor. This function accepts the `ID` number of the cursor and returns the number of rows processed.

182.5.6 FETCH_ROWS or EXECUTE_AND_FETCH

The FETCH_ROWS Function retrieves the rows that satisfy the query. Each successive fetch retrieves another set of rows, until the fetch is unable to retrieve any more rows. Instead of

calling EXECUTE Function and then `FETCH_ROWS`, you may find it more efficient to call EXECUTE_AND_FETCH Function if you are calling `EXECUTE` for a single execution.

Related Topics

- FETCH_ROWS Function
  This function fetches a row from a given cursor.
- EXECUTE Function
  This function executes a given cursor. This function accepts the `ID` number of the cursor and returns the number of rows processed.
- EXECUTE_AND_FETCH Function
  This function executes the given cursor and fetches rows.

182.5.7 VARIABLE_VALUE, VARIABLE_VALUE_PKG, COLUMN_VALUE, or COLUMN_VALUE_LONG

The type of call determines which procedure or function to use.

For queries, call the COLUMN_VALUE Procedure to determine the value of a column retrieved by the FETCH_ROWS Function.

For anonymous blocks containing calls to `PL/SQL` procedures or DML statements with `returning` clause, call the VARIABLE_VALUE Procedures or the VARIABLE_VALUE_PKG Procedure to retrieve the values assigned to the output variables when statements were run.

To fetch only part of a `LONG` database column (which can be up to two gigabytes in size), use the DEFINE_COLUMN_LONG Procedure. You can specify the offset (in bytes) into the column value, and the number of bytes to fetch.

182.5.8 CLOSE_CURSOR

When you no longer need a cursor for a session, close the cursor by calling the CLOSE_CURSOR Procedure. If you are using an Oracle Open Gateway, then you may need to close cursors at other times as well. Consult your *Oracle Open Gateway* documentation for additional information.

Related Topics

- CLOSE_CURSOR Procedure
  This procedure closes a given cursor.

182.6 DBMS_SQL Exceptions

This exception is raised by the `COLUMN_VALUE Procedure` or the `VARIABLE_VALUE Procedures` when the type of the given `OUT` parameter (for where to put the requested value) is different from the type of the value.

```
inconsistent_type EXCEPTION;
  pragma exception_init(inconsistent_type, -6562);
```

Related Topics

- COLUMN_VALUE Procedure
  This procedure returns the value of the cursor element for a given position in a given cursor. This procedure is used to access the data fetched by calling `FETCH_ROWS`.

- VARIABLE_VALUE Procedures
  This procedure returns the value of the named variable for a given cursor. It is used to return the values of bind variables inside PL/SQL blocks or DML statements with `returning` clause.

182.7 DBMS_SQL Examples

These example procedures use the `DBMS_SQL` package.

Example : Using DBMS_SQL Demo

This example does not need dynamic SQL because the text of the statement is known at compile time, but it illustrates the basic concept underlying the package.

The `DEMO` procedure deletes all of the employees from the `EMP` table whose salaries are greater than the salary that you specify when you run `DEMO`.

```
CREATE OR REPLACE PROCEDURE demo(salary IN NUMBER) AS
    cursor_name INTEGER;
    rows_processed INTEGER;
BEGIN
    cursor_name := dbms_sql.open_cursor;
    DBMS_SQL.PARSE(cursor_name, 'DELETE FROM emp WHERE sal > :x',
                   DBMS_SQL.NATIVE);
    DBMS_SQL.BIND_VARIABLE(cursor_name, ':x', salary);
    rows_processed := DBMS_SQL.EXECUTE(cursor_name);
    DBMS_SQL.CLOSE_CURSOR(cursor_name);
EXCEPTION
WHEN OTHERS THEN
    DBMS_SQL.CLOSE_CURSOR(cursor_name);
END;
```

Example 2

The following sample procedure is passed a SQL statement, which it then parses and runs:

```
CREATE OR REPLACE PROCEDURE exec(STRING IN varchar2) AS
    cursor_name INTEGER;
    ret INTEGER;
BEGIN
   cursor_name := DBMS_SQL.OPEN_CURSOR;
```

DDL statements are run by the parse call, which performs the implied commit.

```
   DBMS_SQL.PARSE(cursor_name, string, DBMS_SQL.NATIVE);
   ret := DBMS_SQL.EXECUTE(cursor_name);
   DBMS_SQL.CLOSE_CURSOR(cursor_name);
END;
```

Creating such a procedure enables you to perform the following operations:

- The SQL statement can be dynamically generated at runtime by the calling program.
- The SQL statement can be a DDL statement or a DML without binds.

For example, after creating this procedure, you could make the following call:

```
exec('create table acct(c1 integer)');
```

You could even call this procedure remotely, as shown in the following example. This lets you perform remote DDL.

```
exec@domain.com('CREATE TABLE acct(c1 INTEGER)');
```

Example 3

The following sample procedure is passed the names of a source and a destination table, and copies the rows from the source table to the destination table. This sample procedure assumes that both the source and destination tables have the following columns:

```
id        of type NUMBER
name      of type VARCHAR2(30)
birthdate of type DATE
```

This procedure does not need the use of dynamic SQL; however, it illustrates the concepts of this package.

```
CREATE OR REPLACE PROCEDURE copy (
     source      IN VARCHAR2,
     destination IN VARCHAR2) IS
     id_var             NUMBER;
     name_var           VARCHAR2(30);
     birthdate_var      DATE;
     source_cursor      INTEGER;
     destination_cursor INTEGER;
     ignore             INTEGER;
  BEGIN

  -- Prepare a cursor to select from the source table:
     source_cursor := dbms_sql.open_cursor;
     DBMS_SQL.PARSE(source_cursor,
         'SELECT id, name, birthdate FROM ' || source,
          DBMS_SQL.NATIVE);
     DBMS_SQL.DEFINE_COLUMN(source_cursor, 1, id_var);
     DBMS_SQL.DEFINE_COLUMN(source_cursor, 2, name_var, 30);
     DBMS_SQL.DEFINE_COLUMN(source_cursor, 3, birthdate_var);
     ignore := DBMS_SQL.EXECUTE(source_cursor);

  -- Prepare a cursor to insert into the destination table:
     destination_cursor := DBMS_SQL.OPEN_CURSOR;
     DBMS_SQL.PARSE(destination_cursor,
                  'INSERT INTO ' || destination ||
                  ' VALUES (:id_bind, :name_bind, :birthdate_bind)',
                   DBMS_SQL.NATIVE);

  -- Fetch a row from the source table and insert it into the destination table:
     LOOP
       IF DBMS_SQL.FETCH_ROWS(source_cursor)>0 THEN
         -- get column values of the row
         DBMS_SQL.COLUMN_VALUE(source_cursor, 1, id_var);
         DBMS_SQL.COLUMN_VALUE(source_cursor, 2, name_var);
         DBMS_SQL.COLUMN_VALUE(source_cursor, 3, birthdate_var);

  -- Bind the row into the cursor that inserts into the destination table. You
  -- could alter this example to require the use of dynamic SQL by inserting an
  -- if condition before the bind.
       DBMS_SQL.BIND_VARIABLE(destination_cursor, ':id_bind', id_var);
       DBMS_SQL.BIND_VARIABLE(destination_cursor, ':name_bind', name_var);
       DBMS_SQL.BIND_VARIABLE(destination_cursor, ':birthdate_bind',
                                                              birthdate_var);
       ignore := DBMS_SQL.EXECUTE(destination_cursor);
```

```
      ELSE

  -- No more rows to copy:
        EXIT;
      END IF;
    END LOOP;

  -- Commit and close all cursors:
     COMMIT;
     DBMS_SQL.CLOSE_CURSOR(source_cursor);
     DBMS_SQL.CLOSE_CURSOR(destination_cursor);
   EXCEPTION
     WHEN OTHERS THEN
       IF DBMS_SQL.IS_OPEN(source_cursor) THEN
         DBMS_SQL.CLOSE_CURSOR(source_cursor);
       END IF;
       IF DBMS_SQL.IS_OPEN(destination_cursor) THEN
         DBMS_SQL.CLOSE_CURSOR(destination_cursor);
       END IF;
       RAISE;
  END;
/
```

Example 4: RETURNING clause

With this clause, `INSERT`, `UPDATE`, and `DELETE` statements can return values of expressions in bind variables.

If a single row is inserted, updated, or deleted, then use `DBMS_SQL.BIND_VARIABLE` to bind these outbinds. To get the values in these bind variables, call `DBMS_SQL.VARIABLE_VALUE`

> **Note:**
>
> This process is similar to `DBMS_SQL.VARIABLE_VALUE`, which must be called after running a PL/SQL block with an outbind inside `DBMS_SQL`.

i) Single-row insert

```
      CREATE OR REPLACE PROCEDURE single_Row_insert
           (c1 NUMBER, c2 NUMBER, r OUT NUMBER) is
      c NUMBER;
      n NUMBER;
      begin
        c := DBMS_SQL.OPEN_CURSOR;
        DBMS_SQL.PARSE(c, 'INSERT INTO tab VALUES (:bnd1, :bnd2) ' ||
                          'RETURNING c1*c2 INTO :bnd3', DBMS_SQL.NATIVE);
     DBMS_SQL.BIND_VARIABLE(c, 'bnd1', c1);
        DBMS_SQL.BIND_VARIABLE(c, 'bnd2', c2);
        DBMS_SQL.BIND_VARIABLE(c, 'bnd3', r);
        n := DBMS_SQL.EXECUTE(c);
        DBMS_SQL.VARIABLE_VALUE(c, 'bnd3', r); -- get value of outbind variable
        DBMS_SQL.CLOSE_CURSOR(c);
      END;
      /
```

ii) Single-row update

```
CREATE OR REPLACE PROCEDURE single_Row_update
     (c1 NUMBER, c2 NUMBER, r out NUMBER) IS
c NUMBER;
n NUMBER;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, 'UPDATE tab SET c1 = :bnd1, c2 = :bnd2 ' ||
                    'WHERE rownum < 2 ' ||
                    'RETURNING c1*c2 INTO :bnd3', DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_VARIABLE(c, 'bnd1', c1);
  DBMS_SQL.BIND_VARIABLE(c, 'bnd2', c2);
  DBMS_SQL.BIND_VARIABLE(c, 'bnd3', r);
  n := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.VARIABLE_VALUE(c, 'bnd3', r);-- get value of outbind variable
  DBMS_SQL.CLOSE_CURSOR(c);
END;
/
```

iii) Single-row delete

```
CREATE OR REPLACE PROCEDURE single_Row_Delete
     (c1 NUMBER, r OUT NUMBER) is
c NUMBER;
n number;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, 'DELETE FROM tab WHERE ROWNUM = :bnd1 ' ||
                'RETURNING c1*c2 INTO :bnd2', DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_VARIABLE(c, 'bnd1', c1);
  DBMS_SQL.BIND_VARIABLE(c, 'bnd2', r);
  n := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.VARIABLE_VALUE(c, 'bnd2', r);-- get value of outbind variable
  DBMS_SQL.CLOSE_CURSOR(c);
END;
/
```

iv) Multiple-row insert

```
CREATE OR REPLACE PROCEDURE multi_Row_insert
     (c1 DBMS_SQL.NUMBER_TABLE, c2 DBMS_SQL.NUMBER_TABLE,
      r OUT DBMS_SQL.NUMBER_TABLE) is
c NUMBER;
n NUMBER;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, 'insert into tab VALUES (:bnd1, :bnd2) ' ||
                    'RETURNING c1*c2 INTO :bnd3', DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_ARRAY(c, 'bnd1', c1);
  DBMS_SQL.BIND_ARRAY(c, 'bnd2', c2);
  DBMS_SQL.BIND_ARRAY(c, 'bnd3', r);
  n := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.VARIABLE_VALUE(c, 'bnd3', r);-- get value of outbind variable
  DBMS_SQL.CLOSE_CURSOR(c);
END;
/
```

v) Multiple-row update.

```
CREATE OR REPLACE PROCEDURE multi_Row_update
     (c1 NUMBER, c2 NUMBER, r OUT DBMS_SQL.NUMBER_TABLE) IS
c NUMBER;
n NUMBER;
```

```
BEGIN
   c := DBMS_SQL.OPEN_CURSOR;
   DBMS_SQL.PARSE(c, 'UPDATE tab SET c1 = :bnd1 WHERE c2 = :bnd2 ' ||
                     'RETURNING c1*c2 INTO :bnd3', DBMS_SQL.NATIVE);
   DBMS_SQL.BIND_VARIABLE(c, 'bnd1', c1);
   DBMS_SQL.BIND_VARIABLE(c, 'bnd2', c2);
   DBMS_SQL.BIND_ARRAY(c, 'bnd3', r);
   n := DBMS_SQL.EXECUTE(c);
   DBMS_SQL.VARIABLE_VALUE(c, 'bnd3', r);-- get value of outbind variable
   DBMS_SQL.CLOSE_CURSOR(c);
END;
/
```

Note:

bnd1 and bnd2 can be arrays too. The value of the expression for all the rows updated will be in bnd3. There is no way to determine which rows were updated for each value of bnd1 and bnd2.

vi) Multiple-row delete

```
CREATE OR REPLACE PROCEDURE multi_row_delete
     (c1 DBMS_SQL.NUMBER_TABLE,
      r OUT DBMS_SQL.NUMBER_TABLE) is
c NUMBER;
n NUMBER;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, 'DELETE FROM tab WHERE c1 = :bnd1' ||
                    'RETURNING c1*c2 INTO :bnd2', DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_ARRAY(c, 'bnd1', c1);
  DBMS_SQL.BIND_ARRAY(c, 'bnd2', r);
  n := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.VARIABLE_VALUE(c, 'bnd2', r);-- get value of outbind variable
  DBMS_SQL.CLOSE_CURSOR(c);
END;
/
```

vii) outbind in bulk PL/SQL

```
CREATE OR REPLACE PROCEDURE foo (n NUMBER, square OUT NUMBER) IS
BEGIN square := n * n; END;/

CREATE OR REPLACE PROCEDURE bulk_plsql
   (n DBMS_SQL.NUMBER_TABLE, square OUT DBMS_SQL.NUMBER_TABLE) IS
c NUMBER;
r NUMBER;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, 'BEGIN foo(:bnd1, :bnd2); END;', DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_ARRAY(c, 'bnd1', n);
  DBMS_SQL.BIND_ARRAY(c, 'bnd2', square);
  r := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.VARIABLE_VALUE(c, 'bnd2', square);
END;
/
```

> **Note:**
>
> DBMS_SQL.BIND_ARRAY of number_Table internally binds a number. The number of times statement is run depends on the number of elements in an inbind array.

Example 5: Binds and Defines of User-defined Types in DBMS_SQL

```
CREATE TYPE dnames_var IS VARRAY(7) OF VARCHAR2(30)
/

CREATE TABLE depts (region VARCHAR2(25), dept_names dnames_var)
/
INSERT INTO depts VALUES('Europe', dnames_var('Shipping','Sales','Finance'))
/
INSERT INTO depts VALUES('Americas', dnames_var('Sales','Finance','Shipping'))
/
INSERT INTO depts VALUES('Asia', dnames_var('Finance','Payroll','Shipping','Sales'))
/

CREATE OR REPLACE PROCEDURE update_depts(new_dnames dnames_var, region VARCHAR2) IS
   some_dnames dnames_var;
   c      NUMBER;
   r      NUMBER;
   sql_stmt VARCHAR2(32767) :=
    'UPDATE depts SET dept_names = :b1 WHERE region = :b2 RETURNING dept_names INTO :b3';

BEGIN

   c := DBMS_SQL.OPEN_CURSOR;

   DBMS_SQL.PARSE(c, sql_stmt, dbms_sql.native);

   DBMS_SQL.BIND_VARIABLE(c, 'b1', new_dnames);
   DBMS_SQL.BIND_VARIABLE(c, 'b2', region);
   DBMS_SQL.BIND_VARIABLE(c, 'b3', some_dnames);

   r := DBMS_SQL.EXECUTE(c);

   -- Get value of outbind variable
   DBMS_SQL.VARIABLE_VALUE(c, 'b3', some_dnames);

   DBMS_SQL.CLOSE_CURSOR(c);

   -- select dept_names
   sql_stmt := 'SELECT dept_names FROM depts WHERE region = :b1';

   c := DBMS_SQL.OPEN_CURSOR;
   DBMS_SQL.PARSE(c, sql_stmt, dbms_sql.native);

   DBMS_SQL.DEFINE_COLUMN(c, 1, some_dnames);
   DBMS_SQL.BIND_VARIABLE(c, 'b1', region);

   r := DBMS_SQL.EXECUTE_AND_FETCH(c);

   DBMS_SQL.COLUMN_VALUE(c, 1, some_dnames);

   DBMS_SQL.CLOSE_CURSOR(c);

    -- loop through some_dnames collections
```

```
    FOR i IN some_dnames.FIRST .. some_dnames.LAST  LOOP
        DBMS_OUTPUT.PUT_LINE('Dept. Name = ' || some_dnames(i) || ' Updated!');
    END LOOP;
END;
/

DECLARE
  new_dnames dnames_var;
BEGIN
  new_dnames := dnames_var('Benefits', 'Advertising', 'Contracting',
                           'Executive', 'Marketing');
  update_depts(new_dnames, 'Asia');
END;
/
```

182.8 DBMS_SQL Data Structures

The `DBMS_SQL` package defines `RECORD` type and `TABLE` type data structures.

RECORD Types

- DBMS_SQL DESC_REC Record Type (deprecated)
- DBMS_SQL DESC_REC2 Record Type
- DBMS_SQL DESC_REC3 Record Type
- DBMS_SQL DESC_REC4 Record Type

TABLE Types for DESCRIBE_COLUMNS Procedures

- DBMS_SQL DESC_TAB Table Type
- DBMS_SQL DESC_TAB2 Table Type
- DBMS_SQL DESC_TAB3 Table Type
- DBMS_SQL DESC_TAB4 Table Type

TABLE Types For Scalar and LOB Collections

DBMS_SQL bulk operations are only supported with these predefined DBMS_SQL TABLE types.

- DBMS_SQL BFILE_TABLE Table Type
- DBMS_SQL BINARY_DOUBLE_TABLE Table Type
- DBMS_SQL BINARY_FLOAT_TABLE Table Type
- DBMS_SQL BLOB_TABLE Table Type
- DBMS_SQL CLOB_TABLE Table Type
- DBMS_SQL DATE_TABLE Table Type
- DBMS_SQL INTERVAL_DAY_TO_SECOND_TABLE Table Type
- DBMS_SQL INTERVAL_YEAR_TO_MONTH_TABLE Table Type
- DBMS_SQL JSON_TABLE Table Type
- DBMS_SQL NUMBER_TABLE Table Type
- DBMS_SQL TIME_TABLE Table Type
- DBMS_SQL TIME_WITH_TIME_ZONE_TABLE Table Type

- DBMS_SQL TIMESTAMP_TABLE Table Type
- DBMS_SQL TIMESTAMP_WITH_LTZ_TABLE Table Type
- DBMS_SQL TIMESTAMP_WITH_TIME_ZONE_TABLE Table Type
- DBMS_SQL UROWID_TABLE Table Type
- DBMS_SQL VARCHAR2_TABLE Table Type
- DBMS_SQL VARCHAR2A Table Type
- DBMS_SQL VARCHAR2S Table Type
- DBMS_SQL VECTOR Table Type

182.8.1 DBMS_SQL DESC_REC Record Type

This record type holds the describe information for a single column in a dynamic query.

> **Note:**
>
> This type has been deprecated in favor of the DESC_REC2 Record Type.

It is the element type of the `DESC_TAB` table type and the DESCRIBE_COLUMNS Procedure.

Syntax

```
TYPE desc_rec IS RECORD (
      col_type            BINARY_INTEGER := 0,
      col_max_len         BINARY_INTEGER := 0,
      col_name            VARCHAR2(32)   := '',
      col_name_len        BINARY_INTEGER := 0,
      col_schema_name     VARCHAR2(32)   := '',
      col_schema_name_len BINARY_INTEGER := 0,
      col_precision       BINARY_INTEGER := 0,
      col_scale           BINARY_INTEGER := 0,
      col_charsetid       BINARY_INTEGER := 0,
      col_charsetform     BINARY_INTEGER := 0,
      col_null_ok         BOOLEAN        := TRUE);
TYPE desc_tab IS TABLE OF desc_rec INDEX BY BINARY_INTEGER;
```

Fields

Table 182-2 DESC_REC Fields

Field	Description
col_type	Type of column
col_max_len	Maximum column length
col_name	Name of column
col_name_len	Length of column name
col_schema_name	Column schema name
col_schema_name_len	Length of column schema name
col_precision	Precision of column

Table 182-2 (Cont.) DESC_REC Fields

Field	Description
col_scale	Scale of column
col_charsetid	Column character set id
col_charsetform	Column character set form
col_null_ok	NULL column flag; TRUE, if NULL possible

182.8.2 DBMS_SQL DESC_REC2 Record Type

DESC_REC2 is the element type of the DESC_TAB2 table type and the DESCRIBE_COLUMNS2 Procedure.

This record type is identical to DESC_REC except for the col_name field, which has been expanded to the maximum possible size for VARCHAR2. It is therefore preferred to DESC_REC because column name values can be greater than 32 characters. DESC_REC is deprecated as a result.

Syntax

```
TYPE desc_rec2 IS RECORD (
   col_type            binary_integer := 0,
   col_max_len         binary_integer := 0,
   col_name            varchar2(32767) := '',
   col_name_len        binary_integer := 0,
   col_schema_name     varchar2(32)   := '',
   col_schema_name_len binary_integer := 0,
   col_precision       binary_integer := 0,
   col_scale           binary_integer := 0,
   col_charsetid       binary_integer := 0,
   col_charsetform     binary_integer := 0,
   col_null_ok         boolean        := TRUE);
```

Fields

Table 182-3 DESC_REC2 Fields

Field	Description
col_type	Type of column
col_max_len	Maximum column length
col_name	Name of column
col_name_len	Length of column name
col_schema_name	Column schema name
col_schema_name_len	Length of column schema name
col_precision	Precision of column
col_scale	Scale of column
col_charsetid	Column character set id
col_charsetform	Column character set form
col_null_ok	NULL column flag; TRUE, if NULL possible

Related Topics

- DESCRIBE_COLUMNS2 Procedure
  This procedure describes the specified column. This is an alternative to DESCRIBE_COLUMNS Procedure.

182.8.3 DBMS_SQL DESC_REC3 Record Type

`DESC_REC3` is the element type of the `DESC_TAB3` table type and the `DESCRIBE_COLUMNS3 Procedure`.

`DESC_REC3` is identical to `DESC_REC2` except for two additional fields to hold the type name (`type_name`) and type name len (`type_name_len`) of a column in a dynamic query. These two fields hold the type name and type name length when the column is a user-defined type (a collection or object type). The `col_type_name` and `col_type_name_len` fields are only populated when the `col_type` field's value is 109, the Oracle type number for user-defined types.

Syntax

```
TYPE desc_rec3 IS RECORD (
   col_type              binary_integer := 0,
   col_max_len           binary_integer := 0,
   col_name              varchar2(32767) := '',
   col_name_len          binary_integer := 0,
   col_schema_name       varchar2(32) := '',
   col_schema_name_len   binary_integer := 0,
   col_precision         binary_integer := 0,
   col_scale             binary_integer := 0,
   col_charsetid         binary_integer := 0,
   col_charsetform       binary_integer := 0,
   col_null_ok           boolean := TRUE,
   col_type_name         varchar2(32767)   := '',
   col_type_name_len     binary_integer := 0);
```

Fields

Table 182-4 DESC_REC3 Fields

Field	Description
col_type	Type of column
col_max_len	Maximum column length
col_name	Name of column
col_name_len	Length of column name
col_schema_name	Column schema name
col_schema_name_len	Length of column schema name
col_precision	Precision of column
col_scale	Scale of column
col_charsetid	Column character set ID
col_charsetform	Column character set form
col_null_ok	NULL column flag; TRUE, if NULL possible

Table 182-4 (Cont.) DESC_REC3 Fields

Field	Description
col_type_name	User-define type column type name, this field is valid when col_type is 109
col_type_name_len	Length of user-define type column type name, this field is valid when col_type is 109

Related Topics

- DESCRIBE_COLUMNS3 Procedure
  This procedure describes the specified column. This is an alternative to DESCRIBE_COLUMNS Procedure.

182.8.4 DBMS_SQL DESC_REC4 Record Type

DESC_REC4 is the element type of the DESC_TAB4 table type and the DESCRIBE_COLUMNS3 Procedure.

DESC_REC4 is identical to DESC_REC3 except that it supports longer identifiers in the fields that hold the schema name (col_schema_name) and type name (col_type_name) of a column in a dynamic query.

Syntax

```
TYPE desc_rec4 IS RECORD (
   col_type              binary_integer := 0,
   col_max_len           binary_integer := 0,
   col_name              varchar2(32767) := '',
   col_name_len          binary_integer := 0,
   col_schema_name       DBMS_ID := '',
   col_schema_name_len   binary_integer := 0,
   col_precision         binary_integer := 0,
   col_scale             binary_integer := 0,
   col_charsetid         binary_integer := 0,
   col_charsetform       binary_integer := 0,
   col_null_ok           boolean := TRUE,
   col_type_name         DBMS_ID   := '',
   col_type_name_len     binary_integer := 0);
```

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information about the predefined subtype DBMS_ID.

Fields

Table 182-5 DESC_REC4 Fields

Field	Description
col_type	Type of column
col_max_len	Maximum column length

Table 182-5 (Cont.) DESC_REC4 Fields

Field	Description
col_name	Name of column
col_name_len	Length of column name
col_schema_name	Column schema name
col_schema_name_len	Length of column schema name
col_precision	Precision of column
col_scale	Scale of column
col_charsetid	Column character set ID
col_charsetform	Column character set form
col_null_ok	NULL column flag; TRUE, if NULL possible
col_type_name	User-define type column type name, this field is valid when col_type is 109
col_type_name_len	Length of user-define type column type name, this field is valid when col_type is 109

Related Topics

- DESCRIBE_COLUMNS3 Procedure
  This procedure describes the specified column. This is an alternative to DESCRIBE_COLUMNS Procedure.

182.8.5 DBMS_SQL BFILE_TABLE Table Type

This is a table of BFILE.

Syntax

```
TYPE bfile_table IS TABLE OF BFILE INDEX BY BINARY_INTEGER;
```

182.8.6 DBMS_SQL BINARY_DOUBLE_TABLE Table Type

This is a table of BINARY_DOUBLE.

Syntax

```
TYPE binary_double_table IS TABLE OF BINARY_DOUBLE INDEX BY BINARY_INTEGER;
```

182.8.7 DBMS_SQL BINARY_FLOAT_TABLE Table Type

This is a table of BINARY_FLOAT.

Syntax

```
TYPE binary_float_table IS TABLE OF BINARY_FLOAT INDEX BY BINARY_INTEGER;
```

182.8.8 DBMS_SQL BLOB_TABLE Table Type

This is a table of `BLOB`.

Syntax

```
TYPE blob_table IS TABLE OF BLOB INDEX BY BINARY_INTEGER;
```

182.8.9 DBMS_SQL BOOLEAN_TABLE Table Type

This is a table of `BOOLEAN`.

Syntax

```
TYPE boolean_table IS TABLE OF BOOLEAN INDEX BY BINARY_INTEGER;
```

182.8.10 DBMS_SQL CLOB_TABLE Table Type

This is a table of `CLOB`.

Syntax

```
TYPE clob_table IS TABLE OF CLOB INDEX BY BINARY_INTEGER;
```

182.8.11 DBMS_SQL DATE_TABLE Table Type

This is a table of `DATE`.

Syntax

```
type date_table IS TABLE OF DATE INDEX BY BINARY_INTEGER;
```

182.8.12 DBMS_SQL DESC_TAB Table Type

This is a table of `DESC_REC Record Type`.

Syntax

```
TYPE desc_tab IS TABLE OF desc_rec INDEX BY BINARY_INTEGER;
```

Related Topics

- DBMS_SQL DESC_REC Record Type
  This record type holds the describe information for a single column in a dynamic query.

182.8.13 DBMS_SQL DESC_TAB2 Table Type

This is a table of `DESC_REC2 Record Type`.

Syntax

```
TYPE desc_tab2 IS TABLE OF desc_rec2 INDEX BY BINARY_INTEGER;
```

Related Topics

- DBMS_SQL DESC_REC2 Record Type
  DESC_REC2 is the element type of the DESC_TAB2 table type and the DESCRIBE_COLUMNS2 Procedure.

182.8.14 DBMS_SQL DESC_TAB3 Table Type

This is a table of DESC_REC3 Record Type.

Syntax

```
TYPE desc_tab3 IS TABLE OF desc_rec3 INDEX BY BINARY_INTEGER;
```

Related Topics

- DBMS_SQL DESC_REC3 Record Type
  DESC_REC3 is the element type of the DESC_TAB3 table type and the DESCRIBE_COLUMNS3 Procedure.

182.8.15 DBMS_SQL DESC_TAB4 Table Type

This is a table of DBMS_SQL DESC_REC4 Record Type.

Syntax

```
TYPE DESC_TAB4 IS TABLE OF DESC_REC4 INDEX BY BINARY_INTEGER;
```

Related Topics

- DBMS_SQL DESC_REC4 Record Type
  DESC_REC4 is the element type of the DESC_TAB4 table type and the DESCRIBE_COLUMNS3 Procedure.

182.8.16 DBMS_SQL INTERVAL_DAY_TO_SECOND_TABLE Table Type

This is a table of DSINTERVAL_UNCONSTRAINED.

Syntax

```
 TYPE interval_day_to_second_Table IS TABLE OF
    DSINTERVAL_UNCONSTRAINED INDEX BY binary_integer;
```

182.8.17 DBMS_SQL INTERVAL_YEAR_TO_MONTH_TABLE Table Type

This is a table of YMINTERVAL_UNCONSTRAINED.

Syntax

```
TYPE interval_year_to_month_table IS TABLE OF YMINTERVAL_UNCONSTRAINED
   INDEX BY BINARY_INTEGER;
```

182.8.18 DBMS_SQL JSON_TABLE Table Type

This is a table of JSON.

Syntax

```
TYPE JSON_TABLE IS TABLE OF JSON INDEX BY BINARY_INTEGER;
```

Related Topics

- BIND_ARRAY Procedures
- COLUMN_VALUE Procedure
- DEFINE_ARRAY Procedure
- DBMS_JSON Constants

182.8.19 DBMS_SQL NUMBER_TABLE Table Type

This is a table of `NUMBER`.

Syntax

```
TYPE number_table IS TABLE OF NUMBER INDEX BY BINARY_INTEGER;
```

182.8.20 DBMS_SQL TIME_TABLE Table Type

This is a table of `TIME_UNCONSTRAINED`.

Syntax

```
TYPE time_table IS TABLE OF TIME_UNCONSTRAINED INDEX BY BINARY_INTEGER;
```

182.8.21 DBMS_SQL TIME_WITH_TIME_ZONE_TABLE Table Type

This is a table of `TIME_TZ_UNCONSTRAINED`.

Syntax

```
TYPE time_with_time_zone_table IS TABLE OF TIME_TZ_UNCONSTRAINED
   INDEX BY BINARY_INTEGER;;
```

182.8.22 DBMS_SQL TIMESTAMP_TABLE Table Type

This is a table of `TIMESTAMP_UNCONSTRAINED`.

Syntax

```
TYPE timestamp_table IS TABLE OF TIMESTAMP_UNCONSTRAINED INDEX BY BINARY_INTEGER;
```

182.8.23 DBMS_SQL TIMESTAMP_WITH_LTZ_TABLE Table Type

This is a table of TIMESTAMP_LTZ_UNCONSTRAINED

Syntax

```
TYPE timestamp_with_ltz_table IS TABLE OF
    TIMESTAMP_LTZ_UNCONSTRAINED INDEX BY binary_integer;
```

182.8.24 DBMS_SQL TIMESTAMP_WITH_TIME_ZONE_TABLE Table Type

This is a table of `TIMESTAMP_TZ_UNCONSTRAINED`.

Syntax

```
TYPE timestamp_with_time_zone_Table IS TABLE OF
    TIMESTAMP_TZ_UNCONSTRAINED INDEX BY binary_integer;
```

182.8.25 DBMS_SQL UROWID_TABLE Table Type

This is a table of `UROWID`.

Syntax

```
TYPE urowid_table IS TABLE OF UROWID INDEX BY BINARY_INTEGER;
```

182.8.26 DBMS_SQL VARCHAR2_TABLE Table Type

This is table of `VARCHAR2(4000)`.

Syntax

```
TYPE varchar2_table IS TABLE OF VARCHAR2(4000) INDEX BY BINARY_INTEGER;
```

182.8.27 DBMS_SQL VARCHAR2A Table Type

This is table of `VARCHAR2(32767)`.

Syntax

```
TYPE varchar2a IS TABLE OF VARCHAR2(32767) INDEX BY BINARY_INTEGER;
```

182.8.28 DBMS_SQL VARCHAR2S Table Type

This is table of `VARCHAR2(256)`.

> **Note:**
>
> This type has been superseded by the VARCHAR2A Table Type. Although it is currently retained for backward compatibility of legacy code, it is in the process of deprecation and will be de-supported in a future release.

Syntax

```
TYPE varchar2s IS TABLE OF VARCHAR2(256) INDEX BY BINARY_INTEGER;
```

182.8.29 DBMS_SQL VECTOR Table Type

This is a table of `VECTOR`.

`DBMS_SQL` bulk operations are supported with this predefined `DBMS_SQL TABLE` type.

Syntax

```
TYPE VECTOR_Table IS TABLE OF VECTOR INDEX BY BINARY_INTEGER;
```

182.9 Summary of DBMS_SQL Subprograms

This table lists the `DBMS_SQL` subprograms and briefly describes them.

Table 182-6 DBMS_SQL Package Subprograms

Subprogram	Description
BIND_ARRAY Procedures	Binds a given value to a given collection.
BIND_VARIABLE Procedures	Binds a given value to a given variable.
BIND_VARIABLE_PKG Procedure	Binds a given value to a given package variable.
CLOSE_CURSOR Procedure	Closes given cursor and frees memory.
COLUMN_VALUE Procedure	Returns value of the cursor element for a given position in a cursor.
COLUMN_VALUE_LONG Procedure	Returns a selected part of a `LONG` column, that has been defined using `DEFINE_COLUMN_LONG`.
DEFINE_ARRAY Procedure	Defines a collection to be selected from the given cursor, used only with `SELECT` statements.
DEFINE_COLUMN Procedures	Defines a column to be selected from the given cursor, used only with `SELECT` statements.
DEFINE_COLUMN_CHAR Procedure	Defines a column of type `CHAR` to be selected from the given cursor, used only with `SELECT` statements.
DEFINE_COLUMN_LONG Procedure	Defines a `LONG` column to be selected from the given cursor, used only with `SELECT` statements.

Table 182-6 (Cont.) DBMS_SQL Package Subprograms

Subprogram	Description
DEFINE_COLUMN_RAW Procedure	Defines a column of type `RAW` to be selected from the given cursor, used only with `SELECT` statements.
DEFINE_COLUMN_ROWID Procedure	Defines a column of type `ROWID` to be selected from the given cursor, used only with `SELECT` statements.
DESCRIBE_COLUMNS Procedure	Describes the columns for a cursor opened and parsed through `DBMS_SQL`.
DESCRIBE_COLUMNS2 Procedure	Describes the specified column, an alternative to DESCRIBE_COLUMNS Procedure.
DESCRIBE_COLUMNS3 Procedure	Describes the specified column, an alternative to DESCRIBE_COLUMNS Procedure.
EXECUTE Function	Executes a given cursor.
EXECUTE_AND_FETCH Function	Executes a given cursor and fetch rows.
FETCH_ROWS Function	Fetches a row from a given cursor.
GET_NEXT_RESULT Procedures	Gets the statement of the next result returned to the caller of the recursive statement or, if this caller sets itself as the client for the recursive statement, the next result returned to this caller as client.
IS_OPEN Function	Returns `TRUE` if given cursor is open.
LAST_ERROR_POSITION Function	Returns byte offset in the SQL statement text where the error occurred.
LAST_ROW_COUNT Function	Returns cumulative count of the number of rows fetched
LAST_ROW_ID Function	Returns `ROWID` of last row processed.
LAST_SQL_FUNCTION_CODE Function	Returns SQL function code for statement.
OPEN_CURSOR Functions	Returns cursor ID number of new cursor.
PARSE Procedures	Parses given statement.
RETURN_RESULT Procedures	Returns the result of an executed statement to the client application.
TO_CURSOR_NUMBER Function	Takes an `OPEN`ed strongly or weakly-typed ref cursor and transforms it into a `DBMS_SQL` cursor number.
TO_REFCURSOR Function	Takes an `OPEN`ed, `PARSE`d, and `EXECUTE`d cursor and transforms/migrates it into a PL/SQL manageable `REF CURSOR` (a weakly-typed cursor) that can be consumed by PL/SQL native dynamic SQL switched to use native dynamic SQL.
VARIABLE_VALUE Procedures	Returns value of named variable for given cursor.
VARIABLE_VALUE_PKG Procedure	Returns value of named variable for given cursor. It is used to return the values of bind variables inside PL/SQL blocks or DML statements with returning clause for a declared package. The type of the variable must be declared in the package specification.

182.9.1 BIND_ARRAY Procedures

This procedure binds a given value or set of values to a given variable in a cursor, based on the name of the variable in the statement.

Syntax

```
DBMS_SQL.BIND_ARRAY (
   c                  IN INTEGER,
   name               IN VARCHAR2,
   <variable>         IN <table_type>
 [,index1             IN INTEGER,
   index2             IN INTEGER)] );
```

Where the <variable> and its corresponding <table_type> can be any one of the following matching pairs, with BIND_ARRAY being overloaded to accept different data types.

```
bdbl_tab     Binary_Double_Table
bf_tab       Bfile_Table
bflt_tab     Binary_Float_Table
bl_tab       Blob_Table
bool_tab     Boolean_Table
c_tab        Varchar2_Table
c_tab        Varchar2A
cl_tab       Clob_Table
d_tab        Date_Table
ids_tab      Interval_Day_To_Second_Table
iym_tab      Interval_Year_To_Month_Table
j_tab        Json_Table
n_tab        Number_Table
tm_tab       Time_Table
tms_tab      Timestamp_Table
tstz_tab     Timestamp_With_ltz_Table
tstz_tab     Timestamp_With_Time_Zone_Table
ttz_tab      Time_With_Time_Zone_Table
ur_tab       Urowid_Table
v_tab        Vector_Table
```

Parameters

Table 182-7 BIND_ARRAY Procedure Parameters

Parameter	Description
c	ID number of the cursor to which you want to bind a value.
name	Name of the collection in the statement.
variable	Local variable that has been declared as <table_type>. The table type can be one of the predefined options or a user defined collection type. For a full list of predefined DBMS_SQL table types for scalar and LOB collections, see DBMS_SQL Data Structures.
index1	Index for the table element that marks the lower bound of the range.
index2	Index for the table element that marks the upper bound of the range.

Usage Notes

For binding a range, the table must contain the elements that specify the range — tab(index1) and tab(index2) — but the range does not have to be dense. Index1 must be less than or equal to index2. All elements between tab(index1) and tab(index2) are used in the bind.

If you do not specify indexes in the bind call, and two different binds in a statement specify tables that contain a different number of elements, then the number of elements actually used is the minimum number between all tables. This is also the case if you specify indexes — the minimum range is selected between the two indexes for all tables.

Not all bind variables in a query have to be array binds. Some can be regular binds and the same value are used for each element of the collections in expression evaluations (and so forth).

Bulk Array Binds

Bulk selects, inserts, updates, and deletes can enhance the performance of applications by bundling many calls into one. The `DBMS_SQL` package lets you work on collections of data using the PL/SQL table type.

Table items are unbounded homogeneous collections. In persistent storage, they are like other relational tables and have no intrinsic ordering. But when a table item is brought into the workspace (either by querying or by navigational access of persistent data), or when it is created as the value of a PL/SQL variable or parameter, its elements are given subscripts that can be used with array-style syntax to get and set the values of elements.

The subscripts of these elements need not be dense, and can be any number including negative numbers. For example, a table item can contain elements at locations -10, 2, and 7 only.

When a table item is moved from transient workspace to persistent storage, the subscripts are not stored; the table item is unordered in persistent storage.

At bind time the table is copied out from the PL/SQL buffers into local `DBMS_SQL` buffers (the same as for all scalar types) and then the table is manipulated from the local `DBMS_SQL` buffers. Therefore, if you change the table after the bind call, then that change does not affect the way the execute acts.

Example 182-1 Use DBMS_SQL.BIND_ARRAY with VECTOR

```
CREATE TABLE vec_seq_table(
  col_vector VECTOR(1, float32),
  col_seq NUMBER
);

SET SERVEROUTPUT ON;
DECLARE
  cur NUMBER;
  stmt_1 VARCHAR2(255) :=
    'INSERT INTO vec_seq_table(col_vector, col_seq) VALUES (:1, :2)';
  number_array DBMS_SQL.NUMBER_TABLE;
  v_array DBMS_SQL.VECTOR_TABLE;
  rowsProcessed NUMBER;
BEGIN
  FOR cnt IN 1 .. 5 LOOP
    v_array(cnt) := TO_VECTOR('[' || cnt || ']', 1, float32);
    number_array(cnt) := cnt;
```

```
  END LOOP;

  cur := DBMS_SQL.OPEN_CURSOR();
  DBMS_SQL.PARSE(cur, stmt_1, DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_ARRAY(cur, ':1', v_array);
  DBMS_SQL.BIND_ARRAY(cur, ':2', number_array);
  rowsProcessed := DBMS_SQL.EXECUTE(cur);
  DBMS_SQL.CLOSE_CURSOR(cur);
  COMMIT;
END;
/

SELECT * FROM vec_seq_table ORDER BY col_seq;
```

Result:

```
COL_VECTOR   COL_SEQ
-----------  -------
[1.0E+000]         1
[2.0E+000]         2
[3.0E+000]         3
[4.0E+000]         4
[5.0E+000]         5
```

Example 182-2 Examples Using Bulk DML

This series of examples shows how to use bulk array binds (table items) in the SQL DML statements `INSERT`, `UPDATE` and `DELETE`.

Here is an example of a bulk `INSERT` statement that demonstrates adding seven new employees to the `emp` table:

```
DECLARE
  stmt VARCHAR2(200);
  empno_array      DBMS_SQL.NUMBER_TABLE;
  empname_array    DBMS_SQL.VARCHAR2_TABLE;
  jobs_array       DBMS_SQL.VARCHAR2_TABLE;
  mgr_array        DBMS_SQL.NUMBER_TABLE;
  hiredate_array   DBMS_SQL.VARCHAR2_TABLE;
  sal_array        DBMS_SQL.NUMBER_TABLE;
  comm_array       DBMS_SQL.NUMBER_TABLE;
  deptno_array     DBMS_SQL.NUMBER_TABLE;
  c                NUMBER;
  dummy            NUMBER;
BEGIN
  empno_array(1):= 9001;
  empno_array(2):= 9002;
  empno_array(3):= 9003;
  empno_array(4):= 9004;
  empno_array(5):= 9005;
  empno_array(6):= 9006;
  empno_array(7):= 9007;

  empname_array(1) := 'Dopey';
  empname_array(2) := 'Grumpy';
  empname_array(3) := 'Doc';
  empname_array(4) := 'Happy';
  empname_array(5) := 'Bashful';
```

```
empname_array(6) := 'Sneezy';
empname_array(7) := 'Sleepy';

jobs_array(1) := 'Miner';
jobs_array(2) := 'Miner';
jobs_array(3) := 'Miner';
jobs_array(4) := 'Miner';
jobs_array(5) := 'Miner';
jobs_array(6) := 'Miner';
jobs_array(7) := 'Miner';

mgr_array(1) := 9003;
mgr_array(2) := 9003;
mgr_array(3) := 9003;
mgr_array(4) := 9003;
mgr_array(5) := 9003;
mgr_array(6) := 9003;
mgr_array(7) := 9003;

hiredate_array(1) := '06-DEC-2006';
hiredate_array(2) := '06-DEC-2006';
hiredate_array(3) := '06-DEC-2006';
hiredate_array(4) := '06-DEC-2006';
hiredate_array(5) := '06-DEC-2006';
hiredate_array(6) := '06-DEC-2006';
hiredate_array(7) := '06-DEC-2006';

sal_array(1):= 1000;
sal_array(2):= 1000;
sal_array(3):= 1000;
sal_array(4):= 1000;
sal_array(5):= 1000;
sal_array(6):= 1000;
sal_array(7):= 1000;

comm_array(1):= 0;
comm_array(2):= 0;
comm_array(3):= 0;
comm_array(4):= 0;
comm_array(5):= 0;
comm_array(6):= 0;
comm_array(7):= 0;

deptno_array(1):= 11;
deptno_array(2):= 11;
deptno_array(3):= 11;
deptno_array(4):= 11;
deptno_array(5):= 11;
deptno_array(6):= 11;
deptno_array(7):= 11;

stmt := 'INSERT INTO emp VALUES(
   :num_array, :name_array, :jobs_array, :mgr_array, :hiredate_array,
   :sal_array, :comm_array, :deptno_array)';
c := DBMS_SQL.OPEN_CURSOR;
DBMS_SQL.PARSE(c, stmt, DBMS_SQL.NATIVE);
DBMS_SQL.BIND_ARRAY(c, ':num_array', empno_array);
DBMS_SQL.BIND_ARRAY(c, ':name_array', empname_array);
DBMS_SQL.BIND_ARRAY(c, ':jobs_array', jobs_array);
DBMS_SQL.BIND_ARRAY(c, ':mgr_array', mgr_array);
DBMS_SQL.BIND_ARRAY(c, ':hiredate_array', hiredate_array);
DBMS_SQL.BIND_ARRAY(c, ':sal_array', sal_array);
```

```
  DBMS_SQL.BIND_ARRAY(c, ':comm_array', comm_array);
  DBMS_SQL.BIND_ARRAY(c, ':deptno_array', deptno_array);

  dummy := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.CLOSE_CURSOR(c);
  EXCEPTION WHEN OTHERS THEN
    IF DBMS_SQL.IS_OPEN(c) THEN
      DBMS_SQL.CLOSE_CURSOR(c);
    END IF;
    RAISE;
END;
/
SHOW ERRORS;
```

Here is an example of a bulk `UPDATE` statement that demonstrates updating salaries for four existing employees in the `emp` table:

```
DECLARE
  stmt VARCHAR2(200);
  empno_array     DBMS_SQL.NUMBER_TABLE;
  salary_array    DBMS_SQL.NUMBER_TABLE;
  c               NUMBER;
  dummy           NUMBER;
BEGIN

  empno_array(1):= 7369;
  empno_array(2):= 7876;
  empno_array(3):= 7900;
  empno_array(4):= 7934;

  salary_array(1) := 10000;
  salary_array(2) := 10000;
  salary_array(3) := 10000;
  salary_array(4) := 10000;

  stmt := 'update emp set sal = :salary_array
    WHERE empno = :num_array';
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, stmt, DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_ARRAY(c, ':num_array', empno_array);
  DBMS_SQL.BIND_ARRAY(c, ':salary_array', salary_array);
  dummy := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.CLOSE_CURSOR(c);

  EXCEPTION WHEN OTHERS THEN
    IF DBMS_SQL.IS_OPEN(c) THEN
      DBMS_SQL.CLOSE_CURSOR(c);
    END IF;
    RAISE;
END;
/
```

In a `DELETE` statement, for example, you could bind an array in the `WHERE` clause and have the statement be run for each element in the array:

```
DECLARE
  stmt VARCHAR2(200);
  dept_no_array DBMS_SQL.NUMBER_TABLE;
  c NUMBER;
  dummy NUMBER;
begin
  dept_no_array(1) := 10; dept_no_array(2) := 20;
```

```
  dept_no_array(3) := 30; dept_no_array(4) := 40;
  dept_no_array(5) := 30; dept_no_array(6) := 40;
  stmt := 'delete from emp where deptno = :dept_array';
  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, stmt, DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_ARRAY(c, ':dept_array', dept_no_array, 1, 4);
  dummy := DBMS_SQL.EXECUTE(c);
  DBMS_SQL.CLOSE_CURSOR(c);

  EXCEPTION WHEN OTHERS THEN
    IF DBMS_SQL.IS_OPEN(c) THEN
      DBMS_SQL.CLOSE_CURSOR(c);
    END IF;
    RAISE;
END;
/
```

In the preceding example, only elements 1 through 4 are used as specified by the `BIND_ARRAY` call. Each element of the array potentially deletes a large number of employees from the database.

182.9.2 BIND_VARIABLE Procedures

These procedures bind a given value or set of values to a given variable in a cursor, based on the name of the variable in the statement.

Syntax

```
DBMS_SQL.BIND_VARIABLE (
   c              IN INTEGER,
   name           IN VARCHAR2,
   value          IN <datatype>);
```

Where <datatype> can be any one of the following types:

```
ADT (user-defined object types)
BINARY_DOUBLE
BINARY_FLOAT
BFILE
BLOB
BOOLEAN
CLOB CHARACTER SET ANY_CS
DATE
DSINTERVAL_UNCONSTRAINED
JSON
NESTED table
NUMBER
OPAQUE types
REF
TIME_UNCONSTRAINED
TIME_TZ_UNCONSTRAINED
TIMESTAMP_LTZ_UNCONSTRAINED
TIMESTAMP_TZ_UNCONSTRAINED
TIMESTAMP_UNCONSTRAINED
UROWID
VARCHAR2 CHARACTER SET ANY_CS
VARRAY
VECTOR
YMINTERVAL_UNCONSTRAINED
```

Notice that `BIND_VARIABLE` is overloaded to accept different data types.

The following syntax is also supported for BIND_VARIABLE. The square brackets [] indicate an optional parameter for the BIND_VARIABLE procedure.

```
DBMS_SQL.BIND_VARIABLE (
   c               IN INTEGER,
   name            IN VARCHAR2,
   value           IN VARCHAR2 CHARACTER SET ANY_CS [,out_value_size IN INTEGER]);
```

To bind CHAR, RAW, and ROWID data, you can use the following variations on the syntax:

```
DBMS_SQL.BIND_VARIABLE_CHAR (
   c               IN INTEGER,
   name            IN VARCHAR2,
   value           IN CHAR CHARACTER SET ANY_CS [,out_value_size IN INTEGER]);

DBMS_SQL.BIND_VARIABLE_RAW (
   c               IN INTEGER,
   name            IN VARCHAR2,
   value           IN RAW [,out_value_size IN INTEGER]);

DBMS_SQL.BIND_VARIABLE_ROWID (
   c               IN INTEGER,
   name            IN VARCHAR2,
   value           IN ROWID);
```

Pragmas

```
pragma restrict_references(bind_variable,WNDS);
```

Parameters

Table 182-8 BIND_VARIABLE Procedures Parameters

Parameter	Description
c	ID number of the cursor to which you want to bind a value.
name	Name of the variable in the statement. The length of the bind variable name must be <=30 bytes.
value	Value that you want to bind to the variable in the cursor. For IN and IN/OUT variables, the value has the same type as the type of the value being passed in for this parameter.
out_value_size	Maximum expected OUT value size, in bytes, for the VARCHAR2, RAW, CHAR OUT or IN/OUT variable. If no size is given, then the length of the current value is used. This parameter must be specified if the value parameter is not initialized.

Usage Notes

If the variable is an IN or IN/OUT variable or an IN collection, then the given bind value must be valid for the variable or array type. Bind values for OUT variables are ignored.

The bind variables or collections of a SQL statement are identified by their names. When binding a value to a bind variable or bind array, the string identifying it in the statement must contain a leading colon, as shown in the following example:

```
SELECT emp_name FROM emp WHERE SAL > :X;
```

For this example, the corresponding bind call would look similar to

```
BIND_VARIABLE(cursor_name, ':X', 3500);

or

BIND_VARIABLE (cursor_name, 'X', 3500);
```

182.9.3 BIND_VARIABLE_PKG Procedure

This procedures binds a variable given value or set of values to a given variable in a cursor, based on the name of the variable in the statement. The type of the variable must be declared in the package specification. Bulk operations are not supported for these types.

Syntax

```
DBMS_SQL.BIND_VARIABLE_PKG (
   c              IN INTEGER,
   name           IN VARCHAR2,
   value          IN <datatype>);
```

Where <datatype> can be any one of the following data types:

- RECORD
- VARRAY
- NESTED TABLE
- INDEX BY PLS_INTEGER TABLE
- INDEX BY BINARY_INTEGER TABLE

Table 182-9 BIND_VARIABLE_PKG Parameters

Parameter	Description
c	ID number of the cursor from which to get the values.
name	Name of the variable in the statement for which you are retrieving the value.
value	• Single row option: Returns the value of the variable for the specified position. Oracle raises the exception ORA-06562, inconsistent_type, if the type of this output parameter differs from the actual type of the value, as defined by the call to BIND_VARIABLE_PKG. • Array option: Local variable that has been declared <table_type>

Example 182-3 Dynamic SQL using DBMS_SQL.BIND_VARIABLE_PKG to Bind a Package Variable

The variables types are declared in the package specification. The BIND_VARIABLE_PKG is used to bind the variable v1 in the cursor SQL statement.

```
CREATE OR REPLACE PACKAGE ty_pkg AS
   TYPE rec IS RECORD ( n1 NUMBER, n2 NUMBER);
   TYPE trec IS TABLE OF REC INDEX BY BINARY_INTEGER;
   TYPE trect IS TABLE OF NUMBER;
```

```
   TYPE trecv IS VARRAY(100) OF NUMBER;
END ty_pkg;
/
CREATE OR REPLACE PROCEDURE dyn_sql_ibbi AS
  dummy NUMBER;
  cur   NUMBER;
  v1 ty_pkg.trec;
  str VARCHAR2(3000);
  n1 NUMBER;
  n2 NUMBER;
BEGIN
  FOR i in 1..3 LOOP
     v1(i).n1 := i*10;
     v1(i).n2 := i*20;
  END LOOP;
  str := 'SELECT * FROM TABLE(:v1)' ;
  cur := DBMS_SQL.OPEN_CURSOR();
  DBMS_SQL.PARSE(cur, str, DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_VARIABLE_PKG(cur, ':v1', v1);
  dummy := DBMS_SQL.EXECUTE(cur);
  DBMS_SQL.DEFINE_COLUMN(cur, 1, n1);
  DBMS_SQL.DEFINE_COLUMN(cur, 2, n2);

  LOOP
    IF DBMS_SQL.FETCH_ROWS(cur) > 0 THEN
      -- get column values of the row
       DBMS_SQL.COLUMN_VALUE(cur, 1, n1);
       DBMS_SQL.COLUMN_VALUE(cur, 2, n2);
       DBMS_OUTPUT.PUT_LINE('n1 = '||n1||' n2 = '||n2);
    ELSE
      -- No more rows
       EXIT;
    END IF;
  END LOOP;
  DBMS_SQL.CLOSE_CURSOR(cur);
END dyn_sql_ibbi;
/
EXEC dyn_sql_ibbi;

n1 = 10 n2 = 20
n1 = 20 n2 = 40
n1 = 30 n2 = 60
```

182.9.4 CLOSE_CURSOR Procedure

This procedure closes a given cursor.

Syntax

```
DBMS_SQL.CLOSE_CURSOR (
   c    IN OUT INTEGER);
```

Pragmas

```
pragma restrict_references(close_cursor,RNDS,WNDS);
```

Parameters

Table 182-10 CLOSE_CURSOR Procedure Parameters

Parameter	Mode	Description
c	IN	ID number of the cursor that you want to close.
c	OUT	Cursor is set to null. After you call CLOSE_CURSOR, the memory allocated to the cursor is released and you can no longer fetch from that cursor.

182.9.5 COLUMN_VALUE Procedure

This procedure returns the value of the cursor element for a given position in a given cursor. This procedure is used to access the data fetched by calling FETCH_ROWS.

Syntax

```
DBMS_SQL.COLUMN_VALUE (
   c                IN  INTEGER,
   position         IN  INTEGER,
   value            OUT <datatype>
 [,column_error     OUT NUMBER]
 [,actual_length    OUT INTEGER]);
```

Where square brackets [ ] indicate optional parameters and <datatype> can be any one of the following types:

```
BINARY_DOUBLE
BINARY_FLOAT
BFILE
BLOB
BOOLEAN
CLOB CHARACTER SET ANY_CS
DATE
DSINTERVAL_UNCONSTRAINED
JSON
NUMBER
TIME_TZ_UNCONSTRAINED
TIME_UNCONSTRAINED
TIMESTAMP_LTZ_UNCONSTRAINED
TIMESTAMP_TZ_UNCONSTRAINED
TIMESTAMP_UNCONSTRAINED
UROWID
VARCHAR2 CHARACTER SET ANY_CS
VECTOR
YMINTERVAL_UNCONSTRAINED
user-defined object types
collections (VARRAYs and nested tables)
REFs
Opaque types
```

For variables containing CHAR, RAW, and ROWID data, you can use the following variations on the syntax:

```
DBMS_SQL.COLUMN_VALUE_CHAR (
   c                IN  INTEGER,
   position         IN  INTEGER,
```

```
   value           OUT CHAR CHARACTER SET ANY_CS
 [,column_error    OUT NUMBER]
 [,actual_length   OUT INTEGER]);

DBMS_SQL.COLUMN_VALUE_RAW (
   c               IN  INTEGER,
   position        IN  INTEGER,
   value           OUT RAW
 [,column_error    OUT NUMBER]
 [,actual_length   OUT INTEGER]);

DBMS_SQL.COLUMN_VALUE_ROWID (
   c               IN  INTEGER,
   position        IN  INTEGER,
   value           OUT ROWID
 [,column_error    OUT NUMBER]
 [,actual_length   OUT INTEGER]);
```

The following syntax enables the `COLUMN_VALUE` procedure to accommodate bulk operations:

```
DBMS_SQL.COLUMN_VALUE(
   c                IN             INTEGER,
   position         IN             INTEGER,
   <param_name>     IN OUT NOCOPY  <table_type>);
```

Where the <`param_name`> and its corresponding <`table_type`> can be any one of these matching pairs:

```
bdbl_tab     Binary_Double_Table
bf_tab       Bfile_Table
bflt_tab     Binary_Float_Table
bl_tab       Blob_Table
bool_tab     Boolean_Table
c_tab        Varchar2_Table
c_tab        Varchar2A
cl_tab       Clob_Table
d_tab        Date_Table
ids_tab      Interval_Day_To_Second_Table
iym_tab      Interval_Year_To_Month_Table
j_tab        Json_table
n_tab        Number_Table
tm_tab       Time_Table
tms_tab      Timestamp_Table
tstz_tab     Timestamp_With_ltz_Table
tstz_tab     Timestamp_With_Time_Zone_Table
ttz_tab      Time_With_Time_Zone_Table
ur_tab       Urowid_Table
v_tab        Vector_Table
```

Pragmas

```
pragma restrict_references(column_value,RNDS,WNDS);
```

Parameters

Table 182-11 COLUMN_VALUE Procedure Parameters (Single Row)

Parameter	Description
`c`	ID number of the cursor from which you are fetching the values.

Table 182-11 (Cont.) COLUMN_VALUE Procedure Parameters (Single Row)

Parameter	Description
position	Relative position of the column in the cursor. The first column in a statement has position 1.
value	Returns the value at the specified column. Oracle raises exception ORA-06562, inconsistent_type, if the type of this output parameter differs from the actual type of the value, as defined by the call to DEFINE_COLUMN.
column_error	Returns any error code for the specified column value.
actual_length	The actual length, before any truncation, of the value in the specified column.

Table 182-12 COLUMN_VALUE Procedure Parameters (Bulk)

Parameter	Description
c	ID number of the cursor from which you are fetching the values.
position	Relative position of the column in the cursor. The first column in a statement has position 1.
<param_name>	Local variable that has been declared <table_type>. <param_name> is an IN OUT NOCOPY parameter for bulk operations. For bulk operations, the subprogram appends the new elements at the appropriate (implicitly maintained) index. For instance if on utilizing the DEFINE_ARRAY Procedure a batch size (the cnt parameter) of 10 rows was specified and a start index (lower_bound) of 1 was specified, then the first call to this subprogram after calling the FETCH_ROWS Function will populate elements at index 1..10, and the next call will populate elements 11..20, and so on.

Exceptions

INCONSISTENT_TYPE (ORA-06562) is raised if the type of the given OUT parameter value is different from the actual type of the value. This type was the given type when the column was defined by calling procedure DEFINE_COLUMN.

Example 182-4 Use COLUMN_VALUE Procedure with the VECTOR Data Type

This example demonstrates using the COLUMN_VALUE procedure along with the procedures DEFINE_ARRAY, the table type VECTOR_Table, and the function FROM_VECTOR to interact with a table with a VECTOR column.

```
DROP TABLE dbmsSqlTable;
CREATE TABLE dbmsSqlTable (embedding VECTOR(3, float32), id NUMBER);
INSERT INTO dbmsSqlTable VALUES ('[1.11, 2.22, 3.33]', 1);
INSERT INTO dbmsSqlTable VALUES ('[4.44, 5.55, 6.66]', 2);
INSERT INTO dbmsSqlTable VALUES ('[7.77, 8.88, 9.99]', 3);

SET SERVEROUTPUT ON;

DECLARE
  cur NUMBER;
  stmt_1 VARCHAR2(255) := 'SELECT embedding FROM dbmsSqlTable ORDER BY id';
  vecArray DBMS_SQL.VECTOR_TABLE;
```

```
  rowsProcessed NUMBER;
BEGIN
  cur := DBMS_SQL.OPEN_CURSOR();
  DBMS_SQL.PARSE(cur, stmt_1, DBMS_SQL.NATIVE);
  DBMS_SQL.DEFINE_ARRAY(cur, 1, vecArray, 3, 1);
  rowsProcessed := DBMS_SQL.EXECUTE_AND_FETCH(cur);

  FOR i IN 1..rowsProcessed LOOP
    DBMS_SQL.COLUMN_VALUE(cur, 1, vecArray);
    DBMS_OUTPUT.PUT_LINE('fetched ID ' || i || ': ' ||
FROM_VECTOR(vecArray(i)));
   END LOOP;
  DBMS_SQL.CLOSE_CURSOR(cur);
END;
/
```

Result:

```
fetched ID 1: [1.11000001E+000,2.22000003E+000,3.32999992E+000]
fetched ID 2: [4.44000006E+000,5.55000019E+000,6.65999985E+000]
fetched ID 3: [7.76999998E+000,8.88000011E+000,9.98999977E+000]
```

182.9.6 COLUMN_VALUE_LONG Procedure

This procedure gets part of the value of a long column.

Syntax

```
DBMS_SQL.COLUMN_VALUE_LONG (
   c            IN  INTEGER,
   position     IN  INTEGER,
   length       IN  INTEGER,
   offset       IN  INTEGER,
   value        OUT VARCHAR2,
   value_length OUT INTEGER);
```

Pragmas

```
pragma restrict_references(column_value_long,RNDS,WNDS);
```

Parameters

Table 182-13 COLUMN_VALUE_LONG Procedure Parameters

Parameter	Description
c	Cursor ID number of the cursor from which to get the value.
position	Position of the column of which to get the value.
length	Number of bytes of the long value to fetch.
offset	Offset into the long field for start of fetch.
value	Value of the column as a VARCHAR2.
value_length	Number of bytes actually returned in value.

182.9.7 DEFINE_ARRAY Procedure

This procedure defines the collection for column into which you want to fetch rows (with a FETCH_ROWS call). This procedure lets you do batch fetching of rows from a single SELECT statement. A single fetch call brings over a number of rows into the PL/SQL aggregate object.

When you fetch the rows, they are copied into DBMS_SQL buffers until you run a COLUMN_VALUE call, at which time the rows are copied into the table that was passed as an argument to the COLUMN_VALUE call.

Syntax

```
DBMS_SQL.DEFINE_ARRAY (
   c                  IN INTEGER,
   position           IN INTEGER,
   <variable>         IN <table_type>
   cnt                IN INTEGER,
   lower_bound        IN INTEGER);
```

Where <variable> and its corresponding <table_type> can be any one of the following matching pairs, with DEFINE_ARRAY being overloaded to accept different data types:

```
bdbl_tab     Binary_Double_Table
bf_tab       Bfile_Table
bflt_tab     Binary_Float_Table
bl_tab       Blob_Table
bool_tab     Boolean_Table
c_tab        Varchar2_Table
c_tab        Varchar2A
cl_tab       Clob_Table
d_tab        Date_Table
ids_tab      Interval_Day_To_Second_Table
iym_tab      Interval_Year_To_Month_Table
j_tab        Json_Table
n_tab        Number_Table
tm_tab       Time_Table
tms_tab      Timestamp_Table
tstz_tab     Timestamp_With_ltz_Table
tstz_tab     Timestamp_With_Time_Zone_Table
ttz_tab      Time_With_Time_Zone_Table
ur_tab       Urowid_Table
v_tab        Vector_Table
```

Pragmas

```
pragma restrict_references(define_array,RNDS,WNDS);
```

The subsequent FETCH_ROWS call fetch "count" rows. When the COLUMN_VALUE call is made, these rows are placed in positions lower_bound, lower_bound+1, lower_bound+2, and so on. While there are still rows coming, the user keeps issuing FETCH_ROWS/COLUMN_VALUE calls. The rows keep accumulating in the table specified as an argument in the COLUMN_VALUE call.

Parameters

Table 182-14 DEFINE_ARRAY Procedure Parameters

Parameter	Description
c	ID number of the cursor to which you want to bind an array.
position	Relative position of the column in the array being defined. The first column in a statement has position 1.
variable	Local variable that has been declared as <table_type>. The table type can be one of the predefined options or a user defined collection type. For a full list of predefined DBMS_SQL table types for scalar and LOB collections, see DBMS_SQL Data Structures.
cnt	Number of rows that must be fetched.
lower_bound	Results are copied into the collection, starting at this lower bound index.

Usage Notes

The count (cnt) must be an integer greater than zero; otherwise an exception is raised. The lower_bound can be positive, negative, or zero. A query on which a DEFINE_ARRAY call was issued cannot contain array binds.

Examples

```
PROCEDURE BULK_PLSQL(deptid NUMBER)
    TYPE namelist IS TABLE OF employees.last_name%TYPE;
    TYPE sallist IS TABLE OF employees.salary%TYPE;
    names    namelist;
    sals     sallist;
    c        NUMBER;
    r        NUMBER;
    sql_stmt VARCHAR2(32767) :=
        'SELECT last_name, salary FROM employees WHERE department_id = :b1';

BEGIN
    c := DBMS_SQL.OPEN_CURSOR;
    DBMS_SQL.PARSE(c, sql_stmt, dbms_sql.native);

    DBMS_SQL.BIND_VARIABLE(c, 'b1', deptid);

    DBMS_SQL.DEFINE_ARRAY(c, 1, names, 5);
    DBMS_SQL.DEFINE_ARRAY(c, 2, sals, 5);

    r := DBMS_SQL.EXECUTE(c);

    LOOP
      r := DBMS_SQL.FETCH_ROWS(c);
      DBMS_SQL.COLUMN_VALUE(c, 1, names);
      DBMS_SQL.COLUMN_VALUE(c, 2, sals);
      EXIT WHEN r != 5;
    END LOOP;

    DBMS_SQL.CLOSE_CURSOR(c);

    -- loop through the names and sals collections
    FOR i IN names.FIRST .. names.LAST  LOOP
      DBMS_OUTPUT.PUT_LINE('Name = ' || names(i) || ', salary = ' || sals(i));
```

```
    END LOOP;
END;
/
```

Example 182-5 Example: Defining an Array

The following examples show how to use the `DEFINE_ARRAY` procedure:

```
declare
  c       NUMBER;
  d       NUMBER;
  n_tab   DBMS_SQL.NUMBER_TABLE;
  indx    NUMBER := -10;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
  dBMS_SQL.PARSE(c, 'select n from t order by 1', DBMS_SQL.NATIVE);

  DBMS_SQL.DEFINE_ARRAY(c, 1, n_tab, 10, indx);

  d := DBMS_SQL.EXECUTE(c);
  loop
    d := DBMS_SQL.FETCH_ROWS(c);

    DBMS_SQL.COLUMN_VALUE(c, 1, n_tab);

    EXIT WHEN d != 10;
  END LOOP;

  DBMS_SQL.CLOSE_CURSOR(c);

  EXCEPTION WHEN OTHERS THEN
    IF DBMS_SQL.IS_OPEN(c) THEN
      DBMS_SQL.CLOSE_CURSOR(c);
    END IF;
    RAISE;
END;
/
```

Each time the preceding example calls FETCH_ROWS Function, it fetches 10 rows that are kept in `DBMS_SQL` buffers. When the COLUMN_VALUE Procedure is called, those rows move into the PL/SQL table specified (in this case `n_tab`), at positions -10 to -1, as specified in the `DEFINE` statements. When the second batch is fetched in the loop, the rows go to positions 0 to 9; and so on.

A current index into each array is maintained automatically. This index is initialized to "indx" at `EXECUTE` time and is updated every time `COLUMN_VALUE` is called. If you reexecute at any point, then the current index for each `DEFINE` is reinitialized to "indx".

In this way the entire result of the query is fetched into the table. When `FETCH_ROWS` cannot fetch 10 rows, it returns the number of rows actually fetched (if no rows could be fetched, then it returns zero) and exits the loop.

Here is another example of using the `DEFINE_ARRAY` procedure:

Consider a table `MULTI_TAB` defined as:

```
CREATE TABLE multi_tab (num NUMBER,
                        dat1 DATE,
                        var VARCHAR2(24),
                        dat2 DATE)
```

To select everything from this table and move it into four PL/SQL tables, you could use the following simple program:

```
DECLARE
  c       NUMBER;
  d       NUMBER;
  n_tab  DBMS_SQL.NUMBER_TABLE;
  d_tab1 DBMS_SQL.DATE_TABLE;
  v_tab  DBMS_SQL.VARCHAR2_TABLE;
  d_tab2 DBMS_SQL.DATE_TABLE;
  indx NUMBER := 10;
BEGIN

  c := DBMS_SQL.OPEN_CURSOR;
  DBMS_SQL.PARSE(c, 'select * from multi_tab order by 1', DBMS_SQL.NATIVE);

  DBMS_SQL.DEFINE_ARRAY(c, 1, n_tab,  5, indx);
  DBMS_SQL.DEFINE_ARRAY(c, 2, d_tab1, 5, indx);
  DBMS_SQL.DEFINE_ARRAY(c, 3, v_tab,  5, indx);
  DBMS_SQL.DEFINE_ARRAY(c, 4, d_tab2, 5, indx);

  d := DBMS_SQL.EXECUTE(c);

  LOOP
    d := DBMS_SQL.FETCH_ROWS(c);

    DBMS_SQL.COLUMN_VALUE(c, 1, n_tab);
    DBMS_SQL.COLUMN_VALUE(c, 2, d_tab1);
    DBMS_SQL.COLUMN_VALUE(c, 3, v_tab);
    DBMS_SQL.COLUMN_VALUE(c, 4, d_tab2);

    EXIT WHEN d != 5;
  END LOOP;

  DBMS_SQL.CLOSE_CURSOR(c);

/*
```

The four tables can be used for anything. One usage might be to use `BIND_ARRAY` to move the rows to another table by using a statement such as '`INSERT` into `SOME_T` values (:a, :b, :c, :d);

```
*/

EXCEPTION WHEN OTHERS THEN
    IF DBMS_SQL.IS_OPEN(c) THEN
      DBMS_SQL.CLOSE_CURSOR(c);
    END IF;
    RAISE;
END;
/
```

182.9.8 DEFINE_COLUMN Procedures

This procedure defines a column to be selected from the given cursor. This procedure is only used with `SELECT` cursors.

The column being defined is identified by its relative position in the `SELECT` list of the statement in the given cursor. The type of the `COLUMN` value determines the type of the column being defined.

See also the DEFINE_COLUMN_CHAR Procedure, DEFINE_COLUMN_LONG Procedure, DEFINE_COLUMN_RAW Procedure and DEFINE_COLUMN_ROWID Procedure.

Syntax

```
DBMS_SQL.DEFINE_COLUMN (
   c              IN INTEGER,
   position       IN INTEGER,
   column         IN <datatype>);
```

Where <datatype> can be any one of the following types:

```
BINARY_DOUBLE
BINARY_FLOAT
BFILE
BLOB
BOOLEAN
CLOB CHARACTER SET ANY_CS
DATE
DSINTERVAL_UNCONSTRAINED
JSON
NUMBER
TIME_UNCONSTRAINED
TIME_TZ_UNCONSTRAINED
TIMESTAMP_LTZ_UNCONSTRAINED
TIMESTAMP_TZ_UNCONSTRAINED
TIMESTAMP_UNCONSTRAINED
UROWID
VECTOR
YMINTERVAL_UNCONSTRAINED
user-defined object types
collections (VARRAYs and nested tables)
REFs
Opaque types
```

Note that DEFINE_COLUMN is overloaded to accept different datatypes.

The following syntax is also supported for the DEFINE_COLUMN procedure:

```
DBMS_SQL.DEFINE_COLUMN (
   c              IN INTEGER,
   position       IN INTEGER,
   column         IN VARCHAR2 CHARACTER SET ANY_CS,
   column_size    IN INTEGER);
```

Pragmas

```
pragma restrict_references(define_column,RNDS,WNDS);
```

Parameters

Table 182-15 DEFINE_COLUMN Procedure Parameters

Parameter	Description
c	ID number of the cursor for the row being defined to be selected.
position	Relative position of the column in the row being defined. The first column in a statement has position 1.
column	Value of the column being defined. The type of this value determines the type for the column being defined.

Table 182-15 (Cont.) DEFINE_COLUMN Procedure Parameters

Parameter	Description
column_size	Maximum expected size of the column value in bytes for columns of type VARCHAR2.

Usage Notes

When using character length semantics the maximum number of bytes that can be returned for a column value of type VARCHAR2 is calculated as: column_size * maximum character byte size for the current character set. For example, specifying the column_size as 10 means that a maximum of 30 (10*3) bytes can be returned when using character length semantics with a UTF8 character set regardless of the number of characters this represents.

182.9.9 DEFINE_COLUMN_CHAR Procedure

This procedure defines a column with CHAR data to be selected from the given cursor. This procedure is only used with SELECT cursors.

The column being defined is identified by its relative position in the SELECT list of the statement in the given cursor. The type of the COLUMN value determines the type of the column being defined.

See also the DEFINE_COLUMN Procedures, DEFINE_COLUMN_LONG Procedure, DEFINE_COLUMN_RAW Procedure and DEFINE_COLUMN_ROWID Procedure.

Syntax

```
DBMS_SQL.DEFINE_COLUMN_CHAR (
   c              IN INTEGER,
   position       IN INTEGER,
   column         IN CHAR CHARACTER SET ANY_CS,
   column_size    IN INTEGER);
```

Pragmas

```
pragma restrict_references(define_column,RNDS,WNDS);
```

Parameters

Table 182-16 DEFINE_COLUMN_CHAR Procedure Parameters

Parameter	Description
c	ID number of the cursor for the row being defined to be selected
position	Relative position of the column in the row being defined.The first column in a statement has position 1.
column	Value of the column being defined. The type of this value determines the type for the column being defined.
column_size	Maximum expected size of the column value in characters for columns of type CHAR.

182.9.10 DEFINE_COLUMN_LONG Procedure

This procedure defines a `LONG` column for a `SELECT` cursor. The column being defined is identified by its relative position in the `SELECT` list of the statement for the given cursor. The type of the `COLUMN` value determines the type of the column being defined.

See also the DEFINE_COLUMN Procedures, DEFINE_COLUMN_CHAR Procedure, DEFINE_COLUMN_RAW Procedure and DEFINE_COLUMN_ROWID Procedure.

Syntax

```
DBMS_SQL.DEFINE_COLUMN_LONG (
   c               IN INTEGER,
   position        IN INTEGER);
```

Parameters

Table 182-17 DEFINE_COLUMN_LONG Procedure Parameters

Parameter	Description
`c`	ID number of the cursor for the row being defined to be selected.
`position`	Relative position of the column in the row being defined. The first column in a statement has position 1.

182.9.11 DEFINE_COLUMN_RAW Procedure

This procedure defines a column of type `RAW` to be selected from the given cursor.

This procedure is only used with `SELECT` cursors.

The column being defined is identified by its relative position in the `SELECT` list of the statement in the given cursor. The type of the `COLUMN` value determines the type of the column being defined.

See also the DEFINE_COLUMN Procedures, DEFINE_COLUMN_CHAR Procedure, DEFINE_COLUMN_LONG Procedure and DEFINE_COLUMN_ROWID Procedure.

Syntax

```
DBMS_SQL.DEFINE_COLUMN_RAW (
   c               IN INTEGER,
   position        IN INTEGER,
   column          IN RAW,
   column_size     IN INTEGER);
```

Pragmas

```
pragma restrict_references(define_column,RNDS,WNDS);
```

Parameters

Table 182-18 DEFINE_COLUMN_RAW Procedure Parameters

Parameter	Description
c	ID number of the cursor for the row being defined to be selected.
position	Relative position of the column in the row being defined.The first column in a statement has position 1.
column	Value of the column being defined. The type of this value determines the type for the column being defined.
column_size	Maximum expected size of the column value in bytes for columns of RAW type.

182.9.12 DEFINE_COLUMN_ROWID Procedure

This procedure defines a column of type ROWID to be selected from the given cursor. This procedure is only used with SELECT cursors.

The column being defined is identified by its relative position in the SELECT list of the statement in the given cursor. The type of the COLUMN value determines the type of the column being defined.

See also the DEFINE_COLUMN Procedures, DEFINE_COLUMN_CHAR Procedure, DEFINE_COLUMN_LONG Procedure and DEFINE_COLUMN_RAW Procedure.

Syntax

```
DBMS_SQL.DEFINE_COLUMN_ROWID (
   c              IN INTEGER,
   position       IN INTEGER,
   column         IN ROWID);
```

Pragmas

```
pragma restrict_references(define_column,RNDS,WNDS);
```

Parameters

Table 182-19 DEFINE_COLUMN_ROWID Procedure Parameters

Parameter	Description
c	ID number of the cursor for the row being defined to be selected
position	Relative position of the column in the row being defined.The first column in a statement has position 1.
column	Value of the column being defined. The type of this value determines the type for the column being defined.

182.9.13 DESCRIBE_COLUMNS Procedure

This procedure describes the columns for a cursor opened and parsed through `DBMS_SQL`.

Syntax

```
DBMS_SQL.DESCRIBE_COLUMNS (
   c             IN  INTEGER,
   col_cnt       OUT INTEGER,
   desc_t        OUT DESC_TAB);
```

Parameters

Table 182-20 DESCRIBE_COLUMNS Procedure Parameters

Parameter	Description
`c`	ID number of the cursor for the columns being described
`col_cnt`	Number of columns in the select list of the query
`desc_t`	Describe table to fill in with the description of each of the columns of the query

Example 182-6 Describe Columns

This code can be used as a substitute to the SQL*Plus `DESCRIBE` call by using a `SELECT` * query on the table that you want to describe.

```
DECLARE
  c           NUMBER;
  d           NUMBER;
  col_cnt     INTEGER;
  f           BOOLEAN;
  rec_tab     DBMS_SQL.DESC_TAB;
  col_num    NUMBER;
  PROCEDURE print_rec(rec in DBMS_SQL.DESC_REC) IS
  BEGIN
    DBMS_OUTPUT.NEW_LINE;
    DBMS_OUTPUT.PUT_LINE('col_type            =    ' || rec.col_type);
    DBMS_OUTPUT.PUT_LINE('col_maxlen          =    ' || rec.col_max_len);
    DBMS_OUTPUT.PUT_LINE('col_name            =    ' || rec.col_name);
    DBMS_OUTPUT.PUT_LINE('col_name_len        =    ' || rec.col_name_len);
    DBMS_OUTPUT.PUT_LINE('col_schema_name     =    ' || rec.col_schema_name);
    DBMS_OUTPUT.PUT_LINE('col_schema_name_len =    ' ||
rec.col_schema_name_len);
    DBMS_OUTPUT.PUT_LINE('col_precision       =    ' || rec.col_precision);
    DBMS_OUTPUT.PUT_LINE('col_scale           =    ' || rec.col_scale);
    DBMS_OUTPUT.PUT('col_null_ok         =    ');
    IF (rec.col_null_ok) THEN
      DBMS_OUTPUT.PUT_LINE('true');
    ELSE
      DBMS_OUTPUT.PUT_LINE('false');
    END IF;
  END;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR;
```

```
  DBMS_SQL.PARSE(c, 'SELECT * FROM scott.bonus', DBMS_SQL.NATIVE);

  d := DBMS_SQL.EXECUTE(c);

  DBMS_SQL.DESCRIBE_COLUMNS(c, col_cnt, rec_tab);

/*
 * Following loop could simply be for j in 1..col_cnt loop.
 * Here we are simply illustrating some of the PL/SQL table
 * features.
 */
  col_num := rec_tab.first;
  IF (col_num IS NOT NULL) THEN
    LOOP
      print_rec(rec_tab(col_num));
      col_num := rec_tab.next(col_num);
      EXIT WHEN (col_num IS NULL);
    END LOOP;
  END IF;

  DBMS_SQL.CLOSE_CURSOR(c);
END;
/
```

182.9.14 DESCRIBE_COLUMNS2 Procedure

This procedure describes the specified column. This is an alternative to DESCRIBE_COLUMNS Procedure.

Syntax

```
DBMS_SQL.DESCRIBE_COLUMNS2 (
   c             IN  INTEGER,
   col_cnt       OUT INTEGER,
   desc_t        OUT DESC_TAB2);
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(describe_columns2,WNDS);
```

Parameters

Table 182-21 DESCRIBE_COLUMNS2 Procedure Parameters

Parameter	Description
c	ID number of the cursor for the columns being described.
col_cnt	Number of columns in the select list of the query.
desc_t	Describe table to fill in with the description of each of the columns of the query. This table is indexed from one to the number of elements in the select list of the query.

Related Topics

- DESCRIBE_COLUMNS Procedure
  This procedure describes the columns for a cursor opened and parsed through DBMS_SQL.

182.9.15 DESCRIBE_COLUMNS3 Procedure

This procedure describes the specified column. This is an alternative to DESCRIBE_COLUMNS Procedure.

Syntax

```
DBMS_SQL.DESCRIBE_COLUMNS3 (
   c              IN  INTEGER,
   col_cnt        OUT INTEGER,
   desc_t         OUT DESC_TAB3);

BMS_SQL.DESCRIBE_COLUMNS3 (
   c              IN  INTEGER,
   col_cnt        OUT INTEGER,
   desc_t         OUT DESC_TAB4);
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES(describe_columns3,WNDS);
```

Parameters

Table 182-22 DESCRIBE_COLUMNS3 Procedure Parameters

Parameter	Description
c	ID number of the cursor for the columns being described.
col_cnt	Number of columns in the select list of the query.
desc_t	Describe table to fill in with the description of each of the columns of the query. This table is indexed from one to the number of elements in the select list of the query.

Usage Notes

The cursor passed in by the cursor ID has to be `OPEN`ed and `PARSE`d, otherwise an "invalid cursor id" error is raised.

Examples

```
CREATE TYPE PROJECT_T AS OBJECT
      ( projname          VARCHAR2(20),
        mgr               VARCHAR2(20))
/

CREATE TABLE projecttab(deptno NUMBER, project HR.PROJECT_T)
/

DECLARE
  curid      NUMBER;
  desctab    DBMS_SQL.DESC_TAB3;
  colcnt     NUMBER;
  sql_stmt   VARCHAR2(200) := 'select * from projecttab';
BEGIN

    curid := DBMS_SQL.OPEN_CURSOR;

    DBMS_SQL.PARSE(curid, sql_stmt, DBMS_SQL.NATIVE);
```

```
    DBMS_SQL.DESCRIBE_COLUMNS3(curid, colcnt, desctab);

    FOR i IN 1 .. colcnt LOOP
      IF desctab(i).col_type = 109 THEN
        DBMS_OUTPUT.PUT(desctab(i).col_name || ' is user-defined type: ');
        DBMS_OUTPUT.PUT_LINE(desctab(i).col_schema_name || '.' ||
                             desctab(i).col_type_name);
      END IF;
    END LOOP;

    DBMS_SQL.CLOSE_CURSOR(curid);
END;
/

Output:

PROJECT is user-defined type: HR.PROJECT_T
```

Related Topics

- DESCRIBE_COLUMNS Procedure
  This procedure describes the columns for a cursor opened and parsed through `DBMS_SQL`.

182.9.16 EXECUTE Function

This function executes a given cursor. This function accepts the `ID` number of the cursor and returns the number of rows processed.

The return value is only valid for `INSERT`, `UPDATE`, and `DELETE` statements; for other types of statements, including DDL, the return value is undefined and must be ignored.

Syntax

```
DBMS_SQL.EXECUTE (
   c    IN INTEGER)
  RETURN INTEGER;
```

Parameters

Table 182-23 EXECUTE Function Parameters

Parameter	Description
`c`	Cursor ID number of the cursor to execute.

Return Values

Returns number of rows processed

Usage Notes

The `DBMS_SQL` cursor that is returned by the TO_CURSOR_NUMBER Function performs in the same way as a `DBMS_SQL` cursor that has already been executed. Consequently, calling `EXECUTE` for this cursor will cause an error.

182.9.17 EXECUTE_AND_FETCH Function

This function executes the given cursor and fetches rows.

This function provides the same functionality as calling `EXECUTE` and then calling `FETCH_ROWS`. Calling `EXECUTE_AND_FETCH` instead, however, may reduce the number of network round-trips when used against a remote database.

The `EXECUTE_AND_FETCH` function returns the number of rows actually fetched.

Syntax

```
DBMS_SQL.EXECUTE_AND_FETCH (
   c              IN INTEGER,
   exact          IN BOOLEAN DEFAULT FALSE)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(execute_and_fetch,WNDS);
```

Parameters

Table 182-24 EXECUTE_AND_FETCH Function Parameters

Parameter	Description
`c`	ID number of the cursor to execute and fetch.
`exact`	Set to `TRUE` to raise an exception if the number of rows actually matching the query differs from one. Note: Oracle does not support the exact fetch TRUE option with LONG columns. Even if an exception is raised, the rows are still fetched and available.

Return Values

Returns designated rows

182.9.18 FETCH_ROWS Function

This function fetches a row from a given cursor.

You can call `FETCH_ROWS` repeatedly as long as there are rows remaining to be fetched. These rows are retrieved into a buffer, and must be read by calling `COLUMN_VALUE`, for each column, after each call to `FETCH_ROWS`.

The `FETCH_ROWS` function accepts the ID number of the cursor to fetch, and returns the number of rows actually fetched.

Syntax

```
DBMS_SQL.FETCH_ROWS (
   c              IN INTEGER)
  RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(fetch_rows,WNDS);
```

Parameters

Table 182-25 FETCH_ROWS Function Parameters

Parameter	Description
c	ID number.

Return Values

Returns a row from a given cursor

182.9.19 GET_NEXT_RESULT Procedures

This procedure gets the statement of the next result returned to the caller of the recursive statement or, if this caller sets itself as the client for the recursive statement, the next result returned to this caller as client.

The statements are returned in same order as they are returned by the RETURN_RESULT Procedures.

Syntax

```
DBMS_SQL.GET_NEXT_RESULT(
   c             IN          INTEGER,
   rc            OUT         SYS_REFCURSOR);

DBMS_SQL.GET_NEXT_RESULT(
   c             IN          INTEGER,
   rc            OUT         INTEGER);
```

Parameters

Table 182-26 GET_NEXT_RESULT Procedure Parameters

Parameter	Description
c	Recursive statement cursor
rc	Cursor or ref cursor of the statement of the next returned result

Exceptions

`ORA-01403 no_data_found`: This is raised when there is no further returned statement result.

Usage Notes

- After the cursor of a statement result is retrieved, the caller must close the cursor properly when it is no longer needed.
- The cursors for all unretrieved returned statements will be closed after the cursor of the recursive statement is closed.

Examples

```
DECLARE
  c  INTEGER;
  rc SYS_REFCURSOR;
BEGIN
  c := DBMS_SQL.OPEN_CURSOR(treat_as_client_for_results => TRUE);
  DBMS_SQL.PARSE(c                  => c,
                 statement          => 'begin proc; end;');
  DBMS_SQL.EXECUTE(c);
  LOOP
    BEGIN
      DBMS_SQL.GET_NEXT_RESULT(c, rc);
    EXCEPTIONS
      WHEN no_data_found THEN
        EXIT;
    END;
    LOOP
      FETCH rc INTO ...
      ...
    END LOOP;
  END LOOP;
END;
```

182.9.20 IS_OPEN Function

This function checks to see if the given cursor is currently open.

Syntax

```
DBMS_SQL.IS_OPEN (
   c              IN INTEGER)
  RETURN BOOLEAN;
```

Pragmas

```
pragma restrict_references(is_open,RNDS,WNDS);
```

Parameters

Table 182-27 IS_OPEN Function Parameters

Parameter	Description
c	Cursor ID number of the cursor to check.

Return Values

Returns `TRUE` for any cursor number that has been opened but not closed, and `FALSE` for a `NULL` cursor number. Note that the CLOSE_CURSOR Procedure Procedure `NULL`s out the cursor variable passed to it.

Exceptions

`ORA-29471 DBMS_SQL access denied`: This is raised if an invalid cursor ID number is detected. Once a session has encountered and reported this error, every subsequent DBMS_SQL call in the same session will raise this error, meaning that DBMS_SQL is non-operational for this session.

182.9.21 LAST_ERROR_POSITION Function

This function returns the byte offset in the SQL statement text where the error occurred. The first character in the SQL statement is at position 0.

Syntax

```
DBMS_SQL.LAST_ERROR_POSITION
   RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(last_error_position,RNDS,WNDS);
```

Return Values

Returns the byte offset in the SQL statement text where the error occurred

Usage Notes

Call this function after a `PARSE` call, before any other `DBMS_SQL` procedures or functions are called.

182.9.22 LAST_ROW_COUNT Function

This function returns the cumulative count of the number of rows fetched.

Syntax

```
DBMS_SQL.LAST_ROW_COUNT
   RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(last_row_count,RNDS,WNDS);
```

Return Values

Returns the cumulative count of the number of rows fetched

Usage Notes

Call this function after a `FETCH_ROWS` or an `EXECUTE_AND_FETCH` call. If called after an `EXECUTE` call, then the value returned is zero.

182.9.23 LAST_ROW_ID Function

This function returns the `ROWID` of the last row processed.

Syntax

```
DBMS_SQL.LAST_ROW_ID
   RETURN ROWID;
```

Pragmas

```
pragma restrict_references(last_row_id,RNDS,WNDS);
```

Return Values

Returns the `ROWID` of the last row processed

Usage Notes

Call this function after a `FETCH_ROWS` or an `EXECUTE_AND_FETCH` call.

182.9.24 LAST_SQL_FUNCTION_CODE Function

This function returns the SQL function code for the statement.

These codes are listed in the *Oracle Call Interface Programmer's Guide.*

Syntax

```
DBMS_SQL.LAST_SQL_FUNCTION_CODE
   RETURN INTEGER;
```

Pragmas

```
pragma restrict_references(last_sql_function_code,RNDS,WNDS);
```

Return Values

Returns the SQL function code for the statement

Usage Notes

You must call this function immediately after the SQL statement is run; otherwise, the return value is undefined.

182.9.25 OPEN_CURSOR Functions

This function opens a new cursor.

The `security_level` parameter allows for application of fine-grained control to the security of the opened cursor.

Syntax

```
DBMS_SQL.OPEN_CURSOR (
   treat_as_client_for_results    IN     BOOLEAN    DEFAULT FALSE)
  RETURN INTEGER;

DBMS_SQL.OPEN_CURSOR (
   security_level                 IN     INTEGER,
   treat_as_client_for_results    IN     BOOLEAN    DEFAULT FALSE)
  RETURN INTEGER;
```

Parameters

Table 182-28 OPEN_CURSOR Function Parameters

Parameter	Description
`security_level`	Specifies the level of security protection to enforce on the opened cursor. Valid security level values are `0`, `1`, and `2`. When a `NULL` argument value is provided to this overload, as well as for cursors opened using the overload of open_cursor without the `security_level` parameter, the default security level value `1` will be enforced on the opened cursor. • `Level 0` - allows all `DBMS_SQL` operations on the cursor without any security checks. The cursor may be fetched from, and even re-bound and re-executed, by code running with a different effective userid or roles than those in effect at the time the cursor was parsed. This level of security is off by default. • `Level 1` - requires that the referenced container, effective userid, and roles of the caller to `DBMS_SQL` for bind and execute operations on this cursor must be the same as those of the caller of the most recent parse operation on this cursor. • `Level 2` - requires that the referenced container, effective userid, and roles of the caller to `DBMS_SQL` for all bind, execute, define, describe, and fetch operations on this cursor must be the same as those of the caller of the most recent parse operation on this cursor.
`treat_as_client_for_results`	Allows the caller of the recursive statement to set itself as the client to receive the statement results returned from the recursive statement to client. The statement results returned may be retrieved by the GET_NEXT_RESULT Procedures.

Pragmas

```
pragma restrict_references(open_cursor,RNDS,WNDS);
```

Return Values

Returns the cursor ID number of the new cursor

Usage Notes

- When you no longer need this cursor, you must close it explicitly by calling the CLOSE_CURSOR Procedure.
- You can use cursors to run the same SQL statement repeatedly or to run a new SQL statement. When a cursor is reused, the contents of the corresponding cursor data area are reset when the new SQL statement is parsed. It is never necessary to close and reopen a cursor before reusing it.

182.9.26 PARSE Procedures

This procedure parses the given statement in the given cursor. All statements are parsed immediately. In addition, DDL statements are run immediately when parsed.

There are multiple versions of the `PARSE` procedure:

- Taking a `VARCHAR2` statement as an argument

- Taking a segmented string, one taking `VARCHAR2A`, a `TABLE OF VARCHAR2(32767)`, and another, taking `VARCHAR2S`, a `TABLE OF VARCHAR2(256)`, as argument. These overloads concatenate elements of a PL/SQL table statement and parse the resulting string. You can use these procedures to parse a statement that is longer than the limit for a single `VARCHAR2` variable by splitting up the statement.
- Taking a `CLOB` statement as an argument. You can use the `CLOB` overload version of the parse procedure to parse a SQL statement larger than 32K bytes.

Syntax

Each version has multiple overloads.

```
DBMS_SQL.PARSE (
   c                            IN   INTEGER,
   statement                    IN   VARCHAR2,
   language_flag                IN   INTEGER[
 [,edition                      IN   VARCHAR2 DEFAULT NULL],
   apply_crossedition_trigger   IN   VARCHAR2 DEFAULT NULL,
   fire_apply_trigger           IN   BOOLEAN DEFAULT TRUE]
 [,schema                       IN   VARCHAR2 DEFAULT NULL]
 [,container                    IN   VARCHAR2)];

DBMS_SQL.PARSE (
   c                            IN   INTEGER,
   statement                    IN   CLOB,
   language_flag                IN   INTEGER[
 [,edition                      IN   VARCHAR2 DEFAULT NULL],
   apply_crossedition_trigger   IN   VARCHAR2 DEFAULT NULL,
   fire_apply_trigger           IN   BOOLEAN DEFAULT TRUE]
 [,schema                       IN   VARCHAR2 DEFAULT NULL]
 [,container                    IN   VARCHAR2)];

DBMS_SQL.PARSE (
   c                            IN   INTEGER,
   statement                    IN   VARCHAR2A,
   lb                           IN   INTEGER,
   ub                           IN   INTEGER,
   lfflg                        IN   BOOLEAN,
   language_flag                IN   INTEGER[
 [,edition                      IN   VARCHAR2 DEFAULT NULL],
   apply_crossedition_trigger   IN   VARCHAR2 DEFAULT NULL,
   fire_apply_trigger           IN   BOOLEAN DEFAULT TRUE]
 [,schema                       IN   VARCHAR2 DEFAULT NULL]
 [,container                    IN   VARCHAR2)];

DBMS_SQL.PARSE (
   c                            IN   INTEGER,
   statement                    IN   VARCHAR2s,
   lb                           IN   INTEGER,
   ub                           IN   INTEGER,
   lfflg                        IN   BOOLEAN,
   language_flag                IN   INTEGER[
 [,edition                      IN   VARCHAR2 DEFAULT NULL],
   apply_crossedition_trigger   IN   VARCHAR2 DEFAULT NULL,
   fire_apply_trigger           IN   BOOLEAN DEFAULT TRUE]
 [,schema                       IN   VARCHAR2 DEFAULT NULL]
 [,container                    IN   VARCHAR2)];
```

Parameters

Table 182-29 PARSE Procedure Parameters

Parameter	Description
c	ID number of the cursor in which to parse the statement.
statement	SQL statement to be parsed. SQL statements larger than 32K that may be stored in CLOBs. Unlike a PL/SQL statement, your SQL statement must not include a final semicolon. For example: DBMS_SQL.PARSE(cursor1, 'BEGIN proc; END;', 2); DBMS_SQL.PARSE(cursor1, 'INSERT INTO tab VALUES(1)', 2);
lb	Lower bound for elements in the statement
ub	Upper bound for elements in the statement
lfflg	If TRUE, then insert a linefeed after each element on concatenation.
language_flag	Specifies the behavior for the SQL statement. For more information about the possible values and its corresponding behaviors, see DBMS_SQL Constants
edition	Specifies the edition in which to run the statement under the following conditions: • If NULL and container is NULL, the statement will be run in the current edition. • If a valid container is specified, passing NULL indicates the statement is to run in the target container's default edition. • Given the user and the edition with which the statement is to be executed, the user must have USE privilege on the edition. The following general conditions apply. The contents of the string are processed as a SQL identifier; double quotation marks must surround the remainder of the string if special characters or lowercase characters are present in the edition's actual name, and if double quotation marks are not used the contents will be uppercased.
apply_crossedition_trigger	Specifies the unqualified name of a forward crossedition trigger that is to be applied to the specified SQL. The name is resolved using the edition and current_schema setting in which the statement is to be executed. The trigger must be owned by the user that will execute the statement.If a non-NULL value is specified, the specified crossedition trigger will be executed assuming fire_apply_trigger is TRUE, the trigger is enabled, the trigger is defined on the table which is the target of the statement, the type of the statement matches the trigger's dml_event_clause, any effective WHEN and UPDATE OF restrictions are satisfied, and so on. Other forward crossedition triggers may also be executed, selected using the "crossedition trigger DML rules" applied as if the specified trigger was doing a further DML to the table that is the target of the statement. Non-crossedition triggers and reverse crossedition triggers will not be executed.The contents of the string are processed as a SQL identifier; double quotation marks must surround the remainder of the string if special characters or lowercase characters are present in the trigger's actual name, and if double quotation marks are not used, the contents will be uppercased.

Table 182-29 (Cont.) PARSE Procedure Parameters

Parameter	Description
fire_apply_trigger	Indicates whether the specified apply_crossedition_trigger is itself to be executed, or must only be a guide used in selecting other triggers. This is typically set FALSE when the statement is a replacement for the actions the apply_crossedition_trigger would itself perform. If FALSE, the specified trigger is not executed, but other triggers are still selected for firing as if the specified trigger was doing a DML to the table that is the target of the statement. The apply_crossedition_trigger and fire_apply_trigger parameters are ignored if the statement is not a DML.
schema	Specifies the schema in which to resolve unqualified object names. If NULL, the current schema is the effective user's schema.
container	Name of the target container in which the cursor is to run. If NULL or unspecified, the name of the target container is that of the calling container and no container switch is performed. If a valid container name is specified, the current user must be a common user with SET CONTAINER privilege to switch to the target container. If a container switch completes, the effective user will have its default roles.

Usage Notes

- Using DBMS_SQL to dynamically run DDL statements can cause the program to stop responding. For example, a call to a procedure in a package results in the package being locked until the execution returns to the user side. Any operation that results in a conflicting lock, such as dynamically trying to drop the package before the first lock is released, stops the program from running.
- Because client-side code cannot reference remote package variables or constants, you must explicitly use the values of the constants.

  For example, the following code does *not* compile on the client:

  ```
  DBMS_SQL.PARSE(cur_hdl, stmt_str, DBMS_SQL.NATIVE); -- uses constant DBMS_SQL.NATIVE
  ```

  The following code works on the client, because the argument is explicitly provided:

  ```
  DBMS_SQL.PARSE(cur_hdl, stmt_str, 1); -- compiles on the client
  ```
- The VARCHAR2S type is currently supported for backward compatibility of legacy code. However, you are advised to use VARCHAR2A both for its superior capability and because VARCHAR2S will be deprecated in a future release.
- To parse SQL statements larger than 32 KB, the new CLOB overload version of the PARSE procedure can be used instead of the VARCHAR2A overload.
- If the container parameter value is the same as the calling container, a container switch will not occur. However, the default roles of the current user will be in effect.

Exceptions

If you create a type, procedure, function, or package using DBMS_SQL that has compilation warnings, an ORA-24344 exception is raised, and the PL/SQL unit is still created.

182.9.27 RETURN_RESULT Procedures

This procedure returns the result of an executed statement to the client application.

The result can be retrieved later by the client. Alternatively, it can return the statement result to and be retrieved later by the immediate caller that executes a recursive statement in which this statement result will be returned.

The caller can be:

- A PL/SQL stored procedure executing the recursive statement using `DBMS_SQL`
- A Java stored procedure using JDBC
- A `.NET` stored procedure using `ADO.NET`
- An external procedure using the Oracle Call Interface (OCI)

Syntax

```
DBMS_SQL.RETURN_RESULT(
   rc           IN OUT      SYS_REFCURSOR,
   to_client    IN          BOOLEAN           DEFAULT TRUE);

DBMS_SQL.RETURN_RESULT(
   rc           IN OUT      INTEGER,
   to_client    IN          BOOLEAN           DEFAULT TRUE);
```

Parameters

Table 182-30 RETURN_RESULT Procedure Parameters

Parameter	Description
`rc`	Statement cursor or ref cursor
`to_client`	Returns (or does not return) the statement result to the client. If not, it is returned to the immediate caller.

Usage Notes

- Currently only a SQL query can be returned, and the return of statement results over remote procedure calls is not supported.
- Once the statement is returned, it is no longer accessible except by the client or the immediate caller to which it is returned.
- Statement results cannot be returned when the statement being executed by the client or any intermediate recursive statement is a SQL query and an error is raised.
- A ref cursor being returned can be strongly or weakly-typed.
- A query being returned can be partially fetched.
- Because `EXECUTE IMMEDIATE` statement provides no interface to retrieve the statement results returned from its recursive statement, the cursors of the statement results returned to the caller of the `EXECUTE IMMEDIATE` statement will be closed when the statement completes. To retrieve the returned statement results from a recursive statement in PL/SQL, use `DBMS_SQL` to execute the recursive statement.

Examples

```
CREATE PROCEDURE proc AS
  rc1 sys_refcursor;
  rc2 sys_refcursor;
BEGIN
  OPEN rc1 FOR SELECT * FROM t1;
  DBMS_SQL.RETURN_RESULT(rc1);
  OPEN rc2 FOR SELECT * FROM t2;
  DBMS_SQL.RETURN_RESULT(rc2);
END;
/
```

182.9.28 TO_CURSOR_NUMBER Function

This function takes an OPENed strongly or weakly-typed ref cursor and transforms it into a DBMS_SQL cursor number.

Syntax

```
DBMS_SQL.TO_CURSOR_NUMBER(
   rc IN OUT SYS_REFCURSOR)
  RETURN INTEGER;
```

Parameters

Table 182-31 TO_CURSOR_NUMBER Function Parameters

Parameter	Description
rc	REF CURSOR to be transformed into a cursor number

Return Values

Returns a DBMS_SQL manageable cursor number transformed from a REF CURSOR

Usage Notes

- The REF CURSOR passed in has to be OPENed, otherwise an error is raised.
- Once the REF CURSOR is transformed into a DBMS_SQL cursor number, the REF CURSOR is no longer accessible by any native dynamic SQL operations.
- The DBMS_SQL cursor that is returned by this subprogram performs in the same way as a DBMS_SQL cursor that has already been executed.

Examples

```
CREATE OR REPLACE PROCEDURE DO_QUERY(sql_stmt VARCHAR2) IS
  TYPE CurType IS REF CURSOR;
  src_cur         CurType;
  curid           NUMBER;
  desctab         DBMS_SQL.DESC_TAB;
  colcnt          NUMBER;
  namevar         VARCHAR2(50);
  numvar          NUMBER;
  datevar         DATE;
  empno           NUMBER := 100;
BEGIN
```

```
    -- sql_stmt := 'select ...... from employees where employee_id = :b1';
    OPEN src_cur FOR sql_stmt USING empno;

    -- Switch from native dynamic SQL to DBMS_SQL
    curid := DBMS_SQL.TO_CURSOR_NUMBER (src_cur);

    DBMS_SQL.DESCRIBE_COLUMNS(curid, colcnt, desctab);

    -- Define columns
    FOR i IN 1 .. colcnt LOOP
        IF desctab(i).col_type = 2 THEN
           DBMS_SQL.DEFINE_COLUMN(curid, i, numvar);
        ELSIF desctab(i).col_type = 12 THEN
            DBMS_SQL.DEFINE_COLUMN(curid, i, datevar);
.......
         ELSE
           DBMS_SQL.DEFINE_COLUMN(curid, i, namevar, 25);
         END IF;
    END LOOP;

  -- Fetch Rows
    WHILE DBMS_SQL.FETCH_ROWS(curid) > 0 LOOP
        FOR i IN 1 .. colcnt LOOP
          IF (desctab(i).col_type = 1) THEN
            DBMS_SQL.COLUMN_VALUE(curid, i, namevar);
         ELSIF (desctab(i).col_type = 2) THEN
            DBMS_SQL.COLUMN_VALUE(curid, i, numvar);
          ELSIF (desctab(i).col_type = 12) THEN
            DBMS_SQL.COLUMN_VALUE(curid, i, datevar);
....
          END IF;
        END LOOP;
    END LOOP;

    DBMS_SQL.CLOSE_CURSOR(curid);
END;
/
```

182.9.29 TO_REFCURSOR Function

This function takes an `OPEN`ed, `PARSE`d, and `EXECUTE`d cursor and transforms/migrates it into a PL/SQL manageable `REF CURSOR` (a weakly-typed cursor) that can be consumed by PL/SQL native dynamic SQL switched to use native dynamic SQL.

This subprogram is only used with `SELECT` cursors.

Syntax

```
DBMS_SQL.TO_REFCURSOR(
   cursor_number IN OUT INTEGER)
  RETURN SYS_REFCURSOR;
```

Parameters

Table 182-32 TO_REFCURSOR Function Parameters

Parameter	Description
cursor_number	Cursor number of the cursor to be transformed into REF CURSOR

Return Values

Returns a PL/SQL REF CURSOR transformed from a DBMS_SQL cursor number

Usage Notes

- The cursor passed in by the cursor_number has to be OPENed, PARSEd, and EXECUTEd; otherwise an error is raised.
- Once the cursor_number is transformed into a REF CURSOR, the cursor_number is no longer accessible by any DBMS_SQL operations.
- After a cursor_number is transformed into a REF CURSOR, using DBMS_SQL.IS_OPEN to check to see if the cursor_number is still open results in an error.
- If the cursor number was last parsed with a valid container parameter, it cannot be converted to a REF CURSOR.

Examples

```
CREATE OR REPLACE PROCEDURE DO_QUERY(mgr_id NUMBER) IS
  TYPE CurType IS REF CURSOR;
  src_cur         CurType;
  curid           NUMBER;
  sql_stmt        VARCHAR2(200);
  ret             INTEGER;
  empnos          DBMS_SQL.Number_Table;
  depts           DBMS_SQL.Number_Table;
BEGIN

  -- DBMS_SQL.OPEN_CURSOR
  curid := DBMS_SQL.OPEN_CURSOR;

  sql_stmt :=    'SELECT EMPLOYEE_ID, DEPARTMENT_ID from employees where MANAGER_ID
= :b1';

  DBMS_SQL.PARSE(curid, sql_stmt, DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_VARIABLE(curid, 'b1', mgr_id);
  ret := DBMS_SQL.EXECUTE(curid);

  -- Switch from DBMS_SQL to native dynamic SQL
  src_cur := DBMS_SQL.TO_REFCURSOR(curid);

  -- Fetch with native dynamic SQL
  FETCH src_cur BULK COLLECT INTO empnos, depts;

  IF empnos.COUNT > 0 THEN
    DBMS_OUTPUT.PUT_LINE('EMPNO DEPTNO');
    DBMS_OUTPUT.PUT_LINE('----- ------');
    -- Loop through the empnos and depts collections
    FOR i IN 1 .. empnos.COUNT LOOP
      DBMS_OUTPUT.PUT_LINE(empnos(i) || '   ' || depts(i));
```

```
    END LOOP;
  END IF;
   -- Close cursor
  CLOSE src_cur;
END;
/
```

182.9.30 VARIABLE_VALUE Procedures

This procedure returns the value of the named variable for a given cursor. It is used to return the values of bind variables inside PL/SQL blocks or DML statements with `returning` clause.

Syntax

```
DBMS_SQL.VARIABLE_VALUE (
   c               IN  INTEGER,
   name            IN  VARCHAR2,
   value           OUT NOCOPY <datatype>);
```

Where <datatype> can be any one of the following types:

```
ADT (user-defined object types)
BINARY_DOUBLE
BINARY_FLOAT
BFILE
BLOB
BOOLEAN
CLOB CHARACTER SET ANY_CS
DATE
DSINTERVAL_UNCONSTRAINED
JSON
NESTED table
NUMBER
OPAQUE types
REF
TIME_UNCONSTRAINED
TIME_TZ_UNCONSTRAINED
TIMESTAMP_LTZ_UNCONSTRAINED
TIMESTAMP_TZ_UNCONSTRAINED
TIMESTAMP_UNCONSTRAINED
UROWID
VARCHAR2 CHARACTER SET ANY_CS
VARRAY
VECTOR
YMINTERVAL_UNCONSTRAINED
```

For variables containing `CHAR`, `RAW`, and `ROWID` data, you can use the following variations on the syntax:

```
DBMS_SQL.VARIABLE_VALUE_CHAR (
   c               IN  INTEGER,
   name            IN  VARCHAR2,
   value           OUT CHAR CHARACTER SET ANY_CS);

DBMS_SQL.VARIABLE_VALUE_RAW (
   c               IN  INTEGER,
   name            IN  VARCHAR2,
```

```
   value           OUT RAW);

DBMS_SQL.VARIABLE_VALUE_ROWID (
   c               IN  INTEGER,
   name            IN  VARCHAR2,
   value           OUT ROWID);
```

The following syntax enables the `VARIABLE_VALUE` procedure to accommodate bulk operations:

```
DBMS_SQL.VARIABLE_VALUE (
   c                IN   INTEGER,
   name             IN   VARCHAR2,
   value            OUT NOCOPY <table_type>);
```

For bulk operations, <`table_type`> must be a supported DBMS_SQL predefined TABLE type.

See DBMS_SQL Data Structures

Pragmas

```
pragma restrict_references(variable_value,RNDS,WNDS);
```

Parameters

Table 182-33 VARIABLE_VALUE Procedure Parameters

Parameter	Description
`c`	ID number of the cursor from which to get the values.
`name`	Name of the variable for which you are retrieving the value.
`value`	• Single row option: Returns the value of the variable for the specified position. Oracle raises the exception `ORA-06562`, `inconsistent_type`, if the type of this output parameter differs from the actual type of the value, as defined by the call to `BIND_VARIABLE`. • Array option: Local variable that has been declared <`table_type`>. For bulk operations, `value` is an `OUT NOCOPY` parameter.

182.9.31 VARIABLE_VALUE_PKG Procedure

This procedure returns the value of the named variable for a given cursor.

It is used to return the values of bind variables of collection or record types inside PL/SQL blocks or DML statements with `returning` clause for a declared package. The type of the variable must be declared in the package specification. Bulk operations are not supported for these types.

Syntax

```
DBMS_SQL.VARIABLE_VALUE_PKG (
   c                IN   INTEGER,
   name             IN   VARCHAR2,
   value            OUT NOCOPY <table_type>);
```

Where <datatype> can be any one of the following data types:

- RECORD
- VARRAY

- NESTED TABLE
- INDEX BY PLS_INTEGER TABLE
- INDEX BY BINARY_INTEGER TABLE

Parameters

Table 182-34 VARIABLE_VALUE_PKG Parameters

Parameter	Description
c	ID number of the cursor from which to get the values.
name	Name of the variable for which you are retrieving the value.
value	• Single row option: Returns the value of the variable for the specified position. Oracle raises the exception ORA-06562, inconsistent_type, if the type of this output parameter differs from the actual type of the value, as defined by the call to BIND_VARIABLE_PKG. • Array option: Local variable that has been declared <table_type>.

Example 182-7 Dynamic SQL using DBMS_SQL.VARIABLE_VALUE_PKG to Get the Value of a Bind Variable

The data types are declared in the package specification. The VARIABLE_VALUE_PKG is used to get the value of the bind variable v2 in the cursor SQL statement.

```
CREATE OR REPLACE PACKAGE ty_pkg AS
TYPE rec IS RECORD
   ( n1 NUMBER,
    n2 NUMBER);
TYPE trect IS TABLE OF NUMBER;
END ty_pkg;
/
CREATE OR REPLACE PROCEDURE dyn_sql_nt  AS
  dummy NUMBER;
  cur   NUMBER;
  v1 ty_pkg.trect;
  v2 ty_pkg.trect;
  str VARCHAR2(3000);
BEGIN
  v1 := ty_pkg.trect(1000);
  str := 'declare v1 ty_pkg.trect;  begin v1:=:v1;  v1(1) := 2000; :v2 := v1;
end;' ;
  cur := DBMS_SQL.OPEN_CURSOR();
  DBMS_SQL.PARSE(cur, str, DBMS_SQL.NATIVE);
  DBMS_SQL.BIND_VARIABLE_PKG(cur, ':v1', v1);
  DBMS_SQL.BIND_VARIABLE_PKG(cur, ':v2', v2);
  dummy := DBMS_SQL.EXECUTE(cur);
  DBMS_SQL.VARIABLE_VALUE_PKG(cur, ':v2', v2);
  DBMS_OUTPUT.PUT_LINE('n =  '
  || V2(1));
```

```
  DBMS_SQL.CLOSE_CURSOR(cur);
END dyn_sql_nt;
/
EXEC dyn_sql_nt;

n =  2000
```

183
DBMS_SQL_FIREWALL

The DBMS_SQL_FIREWALL package enables you to monitor users and detect or prevent SQL injection attacks against those users.

This chapter contains the following topics:

- DBMS_SQL_FIREWALL Overview
- DBMS_SQL_FIREWALL Security Model
- DBMS_SQL_FIREWALL Constants
- Summary of DBMS_SQL_FIREWALL Subprograms

183.1 DBMS_SQL_FIREWALL Overview

The DBMS_SQL_FIREWALL PL/SQL package enables you to manage SQL Firewall, which tracks and can block SQL injection attacks.

The DBMS_SQL_FIREWALL package enables you to capture SQL activities of users, create allow-lists (that is, permitted actions) from the captured SQL activities, and then enforce the allow-lists to prevent or detect potential SQL injection attacks. In addition to SQL statements, the allow-list can contain a context list, which is a set of session contexts allowed for database connections. An example of a context can be IP addresses. You can also configure SQL Firewall to not run when Oracle Scheduler is running, because to do so may interfere with Oracle Scheduler operations. After you enable the allow-list, any SQL that the user performs will be monitored by SQL Firewall. SQL that the user performs that is not in the allow-list is considered to be a SQL injection attack. You can configure SQL Firewall to either allow the user to continue performing these SQL operations, or you can block these activities. Note that the SQL operations that violate the allow-list will always be written to a log table that you can query with data dictionary views.

You can configure SQL Firewall in both the root and in individual pluggable databases (PDBs).

Related Topics

- *Oracle Database Security Guides*

183.2 DBMS_SQL_FIREWALL Security Model

Oracle Database protects the administration of SQL Firewall by storing its metadata in tables in the SYS schema.

Hence, these tables rely on dictionary protection, just as other dictionary tables in SYS do. Therefore, users who have the SELECT ANY TABLE system privilege cannot query these tables unless they also have the SELECT ANY DICTIONARY system privilege or are granted the SELECT object privileges on the tables. Only the SYS user can grant these privileges to other users.

Oracle Database stores the SQL Firewall tables in the SYSAUX tablespace by default. If you want to move the SQL Firewall log tables to a different (user-defined) tablespace, then you must first disable SQL Firewall, and then use the MOVE clause of the ALTER TABLE statement to perform the move operation.

To use the procedures in the DBMS_SQL_FIREWALL package, a user must be granted the SQL_FIREWALL_ADMIN role.

Related Topics

- *Oracle Database Security Guide*

183.3 DBMS_SQL_FIREWALL Constants

The DBMS_SQL_FIREWALL package provides constants that are used with several SQL Firewall procedures.

These constants are described in the following table.

Table 183-1 DBMS_SQL_FIREWALL Constants

Name	Type	Value	Description
DBMS_SQL_FIREWALL.ENFORCE_ALL	NUMBER	3	Enforces both allowed SQL and allowed contexts when you run the DBMS_SQL_FIREWALL.ENABLE_ALLOW_LIST procedure
DBMS_SQL_FIREWALL.ENFORCE_CONTEXT	NUMBER	1	Enforces allowed contexts when you run the DBMS_SQL_FIREWALL.ENABLE_ALLOW_LIST procedure.
DBMS_SQL_FIREWALL.ENFORCE_SQL	NUMBER	2	Enforces allowed SQL when you run the DBMS_SQL_FIREWALL.ENABLE_ALLOW_LIST procedure
DBMS_SQL_FIREWALL.ALL_LOGS	NUMBER	3	Purges all logs when you run the DBMS_SQL_FIREWALL.PURGE procedure
DBMS_SQL_FIREWALL.CAPTURE_LOG	NUMBER	1	Purges only capture logs when you run the DBMS_SQL_FIREWALL.PURGE procedure
DBMS_SQL_FIREWALL.IP_ADDRESS	NUMBER	3	Specifies the user's IP address when you run the DBMS_SQL_FIREWALL.ADD_ALLOWED_CONTEXT or DBMS_SQL_FIREWALL.DELETE_ALLOWED_CONTEXT procedure
DBMS_SQL_FIREWALL.OS_PROGRAM	NUMBER	1	Specifies the user's operating system program when you run the DBMS_SQL_FIREWALL.ADD_ALLOWED_CONTEXT or DBMS_SQL_FIREWALL.DELETE_ALLOWED_CONTEXT procedure

Table 183-1 (Cont.) DBMS_SQL_FIREWALL Constants

Name	Type	Value	Description
DBMS_SQL_FIREWALL.OS_USERNAME	NUMBER	2	Specifies an operating system name when you run the DBMS_SQL_FIREWALL.ADD_ALLOWED_CONTEXT or DBMS_SQL_FIREWALL.DELETE_ALLOWED_CONTEXT procedure
DBMS_SQL_FIREWALL.SCHEDULER_JOB	NUMBER	1	Indicates whether SQL Firewall will capture and enforce allow-lists for database connections and SQL executions during Oracle Scheduler operations. Use this constant with the DBMS_SQL_FIREWALL.EXCLUDE and DBMS_SQL_FIREWALL.INCLUDE procedures.
DBMS_SQL_FIREWALL.VIOLATION_LOG	NUMBER	2	Purges only violation logs when you run the DBMS_SQL_FIREWALL.PURGE procedure

183.4 Summary of DBMS_SQL_FIREWALL Subprograms

This table lists and describes the DBMS_SQL_FIREWALL package subprograms.

Table 183-2 DBMS_SQL_FIREWALL Package Subprograms

Subprogram	Description
ADD_ALLOWED_CONTEXT Procedure	Adds a context to the list of allowed contexts for a user who is configured for SQL Firewall
APPEND_ALLOW_LIST Procedure	Appends additional contents to an existing allow-list by using the existing capture logs or violation logs of the user, or both
APPEND_ALLOW_LIST_SINGLE_SQL Procedure	Appends a single SQL record to the violation log or capture log to an existing allow-list
CREATE_CAPTURE Procedure	Creates a SQL Firewall capture for a specified user at a given level
DELETE_ALLOWED_CONTEXT Procedure	Deletes a SQL Firewall context value that had been assigned to a user
DELETE_ALLOWED_SQL Procedure	Deletes a specified entry from the allowed SQL that had been assigned to a user
DISABLE Procedure	Disables SQL Firewall
DISABLE_ALLOW_LIST Procedure	Disables SQL Firewall allow-list enforcement for a given user
DROP_ALLOW_LIST Procedure	Deletes the SQL Firewall allow-list of a specified user

Table 183-2 (Cont.) DBMS_SQL_FIREWALL Package Subprograms

Subprogram	Description
DROP_CAPTURE Procedure	Drops a SQL Firewall capture and deletes all the associated capture logs
ENABLE Procedure	Enables SQL Firewall
ENABLE_ALLOW_LIST Procedure	Enables SQL Firewall allow-list enforcement for a given user
EXCLUDE Procedure	Prevents SQL Firewall from capturing or enforcing allow-lists for database connections and SQL executions during Oracle Scheduler operations
EXPORT_ALLOW_LIST Procedure	Exports the allow-list of the given user in JSON format, into the CLOB provided from the `allow_list` argument
FLUSH_LOGS Procedure	Flushes all the SQL Firewall logs that reside in the memory into the log tables
GENERATE_ALLOW_LIST Procedure	Generates a SQL Firewall allow-list for the specified user by using data from the existing capture logs of the user
IMPORT_ALLOW_LIST Procedure	Imports the allow-list from the specified `CLOB` for the given user, to the target database
INCLUDE Procedure	Enables SQL Firewall to capture and enforce allow-lists for database connections and SQL executions during Oracle Scheduler operations
PURGE_LOG Procedure	Purges SQL Firewall logs
START_CAPTURE Procedure	Starts a SQL Firewall capture for a user
STOP_CAPTURE Procedure	Stops a SQL Firewall capture for a user
UPDATE_ALLOW_LIST_ENFORCEMENT Procedure	Updates the SQL Firewall allow-list enforcement options for the given user

183.4.1 ADD_ALLOWED_CONTEXT Procedure

This procedure adds a context to the list of allowed contexts for a user's SQL Firewall allow-list.

Syntax

```
DBMS_SQL_FIREWALL.ADD_ALLOWED_CONTEXT (
   username       IN  VARCHAR2,
   context_type   IN  NUMBER,
   value          IN  VARCHAR2);
```

Parameters

Table 183-3 ADD_ALLOWED_CONTEXT Procedure Parameters

Parameter	Description
`username`	Specifies the name of the user who has a SQL Firewall allow-list. To find all the users who has an allow-list, query `DBA_SQL_FIREWALL_ALLOW_LISTS`.

Table 183-3 (Cont.) ADD_ALLOWED_CONTEXT Procedure Parameters

Parameter	Description
context_type	Specifies one of the following context types: • DBMS_SQL_FIREWALL.IP_ADDRESS accepts IPv4 and IPv6 addresses and subnets in the CIDR notation. • DBMS_SQL_FIREWALL.OS_USERNAME accepts any valid operating system user name, such as oracle. • DBMS_SQL_FIREWALL.OS_PROGRAM accepts any valid operating system program name, such as sqlplus or SQL Developer.
value	Specifies the value of the context_type constant, such as an IP address for DBMS_SQL_FIREWALL.IP_ADDRESS. To allow a local (bequeathed) connection that does not have an IP address, specify with the value Local for the DBMS_SQL_FIREWALL.IP_ADDRESS type. To specify all values of the context (such as all possible operating system programs), then enter the % wild card character.

Usage Notes

- You can find the user's current context type settings by querying the following data dictionary views:
  - DBA_SQL_FIREWALL_ALLOWED_IP_ADDR
  - DBA_SQL_FIREWALL_ALLOWED_OS_PROG
  - DBA_SQL_FIREWALL_ALLOWED_OS_USER
- Before you can add any contexts for the user, the user's allow-list must be created (using the DBMS_SQL_FIREWALL.GENERATE_ALLOW_LIST procedure).
- This procedure can be run when the allow-list is enabled or disabled, and it takes effects immediately.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.ADD_ALLOWED_CONTEXT (
    username       => 'PFITCH',
    context_type   => DBMS_SQL_FIREWALL.OS_PROGRAM,
    value          => 'SQL Developer'
   );
END;
/
```

183.4.2 APPEND_ALLOW_LIST Procedure

This procedure appends additional contents to an existing allow-list by using the existing capture logs or violation logs of the user, or both.

Syntax

```
DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST (
   username       IN  VARCHAR2,
   source         IN  NUMBER);
```

Parameters

Table 183-4 APPEND_ALLOW_LIST Procedure Parameters

Parameter	Description
username	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query DBA_SQL_FIREWALL_ALLOW_LISTS.
source	Specifies one of the following log types: • DBMS_SQL_FIREWALL.CAPTURE_LOG • DBMS_SQL_FIREWALL.VIOLATION_LOG • DBMS_SQL_FIREWALL.ALL_LOGS

Usage Notes

- DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST processes the specified source logs and identifies contents to be appended to the allow-list. Then it populates the SQL Firewall metadata tables for the allowed SQL and allowed contexts, which will be used during the allow-list enforcement.
- You can run this procedure when the allow-list is either enabled or disabled.
- The change takes effect immediately.
- A new allow-list version number will be associated with all the allowed SQL entries added by the same DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST execution. This new version number will be 1 plus the current maximum allow-list version of the specified user.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST (
    username       => 'PFITCH',
    source         => DBMS_SQL_FIREWALL.CAPTURE_LOG
   );
END;
/
```

183.4.3 APPEND_ALLOW_LIST_SINGLE_SQL Procedure

This procedure appends a single SQL record to the violation log or capture log to an existing allow-list.

This procedure is useful for when you want to individually append SQL commands from the violations log or the capture log to an existing allow-list.

Syntax

```
DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST_SINGLE_SQL (
   username       IN  VARCHAR2,
   sql_signature  IN  VARCHAR2,
   current_user   IN  VARCHAR2,
   top_level      IN  VARCHAR2,
   source         IN  NUMBER DEFAULT);
```

Parameters

Table 183-5 APPEND_ALLOW_LIST_SINGLE_SQL Procedure Parameters

Parameter	Description
username	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query DBA_SQL_FIREWALL_ALLOW_LISTS.
sql_signature	Specifies the signature of the SQL to be added. To find the signature of the SQL for the target record, query the DBA_SQL_FIREWALL_CAPTURE or DBA_SQL_FIREWALL_VIOLATIONS dynamic view.
current_user	Specifies the name of the user who the SQL command was executed as. For example, if user pfitch invokes a definer's rights procedure created in the psmith schema, then all the SQL commands in the procedure are executed as psmith, the current_user. If the procedure is an invoker's rights procedure, then the current_user is the invoker, pfitch.
top_level	Specifies whether the SQL that was executed was top level. Possible values are as follows: • Y (for Yes) means that the target SQL record is top-level (that is, the statement that the user directly runs). • N (for No) means that the target SQL record is not top-level (that is, the SQL command that is issued from PL/SQL units).
source	Specifies the source log to add the SQL record from: • DBMS_SQL_FIREWALL.CAPTURE_LOG • DBMS_SQL_FIREWALL.VIOLATION_LOG (default)

Usage Notes

- DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST_SINGLE_SQL processes the specified source log and identifies the target SQL record to be appended to the allow-list. Then it populates the SQL Firewall metadata tables for the allowed SQL, which will be used during the allow-list enforcement.
- You can run this procedure when the allow-list is either enabled or disabled.
- The change takes effect immediately.
- A new allow-list version number will be associated with the newly added allowed SQL entry.

Example

1. Query the DBA_SQL_FIREWALL_VIOLATIONS or the DBA_SQL_FIREWALL_CAPTURE_LOGS data dictionary view to find the target SQL record that you want to add to the allow-list. Obtain the values for the USERNAME, SQL_SIGNATURE, CURRENT_USER, and TOP_LEVEL columns for the target SQL record.
2. Enter these values in the DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST_SINGLE_SQL SQL procedure to add the target SQL record to the allow-list.
   For example:

```
BEGIN
  DBMS_SQL_FIREWALL.APPEND_ALLOW_LIST_SINGLE_SQL (
    username       => 'PFITCH',
    sql_signature  =>
'7D33A84D0A1B56E382B9A92D01BCD19933969CB16E2AB4934A2258563F5ADB44',
```

```
        current_user    => 'PSMITH',
        top_level       => 'N',
        source          => DBMS_SQL_FIREWALL.CAPTURE_LOG
       );
    END;
    /
```

183.4.4 CREATE_CAPTURE Procedure

This procedure creates a SQL Firewall capture for a specified user at a given level.

Syntax

```
DBMS_SQL_FIREWALL.CREATE_CAPTURE (
   username       IN VARCHAR2,
   top_level_only IN BOOLEAN,
   start_capture  IN BOOLEAN);
```

Parameters

Table 183-6 CREATE_CAPTURE Procedure Parameters

Parameter	Description
username	Specifies the name of the user whose SQL Firewall capture is to be created. To find existing users, query DBA_SQL_FIREWALL_CAPTURES.
top_level_only	• TRUE captures only SQL statements that have been directly issued by the user • FALSE captures both top-level SQL statements and SQL statements that have been issued by PL/SQL units. This setting is the default.
start_capture	• TRUE starts the capture process right away, after you run DBMS_SQL_FIREWALL.CREATE_CAPTURE. This setting is the default. • FALSE does not start the capture process. You can start it later on by using DBMS_SQL_FIREWALL.START_CAPTURE.

Usage Notes

To find the status of existing SQL Firewall captures, including users who have already been configured for SQL Firewall captures, query the DBA_SQL_FIREWALL_CAPTURES data dictionary view.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.CREATE_CAPTURE (
    username         => 'C##HR_ADMIN',
    top_level_only   => TRUE,
    start_capture    => TRUE
  );
END;
/
```

183.4.5 DELETE_ALLOWED_CONTEXT Procedure

This procedure deletes a context from the list of allowed contexts for a user's SQL Firewall allow-list.

Syntax

```
DBMS_SQL_FIREWALL.DELETE_ALLOWED_CONTEXT (
   username       IN  VARCHAR2,
   context_type   IN  NUMBER,
   value          IN  VARCHAR2);
```

Parameters

Table 183-7 DELETE_ALLOWED_CONTEXT Procedure Parameters

Parameter	Description
`username`	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query `DBA_SQL_FIREWALL_ALLOW_LISTS`.
`context_type`	• `DBMS_SQL_FIREWALL.IP_ADDRESS` accepts IPv4 and IPv6 addresses and subnets in the CIDR notation. • `DBMS_SQL_FIREWALL.OS_USERNAME` accepts any valid operating system user name, such as `oracle`. • `DBMS_SQL_FIREWALL.OS_PROGRAM` accepts any valid operating system program name, such as `sqlplus` or `SQL Developer`.
`value`	Specifies the value of the `context_type` constant, such as an IP address for `DBMS_SQL_FIREWALL.IP_ADDRESS`. If you omit this value or specify `NULL`, then all the allowed context values of the specified context type are deleted. This setting is the default.

Usage Notes

- You can find the user's current context type settings by querying the following data dictionary views:
  - `DBA_SQL_FIREWALL_ALLOWED_IP_ADDR`
  - `DBA_SQL_FIREWALL_ALLOWED_OS_PROG`
  - `DBA_SQL_FIREWALL_ALLOWED_OS_USER`
- This procedure can be run when the allow-list is enabled or disabled, and it takes effects immediately.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.DELETE_ALLOWED_CONTEXT, (
    username       => 'PFITCH',
    context_type   => DBMS_SQL_FIREWALL.OS_PROGRAM,
    value          => 'SQL Developer'
   );
END;
/
```

183.4.6 DELETE_ALLOWED_SQL Procedure

This procedure deletes a specified entry from the list of allowed SQL for a user's SQL Firewall allow-list

Syntax

```
DBMS_SQL_FIREWALL.DELETE_ALLOWED_SQL (
   username       IN  VARCHAR2,
   allowed_sql_id IN  NUMBER);
```

Parameters

Table 183-8 DELETE_ALLOWED_SQL Procedure Parameters

Parameter	Description
`username`	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query `DBA_SQL_FIREWALL_ALLOW_LISTS`.
`allowed_sql_id`	Specifies the ID of the allowed SQL entry to be deleted from the allowed SQL of this user.To find this value, query `DBA_SQL_FIREWALL_ALLOWED_SQL`.

Usage Notes

- You can run this procedure when the allow-list is either enabled or disabled.
- The change takes effect immediately.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.DELETE_ALLOWED_SQL (
    username         => 'PFITCH',
    allowed_sql_id   => 1
   );
END;
/
```

183.4.7 DISABLE Procedure

This procedure disables SQL Firewall and stops all the existing captures and allow-lists that are enabled.

Syntax

```
DBMS_SQL_FIREWALL.DISABLE;
```

Parameters

None

Usage Notes

You can find the current status of SQL Firewall by querying the `DBA_SQL_FIREWALL_STATUS` data dictionary view.

Example

```
EXEC DBMS_SQL_FIREWALL.DISABLE;
```

183.4.8 DISABLE_ALLOW_LIST Procedure

This procedure immediately disables SQL Firewall allow-list enforcement for a given user.

Syntax

```
DBMS_SQL_FIREWALL.DISABLE_ALLOW_LIST (
   username       IN  VARCHAR2);
```

Parameters

Table 183-9 DISABLE_ALLOW_LIST Procedure Parameters

Parameter	Description
username	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query DBA_SQL_FIREWALL_ALLOW_LISTS. If you specify NULL, then all allow-lists that are currently enabled will be disabled.

Usage Notes

To find the status of users' allow-lists, query the DBA_SQL_FIREWALL_ALLOW_LISTS data dictionary view.

Example

```
EXEC DBMS_SQL_FIREWALL.DISABLE_ALLOW_LIST ('PFITCH');
```

183.4.9 DROP_ALLOW_LIST Procedure

This procedure deletes the SQL Firewall allow-list of a specified user.

Syntax

```
DBMS_SQL_FIREWALL.DROP_ALLOW_LIST (
   username       IN  VARCHAR2);
```

Parameters

Table 183-10 DROP_ALLOW_LIST Procedure Parameters

Parameter	Description
username	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query DBA_SQL_FIREWALL_ALLOW_LISTS.

Usage Notes

- To find the status of users' allow-lists, query the DBA_SQL_FIREWALL_ALLOW_LISTS data dictionary view.

- You cannot drop an allow-list that is currently enabled. To disable an allow-list, run the `DBMS_SQL_FIREWALL.DISABLE_ALLOW_LIST` procedure.

Example

```
EXEC DBMS_SQL_FIREWALL.DROP_ALLOW_LIST ('PFITCH');
```

183.4.10 DROP_CAPTURE Procedure

This procedure drops a SQL Firewall capture and deletes all the associated capture logs.

Syntax

```
DBMS_SQL_FIREWALL.DROP_CAPTURE (
   username       IN VARCHAR2);
```

Parameters

Table 183-11 DROP_CAPTURE procedure Parameters

Parameter	Description
`username`	Specifies the name of the user whose SQL Firewall capture is to be dropped. To find this user, query `DBA_SQL_FIREWALL_CAPTURES`.

Usage Notes

- To find the status of existing SQL Firewall captures, query the `DBA_SQL_FIREWALL_CAPTURES` data dictionary view.
- You cannot drop a capture that is currently running. To stop the capture, run the `DBMS_SQL_FIREWALL.STOP_CAPTURE` procedure.
- Dropping a capture for a user does not affect the user's allow-list, which can continue to run even if the capture has been dropped. Captures and allow-lists are separate entities.

Example

```
EXEC DBMS_SQL_FIREWALL.DROP_CAPTURE ('C##HR_ADMIN');
```

183.4.11 ENABLE Procedure

This procedure enables SQL Firewall and starts all existing captures and allow-lists that are configured to be enabled.

Syntax

```
DBMS_SQL_FIREWALL.ENABLE;
```

Parameters

None

Usage Notes

You can find the current status of SQL Firewall by querying the `DBA_SQL_FIREWALL_STATUS` data dictionary view.

Example

```
EXEC DBMS_SQL_FIREWALL.ENABLE;
```

183.4.12 ENABLE_ALLOW_LIST Procedure

This procedure immediately enables SQL Firewall allow-list enforcement for a given user.

Syntax

```
DBMS_SQL_FIREWALL.ENABLE_ALLOW_LIST (
   username       IN  VARCHAR2,
   enforce        IN  NUMBER,
   block          IN  BOOLEAN;
```

Parameters

Table 183-12 ENABLE_ALLOW_LIST Procedure Parameters

Parameter	Description
`username`	Specifies the name of the user whose SQL Firewall allow-list is to be enabled. To find this user, query `DBA_SQL_FIREWALL_ALLOW_LISTS`. If you enter `NULL`, then the allow-lists for all users who do not yet have allow-lists enabled are enabled.
`enforce`	• `DBMS_SQL_FIREWALL.ENFORCE_CONTEXT` enforces the allowed contexts that have been configured. • `DBMS_SQL_FIREWALL.ENFORCE_SQL` enforces the allowed SQL that has been configured. • `DBMS_SQL_FIREWALL.ENFORCE_ALL` enforces both allowed contexts and allowed SQL. This setting is the default.
`block`	• `TRUE` blocks user's database connection or the user's SQL execution whenever the user violates the allow-list definition. • `FALSE` allows unmatched user database connections or SQL commands to proceed. This setting is the default.

Usage Notes

- To find the status of users' allow-lists, query the `DBA_SQL_FIREWALL_ALLOW_LISTS` data dictionary view.
- SQL Firewall always generates a violation log for any unmatched database connection or SQL statement regardless of the `block` option setting.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.ENABLE_ALLOW_LIST (
    username       => 'PFITCH',
    enforce        => DBMS_SQL_FIREWALL.ENFORCE_SQL,
```

```
    block          => TRUE
   );
END;
/
```

183.4.13 EXCLUDE Procedure

This procedure prevents SQL Firewall from capturing or enforcing allow-lists for database connections and SQL executions during Oracle Scheduler operations.

Oracle Scheduler jobs are often used in databases for various maintenance purposes. Accidentally interrupting critical jobs can cause undesirable consequences. You can configure SQL Firewall to not capture any SQL statements nor enforce any allow-lists that are run during an Oracle Scheduler job session. This procedure applies to all users that have been configured for SQL Firewall captures and allow-lists. By default, Oracle Scheduler jobs are excluded from SQL Firewall operations.

Syntax

```
DBMS_SQL_FIREWALL.EXCLUDE (
   FEATURE        IN NUMBER);
```

Parameters

Table 183-13 EXCLUDE Procedure Parameters

Parameter	Description
FEATURE	Enter DBMS_SQL_FIREWALL.SCHEDULER_JOB for this value.

Usage Notes

- To find the status of whether SQL Firewall is enforced during Oracle Scheduler operations, query the EXCLUDE_JOBS column of the DBA_SQL_FIREWALL_STATUS data dictionary view. If the output is Y, then Oracle Scheduler jobs are excluded from SQL Firewall operations.
- To enable Oracle Firewall to run during Oracle Scheduler operations, run the DBMS_SQL_FIREWALL.INCLUDE procedure.

Example

```
EXEC DBMS_SQL_FIREWALL.EXCLUDE (DBMS_SQL_FIREWALL.SCHEDULER_JOB);
```

183.4.14 EXPORT_ALLOW_LIST Procedure

This procedure exports the allow-list of the given user in JSON format, into the CLOB provided from the allow_list argument.

Syntax

```
DBMS_SQL_FIREWALL.EXPORT_ALLOW_LIST (
   username       IN      VARCHAR2,
   allow_list     IN/OUT  CLOB;
```

Parameters

Table 183-14 EXPORT_ALLOW_LIST Procedure Parameters

Parameter	Description
username	Specifies the user that the allow-list was created for. To find which user has an allow-list, query DBA_SQL_FIREWALL_ALLOW_LISTS.
allow_list	Specifies the CLOB (which must already exist) into which the exported allow-list must go

Usage Notes

- Before you run this procedure, you must create the CLOB and then pass it to the API (for example, by DBMS_LOB.CREATETEMPORARY for the PL/SQL client, or by OracleConnection.createClob() for JDBC Java client).
- The export operation includes the allow-list's settings (status, enforce, block, top_level_only, generated_on, and status_updated_on timestamp), allowed SQL, and allowed contexts. In addition, the export operation includes all the referenced SQL logs (by the allowed SQL).
- DBMS_SQL_FIREWALL.EXPORT_ALLOW_LIST does not export capture logs or violation logs.
- To find the status of users' allow-lists, query the DBA_SQL_FIREWALL_ALLOW_LISTS data dictionary view.
- If you want to export all the SQL Firewall metadata, which includes captures and allow-lists for all users, then instead of using DBMS_SQL_FIREWALL.EXPORT_ALLOW_LIST, use the include=SQL_FIREWALL clause in the Oracle Data Pump expdp command. See *Oracle Database Security Guide*.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.EXPORT_ALLOW_LIST (
    username       => 'PFITCH',
    allow_list     => ALLOW_LIST_CLOB;
   );
END;
/
```

183.4.15 FLUSH_LOGS Procedure

This procedure flushes all the SQL Firewall logs that reside in the memory into the log tables.

Syntax

```
DBMS_SQL_FIREWALL.FLUSH_LOGS;
```

Parameters

None

Usage Notes

- Usually you do not need to invoke this procedure explicitly, because logs in the memory are flushed to the log tables frequently in the background. But in case if you want to see the capture logs or violation logs immediately after the action during when SQL Firewall is running, you can run this procedure before looking at the logs.
- The `DBMS_SQL_FIREWALL.FLUSH_LOGS` procedure is equivalent to the `DBMS_MEMOPTIMIZE_ADMIN.WRITES_FLUSH` procedure. (See WRITES_FLUSH Procedure.)

Example

```
EXEC DBMS_SQL_FIREWALL.FLUSH_LOGS;
```

183.4.16 GENERATE_ALLOW_LIST Procedure

This procedure generates a SQL Firewall allow-list for the specified user by using the existing capture logs of the user.

Syntax

```
DBMS_SQL_FIREWALL.GENERATE_ALLOW_LIST (
   username       IN  VARCHAR2;
```

Parameters

Table 183-15 GENERATE_ALLOW_LIST Procedure Parameters

Parameter	Description
`username`	Specifies the name of the user who was designated for the SQL Firewall allow-list. To find this user, query `DBA_SQL_FIREWALL_CAPTURES`.

Usage Notes

- To find information about existing generated allow-lists, query the `DBA_SQL_FIREWALL_ALLOW_LISTS` data dictionary view.
- Before you run this procedure, the following components must be in place:
  - The specified user must exist.
  - A capture (using `DBMS_SQL_FIREWALL.CREATE_CAPTURE`) has been created for this user. This capture must be disabled (using `DBMS_SQL_FIREWALL.STOP_CAPTURE`) before you can generate an allow-list for the user.
  - No allow-list exists yet for the user.

Example

```
EXEC DBMS_SQL_FIREWALL.GENERATE_ALLOW_LIST ('PFITCH');
```

183.4.17 IMPORT_ALLOW_LIST Procedure

This procedure imports the allow-list from the specified `CLOB` for the given user, to the target database.

Syntax

```
DBMS_SQL_FIREWALL.IMPORT_ALLOW_LIST (
   username       IN      VARCHAR2,
   allow_list     IN      CLOB;
```

Parameters

Table 183-16 IMPORT_ALLOW_LIST Procedure Parameters

Parameter	Description
`username`	Specifies the user of the exported allow-list. To check whether this user already had an allow-list created in the target database, query `DBA_SQL_FIREWALL_ALLOW_LISTS`.
`allow_list`	Specifies the `CLOB` that was created when the allow-list was exported with `DBMS_SQL_FIREWALL.EXPORT_ALLOW_LIST`.

Usage Notes

- If this user does not have an allow-list in the target database, a new allow-list will be created for this user using the allow-list from the JSON payload. The new allow-list will have the same settings (`status`, `top_level_only`, `enforce`, `block`, `generated_on`, `status_updated_on`), same allowed contexts and same allowed SQL as the one in the JSON. If the specified user already has an allow-list in the target database, then all the settings (`status`, `top_level_only`, `enforce`, `block`, and various timestamps) of the existing allow-list will remain untouched, but only the allowed SQL and allowed contexts from the JSON will be merged into the ones for the existing allow-list.
- In addition, the import operation includes all the referenced SQL logs (by the allowed SQL).
- To find the status of users' allow-lists, query the `DBA_SQL_FIREWALL_ALLOW_LISTS` data dictionary view.
- If you want to import all the SQL Firewall metadata, which includes captures and allow-lists, then instead of using `DBMS_SQL_FIREWALL.IMPORT_ALLOW_LIST`, use the `include=SQL_FIREWALL` clause in the Oracle Data Pump `impdp` command. See *Oracle Database Security Guide*.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.IMPORT_ALLOW_LIST (
    username       => 'PFITCH',
    allow_list     => ALLOW_LIST_CLOB;
   );
END;
/
```

183.4.18 INCLUDE Procedure

This procedure enables SQL Firewall to capture and enforce allow-lists for database connections and SQL executions during Oracle Scheduler operations.

Syntax

```
DBMS_SQL_FIREWALL.INCLUDE (
   FEATURE        IN NUMBER);
```

Parameters

Table 183-17 INCLUDE Procedure Parameters

Parameter	Description
FEATURE	Enter DBMS_SQL_FIREWALL.SCHEDULER_JOB for this value.

Usage Notes

- To find the status of whether SQL Firewall is enforced during Oracle Scheduler operations, query the EXCLUDE_JOBS column of the DBA_SQL_FIREWALL_STATUS data dictionary view. If the output is N, then SQL Firewall can perform during Oracle Scheduler operations.
- To prevent SQL Firewall from running during Oracle Scheduler operations, run the DBMS_SQL_FIREWALL.EXCLUDE procedure.

Example

```
EXEC DBMS_SQL_FIREWALL.INCLUDE (DBMS_SQL_FIREWALL.SCHEDULER_JOB);
```

183.4.19 PURGE_LOG Procedure

This procedure purges SQL Firewall logs that belong to the given user based on the specified purge time (that is, logs that were generated before the specified purge time).

Syntax

```
BEGIN
  DBMS_SQL_FIREWALL.PURGE_LOG (
   username       IN  VARCHAR2,
   purge_time     IN  TIMESTAMP WITH TIME ZONE,
   log_type       IN  NUMBER);
```

Parameters

Table 183-18 PURGE_LOG Procedure Parameters

Parameter	Description
username	Specifies the user whose capture logs or violation logs you want to purge. To see capture logs, query DBA_SQL_FIREWALL_CAPTURE_LOGS; to see violation logs, query DBA_SQL_FIREWALL_VIOLATIONS.

Table 183-18 (Cont.) PURGE_LOG Procedure Parameters

Parameter	Description
purge_time	The timestamp (in TIMESTAMP format) that you can specify to purge only logs that were generated before a certain time. If you omit this value, then Oracle Database purges all logs, regardless of the time when they were generated.
log_type	Specifies the type of the logs to be purged. • DBMS_SQL_FIREWALL.CAPTURE_LOG • DBMS_SQL_FIREWALL.VIOLATION_LOG • DBMS_SQL_FIREWALL.ALL_LOGS (default)

Usage Notes

To find information about SQL Firewall logs, query the DBA_SQL_FIREWALL_VIOLATIONS data dictionary view.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.PURGE_LOG (
    username     => 'PSMITH',
    purge_time   => TO_TIMESTAMP_TZ('23-JAN-22 18.44.42 -07:00', 'DD/MM/YY
HH24:MI:SS TZH:TZM'),
    log_type     => DBMS_SQL_FIREWALL.VIOLATION_LOG
   );
END;
/
```

183.4.20 START_CAPTURE Procedure

This procedure immediately starts a SQL Firewall capture for a user.

Syntax

```
DBMS_SQL_FIREWALL.START_CAPTURE (
   username       IN  VARCHAR2);
```

Parameters

Table 183-19 START_CAPTURE Procedure Parameters

Parameter	Description
username	Specifies the name of the user to be designated for the SQL Firewall capture.

Usage Notes

- A user can only have one SQL Firewall capture. To find if the user already has been configured for a capture, query the DBA_SQL_FIREWALL_CAPTURES data dictionary view.
- After you start the capture process, all SQL the user enters is captured into the SQL Firewall capture log table. You can periodically check the this SQL by querying the DBA_SQL_FIREWALL_CAPTURE_LOGS data dictionary view.

Example

```
EXEC DBMS_SQL_FIREWALL.START_CAPTURE ('PFITCH');
```

183.4.21 STOP_CAPTURE Procedure

This procedure immediately stops a SQL Firewall capture for a given user.

Syntax

```
DBMS_SQL_FIREWALL.STOP_CAPTURE (
   username       IN  VARCHAR2);
```

Parameters

Table 183-20 STOP_CAPTURE Procedure Parameters

Parameter	Description
username	Specifies the name of the user who was designated for the SQL Firewall capture. To find this user, query DBA_SQL_FIREWALL_CAPTURES.

Usage Notes

- The capture process must be currently running before you can run this procedure. You can check its status by querying the DBA_SQL_FIREWALL_CAPTURES data dictionary view.
- After you stop the capture process, you can generate an allow-list for the user by running the DBMS_SQL_FIREWALL.GENERATE_ALLOW_LIST procedure.

Example

```
EXEC DBMS_SQL_FIREWALL.STOP_CAPTURE ('PFITCH');
```

183.4.22 UPDATE_ALLOW_LIST_ENFORCEMENT Procedure

This procedure immediately updates the SQL Firewall allow-list enforcement options for the given user.

Syntax

```
BEGIN
  DBMS_SQL_FIREWALL.UPDATE_ALLOW_LIST_ENFORCEMENT (
   username       IN  VARCHAR2,
   enforce        IN  NUMBER,
   block          IN  BOOLEAN);
```

Parameters

Table 183-21 UPDATE_ALLOW_LIST_ENFORCEMENT Procedure Parameters

Parameter	Description
username	Specifies the name of the user for whom the allow-list was generated. To find this user, query DBA_SQL_FIREWALL_ALLOW_LISTS. If you enter NULL, then the enforcement options of all the existing allow-lists (both enabled or disabled allow-lists) are updated.
enforce	• DBMS_SQL_FIREWALL.ENFORCE_CONTEXT enforces the allowed contexts that have been configured. • DBMS_SQL_FIREWALL.ENFORCE_SQL enforces the allowed SQL that has been configured. • DBMS_SQL_FIREWALL.ENFORCE_ALL enforces both allowed contexts and allowed SQL. This setting is the default.
block	• TRUE blocks user's database connection or the user's SQL execution whenever the user violates the allow-list definition. • FALSE allows unmatched user database connections or SQL commands to proceed. This setting is the default.

Usage Notes

To find the status of users' allow-lists, query the DBA_SQL_FIREWALL_ALLOW_LISTS data dictionary view.

Example

```
BEGIN
  DBMS_SQL_FIREWALL.UPDATE_ALLOW_LIST_ENFORCEMENT (
    username       => 'PFITCH',
    enforce        => DBMS_SQL_FIREWALL.ENFORCE_SQL,
    block          => TRUE
   );
END;
/
```

184
DBMS_SQL_MONITOR

The DBMS_SQL_MONITOR package provides information about Real-Time SQL Monitoring and Real-Time Database Operation Monitoring.

This chapter contains the following topics:

- DBMS_SQL_MONITOR Overview
- DBMS_SQL_MONITOR Security Model
- DBMS_SQL_MONITOR Constants
- Summary of DBMS_SQL_MONITOR Subprograms

> **See Also:**
>
> DBMS_SQLTUNE

184.1 DBMS_SQL_MONITOR Overview

The DBMS_SQL_MONITOR package provides information about Real-Time SQL Monitoring and Real-Time Database Operation Monitoring.

These features provide automatic monitoring of SQL statements, PL/SQL blocks, or composite database operations that are considered high-cost. A simple database operation is a single SQL statement or PL/SQL procedure or function. A composite database operation is activity between two defined points in time in a database session. The monitored data is collected in the V$SQL_MONITOR and V$SQL_PLAN_MONITOR views.

The following subprograms begin and end monitoring of a composite database operation:

- BEGIN_OPERATION Function
- END_OPERATION Procedure

The following subprograms report on monitoring data collected in V$SQL_MONITOR and V$SQL_PLAN_MONITOR:

- REPORT_SQL_MONITOR Function
- REPORT_SQL_MONITOR_XML Function
- REPORT_SQL_MONITOR_LIST Function
- REPORT_SQL_MONITOR_LIST_XML Function

184.2 DBMS_SQL_MONITOR Security Model

This package is available to PUBLIC and executes with invoker's rights privileges. The reporting functions require privileges to select data from the catalog as provided by the role SELECT_CATALOG_ROLE.

184.3 DBMS_SQL_MONITOR Constants

The `DBMS_SQL_MONITOR` package uses the constants shown in the following table.

Table 184-1 DBMS_SQL_MONITOR Constants

Constant	Type	Value	Description
`FORCE_TRACKING`	`VARCHAR2(30)`	`'Y'`	Force track the composite database operation when the operation starts
`NO_FORCE_TRACKING`	`VARCHAR2(30)`	`'N'`	Do not force track the composite database operation when the operation starts. It is only tracked when it has consumed 5 seconds of CPU or I/O time.

184.4 Summary of DBMS_SQL_MONITOR Subprograms

This table lists and describes the `DBMS_SQL_MONITOR` package subprograms.

Table 184-2 DBMS_SQL_MONITOR Package Subprograms

Subprogram	Description
BEGIN_OPERATION Function	This function starts a database operation in the current session.
END_OPERATION Procedure	This function ends a database operation in the current session. If the specified database operation does not exist, then this function has no effect.
REPORT_SQL_MONITOR Function	This function builds a detailed report with monitoring information for a SQL statement, PL/SQL block, or database operation.
REPORT_SQL_MONITOR_XML Function	This function is identical to the `REPORT_SQL_MONITOR` function, except that the return type is `XMLType`.
REPORT_SQL_MONITOR_LIST Function	This function builds a report for all or a subset of database operations that have been monitored by Oracle Database.
REPORT_SQL_MONITOR_LIST_XML Function	This function is identical to the `REPORT_SQL_MONITOR_LIST` function, except that it returns `XMLType`.

184.4.1 BEGIN_OPERATION Function

This function starts a database operation in the current session.

Syntax

```
DBMS_SQL_MONITOR.BEGIN_OPERATION (
   dbop_name       IN VARCHAR2,
```

```
   dbop_eid        IN NUMBER   := NULL,
   forced_tracking IN VARCHAR2 := NO_FORCE_TRACKING,
   attribute_list  IN VARCHAR2 := NULL,
   session_id      IN NUMBER   := NULL,
   session_serial  IN NUMBER   := NULL)
  iRETURN NUMBER;
```

Parameters

Table 184-3 *BEGIN_OPERATION Procedure Parameters*

Parameter	Description
dbop_name	Name for the composite database operation.
dbop_eid	Unique identifier for the current execution of the composite database operation.
forced_tracking	Whether tracking is forced. Possible values are: • FORCE_TRACKING - forces the composite database operation to be tracked when the operation starts. You can also use the string variable Y. • NO_FORCE_TRACKING - tracks the operation only when it has consumed at least 5 seconds of CPU or I/O time. You can also use the string variable N. See "DBMS_SQL_MONITOR Constants".
attribute_list	List of user-created attributes. It is a comma-separated list of name-value pairs (for example, 'table_name=emp, operation=load').
session_id	Session ID of the session to be monitored. If omitted (or null), then the database monitors the current session.
session_serial	Serial number of the session to be monitored. If omitted (or null), then the database uses only the session ID to determine the session.

Return Values

This function returns the database operation execution ID. If the value is null for dbop_eid, then the database generates a unique value.

184.4.2 END_OPERATION Procedure

This procedure ends a database operation in the current session. If the specified database operation does not exist, then this function has no effect.

Syntax

```
DBMS_SQL_MONITOR.END_OPERATION(
   dbop_name       IN VARCHAR2,
   dbop_eid        IN NUMBER);
```

Parameters

Table 184-4 END_OPERATION Procedure Parameters

Parameter	Description
dbop_name	Name of a composite database operation
dbop_eid	Unique identifier for the current execution of the composite database operation

184.4.3 REPORT_SQL_MONITOR Function

This function builds a detailed report with monitoring information for a SQL statement, PL/SQL block, or database operation.

For each operation, it gives key information and associated global statistics. Use this function to get detailed monitoring information for a database operation.

The target database operation for this report can be:

- The last database operation monitored by Oracle Database (default, no parameter).
- The last database operation executed in the specified session and monitored by Oracle Database. The session is identified by its session ID and optionally its serial number (`-1` is current session).
- The last execution of a specific database operation identified by its `sql_id`.
- A specific execution of a database operation identified by the combination `sql_id`, `sql_exec_start`, and `sql_exec_id`.
- The last execution of a specific database operation identified by `dbop_name`.
- The specific execution of a database operation identified by the combination `dbop_name`, `dbop_exec_id`.

Syntax

```
DBMS_SQL_MONITOR.REPORT_SQL_MONITOR (
   sql_id                    IN VARCHAR2 DEFAULT  NULL,
   dbop_name                 IN VARCHAR2 DEFAULT  NULL,
   dbop_exec_id              IN NUMBER   DEFAULT  NULL,
   session_id                IN NUMBER   DEFAULT  NULL,
   session_serial            IN NUMBER   DEFAULT  NULL,
   sql_exec_start            IN DATE     DEFAULT  NULL,
   sql_exec_id               IN NUMBER   DEFAULT  NULL,
   inst_id                   IN NUMBER   DEFAULT  NULL,
   start_time_filter         IN DATE     DEFAULT  NULL,
   end_time_filter           IN DATE     DEFAULT  NULL,
   instance_id_filter        IN NUMBER   DEFAULT  NULL,
   parallel_filter           IN VARCHAR2 DEFAULT  NULL,
   plan_line_filter          IN NUMBER   DEFAULT  NULL,
   event_detail              IN VARCHAR2 DEFAULT  'YES',
   bucket_max_count          IN NUMBER   DEFAULT  128,
   bucket_interval           IN NUMBER   DEFAULT  NULL,
   base_path                 IN VARCHAR2 DEFAULT  NULL,
   last_refresh_time         IN DATE     DEFAULT  NULL,
   report_level              IN VARCHAR2 DEFAULT 'TYPICAL',
   type                      IN VARCHAR2 DEFAULT 'TEXT',
   sql_plan_hash_value       IN NUMBER   DEFAULT  NULL,
```

```
  con_name                  IN VARCHAR2 DEFAULT  NULL)
 RETURN CLOB;
```

Parameters

Table 184-5 *REPORT_SQL_MONITOR Procedure Parameters*

Parameter	Description
sql_id	SQL_ID of the simple database operation for which monitoring information should be displayed. Use NULL (default) to display monitoring information for the last simple database operation monitored by Oracle.
dbop_name	DBOP_NAME for which monitoring information of the composite database operation is displayed
dbop_exec_id	Execution ID for the composite database operation for which monitoring information is displayed
session_id	Targets only the subset of statements executed and monitored on behalf of the specified session. Default is NULL. Use -1 or USERENV('SID') for the current session.
session_serial	In addition to session_id, you can specify the session serial number to ensure the desired session incarnation is targeted. This is ignored when session_id is NULL.
sql_exec_start	Time at which execution of the monitored SQL was started. Only applicable when sql_id is specified. Used to display monitoring information for a particular execution of sql_id. When NULL (default), the last execution of sql_id is shown.
sql_exec_id	A numeric ID generated internally by SQL monitor to identify different executions of the same SQL statement. Thus each execution will have the same sql_id but a different sql_exec_id. Only applicable when sql_id is specified and is used to display monitoring information for a particular execution of sql_id. When NULL (default), the last execution of sql_id is shown.
inst_id	Looks only at queries started on the specified instance. Use -1 to target the current instance. The default, NULL will target all instances.
start_time_filter	If not NULL, the report shows activity from V$ACTIVE_SESSION_HISTORY started after this date. If NULL, the reported activity starts once the targeted database operation has started.
end_time_filter	If not NULL, the report shows activity from V$ACTIVE_SESSION_HISTORY started before this date. If NULL, the reported activity ends when the targeted database operation has ended or SYSDATE if the operation is still executing.
instance_id_filter	Only looks at activity for the specified instance. Use NULL (the default) to target all instances. Only relevant if the query runs in parallel.
parallel_filter	Parallel filter applies only to parallel execution and allows you to select only a subset of the processes involved in the parallel execution. The string parallel_filter can be: • NULL - target all parallel execution servers as wells as the query coordinator • ['qc'][servers(<svr_grp>[,] <svr_set>[,] <srv_num>)] where any NULL value is interpreted as ALL

Table 184-5 (Cont.) *REPORT_SQL_MONITOR Procedure Parameters*

Parameter	Description
`plan_line_filter`	Selects activity and execution statistics for the specified line number in the plan of a SQL.
`event_detail`	When set to `NO`, the activity is aggregated by `wait_class` only. Use `YES` (default) to aggregate by `wait_class`, `event_name`.
`bucket_max_count`	Specifies the maximum number of buckets to create in the report
`bucket_interval`	Represents the exact time interval, in seconds, of all histogram buckets. If specified, `bucket_max_count` is ignored.
`base_path`	URL path for flex HTML resources since flex HTML format requires access to external files (Java scripts and the flash `swf` file).
`last_refresh_time`	If not `NULL` (default), the time when the report was last retrieved (`SYSDATE` attribute of the report tag). Use this option when you want to display the report of an running query and when that report is refreshed on a regular basis. This optimizes the size of the report since only the new changed information will be returned. In particular, the following will be optimized: • SQL text will not be returned when this option is specified • Activity histogram will start at the bucket that intersects that time. The entire content of the bucket is returned, even if `last_refresh_time` is after the start of that bucket

Table 184-5 (Cont.) *REPORT_SQL_MONITOR Procedure Parameters*

Parameter	Description
report_level	Level of detail for the report. Of the following, only one can be specified: • NONE: Minimum possible • BASIC: This is equivalent to sql_text-plan-xplan-sessions-instance-activity_histogram-plan_histogram-metrics where the token "-" implies that report section will not be included in the report. • TYPICAL: Everything but plan_histogram • ALL: Everything In addition, individual report sections can also be enabled or disabled by using a ±*section_name*. Several sections are defined: • XPLAN: Shows explain plan. ON by default. • PLAN: Shows plan monitoring statistics. ON by default. • SESSIONS: Show session details. Applies only to parallel queries. ON by default. • INSTANCE: Shows instance details. Applies only to parallel and cross instance queries. ON by default. • PARALLEL: An umbrella parameter for specifying sessions as well as instance details • ACTIVITY: Shows activity summary at global level, plan line level and session • INSTANCE LEVEL: (If applicable). ON by default. • BINDS: Shows bind information when available. ON by default. • METRICS: Shows metric data (such as CPU and IOs) over time. ON by default • ACTIVITY_HISTOGRAM: Shows a histogram of the overall query activity. ON by default. • PLAN_HISTOGRAM: Shows activity histogram at plan line level. OFF by default. • OTHER: Other information. ON by default. In addition, SQL text can be specified at different levels: • -SQL_TEXT: No SQL text in report • +SQL_TEXT: Alright with partial SQL text, that is, up to the first 2000 chars as stored in GV$SQL_MONITOR • SQL_FULLTEXT: No full SQL text, that is, +sql_text • +SQL_FULLTEXT: Show full SQL text (default)
type	Report type: • TEXT: text report (default) • HTML: simple HTML report • ACTIVE: database active report. Some information (explain plan, activity_histogram, metrics and plan_histogram) is only shown when this type is selected • XML: raw data for the report
sql_plan_hash_value	Targets only those with the specified plan hash value. Default is NULL.
con_name	Container name in a multitenant database.

Return Values

SQL monitor report, an XML document.

Usage Notes

The user invoking this function must have privilege to access the following fixed views:

- GV$SQL_MONITOR
- GV$SQL_PLAN_MONITOR
- GV$ACTIVE_SESSION_HISTORY
- GV$SESSION_LONGOPS
- GV$SQL if SQL full text is requested and its length is greater than 2 KB

184.4.4 REPORT_SQL_MONITOR_XML Function

This function is identical to the REPORT_SQL_MONITOR function, except that the return type is XMLType.

Related Topics

- REPORT_SQL_MONITOR Function
  This function builds a detailed report with monitoring information for a SQL statement, PL/SQL block, or database operation.

184.4.5 REPORT_SQL_MONITOR_LIST Function

This function builds a report for all or a subset of database operations that have been monitored by Oracle Database.

For each database operation, it gives key information and associated global statistics.

Syntax

```
DBMS_SQL_MONITOR.REPORT_SQL_MONITOR_LIST (
   sql_id                    IN VARCHAR2 DEFAULT  NULL,
   dbop_name                 IN VARCHAR2 DEFAULT  NULL,
   monitor_type              IN NUMBER   DEFAULT  MONITOR_TYPE_ALL,
   session_id                IN NUMBER   DEFAULT  NULL,
   session_serial            IN NUMBER   DEFAULT  NULL,
   inst_id                   IN NUMBER   DEFAULT  NULL,
   active_since_date         IN DATE     DEFAULT  NULL,
   active_since_sec          IN NUMBER   DEFAULT  NULL,
   last_refresh_time         IN DATE     DEFAULT  NULL,
   report_level              IN VARCHAR2 DEFAULT 'TYPICAL',
   auto_refresh              IN NUMBER   DEFAULT  NULL,
   base_path                 IN VARCHAR2 DEFAULT  NULL,
   type                      IN VARCHAR2 DEFAULT 'TEXT',
   con_name                  IN VARCHAR2 DEFAULT  NULL)
  RETURN CLOB;
```

Parameters

Table 184-6 *REPORT_SQL_MONITOR_LIST Procedure Parameters*

Parameter	Description
sql_id	SQL_ID of the simple database operation for which monitoring information should be displayed. Use NULL (default) to display monitoring information for the last operation monitored by Oracle Database.
dbop_name	DBOP_NAME for which monitoring information of the composite database operation is displayed.
monitor_type	Monitor type: • MONITOR_TYPE_SQL returns only simple database operations • MONITOR_TYPE_DBOP returns composite database operations • MONITOR_TYPE_ALL returns all types
session_id	Targets only the subset of database operations executed and monitored on behalf of the specified session. Default is NULL. Use -1 or USERENV('SID') for the current session.
session_serial	In addition tosession_id, you can specify the session serial number to ensure the desired session incarnation is targeted. This is ignored when session_id is NULL.
inst_id	Looks only at monitored database operations originating from the specified instance. Use -1 to target the instance where the report executed. To target all instances, use NULL (default).
active_since_date	If not NULL (default), returns monitored database operations that have been active since the specified time. This includes all operations that are executing, as well as all operations that have completed their execution after the specified start time.
active_since_sec	If not NULL (default), returns monitored database operations that have been active since the specified time. This includes all operations that are executing, as well as all operations that have completed their execution after the specified date and time. In this case, the start time is specified relative to the current SYSDATE minus a specified number of seconds. For example, use 3600 to limit the report to all operations that have been active in the past 1 hour.
last_refresh_time	If not NULL (default), the time when the list report was last retrieved. This optimizes the case where an application shows the list and refreshes the report on a regular basis (such as once every 5 seconds). In this case, the report will show details about the execution of monitored queries that have been active since the specified last_refresh_time. For other queries, the report returns the execution key (sql_id, sql_exec_start, and sql_exec_id). Also, for queries that have their first refresh time after the specified date, only the SQL execution key and statistics are returned.
report_level	Level of detail for the report. The level can be BASIC (SQL text up to 200 character), TYPICAL (which include full SQL text assuming that cursor has not aged out, in which case the SQL text is included up to 2000 characters), or ALL which is the same as TYPICAL.
auto_refresh	Specifies the duration in seconds after which report data will be automatically refreshed while the monitored SQL or database operation is still executing. This applies to active report types.

Table 184-6 (Cont.) *REPORT_SQL_MONITOR_LIST Procedure Parameters*

Parameter	Description
`base_path`	URL path for flex HTML resources since flex HTML format requires access to external files (java scripts and the flash `swf` file).
`type`	Report type: • `TEXT`: text report (default) • `HTML`: simple HTML report • `ACTIVE`: database active report. Some information (explain plan, `activity_histogram`, metrics, and `plan_histogram`) is only shown when this type is selected. • `XML`: raw data for the report
`con_name`	Container name in a multitenant database.

Return Values

A report in text, XML, or HTML format that contains the list of the database operations monitored.

Usage Notes

- Use the REPORT_SQL_MONITOR Function to get detailed monitoring information for a single database operation.
- The user invoking this function needs to have the privilege to access the fixed views `GV$SQL_MONITOR` and `GV$SQL`.

184.4.6 REPORT_SQL_MONITOR_LIST_XML Function

This function is identical to the `REPORT_SQL_MONITOR_LIST` function, except that it returns `XMLType`.

Related Topics

- REPORT_SQL_MONITOR_LIST Function
  This function builds a report for all or a subset of database operations that have been monitored by Oracle Database.

185
DBMS_SQL_TRANSLATOR

The `DBMS_SQL_TRANSLATOR` package provides an interface for creating, configuring, and using SQL translation profiles.

This chapter contains the following topics:

- Security Model
- Constants
- Operational Notes
- Exceptions
- Examples
- Summary of DBMS_SQL_TRANSLATOR Subprograms

> **See Also:**
>
> SQL Translation Framework Architecture and Overview in*Oracle Database Migration Guide*

185.1 DBMS_SQL_TRANSLATOR Security Model

`DBMS_SQL_TRANSLATOR` is a invoker's rights package.

When translating a SQL statement or error, the translator package procedure will be invoked with the same current user and current schema as those in which the SQL statement being parsed. The owner of the translator package must be granted the `TRANSLATE SQL` user privilege on the current user.

Additionally, the current user must be granted the `EXECUTE` privilege on the translator package.

185.2 DBMS_SQL_TRANSLATOR Constants

`DBMS_SQL_TRANSLATOR` defines several constants to use when specifying parameter values.

These are shown in the following table.

Table 185-1 DBMS_SQL_TRANSLATOR Constants

Constant	Value	Type	Description
ATTR_EDITIONABLE	'EDITIONABLE'	VARCHAR2(30)	Name of the SQL translation profile attribute that specifies whether the SQL translation profile becomes an editioned or noneditioned object if editioning is later enabled for the schema object type SQL translation profile in the owner's schema (see Operational Notes)
ATTR_FOREIGN_SQL_SYNTAX	'FOREIGN_SQL_SYNTAX'	VARCHAR2(30)	Name of the SQL translation profile attribute that indicates if the profile is for translation of foreign SQL syntax (see Operational Notes)
ATTR_LOG_TRANSLATION_ERROR	'TRANSLATION_ERROR'	VARCHAR2(30)	Name of the SQL translation profile attribute that controls if the profile should log translation error in the database alert log (see Operational Notes)
ATTR_RAISE_TRANSLATION_ERROR	'TRANSLATION_ERROR '	VARCHAR2(30)	Name of the SQL translation profile attribute that controls if the profile should raise translation error if a SQL statement or error fails to be translated (see Operational Notes)
ATTR_TRANSLATE_NEW_SQL	'TRANSLATE_NEW_SQL'	VARCHAR2(30)	Name of the SQL translation profile attribute that controls if the profile should translate new SQL statements and errors (see Operational Notes)
ATTR_TRACE_TRANSLATION	'TRACE_TRANSLATION'	VARCHAR2(30)	Name of the SQL translation profile attribute that controls tracing (see Operational Notes)
ATTR_TRANSLATOR	'TRANSLATOR'	VARCHAR2(30)	Name of the SQL translation profile attribute that specifies the translator package (see Operational Notes)
ATTR_VALUE_TRUE	'TRUE'	VARCHAR2(30)	Value to set a SQL translation profile attribute to true (see Operational Notes)
ATTR_VALUE_FALSE	'FALSE'	VARCHAR2(30)	Value to set a SQL translation profile attribute to false (see Operational Notes)

185.3 DBMS_SQL_TRANSLATOR Operational Notes

The subprograms that modify a profile have DDL transaction semantics and when invoked will commit any open transaction in the session.

ATTR_EDITIONABLE Constant

Editionable is true by default.

ATTR_FOREIGN_SQL_SYNTAX Constant

Foreign SQL syntax is true by default.

ATTR_LOG_TRANSLATION_ERROR Constant

- If log translation is enabled in a SQL translation profile, an alert log is written to the database alert log if no custom translation is found for a SQL statement or error. This allows the user to catch any error in the custom translation in a profile.
- Log translation error is false by default.

ATTR_RAISE_TRANSLATION_ERROR Constant

Raise translation error is false by default.

ATTR_TRANSLATE_NEW_SQL Constant

- The name of the SQL translation profile attribute that controls if the profile should translate new SQL statements and errors. If so, the translator package, if registered, will translate a new SQL statement or error not already translated in custom translations, and also register the new translation as custom translation. If not, any new SQL statement or error encountered will result in a translation error
- Translate new SQL statements and errors is true by default.

ATTR_TRACE_TRANSLATION Constant

- If tracing is enabled in a SQL translation profile, any SQL statement or error translated by the profile in a database session and its translation is written to the database session's trace file.
- Tracing is disabled by default.

ATTR_TRANSLATOR Constant

- The translator package must be a PL/SQL package with the following three procedures. The TRANSLATE_SQL Procedure and the TRANSLATE_ERROR Procedure are called to translate SQL statements and errors. The names of the parameters of the translate procedures must be followed.

```
PROCEDURE TRANSLATE_SQL(
   sql_text        IN  CLOB,
   translated_text OUT CLOB);

PROCEDURE TRANSLATE_ERROR(
   error_code          IN  BINARY_INTEGER,
   translated_code     OUT BINARY_INTEGER,
   translated_sqlstate OUT VARCHAR2);

Parameters:
  profile_name        - profile name
  sql_text            - SQL statement to be translated
  translated_text     - translated SQL statement
  error_code          - Oracle error code
  translated_code     - translated error code
  translated_sqlstate - translated SQLSTATE
```

- When `NULL` is returned in `translated_text`, `translated_code`, or `translated_sqlstate`, it means that no translation is required and the original SQL statement, error code, or `SQLSTATE` is used instead.
- The name of the translator package follows the naming rules for database packages of the form `[schema.]package_name`. When the schema and package names are used, they are set to uppercase by default unless surrounded by double quotation marks. For example, setting a translator package, `translator => 'dbms_tsql_translator'` is the same as `translator => 'Dbms_Tsql_Translator'` and `translator => 'DBMS_TSQL_TRANSLATOR'`, but not the same as `translator => '"dbms_tsql_translator"'`. The default schema name is the profile owner.
- The translator attribute is not set by default.

ATTR_VALUE_TRUE Constant

The value to set a SQL translation profile attribute to true.

ATTR_VALUE_FALSE Constant

The value to set a SQL translation profile attribute to false.

185.4 DBMS_SQL_TRANSLATOR Exceptions

This table lists the exceptions raised by the `DBMS_SQL_TRANSLATOR` package.

Table 185-2 Exceptions Raised by DBMS_SQL_TRANSLATOR

Exception	Error Code	Description
BAD_ARGUMENT	29261	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	1031	User has insufficient privilege for the operation
NO_SUCH_PROFILE	24252	Profile does not exist
NO_SUCH_USER	1918	Profile owner does not exist
NO_TRANSLATION_FOUND	24253	No translation of the SQL statement or error code found
PROFILE_EXISTS	955	Profile already exists

185.5 DBMS_SQL_TRANSLATOR Examples

This is an example of basic SQL translation using `DBMS_SQL_TRANSLATOR`.

Basic SQL Translation

```
BEGIN
  DBMS_SQL_TRANSLATOR.CREATE_PROFILE(
    profile_name     => ' tsql_application');
  DBMS_SQL_TRANSLATOR.SET_ATTRIBUTE(
    profile_name     =>  'tsql_application',
    attribute_name   =>  DBMS_SQL_TRANSLATOR.ATTR_TRANSLATOR,
    attribute_value  =>  'migration_repo.sybase_tsql_translator');
END;
```

185.6 Summary of DBMS_SQL_TRANSLATOR Subprograms

This table lists the `DBMS_SQL_TRANSLATOR` subprograms and briefly describes them.

Table 185-3 DBMS_SQL_TRANSLATOR Package Subprograms

Subprogram	Description
CREATE_PROFILE Procedure	Creates a SQL translation profile
DEREGISTER_SQL_TRANSLATION Procedure	Deregisters the custom translation of a SQL statement in a SQL translation profile

Table 185-3 (Cont.) DBMS_SQL_TRANSLATOR Package Subprograms

Subprogram	Description
DEREGISTER_ERROR_TRANSLATION Procedure	Deregisters the translation of an Oracle error code and SQLSTATE in a SQL translation profile
DROP_PROFILE Procedure	Drops a SQL translation profile and its contents
ENABLE_ERROR_TRANSLATION Procedure	Enables or disables a custom translation of an Oracle error code in a SQL translation profile
ENABLE_SQL_TRANSLATION Procedure	Enables or disables a custom translation of a SQL statement in a SQL translation profile
EXPORT_PROFILE Procedure	Exports the content of a SQL translation profile
IMPORT_PROFILE Procedure	Imports the content of a SQL translation profile
REGISTER_ERROR_TRANSLATION Procedure	Registers a custom translation of an Oracle error code and SQLSTATE in a SQL translation profile
REGISTER_SQL_TRANSLATION Procedure	Registers a custom translation of a SQL statement in a SQL translation profile
SET_ATTRIBUTE Procedure	Sets an attribute of a SQL translation profile
SQL_HASH Function	Computes the hash value of a SQL statement in a SQL translation profile
SQL_ID Function	Computes the SQL identifier of a SQL statement in a SQL translation profile
TRANSLATE_ERROR Procedure	Translates an Oracle error code and an ANSI SQLSTATE using a SQL translation profile
TRANSLATE_SQL Procedure	Translates a SQL statement using a SQL translation profile

185.6.1 CREATE_PROFILE Procedure

This procedure creates a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.CREATE_PROFILE (
   profile_name    IN  VARCHAR2);
```

Parameters

Table 185-4 *CREATE_PROFILE Procedure Parameters*

Parameter	Description
profile_name	Name of profile

Exceptions

Table 185-5 CREATE_PROFILE Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
PROFILE_EXISTS	Profile already exists

Usage Notes

- A SQL translation profile is a database schema object that resides in SQL translation profile namespace. Its name follows the naming rules for database objects of the form `[schema.]name`. When the schema and profile names are used in the `DBMS_SQL_TRANSLATOR` package, they are uppercased unless surrounded by double quotation marks. For example, the translation profile `profile_name => 'tsql_application'` is the same as `profile_name => 'Tsql_Application'` and `profile_name => 'TSQL_APPLICATION'`, but not the same as `profile_name => '"tsql_application"'`.
- A SQL translation profile is an editionable object type.
- A SQL translation profile cannot be created as a common object in a multitenant container database (CDB).
- To destroy a SQL translation profile, use the DROP_PROFILE Procedure.

Examples

```
BEGIN
  DBMS_SQL_TRANSLATOR.CREATE_PROFILE(profile_name => 'tsql_application');
END;
```

185.6.2 DEREGISTER_SQL_TRANSLATION Procedure

This procedure deregisters the custom translation of a SQL statement in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.DEREGISTER_SQL_TRANSLATION (
   profile_name      IN  VARCHAR2,
   sql_text          IN  CLOB);
```

Parameters

Table 185-6 *DEREGISTER_SQL_TRANSLATION Procedure Parameters*

Parameter	Description
profile_name	Name of profile
sql_text	SQL statement

Exceptions

Table 185-7 DEREGISTER_SQL_TRANSLATION Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
PROFILE_EXISTS	Profile already exists

Examples

```
BEGIN
    DBMS_SQL_TRANSLATOR.DEREGISTER_SQL_TRANSLATION(
      profile_name => 'tsql_application',
      sql_text     => 'select top 5 * from emp');
END;
```

185.6.3 DEREGISTER_ERROR_TRANSLATION Procedure

This procedure deregisters the translation of an Oracle error code and SQLSTATE in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.DEREGISTER_ERROR_TRANSLATION (
   profile_name          IN   VARCHAR2,
   error_code            IN   PLS_INTEGER);
```

Parameters

Table 185-8 *DEREGISTER_ERROR_TRANSLATION Procedure Parameters*

Parameter	Description
profile_name	Name of profile
error_code	Oracle error code

Exceptions

Table 185-9 DEREGISTER_ERROR_TRANSLATION Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Examples

```
BEGIN
   DBMS_SQL_TRANSLATOR.DEREGISTER_ERROR_TRANSLATION(
      profile_name    => 'tsql_application',
      error_code      => 1);
END;
```

185.6.4 DROP_PROFILE Procedure

This procedure drops a SQL translation profile and its contents.

Syntax

```
DBMS_SQL_TRANSLATOR.DROP_PROFILE (
   profile_name     IN    VARCHAR2);
```

Parameters

Table 185-10 *DROP_PROFILE Procedure Parameters*

Parameter	Description
profile_name	Name of profile

Exceptions

Table 185-11 DROP_PROFILE Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Examples

```
BEGIN
    DBMS_SQL_TRANSLATOR.DROP_PROFILE(
      profile_name    =>  'tsql_application');
END;
```

185.6.5 ENABLE_ERROR_TRANSLATION Procedure

This procedure enables or disables a custom translation of an Oracle error code in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.ENABLE_ERROR_TRANSLATION (
   profile_name      IN  VARCHAR2,
   sql_text          IN CLOB,
   enable            IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 185-12 *ENABLE_ERROR_TRANSLATION Procedure Parameters*

Parameter	Description
profile_name	Name of profile
sql_text	SQL statement
enable	Enable or disable the translation

Exceptions

Table 185-13 ENABLE_ERROR_TRANSLATION Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Examples

```
BEGIN
  DBMS_SQL_TRANSLATOR.ENABLE_ERROR_TRANSLATION(
      profile_name => 'tsql_application',
      sql_text     => 'SELECT TOP 5 * FROM emp'
      enable       => TRUE);
END;
```

185.6.6 ENABLE_SQL_TRANSLATION Procedure

This procedure enables or disables a custom translation of a SQL statement in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.ENABLE_SQL_TRANSLATION (
   profile_name      IN    VARCHAR2,
   sql_text          IN    CLOB,
   enable            IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 185-14 *ENABLE_SQL_TRANSLATION Procedure Parameters*

Parameter	Description
profile_name	Name of profile
sql_text	SQL statement
enable	Enable or disable the translation

Exceptions

Table 185-15 ENABLE_SQL_TRANSLATION Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Examples

```
BEGIN
  DBMS_SQL_TRANSLATOR.ENABLE_SQL_TRANSLATION(
      profile_name => 'tsql_application',
      sql_text     => 'select top 5 * from emp',
      enable       => TRUE);
END;
```

185.6.7 EXPORT_PROFILE Procedure

This procedure exports the content of a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.EXPORT_PROFILE (
   profile_name     IN          VARCHAR2,
   content          OUT NOCOPY  CLOB);
```

Parameters

Table 185-16 *EXPORT_PROFILE Procedure Parameters*

Parameter	Description
profile_name	Name of profile
content	Content of profile

Exceptions

Table 185-17 EXPORT_PROFILE Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Usage Notes

- The content of the SQL translation profile is exported in XML format as follows. Note that the profile name will not be exported.

```
SQLTranslationProfile Translator="translator package name"
                      ForeignSQLSyntax="TRUE|FALSE"
                      TranslateNewSQL="TRUE|FALSE"
                      RaiseTranslationError="TRUE|FALSE"
                      LogTranslationError="TRUE|FALSE"
                      TraceTranslation="TRUE|FALSE"
                      Editionable="TRUE|FALSE">
  <SQLTranslations>
    <SQLTranslation Enabled="TRUE|FALSE">
      <SQLText>original SQL text</SQLText>
      <TranslatedText>translated SQL text</TranslatedText>
    </SQLTranslation>
    ...
  </SQLTranslations>
  <ErrorTranslations>
    <ErrorTranslation Enabled="TRUE|FALSE">
      <ErrorCode>Oracle error code</ErrorCode>
      <TranslatedCode>translated error code</TranslatedCode>
      <TranslatedSQLSTATE>translated SQLSTATE</TranslatedSQLSTATE>
    </ErrorTranslation>
    ...
  </ErrorTranslations>
</SQLTranslationProfile>
```

- To import the content to a SQL translation profile, use the IMPORT_PROFILE Procedure.

Examples

```
DECLARE
  content CLOB;
BEGIN
    DBMS_SQL_TRANSLATOR.EXPORT_PROFILE(
      profile_name    =>  'tsql_application',
      content         =>  content);
END;
```

185.6.8 IMPORT_PROFILE Procedure

This procedure imports the content of a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.IMPORT_PROFILE (
   profile_name       IN   VARCHAR2,
   content            IN   CLOB);
```

Parameters

Table 185-18 ***IMPORT_PROFILE Procedure Parameters***

Parameter	Description
profile_name	Name of profile
content	Content of profile

Exceptions

Table 185-19 IMPORT_PROFILE Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist

Usage Notes

- The content of the SQL translation profile must be in XML format as used by the EXPORT_PROFILE Procedure. All elements and attributes are optional.
- If the profile does not exist, it is created. If it exists, the content overrides any existing attribute, translator package, SQL or error translation registration.
- To export the content to a SQL translation profile, use the EXPORT_PROFILE Procedure.

Examples

```
DECLARE
  content CLOB;
BEGIN
    DBMS_SQL_TRANSLATOR.IMPORT_PROFILE(
      profile_name    =>  'tsql_application',
      content         =>  content);
END;
```

185.6.9 REGISTER_ERROR_TRANSLATION Procedure

This procedure registers a custom translation of an Oracle error code and SQLSTATE in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.REGISTER_ERROR_TRANSLATION (
   profile_name          IN   VARCHAR2,
   error_code            IN   PLS_INTEGER,
   translated_code       IN   PLS_INTEGER DEFAULT NULL,
   translated_sqlstate   IN   VARCHAR2 DEFAULT NULL,
   enable                IN   BOOLEAN DEFAULT TRUE);
```

Parameters

Table 185-20 *REGISTER_ERROR_TRANSLATION Procedure Parameters*

Parameter	Description
profile_name	Name of profile
error_code	Oracle error code
translated_code	Translated error code
translated_sqlstate	Translated SQLSTATE
enable	Enable or disable the translation

Exceptions

Table 185-21 REGISTER_ERROR_TRANSLATION Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Usage Notes

- When the Oracle Database translates an Oracle error code using a translation profile, it searches for the registered custom translation first, and only invokes the translator package if no match is found.
- When a translation is registered in a profile, it may be disabled. Oracle Database does not search for disabled translations.
- The old translation of the error code and SQLSTATE, if present, is replaced with the new translation.
- To deregister a translation, use the DEREGISTER_ERROR_TRANSLATION Procedure.

Examples

```
BEGIN
      DBMS_SQL_TRANSLATOR.REGISTER_ERROR_TRANSLATION(
      profile_name     => 'tsql_application',
      error_code       => 1,
      translated_code => 2601);
END;
```

185.6.10 REGISTER_SQL_TRANSLATION Procedure

This procedure registers a custom translation of a SQL statement in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.REGISTER_SQL_TRANSLATION (
   profile_name      IN VARCHAR2,
   sql_text          IN CLOB,
   translated_text   IN CLOB DEFAULT NULL,
   enable            IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 185-22 *REGISTER_SQL_TRANSLATION Procedure Parameters*

Parameter	Description
profile_name	Name of profile
sql_text	SQL statement
translated_text	Translated SQL statement

Table 185-22 (Cont.) *REGISTER_SQL_TRANSLATION Procedure Parameters*

Parameter	Description
enable	Enable or disable the translation

Exceptions

Table 185-23 REGISTER_SQL_TRANSLATION Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Usage Notes

- When the Oracle Database translates a statement using a translation profile, it searches for the registered custom translation first, and only invokes the translator package if no match is found.
- When a translation is registered in a profile, it may be disabled. Oracle Database does not search for disabled translations.
- When `translated_text` is `NULL`, no translation is required and the original statement is used.
- The old translation of the SQL statement, if present, is replaced with the new translation.
- To deregister a translation, use the DEREGISTER_SQL_TRANSLATION Procedure.

Examples

```
BEGIN
    DBMS_SQL_TRANSLATOR.REGISTER_SQL_TRANSLATION(
      profile_name    => 'tsql_application',
      sql_text        => 'select top 5 * from emp',
      translated_text => 'SELECT * FROM emp WHERE rownum <= :SYS_N_001');
END;
```

185.6.11 SET_ATTRIBUTE Procedure

This procedure sets an attribute of a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.SET_ATTRIBUTE (
   profile_name      IN  VARCHAR2,
   attribute_name    IN  VARCHAR2,
   attribute_value   IN  VARCHAR2;)
```

Parameters

Table 185-24 *SET_ATTRIBUTE Procedure Parameters*

Parameter	Description
profile_name	Name of profile
attribute_name	Name of attribute
attribute_value	Value of attribute

Exceptions

Table 185-25 SET_ATTRIBUTE Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Usage Notes

See Constants

185.6.12 SQL_HASH Function

This procedure computes the hash value of a SQL statement in the session's SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.SQL_HASH (
   sql_text         IN    CLOB)
RETURN NUMBER DETERMINISTIC;
```

Parameters

Table 185-26 *SQL_HASH Function Parameters*

Parameter	Description
sql_text	SQL statement

Return Values

Returns hash value of the SQL statement in the SQL translation profile

Exceptions

Table 185-27 SQL_HASH Function Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface

Examples

```
DECLARE
  sqltext CLOB;
  txltext CLOB;
  sqlhash NUMBER;
BEGIN
  sqltext  :=  'SELECT TOP 1 * FROM emp';
  sqlhash  :=  DBMS_SQL_TRANSLATOR.SQL_HASH (sqltext);
  SELECT translated_text INTO txltext
    FROM user_sql_translations
    WHERE sql_hash = sqlhash
    AND DBMS_LOB.COMPARE (sql_text, sqltext) = 0;
END;
```

185.6.13 SQL_ID Function

This procedure computes the SQL identifier of a SQL statement in a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.SQL_ID (
   sql_text         IN    CLOB)
RETURN VARCHAR2 DETERMINISTIC;
```

Parameters

Table 185-28 *SQL_ID Function Parameters*

Parameter	Description
sql_text	SQL statement

Return Values

Returns the SQL ID of the SQL statement in the SQL translation profile

Exceptions

Table 185-29 SQL_ID Function Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface

Examples

```
DECLARE
  sqltext CLOB;
```

```
  sqlid   VARCHAR2(13);
BEGIN
  sqltext :=  'SELECT TOP 1 * FROM emp';
  sqlid   :=  DBMS_SQL_TRANSLATOR.SQL_ID (sqltext);
END;
```

185.6.14 TRANSLATE_ERROR Procedure

This procedure translates an Oracle error code and an ANSI SQLSTATE using the session's SQL translation profile

Syntax

```
DBMS_SQL_TRANSLATOR.TRANSLATE_ERROR (
   error_code           IN           PLS_INTEGER,
   translated_code      OUT          PLS_INTEGER,
   translated_sqlstate  OUT NOCOPY   VARCHAR2);
```

Parameters

Table 185-30 *TRANSLATE_ERROR Procedure Parameters*

Parameter	Description
error_code	Oracle error code
translated_code	Translated error code
translated_sqlstate	Translated SQLSTATE

Exceptions

Table 185-31 TRANSLATE_ERROR Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist
NO_TRANSLATION_FOUND	No translation of the SQL statement or error code is found

Examples

```
DECLARE
  translated_code      BINARY_INTEGER;
  translated_sqlstate  VARCHAR2(5);
BEGIN
    DBMS_SQL_TRANSLATOR.TRANSLATE_ERROR(
      error_code           => 1,
      translated_code      => translated_code,
      translated_sqlstate  => translated_sqlstate);
END;
```

185.6.15 TRANSLATE_SQL Procedure

This procedure translates a SQL statement using a SQL translation profile.

Syntax

```
DBMS_SQL_TRANSLATOR.TRANSLATE_SQL (
   sql_text          IN           CLOB,
   translated_text   OUT  NOCOPY  CLOB);
```

Parameters

Table 185-32 *TRANSLATE_SQL Procedure Parameters*

Parameter	Description
sql_text	SQL statement
translated_text	Translated SQL statement

Exceptions

Table 185-33 TRANSLATE_SQL Procedure Exceptions

Exception	Description
BAD_ARGUMENT	Bad argument is passed to the PL/SQL interface
INSUFFICIENT_PRIVILEGE	User has insufficient privilege for the operation
NO_SUCH_USER	Profile owner does not exist
NO_SUCH_PROFILE	Profile does not exist

Examples

```
ALTER SESSION SET SQL_TRANSLATION_PROFILE =  tsql_application;

DECLARE
  translated_text CLOB;
BEGIN
    DBMS_SQL_TRANSLATOR.TRANSLATE_SQL(
      sql_text         => 'select top 5 * from emp',
      translated_text  => translated_text);
END;
```

186
DBMS_SQLDIAG

The DBMS_SQLDIAG package provides an interface to the SQL Diagnosability functionality.

This chapter contains the following topics:

- DBMS_SQLDIAG Overview
- DBMS_SQLDIAG Security Model
- DBMS_SQLDIAG Constants
- Summary of DBMS_SQLDIAG Subprograms

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more information about "Managing Diagnostic Data"

186.1 DBMS_SQLDIAG Overview

In the rare case that a SQL statement fails with a critical error, you can run the SQL Repair Advisor to try to repair the failed statement by using the `DBMS_SQLDIAG` package subprograms.

The SQL Repair Advisor analyzes the statement and in many cases recommends a patch to repair the statement. If you implement the recommendation, the applied SQL patch circumvents the failure by causing the query optimizer to choose an alternate execution plan for future executions.

> **See Also:**
>
> *Oracle Database Administrator's Guide* for more information about how to run the SQL Repair Advisor using the `DBMS_SQLDIAG` package subprograms.

186.2 DBMS_SQLDIAG Security Model

You must have the `ADVISOR` role to execute the `DBMS_SQLDIAG` package.

186.3 DBMS_SQLDIAG Constants

`DBMS_SQLDIAG` defines constants to use when specifying parameter values.

These constants are shown in the following tables:

- Table 186-1 describes the name of SQL repair advisor as seen by the advisor framework

- Table 186-2 describes SQLDIAG advisor task scope parameter values
- Table 186-3 describes SQLDIAG advisor `time_limit` constants
- Table 186-4 describes possible formats for a report
- Table 186-5 describes possible levels of detail in the report
- Table 186-6 describes possible report sections (comma delimited)
- Table 186-7 describes possible values for the `problem_type` parameter of the CREATE_DIAGNOSIS_TASK Functions
- Table 186-8 describes possible values for the `_sql_findings_mode` parameter

Table 186-1 DBMS_SQLDIAG Constants - SQLDIAG Advisor Name

Constant	Type	Value	Description
ADV_SQL_DIAG_NAME	VARCHAR2(18)	SQL Repair Advisor	Name of SQL repair advisor as seen by the advisor framework

Table 186-2 DBMS_SQLDIAG Constants - SQLDIAG Advisor Task Scope Parameter Values

Constant	Type	Value	Description
SCOPE_COMPREHENSIVE	VARCHAR2(13)	COMPREHENSIVE	Detailed analysis of the problem which may take more time to execute
SCOPE_LIMITED	VARCHAR2(7)	LIMITED	Brief analysis of the problem

Table 186-3 DBMS_SQLDIAG Constants - SQLDIAG Advisor time_limit Constants

Constant	Type	Value	Description
TIME_LIMIT_DEFAULT	NUMBER	1800	Default time limit for analysis of the problem

Table 186-4 DBMS_SQLDIAG Constants - Report Type (possible values) Constants

Constant	Type	Value	Description
TYPE_HTML	VARCHAR2(4)	HTML	Report from the REPORT_DIAGNOSIS_TASK Function in HTML form
TYPE_TEXT	VARCHAR2(4)	TEXT	Report from the REPORT_DIAGNOSIS_TASK Function in text form
TYPE_XML	VARCHAR2(3)	XML	Report from the REPORT_DIAGNOSIS_TASK Function in XML form

Table 186-5 DBMS_SQLDIAG Constants - Report Level (possible values) Constants

Constant	Type	Value	Description
LEVEL_ALL	VARCHAR2(3)	ALL	Complete report including annotations about statements skipped over
LEVEL_BASIC	VARCHAR2(5)	BASIC	Shows information about every statement analyzed, including recommendations not implemented
LEVEL_TYPICAL	VARCHAR2(7)	TYPICAL	Simple report shows only information about the actions taken by the advisor.

Table 186-6 DBMS_SQLDIAG Constants - Report Section (possible values) Constants

Constant	Type	Value	Description
SECTION_ALL	VARCHAR2(3)	ALL	All statements
SECTION_ERRORS	VARCHAR2(6)	ERRORS	Statements with errors
SECTION_FINDINGS	VARCHAR2(8)	FINDINGS	Tuning findings
SECTION_INFORMATION	VARCHAR2(11)	INFORMATION	General information
SECTION_PLANS	VARCHAR2(5)	PLANS	Explain plans
SECTION_SUMMARY	VARCHAR2(7)	SUMMARY	Summary information

Table 186-7 DBMS_SQLDIAG Constants - Problem Type Constants

Constant	Type	Value	Description
PROBLEM_TYPE_PERFORMANCE	NUMBER	1	User suspects this is a performance problem
PROBLEM_TYPE_WRONG_RESULTS	NUMBER	2	User suspects the query is giving inconsistent results
PROBLEM_TYPE_COMPILATION_ERROR	NUMBER	3	User sees a crash in compilation
PROBLEM_TYPE_EXECUTION_ERROR	NUMBER	4	User sees a crash in execution
PROBLEM_TYPE_ALT_PLAN_GEN	NUMBER	5	User to explore all alternative plans

Table 186-8 DBMS_SQLDIAG Constants - Findings Filter Constants

Constant	Type	Value	Description
SQLDIAG_FINDINGS_ALL	NUMBER	1	Show all possible findings

Table 186-8 (Cont.) DBMS_SQLDIAG Constants - Findings Filter Constants

Constant	Type	Value	Description
SQLDIAG_FINDINGS_VALIDATION	NUMBER	2	Show status of validation rules over structures
SQLDIAG_FINDINGS_FEATURES	NUMBER	3	Show only features used by the query
SQLDIAG_FINDINGS_FILTER_PLANS	NUMBER	4	Show the alternative plans generated by the advisor
SQLDIAG_FINDINGS_CR_DIFF	NUMBER	5	Show difference between two plans
SQLDIAG_FINDINGS_MASK_VARIANT	NUMBER	6	Mask info for testing
SQLDIAG_FINDINGS_OBJ_FEATURES	NUMBER	7	Show features usage history
SQLDIAG_FINDINGS_BASIC_INFO	NUMBER	8	Show the alternative plans generated by the advisor

186.4 Summary of DBMS_SQLDIAG Subprograms

This table lists the DBMS_SQLDIAG subprograms and briefly describes them.

Table 186-9 DBMS_SQLDIAG Package Subprograms

Subprogram	Description
ACCEPT_SQL_PATCH Function & Procedure	Accepts a recommended SQL patch as recommended by the specified SQL diagnosis task
ALTER_SQL_PATCH Procedure	Alters specific attributes of an existing SQL patch object
CANCEL_DIAGNOSIS_TASK Procedure	Cancels a diagnostic task
CREATE_DIAGNOSIS_TASK Functions	Creates a diagnostic task in order to diagnose a single SQL statement
CREATE_SQL_PATCH Function	Creates an SQL patch based on a set of user specified hints for specific statements identified by SQL text.
CREATE_STGTAB_SQLPATCH Procedure	Creates the staging table used for transporting SQL patches from one system to another
DROP_DIAGNOSIS_TASK Procedure	Drops a diagnostic task
DROP_SQL_PATCH Procedure	Drops the named SQL patch from the database
EXECUTE_DIAGNOSIS_TASK Procedure	Executes a diagnostic task
EXPLAIN_SQL_TESTCASE Function	Explains a SQL test case
EXPORT_SQL_TESTCASE Procedures	Exports a SQL test case to a directory

Table 186-9 (Cont.) DBMS_SQLDIAG Package Subprograms

Subprogram	Description
EXPORT_SQL_TESTCASE_DIR_BY_INC Function	Generates a SQL Test Case corresponding to the incident ID passed as an argument.
EXPORT_SQL_TESTCASE_DIR_BY_TXT Function	Generates a SQL Test Case corresponding to the SQL passed as an argument
GET_FIX_CONTROL Function	Returns the value of fix control for a given bug number
GET_SQL Function	Imports a SQL test case
IMPORT_SQL_TESTCASE Procedures	Imports a SQL test case into a schema
INCIDENTID_2_SQL Procedure	Initializes a `sql_setrow` from an incident ID
INTERRUPT_DIAGNOSIS_TASK Procedure	Interrupts a diagnostic task
LOAD_SQLSET_FROM_TCB Function	Loads a `SQLSET` from Test Case Builder (TCB) file
PACK_STGTAB_SQLPATCH Procedure	SQL patches into the staging table created by the CREATE_STGTAB_SQLPATCH Procedure
REPLAY_SQL_TESTCASE Function	Reports on a diagnostic task
REPORT_SQL Function	Generates a diagnostic report in HTML format for a specific SQL statement.
REPORT_DIAGNOSIS_TASK Function	Reports on a diagnostic task
RESET_DIAGNOSIS_TASK Procedure	Resets a diagnostic task
RESUME_DIAGNOSIS_TASK Procedure	Resumes a diagnostic task
SET_DIAGNOSIS_TASK_PARAMETER Procedure	Sets a diagnosis task parameter
SQL_DIAGNOSE_AND_REPAIR Function	Diagnoses a given SQL statement for a given SQL ID for the given problem type.
UNPACK_STGTAB_SQLPATCH Procedure	Unpacks from the staging table populated by a call to the PACK_STGTAB_SQLPATCH Procedure, using the patch data stored in the staging table to create patches on this system

186.4.1 ACCEPT_SQL_PATCH Function & Procedure

This procedure accepts a recommended SQL patch as recommended by the specified SQL diagnosis task.

Syntax

```
DBMS_SQLDIAG.ACCEPT_SQL_PATCH (
   task_name       IN  VARCHAR2,
   object_id       IN  NUMBER := NULL,
   name            IN  VARCHAR2 := NULL,
   description     IN  VARCHAR2 := NULL,
   category        IN  VARCHAR2 := NULL,
   task_owner      IN  VARCHAR2 := NULL,
   replace         IN  BOOLEAN := FALSE,
```

```
   force_match    IN  BOOLEAN := FALSE)
 RETURN VARCHAR2;

DBMS_SQLDIAG.ACCEPT_SQL_PATCH (
   task_name      IN  VARCHAR2,
   object_id      IN  NUMBER := NULL,
   name           IN  VARCHAR2 := NULL,
   description    IN  VARCHAR2 := NULL,
   category       IN  VARCHAR2 := NULL,
   task_owner     IN  VARCHAR2 := NULL,
   replace        IN  BOOLEAN := FALSE,
   force_match    IN  BOOLEAN := FALSE);
```

Parameters

Table 186-10 ACCEPT_SQL_PATCH Function & Procedure Parameters

Parameter	Description
taskname	Name of the SQL diagnosis task
object_id	Identifier of the advisor framework object representing the SQL statement associated to the diagnosis task
name	Name of the patch. It cannot contain double quotation marks. The name is case sensitive. If not specified, the system will generate a unique name for the SQL patch.
description	User specified string describing the purpose of this SQL patch. Maximum size of description is 500.
category	Category name which must match the value of the SQLDIAGNOSE_CATEGORY parameter in a session for the session to use this patch. It defaults to the value DEFAULT. This is also the default of the SQLDIAGNOSE_CATEGORY parameter. The category must be a valid Oracle identifier. The category name specified is always converted to upper case. The combination of the normalized SQL text and category name create a unique key for a patch. An accept will fail if this combination is duplicated.
task_owner	Owner of the diagnosis task. This is an optional parameter that has to be specified to accept a SQL Patch associated to a diagnosis task owned by another user. The current user is the default value.
replace	If the patch already exists, it will be replaced if this argument is TRUE. It is an error to pass a name that is already being used for another signature/category pair, even with replace set to TRUE.
force_match	If TRUE this causes SQL Patches to target all SQL statements which have the same text after normalizing all literal values into bind variables. (Note that if a combination of literal values and bind values is used in a SQL statement, no bind transformation occurs.) This is analogous to the matching algorithm used by the FORCE option of the CURSOR_SHARING parameter. If FALSE, literals are not transformed. This is analogous to the matching algorithm used by the EXACT option of the CURSOR_SHARING parameter.

Return Values

Name of the SQL patch

Usage Notes

Requires CREATE ANY SQL PROFILE privilege

186.4.2 ALTER_SQL_PATCH Procedure

This procedure alters specific attributes of an existing SQL patch object.

Syntax

```
DBMS_SQLDIAG.ALTER_SQL_PATCH (
   name            IN  VARCHAR2,
   attribute_name  IN  VARCHAR2,
   attribute_value IN  VARCHAR2);
```

Parameters

Table 186-11 ALTER_SQL_PATCH Procedure Parameters

Parameter	Description
`name`	Name of SQL patch to alter.
`attribute_name`	Name of SQL patch to alter. Possible values: • `STATUS` -> can be set to `ENABLED` or `DISABLED` • `NAME` -> can be reset to a valid name (must be a valid Oracle identifier and must be unique). • `DESCRIPTION` -> can be set to any string of size no more than 500 • `CATEGORY` -> can be reset to a valid category name (must be valid Oracle identifier and must be unique when combined with normalized SQL text) This parameter is mandatory and is case sensitive.
`attribute_value`	New value of the attribute. See `attribute_name` for valid attribute values. This parameter is mandatory.

Usage Notes

Requires `ALTER ANY SQL PATCH` privilege

186.4.3 CANCEL_DIAGNOSIS_TASK Procedure

This procedure cancels a diagnostic task.

Syntax

```
DBMS_SQLDIAG.CANCEL_DIAGNOSIS_TASK (
    taskname        IN   VARCHAR2);
```

Parameters

Table 186-12 CANCEL_DIAGNOSIS_TASK Procedure Parameters

Parameter	Description
`taskname`	Name of task

186.4.4 CREATE_DIAGNOSIS_TASK Functions

This function creates a diagnostic task in order to diagnose a single SQL statement. It returns a SQL diagnosis task unique name

Syntax

Prepares the diagnosis of a single statement given its text:

```
DBMS_SQLDIAG.CREATE_DIAGNOSIS_TASK (
    sql_text           IN   CLOB,
    bind_list          IN   sql_binds := NULL,
    user_name          IN   VARCHAR2  := NULL,
    scope              IN   VARCHAR2  := SCOPE_COMPREHENSIVE,
    time_limit         IN   NUMBER    := TIME_LIMIT_DEFAULT,
    task_name          IN   VARCHAR2  := NULL,
    description        IN   VARCHAR2  := NULL,
    problem_type       IN   NUMBER    := PROBLEM_TYPE_PERFORMANCE)
  RETURN VARCHAR2;
```

Prepares the diagnosis of a single statement from the Cursor Cache given its identifier:

```
DBMS_SQLDIAG.CREATE_DIAGNOSIS_TASK (
    sql_id             IN   VARCHAR2,
    plan_hash_value    IN   NUMBER    := NULL,
    scope              IN   VARCHAR2  := SCOPE_COMPREHENSIVE,
    time_limit         IN   NUMBER    := TIME_LIMIT_DEFAULT,
    task_name          IN   VARCHAR2  := NULL,
    description        IN   VARCHAR2  := NULL,
    problem_type       IN   NUMBER    := PROBLEM_TYPE_PERFORMANCE)
  RETURN VARCHAR2;
```

Prepares the diagnosis of a Sqlset:

```
DBMS_SQLDIAG.CREATE_DIAGNOSIS_TASK (
    sqlset_name       IN VARCHAR2,
    basic_filter      IN VARCHAR2 :=  NULL,
    object_filter     IN VARCHAR2 :=  NULL,
    rank1             IN VARCHAR2 :=  NULL,
    rank2             IN VARCHAR2 :=  NULL,
    rank3             IN VARCHAR2 :=  NULL,
    result_percentage IN NUMBER   :=  NULL,
    result_limit      IN NUMBER   :=  NULL,
    scope             IN VARCHAR2 :=  SCOPE_COMPREHENSIVE,
    time_limit        IN NUMBER   :=  TIME_LIMIT_DEFAULT,
    task_name         IN VARCHAR2 :=  NULL,
    description       IN VARCHAR2 :=  NULL,
    plan_filter       IN VARCHAR2 :=  'MAX_ELAPSED_TIME',
    sqlset_owner      IN VARCHAR2 :=  NULL,
    problem_type      IN NUMBER   :=  PROBLEM_TYPE_PERFORMANCE)  RETURN VARCHAR2;
```

Parameters

Table 186-13 CREATE_DIAGNOSIS_TASK Function Parameters

Parameter	Description
sql_text	Text of a SQL statement
bind_list	Set of bind values

Table 186-13 (Cont.) CREATE_DIAGNOSIS_TASK Function Parameters

Parameter	Description
user_name	Username for who the statement/sqlset will be diagnosed
scope	Diagnosis scope (limited/comprehensive)
time_limit	Maximum duration in seconds for the diagnosis session
task_name	Optional diagnosis task name
description	Maximum of 256 SQL diagnosis session description
problem_type	Determines the goal of the task. Possible values are: • PROBLEM_TYPE_WRONG_RESULTS • PROBLEM_TYPE_COMPILATION_ERROR • PROBLEM_TYPE_EXECUTION_ERROR
sql_id	Identifier of the statement
plan_hash_value	Hash value of the SQL execution plan
sqlset_name	Sqlset name
basic_filter	SQL predicate to filter the SQL from the SQL tuning set (STS)
object_filter	Object filter
rank(i)	Order-by clause on the selected SQL
result_percentage	Percentage on the sum of a ranking measure
result_limit	Top L(imit) SQL from (filtered/ranked) SQL
plan_filter	Plan filter. It is applicable in case there are multiple plans (plan_hash_value). This filter allows selecting one plan (plan_hash_value) only. Possible values are: • LAST_GENERATED: plan with most recent timestamp • FIRST_GENERATED: opposite to LAST_GENERATED • LAST_LOADED: plan with most recent first_load_time stat info • FIRST_LOADED: opposite to LAST_LOADED • MAX_ELAPSED_TIME: plan with maximum elapsed time • MAX_BUFFER_GETS: plan with maximum buffer gets • MAX_DISK_READS: plan with maximum disk reads • MAX_DIRECT_WRITES: plan with maximum direct writes • MAX_OPTIMIZER_COST: plan with maximum optimum cost
sqlset_owner	Owner of the sqlset, or null for current schema owner

186.4.5 CREATE_SQL_PATCH Function

This function creates a SQL patch based on a set of user specified hints for specific statements identified by SQL text.

A SQL patch is usually created automatically by the SQL Repair Advisor to prevent any errors during the compilation or execution of a SQL statement. This function provides a way to manually create a SQL patch based on a set of hints that resolves the error.

Syntax

```
DBMS_SQLDIAG.CREATE_SQL_PATCH (
    sql_text        IN   CLOB,
    hint_text       IN   CLOB,
```

```
    name            IN   VARCHAR2   := NULL,
    description     IN   VARCHAR2   := NULL,
    category        IN   VARCHAR2   := NULL,
    validate        IN   BOOLEAN    := TRUE)
RETURN VARCHAR2;

DBMS_SQLDIAG.CREATE_SQL_PATCH (
    sql_id          IN   VARCHAR2,
    hint_text       IN   CLOB,
    name            IN   VARCHAR2   := NULL,
    description     IN   VARCHAR2   := NULL,
    category        IN   VARCHAR2   := NULL,
    validate        IN   BOOLEAN    := TRUE)
RETURN VARCHAR2;
```

Parameters

Table 186-14 CREATE_SQL_PATCH Function Parameters

Parameter	Description
sql_text	Text of the SQL statement
sql_id	The SQL identifier for the SQL statement
hint_text	Hints to include in the SQL patch
name	Optional SQL patch name
description	Description of the SQL patch
category	Category name
validate	Whether to validate the provided hints

Return Values

Both functions return the SQL patch name.

186.4.6 CREATE_STGTAB_SQLPATCH Procedure

This procedure creates the staging table used for transporting SQL patches from one system to another.

Syntax

```
DBMS_SQLDIAG.CREATE_STGTAB_SQLPATCH (
   table_name       IN  VARCHAR2,
   schema_name      IN  VARCHAR2 := NULL,
   tablespace_name  IN  VARCHAR2 := NULL);
```

Parameters

Table 186-15 CREATE_STGTAB_SQLPATCH Procedure Parameters

Parameter	Description
table_name	(Mandatory) Name of the table to create (case-sensitive)
schema_name	Schema to create the table in, or NULL for current schema (case-sensitive)

Table 186-15 (Cont.) CREATE_STGTAB_SQLPATCH Procedure Parameters

Parameter	Description
tablespace_name	Tablespace to store the staging table within, or NULL for current user's default tablespace (case-sensitive)

186.4.7 DROP_DIAGNOSIS_TASK Procedure

This procedure drops a diagnostic task.

Syntax

```
DBMS_SQLDIAG.DROP_DIAGNOSIS_TASK (
    taskname        IN   VARCHAR2);
```

Parameters

Table 186-16 DROP_DIAGNOSIS_TASK Procedure Parameters

Parameter	Description
taskname	Name of task

186.4.8 DROP_SQL_PATCH Procedure

This procedure drops the named SQL patch from the database.

Syntax

```
DBMS_SQLDIAG.DROP_SQL_PATCH (
   name     IN  VARCHAR2,   ignore   IN  BOOLEAN := FALSE);
```

Parameters

Table 186-17 DROP_SQL_PATCH Function & Procedure Parameters

Parameter	Description
name	Name of patch to be dropped. The name is case sensitive.
ignore	Ignore errors due to object not existing.

Usage Notes

Requires DROP ANY SQL PATCH privilege

186.4.9 EXECUTE_DIAGNOSIS_TASK Procedure

This procedure executes a diagnostic task.

Syntax

```
DBMS_SQLDIAG.EXECUTE_DIAGNOSIS_TASK (
    taskname        IN   VARCHAR2);
```

Parameters

Table 186-18 EXECUTE_DIAGNOSIS_TASK Procedure Parameters

Parameter	Description
taskname	Name of task

186.4.10 EXPLAIN_SQL_TESTCASE Function

This procedure explains a SQL test case.

Syntax

```
DBMS_SQLDIAG.EXPLAIN_SQL_TESTCASE (
    sqlTestCase        IN   CLOB)
  RETURN CLOB;
```

Parameters

Table 186-19 EXPLAIN_SQL_TESTCASE Function Parameters

Parameter	Description
sqlTestCase	XML document describing the SQL test case

186.4.11 EXPORT_SQL_TESTCASE Procedures

This procedure exports a SQL test case to a directory.

Syntax

This variant has to be provided with the SQL information.

```
DBMS_SQLDIAG.EXPORT_SQL_TESTCASE (
    directory               IN              VARCHAR2,
    sql_text                IN              CLOB,
    user_name               IN              VARCHAR2  :=  NULL,
    bind_list               IN              sql_binds :=  NULL,
    exportEnvironment       IN              BOOLEAN   :=  TRUE,
    exportMetadata          IN              BOOLEAN   :=  TRUE,
    exportData              IN              BOOLEAN   :=  FALSE,
    exportPkgbody           IN              BOOLEAN   :=  FALSE,
    samplingPercent         IN              NUMBER    :=  100,
    ctrlOptions             IN              VARCHAR2  :=  NULL,
    timeLimit               IN              NUMBER    :=  0,
    testcase_name           IN              VARCHAR2  :=  NULL,
    testcase                IN OUT NOCOPY   CLOB,
    preserveSchemaMapping   IN              BOOLEAN   :=  FALSE,
    version                 IN              VARCHAR2  :=  'COMPATIBLE');
```

This variant extracts the SQL information from an incident file.

```
DBMS_SQLDIAG.EXPORT_SQL_TESTCASE (
    directory               IN              VARCHAR2,
    incident_id             IN              VARCHAR2,
    exportEnvironment       IN              BOOLEAN   :=  TRUE,
```

```
    exportMetadata          IN              BOOLEAN   :=  TRUE,
    exportData              IN              BOOLEAN   :=  FALSE,
    exportPkgbody           IN              BOOLEAN   :=  FALSE,
    samplingPercent         IN              NUMBER    :=  100,
    ctrlOptions             IN              VARCHAR2  :=  NULL,
    timeLimit               IN              NUMBER    :=
                                               DBMS_SQLDIAG.TIME_LIMIT_DEFAULT,
    testcase_name           IN              VARCHAR2  :=  NULL,
    testcase                IN OUT NOCOPY   CLOB,
    preserveSchemaMapping   IN              BOOLEAN   :=  FALSE)
    version                 IN              VARCHAR2  :=  'COMPATIBLE');
```

This variant allow the SQL Test case to be generated from a cursor present in the cursor cache. Use `V$SQL` to get the SQL identifier and the SQL hash value.

```
DBMS_SQLDIAG.EXPORT_SQL_TESTCASE (
    directory               IN              VARCHAR2,
    sql_id                  IN              VARCHAR2,
    plan_hash_value         IN              NUMBER    := NULL,
    exportEnvironment       IN              BOOLEAN   :=  TRUE,
    exportMetadata          IN              BOOLEAN   :=  TRUE,
    exportData              IN              BOOLEAN   :=  FALSE,
    exportPkgbody           IN              BOOLEAN   :=  FALSE,
    samplingPercent         IN              NUMBER    :=  100,
    ctrlOptions             IN              VARCHAR2  :=  NULL,
    timeLimit               IN              NUMBER    :=
                                               DBMS_SQLDIAG.TIME_LIMIT_DEFAULT,
    testcase_name           IN              VARCHAR2  :=  NULL,
    testcase                IN OUT NOCOPY   CLOB,
    preserveSchemaMapping   IN              BOOLEAN   :=  FALSE)
    version                 IN              VARCHAR2  :=  'COMPATIBLE');
```

Parameters

Table 186-20 EXPORT_SQL_TESTCASE Procedure Parameters

Parameter	Description
`directory`	Directory to store the various generated files
`sql_text`	Text of the SQL statement to export
`incident_id`	Incident ID containing the offending SQL
`sql_id`	Identifier of the statement in the cursor cache, automatic workload repository, or the automatic SQL tuning set
`username`	Name of the user schema to use to parse the SQL, defaults to `SYS`
`bind_list`	List of bind values associated to the statement
`exportEnvironment`	`TRUE` if the compilation environment should be exported
`exportMetadata`	`TRUE` if the definition of the objects referenced in the SQL should be exported
`exportData`	`TRUE` if the data of the objects referenced in the SQL should be exported
`exportPkgbody`	`TRUE` if the body of the packages referenced in the SQL are exported
`samplingPercent`	If is `TRUE`, specify the sampling percentage to use to create the dump file

Table 186-20 (Cont.) EXPORT_SQL_TESTCASE Procedure Parameters

Parameter	Description
ctrlOptions	Opaque control parameters. For example, to execute three times, set ctrlOptions with the following string: '<parameter name="mexec_count">3</parameter>'. • name="capture" - BASIC (default) or WITH_RUNTIME_INFO. This parameter defines the mode of TCB capture. BASIC: runs as Oracle release 11*g* TCB and captures all the information that is captured in that release as well as AWR reports, SQL monitor reports and parameter information. WITH_RUNTIME_INFO: TCB captures runtime information for the SQL, such as dynamic sampling data, list of binds, Dynamic Plan info, along with information captured under BASIC mode. Note this must be the same value as used in the IMPORT_SQL_TESTCASE Procedures. • name="mexec_count"–Value is any positive number (N). This parameter tells TCB to execute the statement for N time and capture runtime info at end of each execution. • name="stat_history_since"–Value is date. The object statistics history is exported using this parameter. Statistics history after date specified will be exported. • name="compress"–This option is used to compress the SQL Test Case Builder output files into a zip file by default. The possible values are: – YES – NO The default value is YES. • name="diag_event"–This option is used to specify the level of trace information to include in the SQL Test Case Builder output. The possible values are: – ADS – COMPILER – SQLEXEC_LOW – SQLEXEC_MEDIUM – SQLEXEC_HIGH – SQLEXEC_HIGHEST The default value is ADS + COMPILER. • name="problem_type"–This option is used to assign an issue type for a SQL Test Case Builder test case. For example, if a test case is related to performance regression issue, then you can assign the value of PERFORMANCE to the problem_type option. The possible values are : – PERFORMANCE – WRONG_RESULTS – COMPILATION_ERROR – EXECUTION_ERROR The default value is PERFORMANCE.
timeLimit	How much time should we spend exporting the SQL test case
testcaseName	An optional name for the SQL test case. This is used to prefix all the generated scripts
testcase	Resulting testcase

Table 186-20 (Cont.) EXPORT_SQL_TESTCASE Procedure Parameters

Parameter	Description
preserveSchemaMapping	TRUE if the schema (or schemas) are not re-mapped from the original environment to the test environment
version	Version of database objects to be extracted. This option is only valid for EXPORT. Database objects or attributes incompatible with the version will not be extracted. • COMPATIBLE - (default) the version of the metadata corresponds to the database compatibility level and the compatibility release level for feature (as given in the V$COMPATIBILITY view). Database compatibility must be set to 9.2 or higher. • LATEST - the version of the metadata that specifies the current database version. • A specific database version. For example, if '10.0.0', this cannot be lower than Oracle Database release 10.0.0.

Usage Notes

- A SQL test case generates a set of files needed to help reproduce a SQL failure on a different machine. It contains:
  - a dump file containing schemas objects and statistics (.dmp)
  - the explain plan for the statements (in advanced mode)
  - diagnostic information gathered on the offending statement
  - an import script to execute to reload the objects
  - a SQL script to replay system statistics of the source
  - a table of contents file describing the SQL test case
  - metadata. (xxxxmain.xml)
  - a README.txt file that explain the usage of the TCB
  - the outlines used by the statement (ol.xml)
  - a list of parameters set in the exporting db/env (prmimp.sql)
  - a SQL monitor report, if any (smrpt.html)
  - an AWR report, if any (awrrpt.html)
  - a list of binds used in this statement (bndlst.xml)
- You should not run Test Case Builder (TCB) under user SYS. Instead, use another user who can be granted the DBA privilege.
- The default setting for TCB is that data is not exported. However, in some cases data is required, such as to diagnose an outcome with a result that is not optimal. To export data, call EXPORT_SQL_TESTCASE with exportData=>TRUE and the data will be imported by default, unless turned OFF by importData=>FALSE.
- TCB includes PL/SQL package spec by default, but not the PL/SQL package body. However, you may need to have the package body as well, for example, to invoke the PL/SQL functions, or because you have a Virtual Private Database (VPD) function defined in a package. To export a PL/SQL package body, call EXPORT_SQL_TESTCASE with

`exportPkgbody=>TRUE`. To import a PL/SQL package body, call IMPORT_SQL_TESTCASE Procedures with i`mportPkgbody=>TRUE`.

- To export objects statistics history, the database compatibility should be set to 12.0 or higher.
- This procedure does not export data and statistics on a Global Temporary Table (GTT).

Examples

The user can specify multiple parameters in the `ctrlOptions` encapsulated either by using the `<parameters>` parent tag or without the parent tag.

Using the <parameters> tag

```
<parameters>
<parameter name="capture">with_runtime_info</parameter>
<parameter name="mexec_count">1</parameter>
</parameters>
```

Without the <parameters> tag

```
<parameter name="capture">with_runtime_info</parameter>
<parameter name="mexec_count">1</parameter>'
```

The `compress` option that you can specify in the `ctrlOptions` parameter:

```
opt := '<parameters>
           <parameter name="capture">with_runtime_info</parameter>
           <parameter name="compress">yes</parameter>
       </parameters>';
```

The `diag_event` option that you can specify in the `ctrlOptions` parameter:

```
opt := '<parameters>
           <parameter name="capture">with_runtime_info</parameter>
           <parameter name="compress">yes</parameter>
       </parameters>';
```

The `problem_type` option that you can specify in the `ctrlOptions` parameter:

```
opt := '<parameters>
           <parameter name="capture">with_runtime_info</parameter>
           <parameter name="compress">yes</parameter>
       </parameters>';
```

186.4.12 EXPORT_SQL_TESTCASE_DIR_BY_INC Function

This function generates a SQL test case corresponding to the incident ID passed as an argument. It creates a set of scripts and dump file in the directory passed as an argument.

Syntax

```
DBMS_SQLDIAG.EXPORT_SQL_TESTCASE_DIR_BY_INC (
    incident_id        IN   NUMBER,
    directory          IN   VARCHAR2,
    exportEnvironment  IN   VARCHAR2 := 'TRUE',
    exportMetadata     IN   VARCHAR2 := 'TRUE',
    exportData         IN   VARCHAR2 := 'FALSE',
    samplingPercent    IN   VARCHAR2 := '100',
    ctrlOptions        IN   VARCHAR2 := NULL
```

```
   version            IN   VARCHAR2 := 'COMPATIBLE')
RETURN BOOLEAN;
```

Parameters

Table 186-21 EXPORT_SQL_TESTCASE_DIR_BY_INC Function Parameters

Parameter	Description
incident_id	Incident ID containing the offending SQL. For more information about Incidents, see *Oracle Database Performance Tuning Guide*.
directory	Directory path to the generated files
exportEnvironment	TRUE if the compilation environment should be exported
exportMetadata	TRUE if the definition of the objects referenced in the SQL should be exported
exportData	TRUE if the data of the objects referenced in the SQL should be exported
samplingPercent	If is TRUE, specify the sampling percentage to use to create the dump file
ctrlOptions	Opaque control parameters. For example, to execute three times, set ctrlOptions with the following string: '<parameter name="mexec_count">3</parameter>'. • capture - BASIC (default) or WITH_RUNTIME_INFO. This parameter defines the mode of TCB capture. BASIC: runs as Oracle release 11*g* TCB and captures all the information that is captured in that release as well as AWR reports, SQL monitor reports and parameter information. WITH_RUNTIME_INFO: TCB captures runtime information for the SQL, such as dynamic sampling data, list of binds, Dynamic Plan info, along with information captured under BASIC mode. • name=mexec_count - Value is any positive number (N). This parameter tells TCB to execute the statement for N time and capture runtime info at end of each execution. • name=stat_history_since - Value is date. The object statistics history is exported using this parameter. Statistics history after date specified will be exported.
version	Version of database objects to be extracted. This option is only valid for EXPORT. Database objects or attributes incompatible with the version will not be extracted. • COMPATIBLE - (default) the version of the metadata corresponds to the database compatibility level and the compatibility release level for feature (as given in the V$COMPATIBILITY view). Database compatibility must be set to 9.2 or higher. • LATEST - the version of the metadata that specifies the current database version. • A specific database version. For example, if '10.0.0', this cannot be lower than Oracle Database release 10.0.0.

186.4.13 EXPORT_SQL_TESTCASE_DIR_BY_TXT Function

This function generates a SQL Test Case corresponding to the SQL passed as an argument. It creates a set of scripts and dump files in the directory passed as an argument.

Syntax

```
DBMS_SQLDIAG.EXPORT_SQL_TESTCASE_DIR_BY_TXT (
    incident_id        IN   NUMBER,
    directory          IN   VARCHAR2,
    sql_text           IN   CLOB,
    user_name          IN   VARCHAR2 := 'SYS',
    exportEnvironment  IN   VARCHAR2 := 'TRUE',
    exportMetadata     IN   VARCHAR2 := 'TRUE',
    exportData         IN   VARCHAR2 := 'FALSE',
    samplingPercent    IN   VARCHAR2 := '100',
    ctrlOptions        IN   VARCHAR2 := NULL
    version            IN   VARCHAR2 := 'COMPATIBLE')
  RETURN BOOLEAN;
```

Parameters

Table 186-22 EXPORT_SQL_TESTCASE_DIR_BY_TXT Function Parameters

Parameter	Description
incident_id	Incident ID containing the offending SQL
directory	Directory to store the various generated files
sql_text	Text of the SQL statement to explain
username	Name of the user schema to use to parse the SQL, defaults to SYS
exportEnvironment	TRUE if the compilation environment should be exported
exportMetadata	TRUE if the definition of the objects referenced in the SQL should be exported
exportData	TRUE if the data of the objects referenced in the SQL should be exported
samplingPercent	If is TRUE, specify the sampling percentage to use to create the dump file

Table 186-22 (Cont.) EXPORT_SQL_TESTCASE_DIR_BY_TXT Function Parameters

Parameter	Description
ctrlOptions	Opaque control parameters. For example, to execute three times, set ctrlOptions with the following string: '<parameter name="mexec_count">3</parameter>'. • capture - BASIC (default) or WITH_RUNTIME_INFO. This parameter defines the mode of TCB capture. BASIC: runs as Oracle Release 11*g* TCB and captures all the information that is captured in that release as well as AWR reports, SQL monitor reports and parameter information. WITH_RUNTIME_INFO: TCB captures runtime information for the SQL, such as dynamic sampling data, list of binds, Dynamic Plan info, along with information captured under BASIC mode. • name=mexec_count - Value is any positive number (N). This parameter tells TCB to execute the statement for N time and capture runtime info at end of each execution. • name=stat_history_since - Value is date. The object statistics history is exported using this parameter. Statistics history after date specified will be exported.
version	Version of database objects to be extracted. This option is only valid for EXPORT. Database objects or attributes incompatible with the version will not be extracted. • COMPATIBLE - (default) the version of the metadata corresponds to the database compatibility level and the compatibility release level for feature (as given in the V$COMPATIBILITY view). Database compatibility must be set to 9.2 or higher. • LATEST - the version of the metadata that specifies the current database version. • A specific database version. For example, if '10.0.0', this cannot be lower than Oracle Database Release 10.0.0.

186.4.14 GET_FIX_CONTROL Function

This function returns the value of fix control for a given bug number.

Syntax

```
DBMS_SQLDIAG.GET_FIX_CONTROL (
    bug_number    IN    NUMBER)
  RETURN NUMBER;
```

Parameters

Table 186-23 GET_FIX_CONTROL Function Parameters

Parameter	Description
bug_number	Bug number

186.4.15 GET_SQL Function

This function loads a `sql_setrow` from the trace file associated to an the given incident ID.

Syntax

```
DBMS_SQLDIAG.GET_SQL (
    incident_id  IN     VARCHAR2)
  RETURN SQLSET_ROW;
```

Parameters

Table 186-24 GET_SQL Function Parameters

Parameter	Description
`incident_id`	Identifier of the incident

186.4.16 IMPORT_SQL_TESTCASE Procedures

This procedure imports a SQL test case into a schema.

Syntax

This variant requires a source directory and SQL Testcase metadata object (in XML format).

```
DBMS_SQLDIAG.IMPORT_SQL_TESTCASE (
    directory               IN   VARCHAR2,
    sqlTestCase             IN   CLOB,
    importEnvironment       IN   BOOLEAN   :=  TRUE,
    importMetadata          IN   BOOLEAN   :=  TRUE,
    importData              IN   BOOLEAN   :=  TRUE,
    importPkgbody           IN   BOOLEAN   :=  FALSE,
    importDiagnosis         IN   BOOLEAN   :=  TRUE,
    ignoreStorage           IN   BOOLEAN   :=  TRUE,
    ctrlOptions             IN   VARCHAR2  :=  NULL,
    preserveSchemaMapping   IN   BOOLEAN   :=  FALSE);
```

This variant requires a source directory name of SQL Testcase metadata file.

```
DBMS_SQLDIAG.IMPORT_SQL_TESTCASE (
    directory               IN   VARCHAR2,
    filename                IN   VARCHAR2,
    importEnvironment       IN   BOOLEAN   :=  TRUE,
    importMetadata          IN   BOOLEAN   :=  TRUE,
    importData              IN   BOOLEAN   :=  TRUE,
    importPkgbody           IN   BOOLEAN   :=  FALSE,
    importDiagnosis         IN   BOOLEAN   :=  TRUE,
    ignoreStorage           IN   BOOLEAN   :=  TRUE,
    ctrlOptions             IN   VARCHAR2  :=  NULL,
    preserveSchemaMapping   IN   BOOLEAN   :=  FALSE);
```

Parameters

Table 186-25 IMPORT_SQL_TESTCASE Procedure Parameters

Parameter	Description
directory	Directory containing test case files
filename	Name of a file containing an XML document describing the SQL test case
importEnvironment	TRUE if the compilation environment should be imported
importMetadata	TRUE if the definition of the objects referenced in the SQL should be imported
importData	TRUE if the data of the objects referenced in the SQL should be imported
importPkgbody	TRUE if the body of the packages referenced in the SQL are imported
importDiagnosis	TRUE if the diagnostic information associated to the task should be imported
ignoreStorage	TRUE if the storage attributes should be ignored
ctrlOptions	Opaque control parameters, of which only capture is valid for this subprogam. • capture - BASIC (default) or WITH_RUNTIME_INFO. This parameter defines the mode of TCB capture. BASIC: runs as Oracle Release 11*g* TCB and captures all the information that is captured in that release as well as AWR reports, SQL monitor reports and parameter information. WITH_RUNTIME_INFO: TCB captures runtime information for the SQL, such as dynamic sampling data, list of binds, Dynamic Plan info, along with information captured under BASIC mode.
preserveSchemaMapping	TRUE if the schema (or schemas) are not re-mapped from the original environment to the test environment (schema mapping in the target database will be identical to the source database). Note that when an import is run with preservesSchemaMapping set to TRUE, if the objects in the schemas exists then the import will overwrite the existing objects.

Usage Notes

- A SQL test case generates a set of files needed to help reproduce a SQL failure on a different machine. It contains:
  - a dump file containing schemas objects and statistics (.dmp)
  - the explain plan for the statements (in advanced mode)
  - diagnostic information gathered on the offending statement
  - an import script to execute to reload the objects
  - a SQL script to replay system statistics of the source
  - a table of contents file describing the SQL test case
  - metadata. (xxxxmain.xml)
  - a README.txt file that explain the usage of the TCB
  - the outlines used by the statement (ol.xml)

  - a list of parameters set in the exporting db/env (`prmimp.sql`)
  - a SQL monitor report, if any (`smrpt.html`)
  - an AWR report, if any (`awrrpt.html`)
  - a list of binds used in this statement (`bndlst.xml`)
- You should not run Test Case Builder (TCB) under user `SYS`. Instead, use another user who can be granted the `DBA` privilege
- The default setting for TCB is that data is not exported. However, in some cases data is required, such as to diagnose an outcome with a result that is not optimal. To export data, call EXPORT_SQL_TESTCASE Procedures with `exportData=>TRUE` and the data will be imported by default, unless turned `OFF` by `importData=>FALSE`.
- TCB includes PL/SQL package spec by default, but not the PL/SQL package body. However, you may need to have the package body as well, for example, to invoke the PL/SQL functions, or because you have a Virtual Private Database (VPD) function defined in a package. To export a PL/SQL package body, call EXPORT_SQL_TESTCASE Procedures with `exportPkgbody=>TRUE`. To import a PL/SQL package body, call `IMPORT_SQL_TESTCASE` Procedures with `importPkgbody=>TRUE`.
- The `capture` value used when invoking the EXPORT_SQL_TESTCASE Procedures must be used when calling this procedure.

186.4.17 INCIDENTID_2_SQL Procedure

This procedure initializes a `sql_setrow` from an incident ID.

Syntax

```
DBMS_SQLDIAG.INCIDENTID_2_SQL (
    incident_id   IN     VARCHAR2,
    sql_stmt      OUT    SQLSET_ROW,
    problem_type  OUT    NUMBER,
    err_code      OUT    BINARY_INTEGER,
    err_mesg      OUT    VARCHAR2);
```

Parameters

Table 186-26 INCIDENTID_2_SQL Procedure Parameters

Parameter	Description
`incident_id`	Identifier of the incident
`sql_stmt`	Resulting SQL
`problem_type`	Tentative type of SQL problem (currently among `PROBLEM_TYPE_COMPILATION_ERROR` and `PROBLEM_TYPE_EXECUTION_ERROR`)
`err_code`	Error code if any otherwise it is set to `NULL`
`err_msg`	Error message if any otherwise it is set to `NULL`

186.4.18 INTERRUPT_DIAGNOSIS_TASK Procedure

This procedure interrupts a diagnostic task.

Syntax

```
DBMS_SQLDIAG.INTERRUPT_DIAGNOSIS_TASK (
    taskname        IN   VARCHAR2);
```

Parameters

Table 186-27 INTERRUPT_DIAGNOSIS_TASK Procedure Parameters

Parameter	Description
taskname	Name of task

186.4.19 LOAD_SQLSET_FROM_TCB Function

This function loads a `SQLSET` from a Test Case Builder file.

Syntax

```
DBMS_SQLDIAG.LOAD_SQLSET_FROM_TCB (
    directory        IN     VARCHAR2,
    filename         IN     VARCHAR2,
    sqlset_name      IN     VARCHAR2 DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 186-28 LOAD_SQLSET_FROM_TCB Function Parameters

Parameter	Description
directory	Name of directory
filename	Name of file
sqlset_name	Name of `SQLSET`

186.4.20 PACK_STGTAB_SQLPATCH Procedure

This procedure packs SQL patches into the staging table created by a call to the CREATE_STGTAB_SQLPATCH Procedure.

Syntax

```
DBMS_SQLDIAG.PACK_STGTAB_SQLPATCH (
   patch_name            IN  VARCHAR2 := '%',
   patch_category        IN  VARCHAR2 := 'DEFAULT',
   staging_table_name    IN  VARCHAR2,
   staging_schema_owner  IN  VARCHAR2 := NULL);
```

Parameters

Table 186-29 PACK_STGTAB_SQLPATCH Procedure Parameters

Parameter	Description
patch_name	Name of patch to pack (% wildcards acceptable, case-sensitive)
patch_category	Category to which to pack patches (% wildcards acceptable, case-insensitive)
staging_table_name	(Mandatory) Name of the table to use (case-sensitive)
staging_schema_owner	Schema where the table resides, or NULL for current schema (case-sensitive)

Usage Notes

- Requires: ADMINISTER SQL PLAN MANAGEMENT OBJECT privilege and INSERT privilege on the staging table
- By default, we move all SQL patches in category DEFAULT. Note that the subprogram issues a COMMIT after packing each SQL patch, so if an error is raised in mid-execution, some patches may be in the staging table.

Related Topics

- CREATE_STGTAB_SQLPATCH Procedure
  This procedure creates the staging table used for transporting SQL patches from one system to another.

186.4.21 REPLAY_SQL_TESTCASE Function

This function automates the reproduction of the SQL Test Case.

Syntax

```
DBMS_SQLDIAG.REPLAY_SQL_TESTCASE (
   directory       IN   VARCHAR2,
   filename        IN   VARCHAR2,
   ctrlOptions     IN   VARCHAR2  := NULL,
   format          IN   VARCHAR2  := 'TEXT')
   RETURN CLOB;

DBMS_SQLDIAG.REPLAY_SQL_TESTCASE (
   directory       IN   VARCHAR2,
   sqlTestCase     IN   CLOB,
   ctrlOptions     IN   VARCHAR2  := NULL,
   format          IN   VARCHAR2  := 'TEXT')
   RETURN CLOB;
```

Parameters

Table 186-30 REPLAY_SQL_TESTCASE Function Parameters

Parameter	Description
directory	Directory containing test case files

Table 186-30 (Cont.) REPLAY_SQL_TESTCASE Function Parameters

Parameter	Description
filename	Name of a file containing an XML document describing the SQL test case
ctrlOptions	Opaque control parameters. For example, to execute three times, set ctrlOptions with the following string: '<parameter name="mexec_count">3</parameter>'. • replay - EXPLAIN (default), OUTLINE, EXECUTION or OUTLINE EXECUTION. This parameter defines TCB replay functionality. EXPLAIN: Replay explains the statement without using outlines OUTLINE: Replay uses outlines mode and explains the statement using outlines EXECUTION: Replay executes the statement without using outlines OUTLINE EXECUTION: Replay executes the statement using outlines Note that if the user gives an incorrect parameter value, then the replay runs in default mode and no error is thrown. • name=mexec_count - Value is any positive number (N). This parameter tells TCB to execute the statement for N time and capture runtime info at end of each execution.
sqlTestCase	SQL test case
format	Format of the replay report. Possible formats are: TEXT, XML and HTML..

Examples

```
TCB Replay Mode: Execute
 SELECT /* tcbdynpl_1 */ /*+ gather_plan_statistics */ * FROM (SELECT * FROM emp where
emp.sal > 100) emp, dept WHERE emp.deptno = dept.deptno And emp.sal > 1000 /* tcbdynpl_1
*/

Explain Plan
 Plan Hash Value  : 2219294842

-------------------------------------------------------------------
| Id  | Operation                     | Name           | Rows |
-------------------------------------------------------------------
|   0 | SELECT STATEMENT              |                |   13 |
| * 1 |   HASH JOIN                   |                |   13 |
|   2 |    NESTED LOOPS               |                |      |
|   3 |     NESTED LOOPS              |                |   13 |
|   4 |      STATISTICS COLLECTOR     |                |      |
|   5 |       TABLE ACCESS FULL       | DEPT           |    4 |
| * 6 |      INDEX RANGE SCAN         | EMP_IDX_DEPTNO |      |
| * 7 |     TABLE ACCESS BY INDEX ROWID | EMP          |    3 |
| * 8 |    TABLE ACCESS FULL          | EMP            |   13 |
-------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------
* 1 - access("EMP"."DEPTNO"="DEPT"."DEPTNO")
* 6 - access("EMP"."DEPTNO"="DEPT"."DEPTNO")
* 7 - filter("EMP"."SAL">1000)
* 8 - filter("EMP"."SAL">1000)
```

```
Runtime Plan
 Plan Hash Value  : 2219294842

-----------------------------------------------------
| Id  | Operation            | Name | E-Card | A-Card |
-----------------------------------------------------
|   0 | SELECT STATEMENT     |      |        |      0 |
| * 1 |   HASH JOIN          |      |     13 |      0 |
|   2 |    TABLE ACCESS FULL | DEPT |      4 |      0 |
| * 3 |    TABLE ACCESS FULL | EMP  |     13 |      0 |
-----------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------
* 1 - access("EMP"."DEPTNO"="DEPT"."DEPTNO")
* 3 - filter("EMP"."SAL">1000)

REPLAY Note:
-----------
- Replay used dynamic sampling
- Replay forced Dynamic plan
```

186.4.22 REPORT_SQL Function

Generates a diagnostic report in HTML format for a specific SQL statement.

Syntax

```
DBMS_SQLDIAG.REPORT_SQL (
    sql_id       IN   VARCHAR2,
    directory    IN   VARCHAR2,
    level        IN   VARCHAR2
)
RETURN CLOB;
```

Parameters

Table 186-31 DBMS_SQLDIAG. REPORT_SQL Parameters

Parameter	Description
`sql_id`	ID of the SQL statement.
`directory`	Directory object where the report is written. By default, this parameter is NULL and the report is returned as a CLOB and not written to disk.
`level`	Three options: • `BASIC` – A basic report covering minimal report detail. • `TYPICAL` – A standard (default) report which contains the basic information plus other, more advanced, report sections. • `ALL` – A complete report containing all possible report information.

Usage Notes

By default, the `level` parameter is NULL and the report is returned as a CLOB and not written to disk.

If a directory name is specified for the `DIRECTORY` argument, the file name will be created in the following format: `SQLR_<SQL_ID>_<YYYYMMDDHH24MI>.html`, where `<SQL_ID>` represents the

SQL identifier provided as the `SQL_ID` argument and `<YYYYMMDDHH24MI>` represents the timestamp at which the file was created.

Example 186-1

In this example, `SQL_ID 'gtckcpxmp3ry7'` is passed in, the directory is the standard Data Pump directory, and `ALL` detail is returned in the report.

```
declare my_report clob;

begin
   my_report := dbms_sqldiag.report_sql('gtckcpxmp3ry7',
directory=>'DATA_PUMP_DIR', level=>'ALL');
end;
/
```

186.4.23 REPORT_DIAGNOSIS_TASK Function

This function reports on a diagnostic task. It returns a `CLOB` containing the desired report.

Syntax

```
DBMS_SQLDIAG.REPORT_DIAGNOSIS_TASK (
    taskname           IN   VARCHAR2,
    type               IN   VARCHAR2  := TYPE_TEXT,
    level              IN   VARCHAR2  := LEVEL_TYPICAL,
    section            IN   VARCHAR2  := SECTION_ALL,
    object_id          IN   NUMBER    := NULL,
    result_limit       IN   NUMBER    := NULL,
    owner_name         IN   VARCHAR2  := NULL)
  RETURN CLOB;
```

Parameters

Table 186-32 REPORT_DIAGNOSIS_TASK Function Parameters

Parameter	Description
`taskname`	Name of task to report
`type`	Type of the report. Possible values are: TEXT, HTML, XML (see Table 186-4).
`level`	Format of the recommendations. Possible values are TYPICAL, BASIC, ALL (Table 186-5).
`section`	Particular section in the report. Possible values are: SUMMARY, FINDINGS, PLAN, INFORMATION, ERROR, ALL (Table 186-6).
`object_id`	Identifier of the advisor framework object that represents a given statement in a SQL Tuning Set (STS).
`result_limit`	Number of statements in a STS for which the report is generated
`owner_name`	Name of the task execution to use. If `NULL`, the report will be generated for the last task execution.

186.4.24 RESET_DIAGNOSIS_TASK Procedure

This procedure resets a diagnostic task.

Syntax

```
DBMS_SQLDIAG.RESET_DIAGNOSIS_TASK (
    taskname        IN   VARCHAR2);
```

Parameters

Table 186-33 RESET_DIAGNOSIS_TASK Procedure Parameters

Parameter	Description
taskname	Name of task

186.4.25 RESUME_DIAGNOSIS_TASK Procedure

This procedure resumes a diagnostic path.

Syntax

```
DBMS_SQLDIAG.RESUME_DIAGNOSIS_TASK (
    taskname        IN   VARCHAR2);
```

Parameters

Table 186-34 RESUME_DIAGNOSIS_TASK Procedure Parameters

Parameter	Description
taskname	Name of task

186.4.26 SET_DIAGNOSIS_TASK_PARAMETER Procedure

This procedure is called to update the value of a SQL diagnosis parameter of type `VARCHAR2`.

The task must be set to its initial state before calling this procedure. The diagnosis parameters that can be set by this procedure are:

- `MODE`: diag scope (comprehensive, limited)
- `_SQLDIAG_FINDING_MODE`: findings in the report (see "Table 186-8" for possible values)

Syntax

```
DBMS_SQLDIAG.SET_DIAGNOSIS_TASK_PARAMETER (
    taskname          IN   VARCHAR2,
    parameter         IN   VARCHAR2,    value             IN   NUMBER);
```

Parameters

Table 186-35 SET_DIAGNOSIS_TASK_PARAMETER Procedure Parameters

Parameter	Description
taskname	Identifier of the task to execute
parameter	Name of the parameter to set
value	New value of the specified parameter

186.4.27 SQL_DIAGNOSE_AND_REPAIR Function

Diagnoses a given SQL statement for a given SQL ID for the given problem type. This function creates an incident, populate incident metadata with required information like, SQL ID, SQL text, compilation environment, and so on. It also creates a diagnostic task, executes it and accepts SQL PATCH recommendation for a given SQL ID.

Syntax

```
DBMS_SQLDIAG.SQL_DIAGNOSE_AND_REPAIR (
    sql_text           IN   CLOB,
    bind_list          IN   sql_binds := NULL,
    scope              IN   VARCHAR2  := SCOPE_COMPREHENSIVE,
    time_limit         IN   NUMBER    := TIME_LIMIT_DEFAULT,
    problem_type       IN   NUMBER    := PROBLEM_TYPE_PERFORMANCE,
    auto_apply_patch   IN   VARCHAR2  := YES)
  RETURN NUMBER;


DBMS_SQLDIAG.SQL_DIAGNOSE_AND_REPAIR (
    sql_id             IN   VARCHAR2,
    plan_hash_value    IN   NUMBER   := NULL,
    scope              IN   VARCHAR2 := SCOPE_COMPREHENSIVE,
    time_limit         IN   NUMBER   := TIME_LIMIT_DEFAULT,
    problem_type       IN   NUMBER   := PROBLEM_TYPE_PERFORMANCE,
    auto_apply_patch   IN   VARCHAR2 := YES)
  RETURN NUMBER;

DBMS_SQLDIAG.SQL_DIAGNOSE_AND_REPAIR (
    incident_id        IN   VARCHAR2,
    scope              IN   VARCHAR2 := SCOPE_COMPREHENSIVE,
    time_limit         IN   NUMBER   := TIME_LIMIT_DEFAULT,
    problem_type       IN   NUMBER   := PROBLEM_TYPE_PERFORMANCE,
    auto_apply_patch   IN   VARCHAR2 := YES)
  RETURN NUMBER;
```

Parameters

Table 186-36 SQL_DIAGNOSE_AND_REPAIR Function Parameters

Parameter	Description
sql_text	Text of the SQL statement.
sql_id	SQL ID of the SQL query.

Table 186-36 (Cont.) SQL_DIAGNOSE_AND_REPAIR Function Parameters

Parameter	Description
plan_hash_value	The plan to be used for diagnosis. The default value is NULL.
bind_list	Binds to be used for diagnosis. The default value is NULL.
scope	The scope of diagnostic advisor. Possible values are: • SCOPE_LIMITED—only index and plan analyze are invoked for a given SQL. • SCOPE_COMPREHENSIVE—besides index and plan analyze, auto-tune is called first to tune the statement. The default value is SCOPE_COMPREHENSIVE.
time_limit	Time limit for diagnostic task. The default value is TIME_LIMIT_DEFAULT.
problem_type	Problem type that is being diagnosed. The following problem type are supported: • PROBLEM_TYPE_PERFORMANCE—performance problem. • PROBLEM_TYPE_WRONG_RESULTS— incorrect results. • PROBLEM_TYPE_COMPILATION_ERROR—crash during compilation of the statement. • PROBLEM_TYPE_EXECUTION_ERROR— crash during execution of the statement. The default value is PROBLEM_TYPE_PERFORMANCE.
auto_apply_patch	A value that decides if the recommended SQL patch needs to be accepted. Possible values are: • YES—accepts the recommended SQL patch. • NO—does not accepts recommended SQL patch automatically. User need to manually accept the SQL patch. The default value is YES.

186.4.28 UNPACK_STGTAB_SQLPATCH Procedure

This procedure unpacks from the staging table populated by a call to the PACK_STGTAB_SQLPATCH Procedure. It uses the patch data stored in the staging table to create patches on this system. Users can opt to replace existing patches with patch data when they exist already. In this case, note that it is only possible to replace patches referring to the same statement if the names are the same (see the ACCEPT_SQL_PATCH Function & Procedure).

Syntax

```
DBMS_SQLDIAG.UPPACK_STGTAB_SQLPATCH (
   patch_name            IN  VARCHAR2 := '%',
   patch_category        IN  VARCHAR2 := '%',
   replace               IN  BOOLEAN,
```

```
  staging_table_name    IN  VARCHAR2,
  staging_schema_owner  IN  VARCHAR2 := NULL);
```

Parameters

Table 186-37 UPPACK_STGTAB_SQLPATCH Procedure Parameters

Parameter	Description
patch_name	Name of patch to unpack (% wildcards acceptable, case-sensitive)
patch_category	Category from which to unpack patches (% wildcards acceptable, case-insensitive)
replace	Replace patches if they already exist. Note that patches cannot be replaced if there is one in the staging table with the same name as an active patch on different SQL. The subprogram raises an error if there an attempt to create a patch that already exists.
staging_table_name	(Mandatory) Name of the table to use (case-sensitive)
staging_schema_owner	Schema where the table resides, or NULL for current schema (case-sensitive)

Usage Notes

- Requires: ADMINISTER SQL MANAGEMENT OBJECT privilege and SELECT or READ privilege on the staging table
- By default, all SQL patches in the staging table are moved. The function commits after successfully loading each patch. If it fails in creating an individual patch, it raises an error and does not proceed to those remaining in the staging table.

Related Topics

- PACK_STGTAB_SQLPATCH Procedure
  This procedure packs SQL patches into the staging table created by a call to the CREATE_STGTAB_SQLPATCH Procedure.
- ACCEPT_SQL_PATCH Function & Procedure
  This procedure accepts a recommended SQL patch as recommended by the specified SQL diagnosis task.

187
DBMS_SQLPA

The DBMS_SQLPA package provides the interface to implement the SQL Performance Analyzer.

The chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_SQLPA Subprograms

187.1 DBMS_SQLPA Overview

The `DBMS_SQLPA` package provides a capacity to help users predict the impact of system environment changes on the performance of a SQL workload. The interface lets users build and then compare two different versions of the workload performance, analyze the differences between the two versions, and unmask the SQL statements that might be impacted by the changes.

The package provides a task-oriented interface to implement the SQL Performance Analyzer. For example

1. You use the CREATE_ANALYSIS_TASK Functions to create an analysis task for a single statement or a group of SQL statements.s
2. The EXECUTE_ANALYSIS_TASK Function & Procedure executes a previously created analysis task.
3. The REPORT_ANALYSIS_TASK Function displays the results of an analysis task.

187.2 DBMS_SQLPA Security Model

This package is available to `PUBLIC` and performs its own security checking. All analysis task interfaces (`XXX_ANALYSIS_TASK`) require privilege `ADVISOR`.

187.3 Summary of DBMS_SQLPA Subprograms

This table lists the `DBMS_SQLPA` subprograms and briefly describes them.

Table 187-1 DBMS_SQLPA Package Subprograms

Subprogram	Description
CANCEL_ANALYSIS_TASK Procedure	Cancels the currently executing task analysis of one or more SQL statements
CREATE_ANALYSIS_TASK Functions	Creates an advisor task to process and analyze one or more SQL statements
DROP_ANALYSIS_TASK Procedure	Drops a SQL analysis task

Table 187-1 (Cont.) DBMS_SQLPA Package Subprograms

Subprogram	Description
EXECUTE_ANALYSIS_TASK Function & Procedure	Executes a previously created analysis task
INTERRUPT_ANALYSIS_TASK Procedure	Interrupts the currently executing analysis task
REPORT_ANALYSIS_TASK Function	Displays the results of an analysis task
RESET_ANALYSIS_TASK Procedure	Resets the currently executing analysis task to its initial state
RESUME_ANALYSIS_TASK Procedure	Resumes a previously interrupted analysis task that was created to process a SQL tuning set.
SET_ANALYSIS_TASK_PARAMETER Procedures	Sets the SQL analysis task parameter value
SET_ANALYSIS_DEFAULT_PARAMETER Procedures	Sets the SQL analysis task parameter default value

187.3.1 CANCEL_ANALYSIS_TASK Procedure

This procedure cancels the currently executing analysis task. All intermediate result data is removed from the task.

Syntax

```
DBMS_SQLPA.CANCEL_ANALYSIS_TASK(
 task_name         IN VARCHAR2);
```

Parameters

Table 187-2 CANCEL_ANALYSIS_TASK Procedure Parameters

Parameter	Description
task_name	Name of the task to cancel

Examples

Canceling a task when there is a need to stop it executing and it is not required to view any already-completed results:

```
EXEC DBMS_SQLPA.CANCEL_ANALYSIS_TASK(:my_task);
```

187.3.2 CREATE_ANALYSIS_TASK Functions

These functions create an advisor task to process and analyze one or more SQL statements.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

You can use different forms of this function to:

- Create an analysis task for a single statement given its text.
- Create an analysis task for a single statement from the cursor cache given its identifier.
- Create an analysis task for a single statement from the workload repository given a range of snapshot identifiers.
- Create an analysis task for a SQL tuning set.

In all cases, the function creates an advisor task and sets its parameters.

Syntax

SQL text format. This form of the function is called to prepare the analysis of a single statement given its text.

```
DBMS_SQLPA.CREATE_ANALYSIS_TASK(
  sql_text         IN CLOB,
  bind_list        IN sql_binds := NULL,
  parsing_schema   IN VARCHAR2  := NULL,
  task_name        IN VARCHAR2  := NULL,
  description      IN VARCHAR2  := NULL)
RETURN VARCHAR2;
```

SQL ID format. This form of the function is called to prepare the analysis of a single statement from the cursor cache given its identifier.

```
DBMS_SQLPA.CREATE_ANALYSIS_TASK(
  sql_id           IN VARCHAR2,
  plan_hash_value  IN NUMBER    := NULL,
  task_name        IN VARCHAR2  := NULL,
  con_name         IN VARCHAR2    DEFAULT,
  description      IN VARCHAR2  := NULL)
RETURN VARCHAR2;
```

Workload Repository format. This form of the function is called to prepare the analysis of a single statement from the workload repository given a range of snapshot identifiers.

```
DBMS_SQLPA.CREATE_ANALYSIS_TASK(
  dbid             IN NUMBER       DEFAULT,
  begin_snap       IN NUMBER,
  end_snap         IN NUMBER,
  sql_id           IN VARCHAR2,
  plan_hash_value  IN NUMBER    := NULL,
```

```
  task_name        IN VARCHAR2  := NULL,
  description      IN VARCHAR2  := NULL)
  con_name         IN VARCHAR2     DEFAULT,
RETURN VARCHAR2;
```

SQLSET format. This form of the function is called to prepare the analysis of a SQL tuning set.

```
DBMS_SQLPA.CREATE_ANALYSIS_TASK(
  sqlset_name       IN VARCHAR2,
  basic_filter      IN VARCHAR2 :=  NULL,
  con_name          IN VARCHAR2     DEFAULT,
  order_by          IN VARCHAR2 :=  NULL,
  top_sql           IN VARCHAR2 :=  NULL,
  task_name         IN VARCHAR2 :=  NULL,
  description       IN VARCHAR2 :=  NULL
  sqlset_owner      IN VARCHAR2 :=  NULL)
RETURN VARCHAR2;
```

Parameters

Table 187-3 *CREATE_ANALYSIS_TASK Function Parameters*

Parameter	Description
sql_text	Text of a SQL statement
bind_list	A set of bind values
parsing_schema	Name of the schema where the statement can be compiled
task_name	Optional analysis task name
dbid	The DBID for imported or PDB-level AWR data. If NULL, then the current database DBID is used
con_name	Container for the SPA task. The semantics depend on the function format: For the SQL ID format, this parameter specifies the container from which the database fetches the SQL statement for using with SPA. SPA will analyze the statement in this container. If null, then the database uses the current PDB for SPA analysis. For the AWR format, this parameter specifies the container from whose AWR data the database fetches the SQL statement for using with SPA. SPA will analyze the statement in this container. If null, then the database uses the current PDB for SPA analysis. The following statements are true of all function formats: • In a non-CDB, this parameter is ignored. • In a PDB, this parameter must be null or match the container name of the PDB. Otherwise, error occurs. • In a CDB root, this parameter must be null or match the container name of a container in this CDB. Otherwise, error occurs.
description	Description of the SQL analysis task to a maximum of 256 characters
sql_id	Identifier of a SQL statement
plan_hash_value	Hash value of the SQL execution plan
begin_snap	Begin snapshot identifier
end_snap	End snapshot identifier
sqlset_name	SQL tuning set name
basic_filter	SQL predicate to filter the SQL from the SQL tuning set

Table 187-3 (Cont.) *CREATE_ANALYSIS_TASK Function Parameters*

Parameter	Description
order_by	Order-by clause on the selected SQL
top_sql	Top N SQL after filtering and ranking
sqlset_owner	The owner of the SQL tuning set, or NULL for the current schema owner

Return Values

A SQL analysis task name that is unique by user (two different users can give the same name to their advisor tasks).

Examples

```
variable stmt_task VARCHAR2(64);
variable sts_task  VARCHAR2(64);

-- Sql text format
EXEC :stmt_task := DBMS_SQLPA.CREATE_ANALYSIS_TASK(
     sql_text => 'select quantity_sold from sales s, times t where s.time_id = t.time_id
and s.time_id = TO_DATE(''24-NOV-00'')');

-- Sql id format (cursor cache)
EXEC :stmt_task := DBMS_SQLPA.CREATE_ANALYSIS_TASK(
    sql_id       => 'ay1m3ssvtrh24');

-- Workload repository format
exec :stmt_task := DBMS_SQLPA.CREATE_ANALYSIS_TASK(
    begin_snap  => 1,
    end_snap    => 2,
    sql_id      => 'ay1m3ssvtrh24');

-- Sql tuning set format (first we need to load an STS, then analyze it)
EXEC :sts_task := DBMS_SQLPA.CREATE_ANALYSIS_TASK( -
  sqlset_name    =>   'my_workload', -
  order_by       =>   'BUFFER_GETS', -
  description    =>   'process workload ordered by buffer gets');
```

187.3.3 DROP_ANALYSIS_TASK Procedure

This procedure drops a SQL analysis task.The task and all its result data are deleted.

Syntax

```
DBMS_SQLPA.DROP_ANALYSIS_TASK(
 task_name         IN VARCHAR2);
```

Parameters

Table 187-4 DROP_ANALYSIS_TASK Procedure Parameters

Parameter	Description
task_name	The name of the analysis task to drop

187.3.4 EXECUTE_ANALYSIS_TASK Function & Procedure

This function and procedure executes a previously created analysis task, the function version returning the new execution name.

Syntax

```
DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(
   task_name         IN VARCHAR2,
   execution_type    IN VARCHAR2               := 'test execute',
   execution_name    IN VARCHAR2               := NULL,
   execution_params  IN dbms_advisor.argList   := NULL,
   execution_desc    IN VARCHAR2               := NULL)
 RETURN VARCHAR2;

DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(
   task_name         IN VARCHAR2,
   execution_type    IN VARCHAR2               := 'test execute',
   execution_name    IN VARCHAR2               := NULL,
   execution_params  IN dbms_advisor.argList   := NULL,
   execution_desc    IN VARCHAR2               := NULL);
```

Parameters

Table 187-5 EXECUTE_ANALYSIS_TASK Function & Procedure Parameters

Parameter	Description
task_name	Identifier of the task to execute
execution_type	Type of the action to perform by the function. If NULL it will default to the value of the DEFAULT_EXECUTION_TYPE parameter. Possible values are: • [TEST] EXECUTE - test-execute every SQL statement and collect its execution plans and execution statistics. The resulting plans and statistics will be stored in the advisor framework. This is default. • EXPLAIN PLAN - generate explain plan for every statement in the SQL workload. This is similar to the EXPLAIN PLAN command. The resulting plans will be stored in the advisor framework in association with the task. • COMPARE [PERFORMANCE] - analyze and compare two versions of SQL performance data. The performance data is generated by test-executing or generating explain plan of the SQL statements. Use this option when two executions of type EXPLAIN_PLAN or TEST_EXECUTE already exist in the task • CONVERT SQLSET - used to read the statistics captured in a SQL Tuning Set and model them as a task execution. This can be used when you wish to avoid executing the SQL statements because valid data for the experiment already exists in the SQL Tuning Set.
execution_name	A name to qualify and identify an execution. If not specified, it will be generated by the advisor and returned by function.
execution_params	List of parameters (name, value) for the specified execution. The execution parameters have effect only on the execution for which they are specified. They will override the values for the parameters stored in the task (set through the SET_ANALYSIS_DEFAULT_PARAMETER Procedures).
execution_desc	A 256-length string describing the execution

Usage Notes

SQL performance analyzer task can be executed multiples times without having to reset it. For example, when a task is created to perform a change impact analysis on a SQL workload, the created task has to be executed before making any change in the system environment to build a version of the workload that will be used as a reference for performance analysis. Once the change has been made, a second execution is required to build the post-change version of the workload. Finally, the task has to be executed a third time to let the advisor analyze and compare the performance of the workload in both versions.

Examples

1. Create a task with a purpose of change impact analysis

```
EXEC :tname := DBMS_SQLPA.CREATE_ANALYSIS_TASK(
    sqlset_name     => 'my_sts');
```

2. Make baseline or the before change execution

```
EXEC DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(
    task_name       => :tname,
    execution_type  => 'test execute',
    execution_name  => 'before_change');
```

3. Make change

```
...
```

4. Make the after change version of the workload performance

```
EXEC DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(
   task_name          => :tname, -
   execution_type     => 'test execute',
   execution_name     => 'after_change')
```

5. Compare the two versions of the workload

By default we always compare the results of the two last executions. The SQL Performance Analyzer uses the `elapsed_time` as a default metric for comparison. Here we are changing it to `buffer_gets` instead.

```
EXEC DBMS_SQLPA.SET_ANALYSIS_TASK_PARAMETER(
   :tname,'comparison_metric', 'buffer_gets');
EXEC DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(
   task_name         => :tname, -
   execution_type    => 'compare performance', -
   execution_name    => 'after_change');
```

Use the following call if you would like to explicitly specify the two executions to compare as well as the comparison metric to use.

```
EXEC DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(
   task_name        => :tname, -
   execution_type   => 'compare performance',
   execution_params => dbms_advisor.arglist(
      'execution_name1',
      'before_change',
      'execution_name2',
      'after_change',
      'comparison_metric',
      'buffer_gets'));
```

187.3.5 INTERRUPT_ANALYSIS_TASK Procedure

This procedure interrupts the currently executing analysis task. All intermediate result data will not be removed from the task.

Syntax

```
DBMS_SQLPA.INTERRUPT_ANALYSIS_TASK(
   task_name         IN VARCHAR2);
```

Parameters

Table 187-6 INTERRUPT_ANALYSIS_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of the analysis task to interrupt

Examples

```
EXEC DBMS_SQLPA.INTERRUPT_ANALYSIS_TASK(:my_task);
```

187.3.6 REPORT_ANALYSIS_TASK Function

This procedure displays the results of an analysis task.

Syntax

```
DBMS_SQLPA.REPORT_ANALYSIS_TASK(
   task_name       IN   VARCHAR2,
   type            IN   VARCHAR2   := 'TEXT',
   level           IN   VARCHAR2   := 'TYPICAL',
   section         IN   VARCHAR2   := 'SUMMARY',
   object_id       IN   NUMBER     := NULL,
   top_sql         IN   NUMBER     := 100,
   execution_name  IN   VARCHAR2   := NULL,
   task_owner      IN   VARCHAR2   := NULL,
   order_by        IN   VARCHAR2   := NULL)
RETURN CLOB;
```

Parameters

Table 187-7 REPORT_ANALYSIS_TASK Function Parameters

Parameter	Description
task_name	Name of the task to report
type	Type of the report to produce. Possible values are TEXT (default), HTML, XML and ACTIVE (see Usage Notes).

Table 187-7 (Cont.) REPORT_ANALYSIS_TASK Function Parameters

Parameter	Description
level	Level of detail in the report: • ALL - details of all SQL • BASIC - currently the same as typical • CHANGED - only SQL with changed performance • CHANGED_PLANS - only SQL with plan changes • ERRORS - SQL with errors only • IMPROVED - only improved SQL • REGRESSED - only regressed SQL • TIMEOUT - only SQL which timed-out during execution • TYPICAL (default) - show information about every statement analyzed, including changing and errors • UNCHANGED - only SQL with unchanged performance • UNCHANGED_PLANS - only SQL with unchanged plans • UNSUPPORTED - only SQL not supported by SPAs
section	Optionally limit the report to a single section (ALL for all sections): • SUMMARY (default) - workload summary only • ALL - summary and details on SQL
object_id	Identifier of the advisor framework object that represents a given SQL in a tuning set (STS)
top_sql	Number of SQL statements in a STS for which the report is generated
execution_name	Name of the task execution to use. If NULL, the report will be generated for the last task execution.
task_owner	Owner of the relevant analysis task. Defaults to the current schema owner.
order_by	How to sort SQL statements in the report (summary and body). Possible values: • CHANGE_DIFF - sort SQL statements by change difference in SQL performance in terms of the comparison Metric • NULL (default) - order SQL statement by impact on workload • SQL_IMPACT - order SQL statement by change impact on SQL • WORKLOAD_IMPACT - same as NULL • METRIC_DELTA - same as CHANGE_DIFF

Return Values

A CLOB containing the desired report.

Usage Notes

ACTIVE reports have a rich, interactive user interface similar to Enterprise Manager while not requiring any EM installation. The report file built is in HTML format so it can be interpreted by most modern browsers. The code powering the active report is downloaded transparently by the web browser when the report is first viewed, hence viewing it requires outside connectivity.

Examples

```
-- Get the whole report for the single statement case.
SELECT DBMS_SQLPA.REPORT_ANALYSIS_TASK(:stmt_task) from dual;
```

```
-- Show me the summary for the sts case.
SELECT DBMS_SQLPA.REPORT_ANALYSIS_TASK(:sts_task, 'TEXT', 'TYPICAL', 'SUMMARY')
FROM DUAL;

-- Show me the findings for the statement I'm interested in.
SELECT DBMS_SQLPA.REPORT_ANALYSIS_TASK(:sts_task, 'TEXT', 'TYPICAL', 'ALL', 5) from dual;
```

187.3.7 RESET_ANALYSIS_TASK Procedure

This procedure is called on an analysis task that is not currently executing to prepare it for re-execution.

All intermediate result data will be deleted.

Syntax

```
DBMS_SQLPA.RESET_ANALYSIS_TASK(
 task_name         IN VARCHAR2);
```

Parameters

Table 187-8 RESET_ANALYSIS_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of the analysis task to reset

Examples

```
-- reset and re-execute a task
EXEC DBMS_SQLPA.RESET_ANALYSIS_TASK(:sts_task);

-- re-execute the task
EXEC DBMS_SQLPA.EXECUTE_ANALYSIS_TASK(:sts_task);
```

187.3.8 RESUME_ANALYSIS_TASK Procedure

This procedure resumes a previously interrupted or `FAILED` (with a fatal error) task execution.

Syntax

```
DBMS_SQLPA.RESUME_ANALYSIS_TASK(
 task_name         IN VARCHAR2,
 basic_filter      IN VARCHAR2 := NULL);
```

Parameters

Table 187-9 RESUME_ANALYSIS_TASK Procedure Parameters

Parameter	Description
task_name	Identifier of the analysis task to resume
basic_filter	A SQL predicate to filter the SQL from the SQL tuning set. Note that this filter will be applied in conjunction with the basic filter (parameter basic_filter) that was specified when calling the CREATE_ANALYSIS_TASK Functions.

Usage Notes

Resuming a single SQL analysis task (a task that was created to analyze a single SQL statement as compared to a SQL Tuning Set) is not supported.

Examples

```
-- Interrupt the task
EXEC DBMS_SQLPA.INTERRUPT_ANALYSIS_TASK(:conc_task);

-- Once a task is interrupted, we can elect to reset it, resume it, or check
-- out its results and then decide.  For this example we will just resume.

EXEC DBMS_SQLPA.RESUME_ANALYSIS_TASK(:conc_task);
```

187.3.9 SET_ANALYSIS_TASK_PARAMETER Procedures

This procedure sets the SQL analysis task parameter value.

Syntax

This form of the procedure updates the value of a SQL analysis parameter of type `VARCHAR2`.

```
DBMS_SQLPA.SET_ANALYSIS_TASK_PARAMETER(
  task_name         IN VARCHAR2,
  parameter         IN VARCHAR2,
  value             IN VARCHAR2,
  test_execute_dop  IN NUMBER DEFAULT 0,
  compare_resultset IN BOOLEAN DEFAULT TRUE);
```

This form of the procedure updates the value of a SQL analysis parameter of type `NUMBER`.

```
DBMS_SQLPA.SET_ANALYSIS_TASK_PARAMETER(
  task_name         IN VARCHAR2,
  parameter         IN VARCHAR2,
  value             IN NUMBER
  test_execute_dop  IN NUMBER DEFAULT 0,
  compare_resultset IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 187-10 SET_ANALYSIS_TASK_PARAMETER Procedure Parameters

Parameter	Description
`task_name`	Identifier of the task to execute

Table 187-10 (Cont.) SET_ANALYSIS_TASK_PARAMETER Procedure Parameters

Parameter	Description
parameter	Name of the parameter to set. The possible analysis parameters that can be set by this procedure are: • APPLY_CAPTURED_COMPILENV: indicates whether the advisor could use the compilation environment captured with the SQL statements. The default is 0 (that is, NO). • BASIC_FILTER: basic filter for SQL tuning set • CELL_SIMULATION_ENABLED: for more details, see the helper script tcellsim.sql in the ADMIN directory. • COMPARISON_METRIC: specify an expression of execution statistics to use in performance comparison (Example: buffer_gets, cpu_time + buffer_gets * 10) • DATABASE_LINK: can be set to the global name of a PUBLIC database link. When it is set, SQL Performance Analyzer will use the database link for all TEST EXECUTE and EXPLAIN PLAN operations by sending the SQL statements to the remote database to be processed remotely. The analysis results will still be stored on the local database. • DAYS_TO_EXPIRE: number of days until the task is deleted • DEFAULT_EXECUTION_TYPE: the task will default to this type of execution when none is specified by the EXECUTE_ANALYSIS_TASK Function & Procedure.

Table 187-10 (Cont.) SET_ANALYSIS_TASK_PARAMETER Procedure Parameters

Parameter	Description
parameter (contd.)	• DISABLE_MULTI_EXEC: SQL statements are executed multiple times and runtime statistics are then averaged. Set this parameter to TRUE to disable this capability. In this case, each SQL in the SQL tuning set is executed only once. • EXECUTE_TRIGGERS: Set this parameter to TRUE to execute all statement-level triggers in the FULLDML mode. If the parameter is set to FALSE, then the triggers will not be executed even in FULLDML mode of test execution. Any changes incurred due to potential execution of triggers are always rolled back by SPA. The default value of this parameter is FALSE. • EXECUTION_DAYS_TO_EXPIRE: number of days until the tasks's executions will be deleted (without deleting the task) • EXECUTE_FULLDML: TRUE to execute DML statement fully, including acquiring row locks and modifying rows; FALSE (default) to execute only the query part of the DML without modifying data. When TRUE, SQL Performance Analyzer will issue a rollback following DML execution to prevent persistent changes from being made by the DML. • EXECUTION_NAME1: name of the first task execution to analyze • EXECUTION_NAME2: name of the second task execution to analyze • LOCAL_TIME_LIMIT: per-statement time out (seconds) • METRIC_DELTA_THRESHOLD: threshold of the difference between the SQL performance metric before and after the change. The default value is zero. • NUM_ROWS_TO_FETCH: specifies the number of rows to be fetched for an SQL query. You can use one of the following values: – ALL_ROWS: Fetches all the rows for an SQL query – AVERAGE: Number of result rows is calculated as the ratio of total rows processed and total executions for each SQL in the STS – AUTO: Number of result rows is determined using the value of optimizer_mode parameter of the optimizer environment captured in the STS. If the value of optimizer_mode is ALL_ROWS, then all result rows will be fetched. If its value is FIRST_ROWS_n, then n result rows will be fetched by the SPA. – A valid number: Fetches the exact number of rows specified by in the SQL query The default value is ALL_ROWS. • PLAN_FILTER: plan filter for SQL tuning set (see SELECT_SQLSET for possible values) • PLAN_LINES_COMPARISON: - ALWAYS --line by line comparison of plans in all scenarios.- AUTO - Line by Line comparison of plans only if phv2 is not available and phv1 is different- NONE (default) - line by line comparison of plans only if phv is unknown • RANK_MEASURE1: first ranking measure for SQL tuning set • RANK_MEASURE2: second possible ranking measure for SQL tuning set • RANK_MEASURE3: third possible ranking measure for SQL tuning set

Table 187-10 (Cont.) SET_ANALYSIS_TASK_PARAMETER Procedure Parameters

Parameter	Description
	• REPLACE_SYSDATE_WITH: Returns a fixed date for all calls to SYSDATE within the SPA task execution. You can use one of the following values: – CURRENT_SYSDATE: SYSDATE calls return the current date. – SQLSET_SYSDATE: SYSDATE calls return the value of the column LAST_EXEC_START_TIME in the STS The default value is CURRENT_SYSDATE. • RESUME_FILTER: a extra filter for SQL tuning sets besides BASIC_FILTER • SQL_IMPACT_THRESHOLD: threshold of a change impact on a SQL statement. Same as the previous parameter, but at the level of the SQL statement. • SQL_LIMIT: maximum number of SQL statements to process • SQL_PERCENTAGE: percentage filter of SQL tuning set statements • SQLSET_NAME: name of the SQL tuning set to associate to the specified task or task execution. This parameter is mainly using in comparing two SQL tuning sets using SPA. • SQLSET_OWNER: owner of the SQL tuning set specified using task parameter SQLSET_NAME. • TIME_LIMIT: global time out (seconds)
parameter (contd.)	• WORKLOAD_IMPACT_THRESHOLD: threshold of a SQL statement impact on a workload. Statements which workload change impact is below the absolute value of this threshold will be ignored and not considered for improvement or regression. • CON_DBID_MAPPING: provide a mapping of multitenant container database (CDB) IDs. When it is set, SQL Performance Analyzer uses the new CDB ID when it finds a match for the old CDB ID and executes the SQL in that container.
value	New value of the specified parameter
test_execute_dop	Specifies the requested level of concurrency with which a SPA task should be executed. Values 0 or 1 indicate that the SPA task will run with no additional processes as it used to run in releases prior to Oracle Database 18c Release. A value of n (higher than 1) means that n background SPA processes are being requested to concurrently process the input workload.
compare_resultset	Directs SPA to detect if the result-sets between the two trials being compared are different. If differences are seen in the result-sets of any SQL statement between the two trials being compared, the SPA comparison report will indicate this for every such SQL statement. • If set to TRUE the result set comparison will be performed • If set to FALSE result set comparison will not be performed.

Usage Notes

The actual number of processes granted might be equal to or lower than the number requested using the test_execute_dop parameter. This parameter applies only to test-execute or explain plan type of trials that process a SQL Tuning set.

Examples

To request two concurrent processes to execute the SPA task:

```
    dbms_sqlpa.set_analysis_task_parameter(:tname,'TEST_EXECUTE_DOP',2)
```

To enable result-set validation

```
   exec
dbms_sqlpa.set_analysis_task_parameter(:atname,'COMPARE_RESULTSET','TRUE')
```

To disable result-set validation:

```
    exec
dbms_sqlpa.set_analysis_task_parameter(:atname,'COMPARE_RESULTSET','FALSE')
```

187.3.10 SET_ANALYSIS_DEFAULT_PARAMETER Procedures

This procedure sets the SQL analysis task parameter default value.

Syntax

This form of the procedure updates the default value of an analyzer parameter of type `VARCHAR2`.

```
DBMS_SQLPA.SET_ANALYSIS_DEFAULT_PARAMETER(
   parameter    IN  VARCHAR2,
   value        IN  VARCHAR2);
```

This form of the procedure updates the `default` value of an analyzer parameter of type `NUMBER`.

```
DBMS_SQLPA.SET_ANALYSIS_DEFAULT_PARAMETER(
   parameter    IN  VARCHAR2,
   value        IN  NUMBER);
```

Parameters

Table 187-11 SET_ANALYSIS_DEFAULT_PARAMETER Procedure Parameters

Parameter	Description
parameter	Name of the parameter to set. The possible analysis parameters that can be set by this procedure are: • APPLY_CAPTURED_COMPILENV: indicates whether the advisor could use the compilation environment captured with the SQL statements. The default is 0 (that is, NO). • BASIC_FILTER: basic filter for SQL tuning set • COMPARISON_METRIC: specify an expression of execution statistics to use in performance comparison (Example: buffer_gets, cpu_time + buffer_gets * 10) • DATABASE_LINK: can be set to the global name of a PUBLIC database link. When it is set, SQL Performance Analyzer will use the database link for all TEST EXECUTE and EXPLAIN PLAN operations by sending the SQL statements to the remote database to be processed remotely. The analysis results will still be stored on the local database. • DAYS_TO_EXPIRE: number of days until the task is deleted • DEFAULT_EXECUTION_TYPE: the task will default to this type of execution when none is specified by the EXECUTE_ANALYSIS_TASK Function & Procedure. • EXECUTE_FULLDML: TRUE to execute DML statement fully, including acquiring row locks and modifying rows; FALSE (default) to execute only the query part of the DML without modifying data. When TRUE, SQL Performance Analyzer will issue a rollback following DML execution to prevent persistent changes from being made by the DML. • EXECUTION_DAYS_TO_EXPIRE: number of days until the tasks's executions will be deleted (without deleting the task) • EXECUTION_NAME1: name of the first task execution to analyze • EXECUTION_NAME2: name of the second task execution to analyze • LOCAL_TIME_LIMIT: per-statement time out (seconds)
parameter (contd.)	• PLAN_FILTER: plan filter for SQL tuning set (see SELECT_SQLSET for possible values) • RANK_MEASURE1: first ranking measure for SQL tuning set • RANK_MEASURE2: second possible ranking measure for SQL tuning set • RANK_MEASURE3: third possible ranking measure for SQL tuning set • RESUME_FILTER: a extra filter for SQL tuning sets besides BASIC_FILTER • SQL_IMPACT_THRESHOLD: threshold of a change impact on a SQL statement. Same as the previous parameter, but at the level of the SQL statement. • SQL_LIMIT: maximum number of SQL statements to process • SQL_PERCENTAGE: percentage filter of SQL tuning set statements • TIME_LIMIT: global time out (seconds) • WORKLOAD_IMPACT_THRESHOLD: threshold of a SQL statement impact on a workload. Statements which workload change impact is below the absolute value of this threshold will be ignored and not considered for improvement or regression.

Table 187-11 (Cont.) SET_ANALYSIS_DEFAULT_PARAMETER Procedure Parameters

Parameter	Description
`value`	New value of the specified parameter

188
DBMS_SQLQ

The DBMS_SQLQ package provides the interface for configuring quarantine thresholds for execution plans of SQL statements. If any of the Resource Manager thresholds is equal to or less than the quarantine threshold specified in a SQL statement's quarantine configuration, then the SQL statement is not allowed to run, if it uses the execution plan specified in its quarantine configuration.

This chapter contains the following topics:

- DBMS_SQLQ Overview
- Summary of DBMS_SQLQ Subprograms

188.1 DBMS_SQLQ Overview

The DBMS_SQLQ package provides the interface for configuring quarantine thresholds for execution plans of SQL statements. If any of the Resource Manager thresholds is equal to or less than the quarantine threshold specified in a SQL statement's quarantine configuration, then the SQL statement is not allowed to run, if it uses the execution plan specified in its quarantine configuration.

You can use the DBMS_SQLQ package subprograms to:

- create quarantine configurations for execution plans of SQL statements and specify quarantine thresholds for their resource consumption
- query quarantine thresholds specified in quarantine configurations
- delete quarantine configurations
- transfer quarantine configurations from one database to another

188.2 Summary of DBMS_SQLQ Subprograms

This table lists the DBMS_SQLQ subprograms and briefly describes them.

Table 188-1 DBMS_SQLQ Package Subprograms

Procedure	Description
ALTER_QUARANTINE Procedure	Specifies a quarantine threshold in a quarantine configuration for execution plans of a SQL statement
CREATE_QUARANTINE_BY_SQL_ID Function	Creates a quarantine configuration for execution plans of a SQL statement using SQL ID
CREATE_QUARANTINE_BY_SQL_TEXT Function	Creates a quarantine configuration for execution plans of a SQL statement using SQL text
CREATE_STGTAB_QUARANTINE Procedure	Creates a staging table to store quarantine configurations
DROP_QUARANTINE Procedure	Deletes a quarantine configuration

Table 188-1 (Cont.) DBMS_SQLQ Package Subprograms

Procedure	Description
GET_PARAM_VALUE_QUARANTINE Function	Returns the value for a quarantine threshold specified in a quarantine configuration
PACK_STGTAB_QUARANTINE Function	Adds one or more quarantine configurations to a staging table
UNPACK_STGTAB_QUARANTINE Function	Creates quarantine configurations in a database from a staging table

188.2.1 ALTER_QUARANTINE Procedure

This procedure specifies a quarantine threshold for a resource in a quarantine configuration for execution plans of a SQL statement.

Syntax

```
DBMS_SQLQ.ALTER_QUARANTINE (
   quarantine_name    IN VARCHAR2,
   parameter_name     IN VARCHAR2,
   parameter_value    IN VARCHAR2);
```

Parameters

Table 188-2 ALTER_QUARANTINE Procedure Parameters

Parameter	Description
quarantine_name	Name of the quarantine configuration.
parameter_name	Name of the resource for which quarantine threshold needs to be specified. You can specify any one of the following values: • CPU_TIME: CPU time • ELAPSED_TIME: Elapsed time • IO_MEGABYTES: I/O in megabytes • IO_REQUESTS: Number of physical I/O requests • IO_LOGICAL: Number of logical I/O requests • ENABLED: Flag to enable or disable the quarantine configuration. Specify YES to enable it and NO to disable it. The default value is YES. • AUTOPURGE: Flag to enable or disable automatic purging of the quarantine configuration. If it is set to YES, the quarantine configuration is automatically purged after 53 weeks, if not used. If it is set to NO, the quarantine configuration is never purged. The default value is YES.
parameter_value	Quarantine threshold for the resource specified in parameter_name.

Examples

In the following example, the quarantine threshold specified for CPU time is 5 seconds and elapsed time is 10 seconds for the quarantine configuration `SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4`.

```
BEGIN

  DBMS_SQLQ.ALTER_QUARANTINE(
   QUARANTINE_NAME  =>  'SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4',
   PARAMETER_NAME   =>  'CPU_TIME',
   PARAMETER_VALUE  =>  '5');

  DBMS_SQLQ.ALTER_QUARANTINE(
   QUARANTINE_NAME  =>  'SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4',
   PARAMETER_NAME   =>  'ELAPSED_TIME',
   PARAMETER_VALUE  =>  '10');

END;
/
```

When the SQL statement is executed using the execution plan specified in the quarantine configuration, and if the Resource Manager threshold for CPU time is 5 seconds or less, or elapsed time is 10 seconds or less, then the SQL statement is not allowed to run.

188.2.2 CREATE_QUARANTINE_BY_SQL_ID Function

This function creates a quarantine configuration for execution plans of a SQL statement based on SQL ID.

Syntax

```
DBMS_SQLQ.CREATE_QUARANTINE_BY_SQL_ID (
   sql_id           IN VARCHAR2,
   plan_hash_value  IN NUMBER DEFAULT NULL)
RETURN VARCHAR2;
```

Parameters

Table 188-3 CREATE_QUARANTINE_BY_SQL_ID Function Parameters

Parameter	Description
`sql_id`	SQL ID of the SQL statement.
`plan_hash_value`	Hash value of the execution plan of the SQL statement. Default value is `NULL`. When it is `NULL`, the quarantine configuration applies to all the execution plans of the SQL statement.

Return Value

Name of the quarantine configuration.

Examples

The following example creates a quarantine configuration for the SQL statement having the SQL ID of `8vu7s907prbgr`. The quarantine configuration applies to all the execution plans of the SQL statement.

```
DECLARE
  quarantine_config VARCHAR2(30);
BEGIN
  quarantine_config := DBMS_SQLQ.CREATE_QUARANTINE_BY_SQL_ID(SQL_ID => '8vu7s907prbgr');
END;
/
```

The following example creates a quarantine configuration for the execution plan having the hash value of `3488063716` for the SQL statement having the SQL ID of `8vu7s907prbgr`.

```
DECLARE
  quarantine_config VARCHAR2(30);
BEGIN
  quarantine_config := DBMS_SQLQ.CREATE_QUARANTINE_BY_SQL_ID(SQL_ID => '8vu7s907prbgr',
PLAN_HASH_VALUE => '3488063716');
END;
/
```

188.2.3 CREATE_QUARANTINE_BY_SQL_TEXT Function

This function creates a quarantine configuration for execution plans of a SQL statement based on SQL text.

Syntax

```
DBMS_SQLQ.CREATE_QUARANTINE_BY_SQL_TEXT (
   sql_text         IN CLOB,
   plan_hash_value  IN NUMBER DEFAULT NULL)
RETURN VARCHAR2;
```

Parameters

Table 188-4 CREATE_QUARANTINE_BY_SQL_TEXT Function Parameters

Parameter	Description
`sql_text`	SQL statement.
`plan_hash_value`	Hash value of the execution plan of the SQL statement. Default value is `NULL`. When it is `NULL`, the quarantine configuration applies to all the execution plans of the SQL statement.

Return Value

Name of the quarantine configuration.

Examples

The following example creates a quarantine configuration that applies to all the execution plans of the SQL statement `'select count(*) from emp'`.

```
DECLARE
  quarantine_config VARCHAR2(30);
BEGIN
  quarantine_config := DBMS_SQLQ.CREATE_QUARANTINE_BY_SQL_TEXT(SQL_TEXT =>
to_clob('select count(*) from emp'));
END;
/
```

The following example creates a quarantine configuration for the execution plan having the hash value of `3488063716` for the SQL statement having the SQL text of `'select count(*) from emp'`.

```
DECLARE
  quarantine_config VARCHAR2(30);
BEGIN
  quarantine_config := DBMS_SQLQ.CREATE_QUARANTINE_BY_SQL_TEXT(SQL_TEXT =>
to_clob('select count(*) from emp'), PLAN_HASH_VALUE => '3488063716');
END;
/
```

188.2.4 CREATE_STGTAB_QUARANTINE Procedure

This procedure creates a staging table to store the quarantine configurations, so that the staging table can be exported from the current database and imported into another database, thus enabling the quarantine configurations to be used across databases.

Syntax

```
DBMS_SQLQ.CREATE_STGTAB_QUARANTINE (
   staging_table_name   IN VARCHAR2,
   staging_table_owner  IN VARCHAR2 DEFAULT NULL,
   tablespace_name      IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 188-5 CREATE_STGTAB_QUARANTINE Procedure Parameters

Parameter	Description
`staging_table_name`	Name of the staging table.
`staging_table_owner`	Name of the schema owner of the staging table. Default value is `NULL`, which means the database user executing this procedure is set as the staging table owner.
`tablespace_name`	Name of the tablespace in which the staging table needs to be created. Default value is `NULL`, which means the staging table is created in the default tablespace of the database.

Examples

The following example creates the staging table `TBL_STG_QUARANTINE` in the default tablespace of the database and sets its table owner to the database user executing this procedure.

```
BEGIN
  DBMS_SQLQ.CREATE_STGTAB_QUARANTINE(STAGING_TABLE_NAME => 'TBL_STG_QUARANTINE');
END;
/
```

188.2.5 DROP_QUARANTINE Procedure

This procedure deletes a quarantine configuration.

Syntax

```
DBMS_SQLQ.DROP_QUARANTINE(quarantine_name IN VARCHAR2);
```

Parameters

Table 188-6 DROP_QUARANTINE Procedure Parameters

Parameter	Description
quarantine_name	Name of the quarantine configuration to delete.

Examples

The following example deletes the quarantine configuration having the name `SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4`.

```
BEGIN
  DBMS_SQLQ.DROP_QUARANTINE('SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4');
END;
/
```

188.2.6 GET_PARAM_VALUE_QUARANTINE Function

This function returns the quarantine threshold for a resource specified in a quarantine configuration.

Syntax

```
DBMS_SQLQ.GET_PARAM_VALUE_QUARANTINE (
   quarantine_name    IN VARCHAR2,
   parameter_name     IN VARCHAR2)
RETURN VARCHAR2;
```

Parameters

Table 188-7 GET_PARAM_VALUE_QUARANTINE Function Parameters

Parameter	Description
quarantine_name	Name of the quarantine configuration.
parameter_name	Resource for which the quarantine threshold needs to be retrieved.

Return Value

Returns the quarantine threshold for a resource specified in a quarantine configuration.

Examples

The following example returns the quarantine threshold for CPU time specified in the quarantine configuration having the name SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4.

```
DECLARE
  quarantine_config_setting_value VARCHAR2(30);
BEGIN
  quarantine_config_setting_value := DBMS_SQLQ.GET_PARAM_VALUE_QUARANTINE(
                                       QUARANTINE_NAME  =>
'SQL_QUARANTINE_3z0mwuq3aqsm8cfe7a0e4',
                                       PARAMETER_NAME   =>  'CPU_TIME');
END;
/
```

188.2.7 PACK_STGTAB_QUARANTINE Function

This function adds one or more quarantine configurations into a staging table.

Syntax

```
DBMS_SQLQ.PACK_STGTAB_QUARANTINE (
   staging_table_name    IN VARCHAR2,
   staging_table_owner   IN VARCHAR2 DEFAULT NULL,
   name                  IN VARCHAR2 DEFAULT '%',
   sql_text              IN VARCHAR2 DEFAULT '%',
   enabled               IN VARCHAR2 DEFAULT NULL)
RETURN NUMBER;
```

Parameters

Table 188-8 PACK_STGTAB_QUARANTINE Function Parameters

Parameter	Description
staging_table_name	Name of the staging table in which the quarantine configurations need to be added.
staging_table_owner	Name of the schema owner of the staging table. Default value is NULL, which means the database user executing this procedure is set as the staging table owner.

Table 188-8 (Cont.) PACK_STGTAB_QUARANTINE Function Parameters

Parameter	Description
name	Name of the quarantine configuration. Its value is case-sensitive and it accepts wildcard characters.
sql_text	SQL statement text. Its value is case-sensitive and it accepts wildcard characters.
enabled	Flag indicating whether the quarantine configuration should be enabled or disabled. If it is set to YES, then the quarantine configuration is enabled, else it is disabled. Default value is NULL, which means the quarantine configuration is disabled by default.

Return Value

Number of quarantine configurations added to the staging table.

Examples

The following example adds all the quarantine configurations having the names starting with SQL_QUARANTINE_ into the staging table TBL_STG_QUARANTINE.

```
DECLARE
  quarantine_configs NUMBER;
BEGIN
  quarantine_configs := DBMS_SQLQ.PACK_STGTAB_QUARANTINE(
                          STAGING_TABLE_NAME => 'TBL_STG_QUARANTINE',
                          NAME => 'SQL_QUARANTINE_%');
END;
/
```

188.2.8 UNPACK_STGTAB_QUARANTINE Function

This function creates quarantine configurations in a database from a staging table.

Syntax

```
DBMS_SQLQ.UNPACK_STGTAB_QUARANTINE (
   staging_table_name    IN VARCHAR2,
   staging_table_owner   IN VARCHAR2 DEFAULT NULL,
   name                  IN VARCHAR2 DEFAULT '%',
   sql_text              IN VARCHAR2 DEFAULT '%',
   enabled               IN VARCHAR2 DEFAULT NULL)
RETURN NUMBER;
```

Parameters

Table 188-9 UNPACK_STGTAB_QUARANTINE Function Parameters

Parameter	Description
staging_table_name	Name of the staging table from which the quarantine configurations need to be created in the database.
staging_table_owner	Name of the schema owner of the staging table. Default value is NULL, which means the database user executing this procedure is set as the staging table owner.
name	Name of the quarantine configuration. Its value is case-sensitive and it accepts wildcard characters.
sql_text	SQL statement text. Its value is case-sensitive and it accepts wildcard characters.
enabled	Flag indicating whether the quarantine configuration should be enabled or disabled. If its value is YES, then the quarantine configuration is enabled, else it is disabled. Default value is NULL, which means the quarantine configuration is disabled by default.

Return Value

Number of quarantine configurations created in the database from the staging table.

Examples

The following example creates the quarantine configurations in the database from all the quarantine configurations stored in the staging table TBL_STG_QUARANTINE.

```
DECLARE
  quarantine_configs NUMBER;
BEGIN
  quarantine_configs := DBMS_SQLQ.UNPACK_STGTAB_QUARANTINE(
                                    STAGING_TABLE_NAME => 'TBL_STG_QUARANTINE');
END;
/
```

189
DBMS_SQLSET

The `DBMS_SQLSET` package provides an interface to manage SQL tuning sets.

This package provides the same subprograms, although in some cases with slightly different names, as the SQL tuning set subprograms in `DBMS_SQLTUNE`. The difference is that `DBMS_SQLSET` does not require the Oracle Tuning Pack.

This chapter contains the following topics:

- DBMS_SQLSET Overview
- DBMS_SQLSET Security Model
- DBMS_SQLSET Data Structures
- Summary of DBMS_SQLSET Subprograms

189.1 DBMS_SQLSET Overview

Use this package to manage SQL tuning sets.

> **Note:**
>
> All `DBMS_SQLSET` subprograms have equivalents in the `DBMS_SQLTUNE` package.

SQL tuning sets store SQL statements along with the following information:

- The execution context, such as the parsing schema name and bind values
- Execution statistics such as average elapsed time and execution count
- Execution plans, which are the sequence of operations that the database performs to run SQL statements
- Row source statistics such as the number of rows processed for each operation executed within the plan

You can create SQL tuning sets by filtering or ranking SQL statements from several sources:

- The shared SQL area using the SELECT_CURSOR_CACHE Function
- Top SQL statements from the Automatic Workload Repository using the SELECT_WORKLOAD_REPOSITORY Function
- Other SQL tuning sets using the SELECT_SQLSET Function
- SQL Performance Analyzer task comparison results using the SELECT_SQLPA_TASK Function
- SQL Trace files using the SELECT_SQL_TRACE Function
- A user-defined workload

The complete group of subprograms listed in Summary of DBMS_SQLSET Subprograms facilitates this functionality. As examples:

- The CREATE_SQLSET Procedure and Function creates a SQL tuning set object in the database.
- The LOAD_SQLSET Procedure populates the SQL tuning set with a set of selected SQL.
- The CAPTURE_CURSOR_CACHE Procedure collects SQL statements from the shared SQL area over a specified time interval, attempting to build a realistic picture of database workload.

189.2 DBMS_SQLSET Security Model

This package is available to `PUBLIC` and performs its own security checking.

SQL tuning set subprograms require either the `ADMINISTER SQL TUNING SET` or the `ADMINISTER ANY SQL TUNING SET` privilege. Users having the `ADMINISTER SQL TUNING SET` privilege can only create and modify a SQL tuning set that they own. The `ADMINISTER ANY SQL TUNING SET` privilege allows users to manipulate all SQL tuning sets, even those owned by other users. For example, you can use CREATE_SQLSET Procedure and Function to create a SQL tuning set to be owned by a different user. In this case, the different user need not have the `ADMINISTER SQL TUNING SET` privilege to manipulate the SQL tuning set.

189.3 DBMS_SQLSET Data Structures

The `SELECT_*` subprograms in the `DBMS_SQLTUNE` package return objects of the `SQLSET_ROW` type.

Object Types

- SQLSET_ROW Object Type

189.3.1 SQLSET_ROW Object Type

The `SQLSET_ROW` object models the content of a SQL tuning set for the user.

Logically, a SQL tuning set is a collection of `SQLSET_ROW` objects. Each `SQLSET_ROW` contains a single SQL statement along with its execution context, statistics, binds, and plan. The `SELECT_*` subprograms each model a data source as a collection of `SQLSET_ROW` objects, with each object uniquely identified by (`sql_id`, `plan_hash_value`). Similarly, the `LOAD_SQLSET` procedure takes as input a cursor whose row type is `SQLSET_ROW`, treating each `SQLSET_ROW` in isolation according to the policies requested by the user.

Several subprograms package accept basic filters on the content of a SQL tuning set or data source. These filters are expressed in terms of the attributes within the `SQLSET_ROW` as defined.

Syntax

```
CREATE TYPE sqlset_row AS object (
  sql_id                    VARCHAR(13),
  force_matching_signature  NUMBER,
  sql_text                  CLOB,
  object_list               sql_objects,
  bind_data                 RAW(2000),
  parsing_schema_name       VARCHAR2(30),
```

```
 module                  VARCHAR2(48),
 action                  VARCHAR2(32),
 elapsed_time            NUMBER,
 cpu_time                NUMBER,
 buffer_gets             NUMBER,
 disk_reads              NUMBER,
 direct_writes           NUMBER,
 rows_processed          NUMBER,
 fetches                 NUMBER,
 executions              NUMBER,
 end_of_fetch_count      NUMBER,
 optimizer_cost          NUMBER,
 optimizer_env           RAW(2000),
 priority                NUMBER,
 command_type            NUMBER,
 first_load_time         VARCHAR2(19),
 stat_period             NUMBER,
 active_stat_period      NUMBER,
 other                   CLOB,
 plan_hash_value         NUMBER,
 sql_plan                sql_plan_table_type,
 bind_list               sql_binds,
 con_dbid                NUMBER,
 last_exec_start_time    VARCHAR2(19))
```

Attributes

Table 189-1 SQLSET_ROW Attributes

Attribute	Description
sql_id	Unique SQL ID.
forcing_matching_signature	Signature with literals, case, and whitespace removed.
sql_text	Full text for the SQL statement.
object_list	Currently not implemented.
bind_data	Bind data as captured for this SQL. Note that you cannot stipulate an argument for this parameter and also for bind_list - they are mutually exclusive.
parsing_schema_name	Schema where the SQL is parsed.
module	Last application module for the SQL.
action	Last application action for the SQL.
elapsed_time	Sum total elapsed time for this SQL statement.
cpu_time	Sum total CPU time for this SQL statement.
buffer_gets	Sum total number of buffer gets.
disk_reads	Sum total number of disk reads.
direct_writes	Sum total number of direct path writes.
rows_processed	Sum total number of rows processed by this SQL.
fetches	Sum total number of fetches.
executions	Total executions of this SQL statement.

Table 189-1 (Cont.) SQLSET_ROW Attributes

Attribute	Description
end_of_fetch_count	Number of times the SQL statement was fully executed with all of its rows fetched.
optimizer_cost	Optimizer cost for this SQL.
optimizer_env	Optimizer environment for this SQL statement.
priority	User-defined priority (1,2,3).
command_type	Statement type, such as INSERT or SELECT.
first_load_time	Load time of the parent cursor.
stat_period	Period of time (seconds) when the statistics of this SQL statement were collected.
active_stat_period	Effective period of time (in seconds) during which the SQL statement was active.
other	Other column for user-defined attributes.
plan_hash_value	Plan hash value of the plan.
sql_plan	Execution plan for the SQL statement.
bind_list	List of user-specified binds for the SQL statement. This is used for user-specified workloads. Note that you cannot stipulate an argument for this parameter and also for bind_data: they are mutually exclusive.
con_dbid	DBID of the PDB or CDB root.
last_exec_start_time	Most recent execution start time of this SQL statement.

189.4 Summary of DBMS_SQLSET Subprograms

This table lists the DBMS_SQLSET subprograms and briefly describes them.

Table 189-2 DBMS_SQLSET Package Subprograms

Subprogram	Description
ADD_REFERENCE Function	This procedure adds a new reference to an existing SQL tuning set to indicate its use by a client.
CAPTURE_CURSOR_CACHE Procedure	This procedure captures a workload from the shared SQL area into a SQL tuning set.
CREATE_SQLSET Procedure and Function	This procedure or function creates a SQL tuning set object in the database.
CREATE_STGTAB Procedure	This procedure creates a staging table through which SQL tuning sets are imported and exported.
DELETE_SQLSET Procedure	This procedure deletes a set of SQL statements from a SQL tuning set.
DROP_SQLSET Procedure	This procedure drops a SQL tuning set if it is not active.

Table 189-2 (Cont.) DBMS_SQLSET Package Subprograms

Subprogram	Description
LOAD_SQLSET Procedure	This procedure populates the SQL tuning set with a set of selected SQL statements. You can call the procedure multiple times to add new SQL statements or replace attributes of existing statements.
PACK_STGTAB Procedure	This procedure copies one or more SQL tuning sets from their location in the `SYS` schema to a staging table created by the `CREATE_STGTAB` procedure.
REMAP_STGTAB Procedure	This procedure changes the tuning set names and owners in the staging table so that they can be unpacked with different values.
REMOVE_REFERENCE Procedure	This procedure deactivates a SQL tuning set to indicate that it is no longer used by the client.
SELECT_CURSOR_CACHE Function	This function collects SQL statements from the workload repository.
SELECT_SQL_TRACE Function	This table function reads the content of one or more trace files and returns the SQL statements it finds in the format of `sqlset_row`.
SELECT_SQLPA_TASK Function	This function collects SQL statements from a SQL Performance Analyzer comparison task.
SELECT_SQLSET Function	This is a table function that reads the contents of a SQL tuning set.
SELECT_WORKLOAD_REPOSITORY Function	This function collects SQL statements from the workload repository.
UNPACK_STGTAB Procedure	This procedure copies one or more SQL tuning sets from their location in the staging table into the SQL tuning sets schema, making them proper SQL tuning sets.
UPDATE_SQLSET Procedures	This overloaded procedure updates selected fields for SQL statements in a SQL tuning set.

189.4.1 ADD_REFERENCE Function

This procedure adds a new reference to an existing SQL tuning set to indicate its use by a client.

Syntax

```
DBMS_SQLSET.ADD_REFERENCE (
   sqlset_name  IN  VARCHAR2,
   description  IN  VARCHAR2 := NULL,
   sqlset_owner IN  VARCHAR2 :=NULL)
 RETURN NUMBER;
```

Parameters

The parameters are identical for `DBMS_SQLTUNE.ADD_SQLSET_REFERENCE` and `DBMS_SQLSET.ADD_REFERENCE`.

Table 189-3 ADD_SQLSET_REFERENCE and ADD_REFERENCE Function Parameters

Parameter	Description
sqlset_name	Specifies the name of the SQL tuning set.
description	Provides an optional description of the usage of SQL tuning set. The description is truncated if longer than 256 characters.
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for the current schema owner.

Return Values

The identifier of the added reference.

Usage Notes

Adding a reference to a SQL tuning set prevents the tuning set from being modified while it is being used. Invoking SQL Tuning Advisor on the SQL tuning set adds a reference automatically, so use ADD_REFERENCE only when the automatically generated reference is not sufficient. The ADD_REFERENCE function returns a reference ID that you can later supply to the REMOVE_SQLSET_REFERENCE procedure. Query the DBA_SQLSET_REFERENCES view to find all references to a specified SQL tuning set.

Examples

This example generates a reference to the SQL tuning set named my_workload and stores it in the b_rid variable.

```
VARIABLE b_rid NUMBER;
EXEC :b_rid := DBMS_SQLSET.ADD_REFERENCE(sqlset_name => 'my_workload',
description => 'my sts ref');
```

189.4.2 CAPTURE_CURSOR_CACHE Procedure

This procedure captures a workload from the shared SQL area into a SQL tuning set.

The procedure polls the cache multiple times over a time period, and updates the workload data stored there. It can execute over as long a period as required to capture an entire system workload.

Syntax

```
DBMS_SQLSET.CAPTURE_CURSOR_CACHE (
    sqlset_name         IN VARCHAR2,
    time_limit          IN POSITIVE := 1800,
    repeat_interval     IN POSITIVE := 300,
    capture_option      IN VARCHAR2 := 'MERGE',
    capture_mode        IN NUMBER   := MODE_REPLACE_OLD_STATS,
    basic_filter        IN VARCHAR2 := NULL,
    sqlset_owner        IN VARCHAR2 := NULL,
    recursive_sql       IN VARCHAR2 := HAS_RECURSIVE_SQL);
```

Parameters

The parameters are the same for both `DBMS_SQLTUNE.CAPTURE_CURSOR_CACHE_SQLSET` and `DBMS_SQLSET.CAPTURE_CURSOR_CACHE`.

Table 189-4 CAPTURE_CURSOR_CACHE_SQLSET and CAPTURE_CURSOR_CACHE Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the SQL tuning set name
`time_limit`	Defines the total amount of time, in seconds, to execute.
`repeat_interval`	Defines the amount of time, in seconds, to pause between sampling.
`capture_option`	Specifies whether to insert new statements, update existing statements, or both. Values are `INSERT`, `UPDATE`, or `MERGE`. The values are the same as for `load_option` in `load_sqlset`.
`capture_mode`	Specifies the capture mode (`UPDATE` and `MERGE` capture options). Possible values: • `MODE_REPLACE_OLD_STATS` — Replaces statistics when the number of executions is greater than the number stored in the SQL tuning set • `MODE_ACCUMULATE_STATS` — Adds new values to current values for SQL that is already stored. Note that this mode detects if a statement has been aged out, so the final value for a statistics is the sum of the statistics of all cursors that statement existed under.
`basic_filter`	Defines a filter to apply to the shared SQL area for each sample. If `basic_filter` is not set by the caller, then the subprogram captures only statements of type `CREATE TABLE`, `INSERT`, `SELECT`, `UPDATE`, `DELETE`, and `MERGE`.
`sqlset_owner`	Specifies the owner of the SQL tuning set or `NULL` for current schema owner
`recursive_sql`	Defines a filter that includes recursive SQL in the SQL tuning set (`HAS_RECURSIVE_SQL`) or excludes it (`NO_RECURSIVE_SQL`).

Examples

In this example capture takes place over a 30-second period, polling the cache once every five seconds. This captures all statements run during that period but not before or after. If the same statement appears a second time, the process replaces the stored statement with the new occurrence.

Note that in production systems the time limit and repeat interval would be set much higher. You should tune the `time_limit` and `repeat_interval` parameters based on the workload time and shared SQL area turnover properties of your system.

```
EXEC DBMS_SQLSET.CAPTURE_CURSOR_CACHE( -
                                       sqlset_name     => 'my_workload', -
                                       time_limit      =>  30, -
                                       repeat_interval =>  5);
```

In the following call you accumulate execution statistics as you go. This option produces an accurate picture of the cumulative activity of each cursor, even across age-outs, but it is more expensive than the previous example.

```
EXEC DBMS_SQLSET.CAPTURE_CURSOR_CACHE( -
                         sqlset_name     => 'my_workload', -
                         time_limit      => 30, -
                         repeat_interval => 5, -
                         capture_mode    =>
DBMS_SQLSET.MODE_ACCUMULATE_STATS);
```

This call performs a very inexpensive capture where you only insert new statements and do not update their statistics once they have been inserted into the SQL tuning set

```
EXEC DBMS_SQLSET.CAPTURE_CURSOR_CACHE( -
                                    sqlset_name     => 'my_workload', -
                                    time_limit      => 30, -
                                    repeat_interval => 5, -
                                    capture_option  => 'INSERT');
```

189.4.3 CREATE_SQLSET Procedure and Function

This procedure or function creates a SQL tuning set object in the database.

Syntax

```
DBMS_SQLSET.CREATE_SQLSET (
   sqlset_name  IN  VARCHAR2,
   description  IN  VARCHAR2 := NULL
   sqlset_owner IN  VARCHAR2 := NULL);


DBMS_SQLSET.CREATE_SQLSET (
   sqlset_name  IN  VARCHAR2 := NULL,
   description  IN  VARCHAR2 := NULL,
   sqlset_owner IN  VARCHAR2 := NULL)
 RETURN VARCHAR2;
```

Parameters

Table 189-5 CREATE_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the created SQL tuning set. The name is the name passed to the function. If no name is passed to the function, then the function generates an automatic name.
description	Provides an optional description of the SQL tuning set.
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for the current schema owner.

Return Values

Specifies the name of the created SQL tuning set. The name is the name passed to the function. If no name is passed to the function, then the function generates an automatic name.

Examples

```
EXEC DBMS_SQLSET.CREATE_SQLSET(-
  sqlset_name => 'my_workload', -
  description => 'complete application workload');
```

189.4.4 CREATE_STGTAB Procedure

This procedure creates a staging table through which SQL tuning sets are imported and exported.

Syntax

```
DBMS_SQLSET.CREATE_STGTAB (
   table_name        IN VARCHAR2,
   schema_name       IN VARCHAR2 := NULL,
   tablespace_name   IN VARCHAR2 := NULL,
   db_version        IN NUMBER   := NULL);
```

Parameters

Table 189-6 CREATE_STGTAB_SQLSET and CREATE_STGTAB Procedure Parameters

Parameter	Description
table_name	Specifies the of the table to create. The name is case sensitive.
schema_name	Defines the schema in which to create the table, or NULL for the current schema. The name is case sensitive.
tablespace_name	Specifies the tablespace in which to store the staging table, or NULL for the default tablespace of the current user. The name is case sensitive.
db_version	Specifies the database version that determines the format of the staging table. You can also create an older database version staging table to export an STS to an older database version. Use one of the following values: • NULL (default) — Specifies the current database version. • STS_STGTAB_10_2_VERSION — Specifies the 10.2 database version. • STS_STGTAB_11_1_VERSION — Specifies the 11.1 database version. • STS_STGTAB_11_2_VERSION — Specifies the 11.2 database version. • STS_STGTAB_12_1_VERSION — Specifies the 12.1 database version. • STS_STGTAB_12_2_VERSION — Specifies the 12.2 database version.

Examples

Create a staging table for packing and eventually exporting a SQL tuning sets:

```
EXEC DBMS_SQLSET.CREATE_STGTAB(table_name => 'STGTAB_SQLSET');
```

Create a staging table to pack a SQL tuning set in Oracle Database 12*c* Release 1 (12.1.0.2) format:

```
BEGIN
  DBMS_SQLSET.CREATE_STGTAB(
      table_name => 'STGTAB_SQLSET_121'
  ,   db_version => DBMS_SQLSET.STS_STGTAB_12_1_VERSION );
END;
```

189.4.5 DELETE_SQLSET Procedure

This procedure deletes a set of SQL statements from a SQL tuning set.

Syntax

```
DBMS_SQLSET.DELETE_SQLSET (
   sqlset_name    IN  VARCHAR2,
   basic_filter   IN  VARCHAR2 := NULL,
   sqlset_owner   IN  VARCHAR2 := NULL);
```

Parameters

Table 189-7 DELETE_SQLSET Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set.
`basic_filter`	Specifies the SQL predicate to filter the SQL from the SQL tuning set. This basic filter is used as a where clause on the SQL tuning set content to select a desired subset of SQL from the SQL tuning set.
`sqlset_owner`	Specifies the owner of the SQL tuning set, or `NULL` for current schema owner.

Examples

```
-- Delete all statements in a sql tuning set.
EXEC DBMS_SQLSET.DELETE_SQLSET(sqlset_name   => 'my_workload');

-- Delete all statements in a sql tuning set which ran for less than a second
EXEC DBMS_SQLSET.DELETE_SQLSET(sqlset_name   => 'my_workload', -
                               basic_filter  => 'elapsed_time < 1000000');
```

189.4.6 DROP_SQLSET Procedure

This procedure drops a SQL tuning set if it is not active.

Syntax

```
DBMS_SQLSET.DROP_SQLSET (
   sqlset_name    IN  VARCHAR2,
   sqlset_owner   IN  VARCHAR2 := NULL);
```

Parameters

Table 189-8 DROP_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the SQL tuning set.
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for current schema owner.

Usage Notes

You cannot drop a SQL tuning set when it is referenced by one or more clients.

Examples

```
-- Drop the sqlset.
EXEC DBMS_SQLSET.DROP_SQLSET ('my_workload');
```

189.4.7 LOAD_SQLSET Procedure

This procedure populates the SQL tuning set with a set of selected SQL statements. You can call the procedure multiple times to add new SQL statements or replace attributes of existing statements.

Syntax

```
DBMS_SQLSET.LOAD_SQLSET (
   sqlset_name       IN  VARCHAR2,
   populate_cursor   IN  sqlset_cursor,
   load_option       IN VARCHAR2 := 'INSERT',
   update_option     IN VARCHAR2 := 'REPLACE',
   update_condition  IN VARCHAR2 :=  NULL,
   update_attributes IN VARCHAR2 :=  NULL,
   ignore_null       IN BOOLEAN  :=  TRUE,
   commit_rows       IN POSITIVE :=  NULL,
   sqlset_owner      IN VARCHAR2 := NULL);
```

Parameters

Table 189-9 LOAD_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of SQL tuning set to be loaded.
populate_cursor	Specifies the cursor reference to the SQL tuning set to be loaded.
load_option	Specifies which statements are loaded into the SQL tuning set. The possible values are: • INSERT (default) — Adds only new statements. • UPDATE — Updates existing the SQL statements and ignores any new statements. • MERGE — Inserts new statements and updates the information of the existing ones.

Table 189-9 (Cont.) LOAD_SQLSET Procedure Parameters

Parameter	Description
update_option	Specifies how existing SQL statements are updated. This parameter is considered only if load_option is specified with UPDATE or MERGE as an option. The possible values are: • REPLACE (default) — Updates the statement using the new statistics, bind list, object list, and so on. • ACCUMULATE — Combines attributes when possible (for example, statistics such as elapsed_time), and otherwise replaces the existing values (for example, module and action) with the provided values. The SQL statement attributes that can be accumulated are: elapsed_time, buffer_gets, direct_writes, disk_reads, row_processed, fetches, executions, end_of_fetch_count, stat_period and active_stat_period.
update_condition	Specifies when to perform the update. The procedure only performs the update when the specified condition is satisfied. The condition can refer to either the data source or destination. The condition must use the following prefixes to refer to attributes from the source or the destination: • OLD — Refers to statement attributes from the SQL tuning set (destination). • NEW — Refers to statement attributes from the input statements (source).
update_attributes	Specifies the list of SQL statement attributes to update during a merge or update. The possible values are: • NULL (default) — Specifies the content of the input cursor except the execution context. On other terms, it is equivalent to ALL without execution contexts such as module and action. • BASIC — Specifies statistics and binds only. • TYPICAL — Specifies BASIC with SQL plans (without row source statistics) and without an object reference list. • ALL — Specifies all attributes, including the execution context attributes such as module and action. • List of comma separated attribute names to update: – EXECUTION_CONTEXT – EXECUTION_STATISTICS – SQL_BINDS – SQL_PLAN – SQL_PLAN_STATISTICS (similar to SQL_PLAN with added row source statistics)
ignore_null	Specifies whether to update attributes when the new value is NULL. If TRUE, then the procedure does not update an attribute when the new value is NULL. That is, do not override with NULL values unless intentional.
commit_rows	Specifies whether to commit statements after DML. If a value is provided, then the load commits after each specified number of statements is inserted. If NULL is provided, then the load commits only once, at the end of the operation. Providing a value for this argument enables you to monitor the progress of a SQL tuning set load operation in the DBA_SQLSET views. The STATEMENT_COUNT value increases as new SQL statements are loaded.

Table 189-9 (Cont.) LOAD_SQLSET Procedure Parameters

Parameter	Description
sqlset_owner	Defines the owner of the SQL tuning set, or the current schema owner (or NULL for the current owner).

Exceptions

- This procedure returns an error when sqlset_name is invalid, or a corresponding SQL tuning set does not exist, or the populate_cursor is incorrect and cannot be executed.
- Exceptions are also raised when invalid filters are provided. Filters can be invalid either because they don't parse (for example, they refer to attributes not in sqlset_row), or because they violate the user's privileges.

Usage Notes

Rows in the input populate_cursor must be of type SQLSET_ROW.

Examples

In this example, you create and populate a SQL tuning set with all shared SQL area statements with an elapsed time of 5 seconds or more, excluding statements that belong to SYS schema. You select all attributes of the SQL statements and load them in the tuning set using the default mode. The default mode loads only new statements because the SQL tuning set is empty.

```
-- create the tuning set
EXEC DBMS_SQLSET.CREATE_SQLSET('my_workload');

-- populate the tuning set from the shared SQL area
DECLARE
 cur DBMS_SQLSET.SQLSET_CURSOR;
BEGIN
 OPEN cur FOR
   SELECT VALUE(P)
     FROM table(
       DBMS_SQLSET.SELECT_CURSOR_CACHE(
         'parsing_schema_name <> ''SYS'' AND elapsed_time > 5000000',
          NULL, NULL, NULL, NULL, 1, NULL,
         'ALL')) P;

DBMS_SQLSET.LOAD_SQLSET(sqlset_name     => 'my_workload',
                        populate_cursor => cur);

END;
/
```

Now you want to augment this information with what is stored in the workload repository (AWR). You populate the tuning set with 'ACCUMULATE' as your update_option because it is assumed the cursors currently in the cache have aged out since the snapshot was taken.

You omit the `elapsed_time` filter because it is assumed that any statement captured in AWR is important, but still you throw away the `SYS`-parsed cursors to avoid recursive SQL.

```
DECLARE
  cur DBMS_SQLSET.SQLSET_CURSOR;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
      FROM table(
        DBMS_SQLSET.SELECT_WORKLOAD_REPOSITORY(1,2,
                                               'parsing_schema_name <>
''SYS''',
                                               NULL, NULL,NULL,NULL,
                                               1,
                                               NULL,
                                               'ALL')) P;

  DBMS_SQLSET.LOAD_SQLSET(sqlset_name     => 'my_workload',
                          populate_cursor => cur,
                          load_option     => 'MERGE',
                          update_option   => 'ACCUMULATE');
END;
```

The following example is a simple load that only inserts new statements from the workload repository, skipping existing ones (in the SQL tuning set). Note that `'INSERT'` is the default value for the `load_option` argument of the `LOAD_SQLSET` procedure.

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
  SELECT VALUE(P)
  FROM table(DBMS_SQLSET.SELECT_WORKLOAD_REPOSITORY(1,2)) P;
  DBMS_SQLSET.LOAD_SQLSET(sqlset_name => 'my_workload', populate_cursor =>
cur);
END;
/
```

The next example demonstrates a load with `UPDATE` option. This updates statements that already exist in the SQL tuning set but does not add new ones. By default, old statistics are replaced by their new values.

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLSET.SELECT_CURSOR_CACHE) P;

  DBMS_SQLSET.LOAD_SQLSET(sqlset_name     => 'my_workload',
                          populate_cursor => cur,
                          load_option     => 'UPDATE');
END;
/
```

189.4.8 PACK_STGTAB Procedure

This procedure copies one or more SQL tuning sets from their location in the `SYS` schema to a staging table created by the `CREATE_STGTAB` procedure.

Syntax

```
DBMS_SQLSET.PACK_STGTAB (
   sqlset_name          IN VARCHAR2,
   sqlset_owner         IN VARCHAR2 := NULL,
   staging_table_name   IN VARCHAR2,
   staging_schema_owner IN VARCHAR2 := NULL,
   db_version           IN NUMBER   := NULL);
```

Parameters

The parameters are identical for the `DBMS_SQLTUNE.PACK_STGTAB_SQLSET` and `DBMS_SQLSET.PACK_STGTAB` procedures.

Table 189-10 PACK_STGTAB_SQLSET and PACK_STGTAB Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set to pack. The name is case sensitive. Wildcard characters (`%`) are permitted.
`sqlset_owner`	Specifies the category from which to pack SQL tuning sets. The name is case sensitive. Wildcard characters (`%`) are permitted.
`staging_table_name`	Specifies the name of the table to use. The value is case sensitive.
`staging_schema_owner`	Specifies the schema where the table resides, or `NULL` for the current schema. The value is case sensitive.
`db_version`	Specifies the database version that determines the format of the staging table. You can also create an older database version staging table to export an STS to an older database version. Use any of the following values: • `NULL` (default) — Specifies the current database version. • `STS_STGTAB_10_2_VERSION` — Specifies the 10.2 database version. • `STS_STGTAB_11_1_VERSION` — Specifies the 11.1 database version. • `STS_STGTAB_11_2_VERSION` — Specifies the 11.2 database version.

Usage Notes

- To move more than one SQL tuning set, call this procedure multiple times. You can then move the populated staging table to a destination database using any method, such as a database link or Oracle Data Pump, and then unpack the SQL tuning set in the destination database.
- This function issues a `COMMIT` after packing each SQL tuning set. If an error is raised mid-execution, then clear the staging table by deleting its rows.

Examples

Put all SQL tuning sets on the database in the staging table:

```
BEGIN
  DBMS_SQLSET.PACK_STGTAB(
      sqlset_name        => '%'
  ,   sqlset_owner       => '%'
  ,   staging_table_name => 'STGTAB_SQLSET');
END;
```

Put only those SQL tuning sets owned by the current user in the staging table:

```
BEGIN
  DBMS_SQLSET.PACK_STGTAB(
      sqlset_name        => '%'
  ,   staging_table_name => 'STGTAB_SQLSET');
END;
```

Pack a specific SQL tuning set:

```
BEGIN
  DBMS_SQLSET.PACK_STGTAB(
      sqlset_name         => 'my_workload'
  ,   staging_table_name  => 'STGTAB_SQLSET');
END;
```

Pack a second SQL tuning set:

```
BEGIN
  DBMS_SQLSET.PACK_STGTAB(
      sqlset_name         => 'workload_subset'
  ,   staging_table_name  => 'STGTAB_SQLSET');
END;
```

Pack the STS `my_workload_subset` into a staging table `stgtab_sqlset` created for Oracle Database 11*g* Release 2 (11.2):

```
BEGIN
  DBMS_SQLSET.PACK_STGTAB(
      sqlset_name          => 'workload_subset'
  ,   staging_table_name   => 'STGTAB_SQLSET'
  ,   db_version           => DBMS_SQLSET.STS_STGTAB_11_2_VERSION);
END;
```

189.4.9 REMAP_STGTAB Procedure

This procedure changes the tuning set names and owners in the staging table so that they can be unpacked with different values.

Syntax

```
DBMS_SQLSET.REMAP_STGTAB (
   old_sqlset_name         IN VARCHAR2,
   old_sqlset_owner        IN VARCHAR2 := NULL,
   new_sqlset_name         IN VARCHAR2 := NULL,
   new_sqlset_owner        IN VARCHAR2 := NULL,
   staging_table_name      IN VARCHAR2,
   staging_schema_owner    IN VARCHAR2 := NULL
   old_con_dbid            IN NUMBER   := NULL,
   new_con_dbid            IN NUMBER   := NULL);
);
```

Parameters

The parameters are identical for the DBMS_SQLTUNE.REMAP_STGTAB_SQLSET and DBMS_SQLSET.REMAP_SQLSET procedures.

Table 189-11 REMAP_STGTAB_SQLSET and REMAP_SQLSET Procedure Parameters

Parameter	Description
`old_sqlset_name`	Specifies the name of the tuning set to target for a remap operation. Wildcard characters (`%`) are not supported.
`old_sqlset_owner`	Specifies the new name of the tuning set owner to target for a remap operation. `NULL` for current schema owner
`new_sqlset_name`	Specifies the new name for the tuning set, or `NULL` to keep the same tuning set name.
`new_sqlset_owner`	Specifies the new owner for the tuning set, or `NULL` to keep the same owner name.
`staging_table_name`	Specifies the name of the table on which to perform the remap operation. The value is case sensitive.
`staging_schema_owner`	Specifies the name of staging table owner, or `NULL` for the current schema owner. The value is case sensitive.
`old_con_dbid`	Specifies the old container DBID to be remapped to a new container DBID. Specify `NULL` to use the same container DBID. You must provide both `old_con_dbid` and `new_con_dbid` for the remap to succeed.
`new_con_dbid`	Specifies the new container DBID to replace with the old container DBID. Specify `NULL` to use the same container DBID. You must provide both `old_con_dbid` and `new_con_dbid` for the remap to succeed.

Usage Notes

Call this procedure multiple times to remap more than one tuning set name or owner. This procedure only handles one tuning set per call.

Examples

```
-- Change the name of an STS in the staging table before unpacking it.
BEGIN
  DBMS_SQLSET.REMAP_STGTAB(
      old_sqlset_name    =>
'my_workload'
  ,   old_sqlset_owner   => 'SH'
  ,   new_sqlset_name    =>
'imp_workload'
  ,   staging_table_name => 'STGTAB_SQLSET');

-- Change the owner of an STS in the staging table before unpacking it.
  DBMS_SQLSET.REMAP_STGTAB(
      old_sqlset_name    => 'imp_workload'
  ,   old_sqlset_owner   => 'SH'
  ,   new_sqlset_owner   => 'SYS'
  ,   staging_table_name => 'STGTAB_SQLSET');
END;
```

189.4.10 REMOVE_REFERENCE Procedure

This procedure deactivates a SQL tuning set to indicate that it is no longer used by the client.

Syntax

```
DBMS_SQLSET.REMOVE_REFERENCE (
   sqlset_name    IN  VARCHAR2,
   reference_id   IN  NUMBER,
   sqlset_owner   IN  VARCHAR2 := NULL,
   force_remove   IN  NUMBER   := 0);
```

Parameters

The parameters are identical for the `DBMS_SQLTUNE.REMOVE_SQLSET_REFERENCE` and `DBMS_SQLSET.REMOVE_REFERENCE` procedures.

Table 189-12 REMOVE_SQLSET_REFERENCE and REMOVE_REFERENCE Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set.
`reference_id`	Specifies the identifier of the reference to remove.
`sqlset_owner`	Specifies the owner of the SQL tuning set (or `NULL` for the current schema owner).
`force_remove`	Specifies whether references can be removed for other users (`1`) or whether they cannot be removed (`0`). Setting this parameter to `1` only takes effect when the user has the `ADMINISTER ANY SQL TUNING SET` privilege. Otherwise, the database only removes references owned by the user.

Examples

You can remove references on a given SQL tuning set when you finish using it and want to make it writable again. The following example removes the reference to `my_workload`:

```
EXEC DBMS_SQLSET.REMOVE_REFERENCE(sqlset_name   => 'my_workload', -
                                  reference_id  => :rid);
```

Use the `DBA_SQLSET_REFERENCES` view to find all references to a given SQL tuning set.

189.4.11 SELECT_CURSOR_CACHE Function

This function collects SQL statements from the shared SQL area.

Syntax

```
DBMS_SQLSET.SELECT_CURSOR_CACHE (
  basic_filter         IN   VARCHAR2 := NULL,
  object_filter        IN   VARCHAR2 := NULL,
  ranking_measure1     IN   VARCHAR2 := NULL,
  ranking_measure2     IN   VARCHAR2 := NULL,
  ranking_measure3     IN   VARCHAR2 := NULL,
  result_percentage    IN   NUMBER   := 1,
  result_limit         IN   NUMBER   := NULL,
  attribute_list       IN   VARCHAR2 := 'TYPICAL',
  recursive_sql        IN   VARCHAR2 := HAS_RECURSIVE_SQL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 189-13 SELECT_CURSOR_CACHE Function Parameters

Parameter	Description
`basic_filter`	Specifies the SQL predicate that filters the SQL from the shared SQL area defined on attributes of the `SQLSET_ROW`. If `basic_filter` is not set by the caller, then the subprogram captures only statements of the type `CREATE TABLE`, `INSERT`, `SELECT`, `UPDATE`, `DELETE`, and `MERGE`.
`object_filter`	Currently not supported.
`ranking_measure(n)`	Defines an `ORDER BY` clause on the selected SQL.
`result_percentage`	Specifies a filter that picks the top *n*% according to the supplied ranking measure. The value applies only if one ranking measure is supplied.
`result_limit`	Defines the top limit SQL from the filtered source ranked by the ranking measure.

Table 189-13 (Cont.) SELECT_CURSOR_CACHE Function Parameters

Parameter	Description
attribute_list	Specifies the list of SQL statement attributes to return in the result. Possible values are: • TYPICAL — Specifies BASIC plus SQL plan (without row source statistics) and without object reference list (default). • BASIC — Specifies all attributes (such as execution statistics and binds) except the plans. The execution context is always part of the result. • ALL — Specifies all attributes. • Comma-separated list of attribute names. This values returns only a subset of SQL attributes: – EXECUTION_STATISTICS – BIND_LIST – OBJECT_LIST – SQL_PLAN – SQL_PLAN_STATISTICS — Similar to SQL_PLAN plus row source statistics
recursive_sql	Specifies that the filter must include recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL, which is the default) or exclude it (NO_RECURSIVE_SQL).

Return Values

This function returns a one SQLSET_ROW per SQL_ID or PLAN_HASH_VALUE pair found in each data source.

Usage Notes

- Filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.
- Users need privileges on the shared SQL area views.

Example 189-1 Statements with 500 or More Buffer Gets

This query obtains the SQL IDs and SQL text for statements with 500 buffer gets:

```
SELECT SQL_ID, SQL_TEXT
FROM   TABLE(DBMS_SQLSET.SELECT_CURSOR_CACHE('buffer_gets > 500'))
ORDER BY sql_id;
```

Example 189-2 All Information About a Statement

The following query obtains all information about the SQL statement with the SQL ID 4rm4183czbs7j:

```
SELECT * FROM TABLE(DBMS_SQLSET.SELECT_CURSOR_CACHE('sql_id =
''4rm4183czbs7j'''));
```

Example 189-3 Multiple Plans for a SQL Statement

A data source may store multiple plans for each SQL statement. The output of the `SELECT_CURSOR_CACHE` function is a SQL row set object that is uniquely identified by SQL ID and plan hash value. This example queries the plan hash values for the statement with the SQL ID `ay1m3ssvtrh24`:

```
SELECT sql_id, plan_hash_value
FROM table(DBMS_SQLSET.select_cursor_cache('sql_id = ''ay1m3ssvtrh24'''))
ORDER BY sql_id, plan_hash_value;
```

Example 189-4 Processing All Statements in the Shared SQL Area

This example processes all statements in the shared SQL area:

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(p)
    FROM TABLE(DBMS_SQLSET.SELECT_CURSOR_CACHE) p;

  -- Process each statement in cursor (or pass cursor to load_sqlset).

  CLOSE cur;
END;
/
```

Example 189-5 Process Statements Not Parsed by SYS

This example processes all statements not parsed in the `SYS` schema:

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur for
    SELECT VALUE(p)
    FROM TABLE(
     DBMS_SQLSET.SELECT_CURSOR_CACHE('parsing_schema_name <> ''SYS''')) p;

  -- Process each statement (or pass cursor to load_sqlset).

  CLOSE cur;
end;
/
```

Example 189-6 All Statements from an Application Module and Action

This example processes all statements from a specified application module and action:

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(p)
```

```
    FROM TABLE(
      DBMS_SQLSET.SELECT_CURSOR_CACHE(
         'module = ''MY_APPLICATION'' and action = ''MY_ACTION''')) p;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;/
```

Example 189-7 All Statements Whose Elapsed Time Is At Least Five Seconds

This example processes all statements that ran for at least five seconds:

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLSET.SELECT_CURSOR_CACHE('elapsed_time > 5000000')) P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

Example 189-8 Statements Parsed in the APPS Schema

This example processes all SQL statements that were parsed in the `APPS` schema and had more than 100 buffer gets:

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(p)
    FROM TABLE(
      DBMS_SQLSET.SELECT_CURSOR_CACHE(
        'buffer_gets > 100 and parsing_schema_name = ''APPS'''))p;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

Example 189-9 Plans and SQL Statements

This example processes all SQL statements exceeding 5 seconds. It also selects the plans for these statements. For performance reasons, the example selects execution statistics and SQL binds. The `SQL_PLAN` attribute of `sqlset_row` is `NULL`.

```
-- select all statements exceeding 5 seconds in elapsed time, but also
-- select the plans (by default we only select execution stats and binds
-- for performance reasons - in this case the SQL_PLAN attribute of sqlset_row
```

```
-- is NULL)
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(p)
    FROM TABLE(DBMS_SQLSET.SELECT_CURSOR_CACHE(
      basic_filter      => 'elapsed_time > 5000000',
      object_filter     => NULL,
      ranking_measure1  => NULL,
      ranking_measure2  => NULL,
      ranking_measure3  => NULL,
      result_percentage => 1,
      result_limit      => NULL,
      attribute_list    => 'EXECUTION_STATISTICS, SQL_BINDS, SQL_PLAN',
      recursive_sql     => HAS_RECURSIVE_SQL)) p;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;/
```

Example 189-10 Top 100 Statements Ordered by Elapsed Time

This example selects the top 100 statements in the shared SQL area, ordered by elapsed time:

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(p)
    FROM TABLE(DBMS_SQLSET.SELECT_CURSOR_CACHE(
      basic_filter      => NULL,
      object_filter     => NULL,
      ranking_measure_1 => 'ELAPSED_TIME',
      ranking_measure_2 => NULL,
      ranking_measure_3 => NULL,
      result_percentage => 1,
      result_limit      => 100,
      attribute_list    => 'TYPICAL',
      recursive_sql     => HAS_RECURSIVE_SQL))) p;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

Example 189-11 Statements Responsible for Most Buffer Gets

This example processes statements that cumulatively account for 90% of the buffer gets in the shared SQL area. The buffer gets of all statements added together is approximately 90% of the sum of all statements currently in the shared SQL area.

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLSET.SELECT_CURSOR_CACHE(
      basic_filter      => NULL,
      object_filter     => NULL,
      ranking_measure_1 => 'BUFFER_GETS',
      ranking_measure_2 => NULL,
      ranking_measure_3 => NULL,
      result_percentage => .9,
      result_limit      => NULL,
      attribute_list    => 'TYPICAL',
      recursive_sql     => HAS_RECURSIVE_SQL))) p;

  -- Process each statement (or pass cursor to load_sqlset).

  CLOSE cur;
END;
/
```

189.4.12 SELECT_SQL_TRACE Function

This table function reads the content of one or more trace files and returns the SQL statements it finds in the format of `sqlset_row`.

Syntax

```
DBMS_SQLSET.SELECT_SQL_TRACE (
  directory              IN VARCHAR2,
  file_name              IN VARCHAR2 := NULL,
  mapping_table_name     IN VARCHAR2 := NULL,
  mapping_table_owner    IN VARCHAR2 := NULL,,
  select_mode            IN POSITIVE := SINGLE_EXECUTION,
  options                IN BINARY_INTEGER := LIMITED_COMMAND_TYPE,
  pattern_start          IN VARCHAR2 := NULL,
  pattern_end            IN VARCHAR2 := NULL,
  result_limit           IN POSITIVE := NULL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 189-14 SELECT_SQL_TRACE Function Parameters

Parameter	Description
`directory`	Defines the directory object containing the trace files. This field is mandatory.

Table 189-14 (Cont.) SELECT_SQL_TRACE Function Parameters

Parameter	Description
file_name	Specifies all or part of the name of the trace files. If NULL, then the function uses the current or most recent file in the specified location or path. '%' wildcards are supported for matching trace file names.
mapping_table_name	Specifies the mapping table name. Note that the mapping table name is case insensitive. If the mapping table name is NULL, then the function uses the mappings in the current database.
mapping_table_owner	Specifies the mapping table owner. If it is NULL, then the function uses the current user.
select_mode	Specifies the mode for selecting SQL from the trace. Possible values are: • SINGLE_EXECUTION — Returns one execution of a SQL. This is the default. • ALL_EXECUTIONS — Returns all executions.
options	Specifies which types of SQL statements are returned. • LIMITED_COMMAND_TYPE — Returns the SQL statements with the command types CREATE, INSERT, SELECT, UPDATE, DELETE, and MERGE. This value is the default. • ALL_COMMAND_TYPE — Returns the SQL statements with all command types.
pattern_start	Specifies the delimiting pattern of the trace file sections to consider. CURRENTLY INOPERABLE.
pattern_end	Specifies the closing delimiting pattern of the trace file sections to process. CURRENTLY INOPERABLE.
result_limit	Specifies the top SQL from the filtered source. Default to MAXSB4 if NULL.

Return Values

This function returns a SQLSET_ROW object.

Usage Notes

The ability to create a directory object for the system directory creates a potential security issue. For example, in a CDB, all containers write trace files to the same directory. A local user with SELECT privileges on this directory can read the contents of trace files belonging to any container.

To prevent this type of unauthorized access, copy the files from the default SQL trace directory into a different directory, and then create a directory object. Use the PATH_PREFIX clause of the CREATE PLUGGABLE DATABASE statement to ensure that all directory object paths associated with the PDB are restricted to the specified directory or its subdirectories.

Examples

The following code shows how to enable SQL trace for a few SQL statements and load the results into a SQL tuning set:

```
-- turn on the SQL trace in the capture database
ALTER SESSION SET EVENTS '10046 TRACE NAME CONTEXT FOREVER, LEVEL 4'
```

```
-- run sql statements
SELECT 1 FROM DUAL;
SELECT COUNT(*) FROM dba_tables WHERE table_name = :mytab;

ALTER SESSION SET EVENTS '10046 TRACE NAME CONTEXT OFF';

-- create mapping table from the capture database
CREATE TABLE mapping AS
SELECT object_id id, owner, substr(object_name, 1, 30) name
   FROM dba_objects
   WHERE object_type NOT IN ('CONSUMER GROUP', 'EVALUATION CONTEXT',
                             'FUNCTION', 'INDEXTYPE', 'JAVA CLASS',
                             'JAVA DATA', 'JAVA RESOURCE', 'LIBRARY',
                             'LOB', 'OPERATOR', 'PACKAGE',
                             'PACKAGE BODY', 'PROCEDURE', 'QUEUE',
                             'RESOURCE PLAN', 'TRIGGER', 'TYPE',
                             'TYPE BODY')
UNION ALL
SELECT user_id id, username owner, NULL name
   FROM dba_users;

-- create the directory object where the SQL traces are stored
CREATE DIRECTORY SQL_TRACE_DIR as '/home/foo/trace';

-- create the STS
EXEC DBMS_SQLSET.CREATE_SQLSET('my_sts', 'test purpose');

-- load the SQL statements into STS from SQL TRACE
DECLARE
   cur sys_refcursor;
BEGIN
   OPEN cur FOR
   SELECT value(p)
     FROM TABLE(
        DBMS_SQLSET.SELECT_SQL_TRACE(
           directory=>'SQL_TRACE_DIR',
           file_name=>'%trc',
           mapping_table_name=>'mapping')) p;
   DBMS_SQLSET.LOAD_SQLSET('my_sts', cur);
   CLOSE cur;
END;
/
```

See Also:

Oracle Database SQL Language Reference to learn more about the `PATH_PREFIX` clause

189.4.13 SELECT_SQLPA_TASK Function

This function collects SQL statements from a SQL Performance Analyzer comparison task.

> **See Also:**
>
> *Oracle Database Testing Guide* for a `SELECT_SQLPA_TASK` example

Syntax

```
DBMS_SQLSET.SELECT_SQLPA_TASK(
    task_name          IN VARCHAR2,
    task_owner         IN VARCHAR2 := NULL,
    execution_name     IN VARCHAR2 := NULL,
    level_filter       IN VARCHAR2 := 'REGRESSED',
    basic_filter       IN VARCHAR2 := NULL,
    object_filter      IN VARCHAR2 := NULL,
    attribute_list     IN VARCHAR2 := 'TYPICAL')
RETURN sys.sqlset PIPELINED;
```

Parameters

Table 189-15 SELECT_SQLPA_TASK Function Parameters

Parameter	Description
`task_name`	Specifies the name of the SQL Performance Analyzer task.
`task_owner`	Specifies the owner of the SQL Performance Analyzer task. If `NULL`, then assume the current user.
`execution_name`	Specifies the name of the SQL Performance Analyzer task execution (type `COMPARE PERFORMANCE`) from which the provided filters will be applied. If `NULL`, then assume the most recent `COMPARE PERFORMANCE` execution.
`level_filter`	Specifies which subset of SQL statements to include. Same format as `DBMS_SQLPA.REPORT_ANALYSIS_TASK.LEVEL`, with some possible strings removed. • `IMPROVED` includes only improved SQL. • `REGRESSED` includes only regressed SQL (default). • `CHANGED` includes only SQL with changed performance. • `UNCHANGED` includes only SQL with unchanged performance. • `CHANGED_PLANS` includes only SQL with plan changes. • `UNCHANGED_PLANS` includes only SQL with unchanged plans. • `ERRORS` includes only SQL with errors only. • `MISSING_SQL` includes only missing SQL statements (across STS). • `NEW_SQL` includes only new SQL statements (across STS).
`basic filter`	Specifies the SQL predicate to filter the SQL in addition to the level filters.
`object_filter`	Currently not supported.

Table 189-15 (Cont.) SELECT_SQLPA_TASK Function Parameters

Parameter	Description
attribute_list	Defines the SQL statement attributes to return in the result. Possible values are: • TYPICAL — Returns BASIC plus the SQL plan (without row source statistics) and without an object reference list. This is the default. • BASIC — Returns all attributes (such as execution statistics and binds) except the plans. The execution context is always part of the result. • ALL — Returns all attributes. • Comma-separated list of attribute names this allows to return only a subset of SQL attributes: EXECUTION_STATISTICS, SQL_BINDS, SQL_PLAN_STATISTICS (similar to SQL_PLAN + row source statistics).

Return Values

This function returns a SQL tuning set object.

Usage Notes

For example, you can use this function to create a SQL tuning set containing the subset of SQL statements that regressed during a SQL Performance Analyzer (SPA) experiment. You can also specify other arbitrary filters.

189.4.14 SELECT_SQLSET Function

This is a table function that reads the contents of a SQL tuning set.

Syntax

```
DBMS_SQLSET.SELECT_SQLSET (
  sqlset_name         IN   VARCHAR2,
  basic_filter        IN   VARCHAR2 := NULL,
  object_filter       IN   VARCHAR2 := NULL,
  ranking_measure1    IN   VARCHAR2 := NULL,
  ranking_measure2    IN   VARCHAR2 := NULL,
  ranking_measure3    IN   VARCHAR2 := NULL,
  result_percentage   IN   NUMBER   := 1,
  result_limit        IN   NUMBER   := NULL)
  attribute_list      IN   VARCHAR2 := 'TYPICAL',
  plan_filter         IN   VARCHAR2 := NULL,
  sqlset_owner        IN   VARCHAR2 := NULL,
  recursive_sql       IN   VARCHAR2 := HAS_RECURSIVE_SQL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 189-16 SELECT_SQLSET Function Parameters

Parameter	Description
sqlset_name	Specifies the name of the SQL tuning set to query.

Table 189-16 (Cont.) SELECT_SQLSET Function Parameters

Parameter	Description
basic_filter	Specifies the SQL predicate to filter the SQL from the SQL tuning set defined on attributes of the SQLSET_ROW.
object_filter	Currently not supported.
ranking_measure(n)	Specifies an ORDER BY clause on the selected SQL.
result_percentage	Specifies a filter that picks the top *n*% according to the supplied ranking measure. Note that this parameter applies only if one ranking measure is supplied.
result_limit	The top limit SQL from the filtered source, ranked by the ranking measure.
attribute_list	Defines the SQL statement attributes to return in the result. The possible values are: • BASIC — Returns all attributes (such as execution statistics and binds) except the plans. The execution context is included in the result. • TYPICAL — Returns BASIC plus the SQL plan, but without row source statistics and without the object reference list. This is the default. • ALL — Returns all attributes. • Comma-separated list of attribute names. This value enables the function to return only a subset of SQL attributes: – EXECUTION_STATISTICS – SQL_BINDS – SQL_PLAN_STATISTICS (similar to SQL_PLAN plus row source statistics)
plan_filter	Specifies the plan filter. This parameter enables you to select a single plan when a statement has multiple plans. Possible values are: • LAST_GENERATED — Returns the plan with the most recent timestamp. • FIRST_GENERATED — Returns the plan with the least recent timestamp. • LAST_LOADED — Returns the plan with the most recent FIRST_LOAD_TIME statistical information. • FIRST_LOADED — Returns the plan with the least recent FIRST_LOAD_TIME statistical information. • MAX_ELAPSED TIME — Returns the plan with the maximum elapsed time. • MAX_BUFFER_GETS — Returns the plan with the maximum buffer gets. • MAX_DISK_READS — Returns the plan with the maximum disk reads. • MAX_DIRECT_WRITES — Returns the plan with the maximum direct writes. • MAX_OPTIMIZER_COST — Returns the plan with the maximum optimizer cost value.
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for the current schema owner.
recursive_sql	Specifies that the filter must include recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL, which is the default) or exclude it (NO_RECURSIVE_SQL).

Return Values

This function returns one `SQLSET_ROW` per `SQL_ID` or `PLAN_HASH_VALUE` pair found in each data source.

Usage Notes

Filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.

Examples

```
-- select from a sql tuning set
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE (P)
    FROM table(DBMS_SQLSET.SELECT_SQLSET('my_workload')) P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

189.4.15 SELECT_WORKLOAD_REPOSITORY Function

This function collects SQL statements from the workload repository.

The overloaded forms enable you to collect SQL statements from the following sources:

- Snapshots between `begin_snap` and `end_snap`
- A workload repository baseline

Syntax

```
DBMS_SQLSET.SELECT_WORKLOAD_REPOSITORY (
  begin_snap        IN NUMBER,
  end_snap          IN NUMBER,
  basic_filter      IN VARCHAR2 := NULL,
  object_filter     IN VARCHAR2 := NULL,
  ranking_measure1  IN VARCHAR2 := NULL,
  ranking_measure2  IN VARCHAR2 := NULL,
  ranking_measure3  IN VARCHAR2 := NULL,
  result_percentage IN NUMBER   := 1,
  result_limit      IN NUMBER   := NULL,
  attribute_list    IN VARCHAR2 := 'TYPICAL',
  recursive_sql     IN VARCHAR2 := HAS_RECURSIVE_SQL,
  dbid              IN NUMBER   := NULL)
 RETURN sys.sqlset PIPELINED;

DBMS_SQLSET.SELECT_WORKLOAD REPOSITORY (
  baseline_name     IN VARCHAR2,
```

```
  basic_filter      IN VARCHAR2 := NULL,
  object_filter     IN VARCHAR2 := NULL,
  ranking_measure1  IN VARCHAR2 := NULL,
  ranking_measure2  IN VARCHAR2 := NULL,
  ranking_measure3  IN VARCHAR2 := NULL,
  result_percentage IN NUMBER   := 1,
  result_limit      IN NUMBER   := NULL,
  attribute_list    IN VARCHAR2 := 'TYPICAL',
  recursive_sql     IN VARCHAR2 := HAS_RECURSIVE_SQL,
  dbid              IN NUMBER   := NULL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 189-17 SELECT_WORKLOAD_REPOSITORY Function Parameters

Parameter	Description
begin_snap	Defines the beginning AWR snapshot (non-inclusive).
end_snap	Defines the ending AWR snapshot (inclusive).
baseline_name	Specifies the name of the AWR baseline period.
basic_filter	Specifies the SQL predicate to filter the SQL from the workload repository. The filter is defined on attributes of the SQLSET_ROW. If basic_filter is not set by the caller, then the subprogram captures only statements of type CREATE TABLE, INSERT, SELECT, UPDATE, DELETE, and MERGE.
object_filter	Currently not supported.
ranking_measure(n)	Defines an ORDER BY clause on the selected SQL.
result_percentage	Specifies a filter that picks the top *n*% according to the supplied ranking measure. Note that this percentage applies only if one ranking measure is given.
result_limit	Specifies the top limit SQL from the source according to the supplied ranking measure.
attribute_list	Specifies the SQL statement attributes to return in the result. The possible values are: • TYPICAL — Returns BASIC plus SQL plan (without row source statistics) and without object reference list. This is the default. • BASIC — Returns all attributes (such as execution statistics and binds) are returned except the plans. The execution context is always part of the result. • ALL — Returns all attributes • Comma-separated list of attribute names this allows to return only a subset of SQL attributes: EXECUTION_STATISTICS, SQL_BINDS, SQL_PLAN_STATISTICS (similar to SQL_PLAN plus row source statistics).
recursive_sql	Specifies the filter that includes recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL) or excludes it (NO_RECURSIVE_SQL).
dbid	Specifies the DBID for imported or PDB-level AWR data. If NULL, then the function uses the current database DBID.

Return Values

This function returns one SQLSET_ROW per SQL_ID or PLAN_HASH_VALUE pair found in each data source.

Usage Notes

Filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.

Examples

```
-- select statements from snapshots 1-2
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE (P)
    FROM table(DBMS_SQLSET.SELECT_WORKLOAD_REPOSITORY(1,2)) P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

189.4.16 UNPACK_STGTAB Procedure

This procedure copies one or more SQL tuning sets from their location in the staging table into the SQL tuning sets schema, making them proper SQL tuning sets.

Syntax

```
DBMS_SQLSET.UNPACK_STGTAB (
   sqlset_name          IN VARCHAR2 := '%',
   sqlset_owner         IN VARCHAR2 := NULL,
   replace              IN BOOLEAN,
   staging_table_name   IN VARCHAR2,
   staging_schema_owner IN VARCHAR2 := NULL);
```

Parameters

The parameters are identical for DBMS_SQLTUNE.UNPACK_STGTAB_SQLSET and DBMS_SQLSET.UNPACK_STGTAB.

Table 189-18 UNPACK_STGTAB_SQLSET and UNPACK_STGTAB Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the tuning set to unpack (not null). Wildcard characters (%) are supported to unpack multiple tuning sets in a single call. For example, specify % to unpack all tuning sets from the staging table.

Table 189-18 (Cont.) UNPACK_STGTAB_SQLSET and UNPACK_STGTAB Procedure Parameters

Parameter	Description
sqlset_owner	Specifies the name of tuning set owner, or NULL for the current schema owner. Wildcard characters (%) are supported.
replace	Specifies whether to replace an existing SQL tuning set. If FALSE, then this procedure raises errors when you try to create a tuning set that already exists.
staging_table_name	Specifies the name of the staging table, moved after a call to the DBMS_SQLTUNE.PACK_STGTAB_SQLSET or DBMS_SQLSET.PACK_STGTAB procedure (case-sensitive).
staging_schema_owner	Specifies the name of staging table owner, or NULL for the current schema owner (case-sensitive).

Examples

```
 -- unpack all STS in the staging table
EXEC DBMS_SQLSET.UNPACK_STGTAB(sqlset_name         => '%', -
                               sqlset_owner        => '%', -
                               replace             => FALSE, -
                               staging_table_name  => 'STGTAB_SQLSET');

-- errors can arise during STS unpack when a STS in the staging table has the
-- same name/owner as STS on the system.  In this case, users should call
-- remap_stgtab_sqlset to patch the staging table and with which to call
unpack
-- Replace set to TRUE.
EXEC DBMS_SQLSET.UNPACK_STGTAB(sqlset_name         => '%', -
                               sqlset_owner        => '%', -
                               replace             => TRUE, -
                               staging_table_name  => 'STGTAB_SQLSET');
```

189.4.17 UPDATE_SQLSET Procedures

This overloaded procedure updates selected fields for SQL statements in a SQL tuning set.

Syntax

```
DBMS_SQLSET.UPDATE_SQLSET (
   sqlset_name      IN  VARCHAR2,
   sql_id           IN  VARCHAR2,
   plan_hash_value  IN  NUMBER := NULL,
   attribute_name   IN  VARCHAR2,
   attribute_value  IN  VARCHAR2 := NULL,
   sqlset_owner     IN  VARCHAR2 := NULL);

DBMS_SQLSET.UPDATE_SQLSET (
   sqlset_name      IN  VARCHAR2,
   sql_id           IN  VARCHAR2,
   plan_hash_value  IN  NUMBER := NULL,
   attribute_name   IN  VARCHAR2,
```

```
   attribute_value  IN  NUMBER := NULL,
   sqlset_owner     IN  VARCHAR2 := NULL);
```

Parameters

Table 189-19 UPDATE_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the SQL tuning set.
sql_id	Specifies the identifier of the SQL statement to be updated.
plan_hash value	Specifies the hash value of the execution plan for a SQL statement. Use this parameter when you want to update the attribute for a specific plan for a statement, but not all plans for the statement.
attribute_name	Specifies the name of the attribute to be modified. You can update the text field for MODULE, ACTION, PARSING_SCHEMA_NAME, and OTHER. The only numerical field that you can update is PRIORITY. If a statement has multiple plans, then the procedure changes the attribute value for all plans.
attribute_value	Specifies the new value of the attribute.

190
DBMS_SQLTUNE

The `DBMS_SQLTUNE` package is the interface for tuning SQL on demand. The related package `DBMS_AUTO_SQLTUNE` package provides the interface for SQL Tuning Advisor run as an automated task.

The chapter contains the following topics:

- DBMS_SQLTUNE Overview
- DBMS_SQLTUNE Security Model
- DBMS_SQLTUNE Data Structures
- DBMS_SQLTUNE Subprogram Groups
- Summary of DBMS_SQLTUNE Subprograms

> **See Also:**
>
> "DBMS_AUTO_SQLTUNE Overview"

190.1 DBMS_SQLTUNE Overview

The `DBMS_SQLTUNE` package provides a number of interrelated areas of functionality.

This section contains the following topics:

- DBMS_SQLTUNE SQL Tuning Advisor Subprograms
- DBMS_SQLTUNE SQL Profile Subprograms
- DBMS_SQLTUNE SQL Tuning Set Subprograms

SQL Tuning Advisor

SQL Tuning Advisor is one of a suite of advisors, a set of expert systems that identifies and helps resolve database performance problems. Specifically, SQL Tuning Advisor automates tuning of problematic SQL statements. It takes one or more SQL statements as input and gives precise advice on how to tune the statements. The advisor provides the advice in the form of SQL actions for tuning the SQL along with their expected performance benefit.

The group of DBMS_SQLTUNE SQL Tuning Advisor Subprograms provide a task-oriented interface that enables you to access the advisor. You can call the following subprograms in the order given to use some of SQL Tuning Advisor's features:

1. CREATE_TUNING_TASK Functions creates a tuning task for tuning one or more SQL statements.
2. The EXECUTE_TUNING_TASK Function and Procedure executes a previously created tuning task.
3. The REPORT_TUNING_TASK Function displays the results of a tuning task.

4. You use the SCRIPT_TUNING_TASK Function to create a SQL*Plus script which can then be executed to implement a set of Advisor recommendations

SQL Profile Subprograms

SQL Tuning Advisor may recommend the creation of a SQL profile to improve the performance of a statement. SQL profiles consist of auxiliary statistics specific to the statement. The query optimizer makes estimates about cardinality, selectivity, and cost that can sometimes be off by a significant amount, resulting in poor execution plans. The SQL profile addresses this problem by collecting additional information using sampling and partial execution techniques to adjust these estimates.

The group of DBMS_SQLTUNE SQL Profile Subprograms provides a mechanism for delivering statistics to the optimizer that targets one particular SQL statement, and helps the optimizer make good decisions for that statement by giving it the most accurate statistical information possible. For example:

- You can use the ACCEPT_SQL_PROFILE Procedure and Function to accept a SQL profile recommended by SQL Tuning Advisor.
- You can alter the `STATUS`, `NAME`, `DESCRIPTION`, and `CATEGORY` attributes of an existing SQL profile with the ALTER_SQL_PROFILE Procedure.
- You can drop a SQL profile with the DROP_SQL_PROFILE Procedure.

SQL Tuning Sets

SQL tuning sets store SQL statements along with the following information:

- The execution context, such as the parsing schema name and bind values
- Execution statistics such as average elapsed time and execution count
- Execution plans, which are the sequence of operations that the database performs to run SQL statements
- Row source statistics such as the number of rows processed for each operation executed within the plan

You can create SQL tuning sets by filtering or ranking SQL statements from several sources:

- The shared SQL area using the SELECT_CURSOR_CACHE Function
- Top SQL statements from the Automatic Workload Repository using the SELECT_WORKLOAD_REPOSITORY Function
- Other SQL tuning sets using the SELECT_SQLSET Function
- SQL Performance Analyzer task comparison results using the SELECT_SQLPA_TASK Function
- SQL Trace files using the SELECT_SQL_TRACE Function
- A user-defined workload

The complete group of DBMS_SQLTUNE SQL Tuning Set Subprograms facilitates this functionality. As examples:

- The CREATE_SQLSET Procedure and Function creates a SQL tuning set object in the database.
- The LOAD_SQLSET Procedure populates the SQL tuning set with a set of selected SQL.

- The CAPTURE_CURSOR_CACHE_SQLSET Procedure collects SQL statements from the shared SQL area over a specified time interval, attempting to build a realistic picture of database workload.

> **Note:**
>
> When manipulating SQL tuning sets, you can use DBMS_SQLSET as an alternative to `DBMS_SQLTUNE`.

Import and Export of SQL Tuning Sets and SQL Profiles

Use `DBMS_SQLTUNE` subprograms to move SQL profiles and SQL tuning sets from one system to another using a common programmatic model. In both cases, you create a staging table on the source database and populate this staging table with the relevant data. You then move that staging table to the destination system following the method of your choice (such as Oracle Data Pump, or a database link), where it is used to reconstitute the objects in their original form. The following steps are implemented by means of subprograms included in this package:

1. To create the staging table on the source system, call the CREATE_STGTAB_SQLPROF Procedure or the CREATE_STGTAB_SQLSET Procedure.
2. To populate the staging table with information from the source system, call the PACK_STGTAB_SQLPROF Procedure or PACK_STGTAB_SQLSET Procedure.
3. Move the staging table to the destination system.
4. To re-create the object on the new system, call the UNPACK_STGTAB_SQLPROF Procedure or the UNPACK_STGTAB_SQLSET Procedure.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* for more information about programmatic flow

Automatic Tuning Task Functions

The automated system task `SYS_AUTO_SQL_TUNING_TASK` is created by the database as part of the catalog scripts. This task automatically chooses a set of high-load SQL from AWR and runs SQL Tuning Advisor on this SQL. The automated task performs the same comprehensive analysis as any other SQL Tuning task.

You can obtain a report on the activity of the Automatic SQL Tuning task through the `DBMS_AUTO_SQLTUNE.REPORT_AUTO_TUNING_TASK` API.

> **See Also:**
>
> DBMS_AUTO_SQLTUNE for the list of subprograms that you can use to manage the automated SQL tuning task.

Real-Time SQL Monitoring

Real-time SQL Monitoring enables DBAs or performance analysts to monitor the execution of long-running SQL statements while they are executing. Both cursor statistics (such as CPU times and IO times) and execution plan statistics (such as number of output rows, memory and temp space used) are updated in almost real time during statement execution. The `V$SQL_MONITOR` and `V$SQL_PLAN_MONITOR` views expose these statistics. In addition, `DBMS_SQLTUNE` provides the `REPORT_SQL_MONITOR` and `REPORT_SQL_MONITOR_LIST` functions to report monitoring information.

> **Note:**
>
> `DBMS_SQL_MONITOR` also contains the `REPORT_SQL_MONITOR` and `REPORT_SQL_MONITOR_LIST` functions.

Tuning a Standby Database Workload

In some cases, a standby database can assume a reporting role in addition to its data protection role. The standby database can have its own workload of queries, some of which may require tuning. You can issue SQL Tuning Advisor statements on a standby database, which is read-only. A standby-to-primary database link enables `DBMS_SQLTUNE` to write data to and read data from the primary database. The procedures that are eligible for tuning standby workloads include the `database_link_to` parameter.

190.2 DBMS_SQLTUNE Security Model

This package is available to `PUBLIC` and performs its own security checking.

Note the following:

- Because SQL Tuning Advisor relies on the advisor framework, all tuning task interfaces (`*_TUNING_TASK`) require the `ADVISOR` privilege.
- SQL tuning set subprograms (`*_SQLSET`) require either of the following privileges:
  - `ADMINISTER SQL TUNING SET`

    You can only create and modify a SQL tuning set that you own.
  - `ADMINISTER ANY SQL TUNING SET`

    You can operate on all SQL tuning sets, even those owned by other users.
- In earlier releases, three different privileges were needed to invoke subprograms involving SQL profiles:
  - `CREATE ANY SQL PROFILE`
  - `ALTER ANY SQL PROFILE`
  - `DROP ANY SQL PROFILE`

  The preceding privileges have been deprecated in favor of `ADMINISTER SQL MANAGEMENT OBJECT` .

190.3 DBMS_SQLTUNE Data Structures

The `SELECT_*` subprograms in the `DBMS_SQLTUNE` package return objects of the `SQLSET_ROW` type.

Object Types

- SQLSET_ROW Object Type

190.3.1 SQLSET_ROW Object Type

The `SQLSET_ROW` object models the content of a SQL tuning set for the user.

Logically, a SQL tuning set is a collection of `SQLSET_ROW` objects. Each `SQLSET_ROW` contains a single SQL statement along with its execution context, statistics, binds, and plan. The `SELECT_*` subprograms each model a data source as a collection of `SQLSET_ROW` objects, with each object uniquely identified by (`sql_id`, `plan_hash_value`). Similarly, the `LOAD_SQLSET` procedure takes as input a cursor whose row type is `SQLSET_ROW`, treating each `SQLSET_ROW` in isolation according to the policies requested by the user.

Several subprograms package accept basic filters on the content of a SQL tuning set or data source. These filters are expressed in terms of the attributes within the `SQLSET_ROW` as defined.

Syntax

```
CREATE TYPE sqlset_row AS object (
  sql_id                    VARCHAR(13),
  force_matching_signature  NUMBER,
  sql_text                  CLOB,
  object_list               sql_objects,
  bind_data                 RAW(2000),
  parsing_schema_name       VARCHAR2(30),
  module                    VARCHAR2(48),
  action                    VARCHAR2(32),
  elapsed_time              NUMBER,
  cpu_time                  NUMBER,
  buffer_gets               NUMBER,
  disk_reads                NUMBER,
  direct_writes             NUMBER,
  rows_processed            NUMBER,
  fetches                   NUMBER,
  executions                NUMBER,
  end_of_fetch_count        NUMBER,
  optimizer_cost            NUMBER,
  optimizer_env             RAW(2000),
  priority                  NUMBER,
  command_type              NUMBER,
  first_load_time           VARCHAR2(19),
  stat_period               NUMBER,
  active_stat_period        NUMBER,
  other                     CLOB,
  plan_hash_value           NUMBER,
  sql_plan                  sql_plan_table_type,
  bind_list                 sql_binds,
```

```
  con_dbid                 NUMBER,
  last_exec_start_time     VARCHAR2(19))
```

Attributes

Table 190-1 SQLSET_ROW Attributes

Attribute	Description
sql_id	Unique SQL ID.
forcing_matching_signature	Signature with literals, case, and whitespace removed.
sql_text	Full text for the SQL statement.
object_list	Currently not implemented.
bind_data	Bind data as captured for this SQL. Note that you cannot stipulate an argument for this parameter and also for bind_list - they are mutually exclusive.
parsing_schema_name	Schema where the SQL is parsed.
module	Last application module for the SQL.
action	Last application action for the SQL.
elapsed_time	Sum total elapsed time for this SQL statement.
cpu_time	Sum total CPU time for this SQL statement.
buffer_gets	Sum total number of buffer gets.
disk_reads	Sum total number of disk reads.
direct_writes	Sum total number of direct path writes.
rows_processed	Sum total number of rows processed by this SQL.
fetches	Sum total number of fetches.
executions	Total executions of this SQL statement.
end_of_fetch_count	Number of times the SQL statement was fully executed with all of its rows fetched.
optimizer_cost	Optimizer cost for this SQL.
optimizer_env	Optimizer environment for this SQL statement.
priority	User-defined priority (1,2,3).
command_type	Statement type, such as INSERT or SELECT.
first_load_time	Load time of the parent cursor.
stat_period	Period of time (seconds) when the statistics of this SQL statement were collected.
active_stat_period	Effective period of time (in seconds) during which the SQL statement was active.
other	Other column for user-defined attributes.
plan_hash_value	Plan hash value of the plan.
sql_plan	Execution plan for the SQL statement.
bind_list	List of user-specified binds for the SQL statement. This is used for user-specified workloads. Note that you cannot stipulate an argument for this parameter and also for bind_data: they are mutually exclusive.

Table 190-1 (Cont.) SQLSET_ROW Attributes

Attribute	Description
con_dbid	DBID of the PDB or CDB root.
last_exec_start_time	Most recent execution start time of this SQL statement.

190.4 DBMS_SQLTUNE Subprogram Groups

DBMS_SQLTUNE subprograms are grouped by function.

- DBMS_SQLTUNE SQL Tuning Advisor Subprograms
- DBMS_SQLTUNE SQL Profile Subprograms
- DBMS_SQLTUNE SQL Tuning Set Subprograms
- DBMS_SQLTUNE Real-Time SQL Monitoring Subprograms
- DBMS_SQLTUNE SQL Performance Reporting Subprograms

190.4.1 DBMS_SQLTUNE SQL Tuning Advisor Subprograms

This subprogram group provides an interface to manage SQL tuning tasks.

Table 190-2 SQL Tuning Task Subprograms

Subprogram	Description
"CANCEL_TUNING_TASK Procedure"	Cancels the currently executing tuning task
"CREATE_SQL_PLAN_BASELINE Procedure"	Creates a SQL plan baseline for an existing plan
"CREATE_TUNING_TASK Functions"	Creates a tuning of a single statement or SQL tuning set for either SQL Tuning Advisor
"DROP_TUNING_TASK Procedure"	Drops a SQL tuning task
"EXECUTE_TUNING_TASK Function and Procedure"	Executes a previously created tuning task
"IMPLEMENT_TUNING_TASK Procedure"	Implements a set of SQL profile recommendations made by SQL Tuning Advisor
"INTERRUPT_TUNING_TASK Procedure"	Interrupts the currently executing tuning task
"REPORT_AUTO_TUNING_TASK Function"	Displays a report from the automatic tuning task, reporting on a range of executions
"REPORT_TUNING_TASK Function"	Displays the results of a tuning task
"RESET_TUNING_TASK Procedure"	Resets the currently executing tuning task to its initial state
"RESUME_TUNING_TASK Procedure"	Resumes a previously interrupted task that was created to process a SQL tuning set
"SCHEDULE_TUNING_TASK Function"	Creates a tuning task and schedules its execution as a scheduler job
"SCRIPT_TUNING_TASK Function"	Creates a SQL*Plus script which can then be executed to implement a set of SQL Tuning Advisor recommendations

Table 190-2 (Cont.) SQL Tuning Task Subprograms

Subprogram	Description
"SET_TUNING_TASK_PARAMETER Procedures"	Updates the value of a SQL tuning parameter of type `VARCHAR2` or `NUMBER`

"Summary of DBMS_SQLTUNE Subprograms" contains a complete listing of all subprograms in the package.

190.4.2 DBMS_SQLTUNE SQL Profile Subprograms

This subprogram group provides an interface to manage SQL profiles.

Table 190-3 SQL Profile Subprograms

Subprogram	Description
ACCEPT_ALL_SQL_PROFILES Procedure	Accepts all SQL profiles recommended by a specific execution of a tuning task
ACCEPT_SQL_PROFILE Procedure and Function	Creates a SQL profile for the specified tuning task
ALTER_SQL_PROFILE Procedure	Alters specific attributes of an existing SQL profile object
CREATE_STGTAB_SQLPROF Procedure	Creates the staging table used for copying SQL profiles from one system to another
DROP_SQL_PROFILE Procedure	Drops the named SQL profile from the database
PACK_STGTAB_SQLPROF Procedure	Moves profile data out of the `SYS` schema into the staging table
REMAP_STGTAB_SQLPROF Procedure	Changes the profile data values kept in the staging table prior to performing an unpack operation
SQLTEXT_TO_SIGNATURE Function	Returns a SQL text's signature
UNPACK_STGTAB_SQLPROF Procedure	Uses the profile data stored in the staging table to create profiles on this system

"Summary of DBMS_SQLTUNE Subprograms" contains a complete listing of all subprograms in the package.

190.4.3 DBMS_SQLTUNE SQL Tuning Set Subprograms

This subprogram group provides an interface to manage SQL tuning sets.

Table 190-4 SQL Tuning Set Subprograms

Subprogram	Description
ADD_SQLSET_REFERENCE Function	Adds a new reference to an existing SQL tuning set to indicate its use by a client
CAPTURE_CURSOR_CACHE_SQLSET Procedure	Over a specified time interval incrementally captures a workload from the shared SQL area into a SQL tuning set
CREATE_SQLSET Procedure and Function	Creates a SQL tuning set object in the database

Table 190-4 (Cont.) SQL Tuning Set Subprograms

Subprogram	Description
CREATE_STGTAB_SQLSET Procedure	Creates a staging table through which SQL Tuning Sets are imported and exported
DELETE_SQLSET Procedure	Deletes a set of SQL statements from a SQL tuning set
DROP_SQLSET Procedure	Drops a SQL tuning set if it is not active
LOAD_SQLSET Procedure	Populates the SQL tuning set with a set of selected SQL
PACK_STGTAB_SQLSET Procedure	Copies tuning sets out of the `SYS` schema into the staging table
REMAP_STGTAB_SQLSET Procedure	Changes the tuning set names and owners in the staging table so that they can be unpacked with different values than they had on the host system
REMOVE_SQLSET_REFERENCE Procedure	Deactivates a SQL tuning set to indicate it is no longer used by the client
SELECT_CURSOR_CACHE Function	Collects SQL statements from the shared SQL area
SELECT_SQL_TRACE Function	Reads the content of one or more trace files and returns the SQL statements it finds in the format of `sqlset_row`
SELECT_SQLPA_TASK Function	Collects SQL statements from a SQL performance analyzer comparison task
SELECT_SQLSET Function	Collects SQL statements from an existing SQL tuning set
SELECT_WORKLOAD_REPOSITORY Function	Collects SQL statements from the workload repository
UNPACK_STGTAB_SQLSET Procedure	Copies one or more SQL tuning sets from the staging table
UPDATE_SQLSET Procedures	Updates whether selected string fields for a SQL statement in a SQL tuning set or the set numerical attributes of a SQL in a SQL tuning set

The Summary of DBMS_SQLTUNE Subprograms contains a complete listing of all subprograms in the package.

190.4.4 DBMS_SQLTUNE Real-Time SQL Monitoring Subprograms

This subprogram group provides function to report on monitoring data collected in `V$SQL_MONITOR` and `V$SQL_PLAN_MONITOR`.

Table 190-5 Real-Time SQL Monitoring Subprograms

Subprogram	Description
REPORT_SQL_MONITOR Function	Reports on Real-Time SQL Monitoring
REPORT_SQL_MONITOR_LIST Function	Builds a report for all or a subset of statements monitored by Oracle Database
REPORT_SQL_MONITOR_LIST_XML Function	Builds an XML report for all or a subset of statements monitored by Oracle Database

190.4.5 DBMS_SQLTUNE SQL Performance Reporting Subprograms

This subprogram group provides detailed reports on SQL performance using statistics from the shared SQL area and automatic workload repository (AWR).

Table 190-6 SQL Performance Reporting Subprograms

Subprogram	Description
REPORT_SQL_DETAIL Function	This function reports on a specific SQL ID.
REPORT_SQL_MONITOR Function	This function builds a report (text, simple HTML, active HTML, XML) for the monitoring information collected on behalf of the targeted statement execution.
REPORT_SQL_MONITOR_LIST Function	This function builds a report for all or a sub-set of statements monitored by Oracle. For each statement, the subprogram gives key information and associated global statistics.
REPORT_TUNING_TASK Function	This function displays the results of a tuning task.
REPORT_TUNING_TASK_XML Function	This function displays an XML report of a tuning task.

190.5 Summary of DBMS_SQLTUNE Subprograms

This table lists the `DBMS_SQLTUNE` subprograms and briefly describes them.

Table 190-7 DBMS_SQLTUNE Package Subprograms

Subprogram	Description	Group
ACCEPT_ALL_SQL_PROFILES Procedure	Accepts all SQL profiles recommended by a particular execution of a particular tuning task	DBMS_SQLTUNE SQL Profile Subprograms
ACCEPT_SQL_PROFILE Procedure and Function	Creates a SQL profile for the specified tuning task	DBMS_SQLTUNE SQL Profile Subprograms
ADD_SQLSET_REFERENCE Function	Adds a new reference to an existing SQL tuning set to indicate its use by a client	DBMS_SQLTUNE SQL Tuning Set Subprograms
ALTER_SQL_PROFILE Procedure	Alters specific attributes of an existing SQL profile object	DBMS_SQLTUNE SQL Profile Subprograms
CANCEL_TUNING_TASK Procedure	Cancels the currently executing tuning task	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
CAPTURE_CURSOR_CACHE_SQLSET Procedure	Over a specified time interval incrementally captures a workload from the shared SQL area into a SQL tuning set	DBMS_SQLTUNE SQL Tuning Set Subprograms
CREATE_SQL_PLAN_BASELINE Procedure	Creates a SQL plan baseline for an existing plan	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
CREATE_SQLSET Procedure and Function	Creates a SQL tuning set object in the database	DBMS_SQLTUNE SQL Tuning Set Subprograms

Table 190-7 (Cont.) DBMS_SQLTUNE Package Subprograms

Subprogram	Description	Group
CREATE_STGTAB_SQLPROF Procedure	Creates the staging table used for copying SQL profiles from one system to another	DBMS_SQLTUNE SQL Profile Subprograms
CREATE_STGTAB_SQLSET Procedure	Creates a staging table through which SQL tuning sets are imported and exported	DBMS_SQLTUNE SQL Tuning Set Subprograms
CREATE_TUNING_TASK Functions	Creates a tuning of a single statement or SQL tuning set for either SQL Tuning Advisor	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
DELETE_SQLSET Procedure	Deletes a set of SQL statements from a SQL tuning set	DBMS_SQLTUNE SQL Tuning Set Subprograms
DROP_SQL_PROFILE Procedure	Drops the named SQL profile from the database	DBMS_SQLTUNE SQL Profile Subprograms
DROP_SQLSET Procedure	Drops a SQL tuning set if it is not active	DBMS_SQLTUNE SQL Tuning Set Subprograms
DROP_TUNING_TASK Procedure	Drops a SQL tuning task	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
EXECUTE_TUNING_TASK Function and Procedure	Executes a previously created tuning task	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
IMPLEMENT_TUNING_TASK Procedure	implements a set of SQL profile recommendations made by SQL Tuning Advisor	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
INTERRUPT_TUNING_TASK Procedure	Interrupts the currently executing tuning task	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
LOAD_SQLSET Procedure	Populates the SQL tuning set with a set of selected SQL	DBMS_SQLTUNE SQL Tuning Set Subprograms
PACK_STGTAB_SQLPROF Procedure	Moves profile data out of the `SYS` schema into the staging table	DBMS_SQLTUNE SQL Profile Subprograms
PACK_STGTAB_SQLSET Procedure	Moves tuning sets out of the `SYS` schema into the staging table	DBMS_SQLTUNE SQL Tuning Set Subprograms
REMAP_STGTAB_SQLPROF Procedure	Changes the profile data values kept in the staging table prior to performing an unpack operation	DBMS_SQLTUNE SQL Profile Subprograms
REMAP_STGTAB_SQLSET Procedure	Changes the tuning set names and owners in the staging table so that they can be unpacked with different values than they had on the host system	DBMS_SQLTUNE SQL Tuning Set Subprograms
REMOVE_SQLSET_REFERENCE Procedure	Deactivates a SQL tuning set to indicate it is no longer used by the client	DBMS_SQLTUNE SQL Tuning Set Subprograms
REPORT_AUTO_TUNING_TASK Function	Displays a report from the automatic tuning task, reporting on a range of subtasks	DBMS_SQLTUNE SQL Tuning Set Subprograms
REPORT_SQL_DETAIL Function	Reports on a specific SQL ID	DBMS_SQLTUNE SQL Performance Reporting Subprograms

Table 190-7 (Cont.) DBMS_SQLTUNE Package Subprograms

Subprogram	Description	Group
REPORT_SQL_MONITOR Function	Builds a report (text, simple HTML, active HTML, XML) for the monitoring information collected on behalf of the targeted statement execution	DBMS_SQLTUNE Real-Time SQL Monitoring Subprograms
REPORT_SQL_MONITOR_LIST Function	Builds a report for all or a subset of statements monitored by Oracle Database. For each statement, the subprogram gives key information and associated global statistics	DBMS_SQLTUNE Real-Time SQL Monitoring Subprograms
REPORT_SQL_MONITOR_LIST_XML Function	Equivalent to the `REPORT_SQL_MONITOR_LIST` function, except that it returns `XMLType`	DBMS_SQLTUNE Real-Time SQL Monitoring Subprograms
REPORT_TUNING_TASK Function	Displays the results of a tuning task	DBMS_SQLTUNE SQL Performance Reporting Subprograms
REPORT_TUNING_TASK_XML Function	Displays an XML report of a tuning task	DBMS_SQLTUNE SQL Performance Reporting Subprograms
RESET_TUNING_TASK Procedure	Resets the currently executing tuning task to its initial state	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
RESUME_TUNING_TASK Procedure	Resumes a previously interrupted task that was created to process a SQL tuning set	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
SCHEDULE_TUNING_TASK Function	Creates a SQL tuning task and schedule its execution as a scheduler job	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
SCRIPT_TUNING_TASK Function	Creates a SQL*Plus script which can then be executed to implement a set of SQL Tuning Advisor recommendations	DBMS_SQLTUNE SQL Tuning Advisor Subprograms
SELECT_CURSOR_CACHE Function	Collects SQL statements from the shared SQL area	DBMS_SQLTUNE SQL Tuning Set Subprograms
SELECT_SQL_TRACE Function	Reads the content of one or more trace files and returns the SQL statements it finds in the format of `sqlset_row`	DBMS_SQLTUNE SQL Tuning Set Subprograms
SELECT_SQLPA_TASK Function	Collects SQL statements from a SQL Performance Analyzer comparison task	DBMS_SQLTUNE SQL Tuning Set Subprograms
SELECT_SQLSET Function	Collects SQL statements from an existing SQL tuning set	DBMS_SQLTUNE SQL Tuning Set Subprograms
SELECT_WORKLOAD_REPOSITORY Function	Collects SQL statements from the workload repository	DBMS_SQLTUNE SQL Tuning Set Subprograms
SET_TUNING_TASK_PARAMETER Procedures	Updates the value of a SQL tuning parameter of type `VARCHAR2` or `NUMBER`	DBMS_SQLTUNE SQL Tuning Advisor Subprograms

Table 190-7 (Cont.) DBMS_SQLTUNE Package Subprograms

Subprogram	Description	Group
SQLTEXT_TO_SIGNATURE Function	Returns a SQL text's signature	DBMS_SQLTUNE SQL Profile Subprograms
UNPACK_STGTAB_SQLPROF Procedure	Uses the profile data stored in the staging table to create profiles on this system	DBMS_SQLTUNE SQL Profile Subprograms
UNPACK_STGTAB_SQLSET Procedure	Moves one or more SQL tuning sets from the staging table	DBMS_SQLTUNE SQL Tuning Set Subprograms
UPDATE_SQLSET Procedures	Updates selected fields for a SQL statement in a SQL tuning set	DBMS_SQLTUNE SQL Tuning Set Subprograms

190.5.1 ACCEPT_ALL_SQL_PROFILES Procedure

This procedure accepts all SQL profiles recommended by a specific execution of a tuning task, and sets the attributes of the SQL profiles according to the parameter values passed by the user.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.ACCEPT_ALL_SQL_PROFILES (
   task_name         IN VARCHAR2,
   category          IN VARCHAR2 := NULL,
   replace           IN BOOLEAN  := FALSE,
   force_match       IN BOOLEAN  := FALSE,
   profile_type      IN VARCHAR2 := REGULAR_PROFILE,
   autotune_period   IN NUMBER   := NULL,
   execution_name    IN VARCHAR2 := NULL,
   task_owner        IN VARCHAR2 := NULL,
   description       IN VARCHAR2 := NULL,
   database_link_to  IN VARCHAR2 := NULL);
```

Parameters

Table 190-8 ACCEPT_ALL_SQL_PROFILES Procedure Parameters

Parameter	Description
task_name	The (mandatory) name of the SQL tuning task

Table 190-8 (Cont.) ACCEPT_ALL_SQL_PROFILES Procedure Parameters

Parameter	Description
category	This is the category name which must match the value of the SQLTUNE_CATEGORY parameter in a session for the session to use this SQL profile. It defaults to the value "DEFAULT". This is also the default of the SQLTUNE_CATEGORY parameter. The category must be a valid Oracle identifier. The category name specified is always converted to upper case. The combination of the normalized SQL text and category name creates a unique key for a SQL profile. An ACCEPT_SQL_PROFILE fails if this combination is duplicated.
replace	If the profile already exists, it is replaced if this argument is TRUE. It is an error to pass a name that is already being used for another signature/ category pair, even with replace set to TRUE.
force_match	If TRUE this causes SQL profiles to target all SQL statements which have the same text after normalizing all literal values into bind variables. (Note that if a combination of literal values and bind values is used in a SQL statement, no bind transformation occurs.) This is analogous to the matching algorithm used by the FORCE option of the cursor_sharing parameter. If FALSE, literals are not transformed. This is analogous to the matching algorithm used by the EXACT option of the cursor_sharing parameter.
profile_type	Options: • REGULAR_PROFILE - profile without a change to parallel execution (Default, equivalent to NULL). Note that if the SQL statement currently has a parallel execution plan, the regular profile will cause the optimizer to choose a different, but still parallel, execution plan. • PX_PROFILE - regular profile with a change to parallel execution
autotune_period	The time period for the automatic SQL tuning. This setting applies only to the automatic SQL Tuning Advisor task. Possible values are as follows: • null or negative value (default) - all or full. The result includes all task executions. • 0 - result of the current or most recent task execution. • 1 - result for the most recent 24-hour period. • 7 - result for the most recent 7-day period. The procedure interprets any other value as the time of the most recent task execution minus the value of this argument.
execution_name	Name of the task execution to use. If null, then the procedure generates the report for the most recent task execution.
task_owner	Owner of the tuning task. This is an optional parameter that must be specified to accept a SQL profile associated to a tuning task owned by another user. The current user is the default value.
description	A user specified string describing the purpose of the SQL profile. The description is truncated if longer than 256 characters. The maximum size is 500 characters.

Table 190-8 (Cont.) ACCEPT_ALL_SQL_PROFILES Procedure Parameters

Parameter	Description
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY *password* USING 'inst1';

Security Model

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required. The CREATE ANY SQL PROFILE privilege is deprecated.

190.5.2 ACCEPT_SQL_PROFILE Procedure and Function

This subprogram creates a SQL profile recommended by SQL Tuning Advisor.

The SQL text is normalized for matching purposes although it is stored in the data dictionary in denormalized form for readability. SQL text is provided through a reference to the SQL Tuning task. If the referenced SQL statement does not exist, then the database reports an error.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.ACCEPT_SQL_PROFILE (
   task_name     IN  VARCHAR2,
   object_id     IN  NUMBER   := NULL,
   name          IN  VARCHAR2 := NULL,
   description   IN  VARCHAR2 := NULL,
   category      IN  VARCHAR2 := NULL);
   task_owner    IN VARCHAR2  := NULL,
   replace       IN BOOLEAN   := FALSE,
```

```
   force_match  IN BOOLEAN   := FALSE,
   profile_type IN VARCHAR2  := REGULAR_PROFILE);


DBMS_SQLTUNE.ACCEPT_SQL_PROFILE (
   task_name         IN  VARCHAR2,
   object_id         IN  NUMBER    := NULL,
   name              IN  VARCHAR2  := NULL,
   description       IN  VARCHAR2  := NULL,
   category          IN  VARCHAR2  := NULL;
   task_owner        IN  VARCHAR2  := NULL,
   replace           IN  BOOLEAN   := FALSE,
   force_match       IN  BOOLEAN   := FALSE,
   profile_type      IN  VARCHAR2  := REGULAR_PROFILE,
   database_link_to  IN  VARCHAR2  := NULL)
 RETURN VARCHAR2;
```

Parameters

Table 190-9 ACCEPT_SQL_PROFILE Procedure and Function Parameters

Parameter	Description
task_name	The (mandatory) name of the SQL tuning task
object_id	The identifier of the advisor framework object representing the SQL statement associated with the tuning task
name	The name of the SQL profile. It cannot contain double quotation marks. The name is case sensitive. If not specified, the system generates a unique name for the SQL profile.
description	A user specified string describing the purpose of the SQL profile. The description is truncated if longer than 256 characters. The maximum size is 500 characters.
category	The category name. This name must match the value of the SQLTUNE_CATEGORY parameter in a session for the session to use this SQL profile. It defaults to the value "DEFAULT". This is also the default of the SQLTUNE_CATEGORY parameter. The category must be a valid Oracle identifier. The category name specified is always converted to upper case. The combination of the normalized SQL text and category name creates a unique key for a SQL profile. An ACCEPT_SQL_PROFILE fails if this combination is duplicated.
task_owner	Owner of the tuning task. This is an optional parameter that has to be specified to accept a SQL profile associated to a tuning task owned by another user. The current user is the default value.
replace	If the profile already exists, it is replaced if this argument is TRUE. It is an error to pass a name that is already being used for another signature/ category pair, even with replace set to TRUE.
force_match	If TRUE this causes SQL profiles to target all SQL statements which have the same text after normalizing all literal values into bind variables. (Note that if a combination of literal values and bind values is used in a SQL statement, no bind transformation occurs.) This is analogous to the matching algorithm used by the FORCE option of the cursor_sharing parameter. If FALSE, literals are not transformed. This is analogous to the matching algorithm used by the EXACT option of the cursor_sharing parameter.

Table 190-9 (Cont.) ACCEPT_SQL_PROFILE Procedure and Function Parameters

Parameter	Description
profile_type	Options: • REGULAR_PROFILE - profile without a change to parallel execution (Default, equivalent to NULL). Note that if the SQL statement currently has a parallel execution plan, the regular profile will cause the optimizer to choose a different, but still parallel, execution plan. • PX_PROFILE - regular profile with a change to parallel execution
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY *password* USING 'inst1';

Return Values

The name of the SQL profile.

Usage Notes

The ADMINISTER SQL MANAGEMENT OBJECT privilege is required. The CREATE ANY SQL PROFILE privilege is deprecated.

Examples

You use both the procedure and the function versions of the subprogram in the same way except you must specify a return value to invoke the function. Here we give examples of the procedure only.

In this example, you tune a single SQL statement form the workload repository and you create the SQL profile recommended by SQL Tuning Advisor.

```
VARIABLE stmt_task VARCHAR2(64);
VARIABLE sts_task VARCHAR2(64);

-- create a tuning task tune the statement
EXEC :stmt_task := DBMS_SQLTUNE.CREATE_TUNING_TASK(
   begin_snap  => 1, -
   end_snap    => 2, -
   sql_id      => 'ay1m3ssvtrh24');

-- execute the resulting task
```

```
EXEC DBMS_SQLTUNE.EXECUTE_TUNING_TASK(:stmt_task);

EXEC DBMS_SQLTUNE.ACCEPT_SQL_PROFILE(:stmt_task);
```

Note that you do not have to specify the ID (that is, `object_id`) for the advisor framework object created by SQL Tuning Advisor to represent the tuned SQL statement.

You might also want to accept the recommended SQL profile in a different category, (for example, `TEST`), so that it is not used by default.

```
EXEC DBMS_SQLTUNE.ACCEPT_SQL_PROFILE (
   task_name  =>  :stmt_task, -
   category   =>  'TEST');
```

You can use command `ALTER SESSION SET SQLTUNE_CATEGORY = 'TEST'` to see how this profile behaves.

The following call creates a SQL profile that targets any SQL statement with the same `force_matching_signature` as the tuned statement.

```
EXEC DBMS_SQLTUNE.ACCEPT_SQL_PROFILE (task_name   =>  :stmt_task, -
                                      force_match =>  TRUE);
```

In the following example, you tune a SQL tuning set, and you create a SQL profile for only one of the SQL statements in the SQL tuning set. The SQL statement is represented by an advisor framework object with ID equal to `5`. You must pass an object ID to the `ACCEPT_SQL_PROFILE` procedure because there are potentially many SQL profiles for the tuning task. This object ID is given along with the report.

```
EXEC :sts_task := DBMS_SQLTUNE.CREATE_TUNING_TASK ( -
  sqlset_name   => 'my_workload',  -
  rank1         => 'ELAPSED_TIME', -
  time_limit    => 3600,           -
  description   => 'my workload ordered by elapsed time');

-- execute the resulting task
EXEC DBMS_SQLTUNE.EXECUTE_TUNING_TASK(:sts_task);

 -- create the profile for the sql statement corresponding to object_id = 5.
EXEC DBMS_SQLTUNE.ACCEPT_SQL_PROFILE (
   task_name  =>  :sts_task, -
   object_id  =>  5);
```

190.5.3 ADD_SQLSET_REFERENCE Function

This procedure adds a new reference to an existing SQL tuning set to indicate its use by a client.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.ADD_SQLSET_REFERENCE (
   sqlset_name  IN  VARCHAR2,
   description  IN  VARCHAR2 := NULL)
 RETURN NUMBER;
```

Parameters

The parameters are identical for `DBMS_SQLTUNE.ADD_SQLSET_REFERENCE` and `DBMS_SQLSET.ADD_REFERENCE`.

Table 190-10 ADD_SQLSET_REFERENCE and ADD_REFERENCE Function Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set.
`description`	Provides an optional description of the usage of SQL tuning set. The description is truncated if longer than 256 characters.
`sqlset_owner`	Specifies the owner of the SQL tuning set, or `NULL` for the current schema owner.

Return Values

The identifier of the added reference.

Examples

You can add reference to a SQL tuning set. This prevents the tuning set from being modified while it is being used. References are automatically added when you invoke SQL Tuning Advisor on the SQL tuning set, so you should use this function for custom purposes only. The function returns a reference ID that is used to remove it later. You use the `REMOVE_SQLSET_REFERENCE` procedure to delete references to a SQL tuning set.

```
VARIABLE rid NUMBER;
EXEC :rid := DBMS_SQLTUNE.ADD_SQLSET_REFERENCE( -
                                sqlset_name => 'my_workload', -
                                description => 'my sts reference');
```

You can use the `DBA_SQLSET_REFERENCES` view to find all references on a given SQL tuning set.

190.5.4 ALTER_SQL_PROFILE Procedure

This procedure alters specific attributes of an existing SQL profile object.

The following attributes can be altered (using these attribute names):

- `STATUS` can be set to `ENABLED` or `DISABLED`.
- `NAME` can be reset to a valid name which must be a valid Oracle identifier and must be unique.
- `DESCRIPTION` can be set to any string of size no more than 500 characters.
- `CATEGORY` can be reset to a valid category name which must be a valid Oracle identifier and must be unique when combined with normalized SQL text).

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.ALTER_SQL_PROFILE (
   name                 IN  VARCHAR2,
   attribute_name       IN  VARCHAR2,
   value                IN  VARCHAR2);
```

Parameters

Table 190-11 ALTER_SQL_PROFILE Procedure Parameters

Parameter	Description
`name`	The (mandatory) name of the existing SQL profile to alter
`attribute_name`	The (mandatory) attribute name to alter (case insensitive) using valid attribute names
`value`	The (mandatory) new value of the attribute using valid attribute values

Usage Notes

Requires the `ALTER ANY SQL PROFILE` privilege.

Examples

```
-- Disable a profile, so it is not be used by any sessions.
EXEC DBMS_SQLTUNE.ALTER_SQL_PROFILE ( name            =>  :pname,   -
                                      attribute_name  =>  'STATUS', -
                                      value           =>  'DISABLED');

-- Enable it back:
EXEC DBMS_SQLTUNE.ALTER_SQL_PROFILE ( name            =>   :pname,   -
                                      attribute_name  =>   'STATUS', -
```

```
                                        value           =>   'ENABLED');

-- Change the category of the profile so it is used only by sessions
-- with category set to TEST.
-- Use ALTER SESSION SET SQLTUNE_CATEGORY = 'TEST' to see how this profile
-- behaves.
EXEC DBMS_SQLTUNE.ALTER_SQL_PROFILE ( name            =>  :pname,   -
                                      attribute_name  =>  'CATEGORY', -
                                      value           =>  'TEST');

-- Change it back:
EXEC DBMS_SQLTUNE.ALTER_SQL_PROFILE ( name            =>  :pname,   -
                                      attribute_name  =>  'CATEGORY', -
                                      value           =>  'DEFAULT');
```

190.5.5 CANCEL_TUNING_TASK Procedure

This procedure cancels the currently executing tuning task. All intermediate result data is deleted.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.CANCEL_TUNING_TASK (
 task_name          IN VARCHAR2);
```

Parameters

Table 190-12 CANCEL_TUNING_TASK Procedure Parameters

Parameter	Description
task_name	Specifies the name of the task to cancel

Examples

You cancel a task when you need to stop it executing and do not require to view any already-completed results.

```
EXEC DBMS_SQLTUNE.CANCEL_TUNING_TASK(:my_task);
```

190.5.6 CAPTURE_CURSOR_CACHE_SQLSET Procedure

This procedure captures a workload from the shared SQL area into a SQL tuning set.

The procedure polls the cache multiple times over a time period, and updates the workload data stored there. It can execute over as long a period as required to capture an entire system workload.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.CAPTURE_CURSOR_CACHE_SQLSET (
    sqlset_name         IN VARCHAR2,
    time_limit          IN POSITIVE := 1800,
    repeat_interval     IN POSITIVE := 300,
    capture_option      IN VARCHAR2 := 'MERGE',
    capture_mode        IN NUMBER   := MODE_REPLACE_OLD_STATS,
    basic_filter        IN VARCHAR2 := NULL,
    sqlset_owner        IN VARCHAR2 := NULL,
    recursive_sql       IN VARCHAR2 := HAS_RECURSIVE_SQL);
```

Parameters

The parameters are the same for both `DBMS_SQLTUNE.CAPTURE_CURSOR_CACHE_SQLSET` and `DBMS_SQLSET.CAPTURE_CURSOR_CACHE`.

Table 190-13 CAPTURE_CURSOR_CACHE_SQLSET and CAPTURE_CURSOR_CACHE Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the SQL tuning set name
`time_limit`	Defines the total amount of time, in seconds, to execute.
`repeat_interval`	Defines the amount of time, in seconds, to pause between sampling.
`capture_option`	Specifies whether to insert new statements, update existing statements, or both. Values are `INSERT`, `UPDATE`, or `MERGE`. The values are the same as for `load_option` in `load_sqlset`.
`capture_mode`	Specifies the capture mode (`UPDATE` and `MERGE` capture options). Possible values: • `MODE_REPLACE_OLD_STATS` — Replaces statistics when the number of executions is greater than the number stored in the SQL tuning set • `MODE_ACCUMULATE_STATS` — Adds new values to current values for SQL that is already stored. Note that this mode detects if a statement has been aged out, so the final value for a statistics is the sum of the statistics of all cursors that statement existed under.

Table 190-13 (Cont.) CAPTURE_CURSOR_CACHE_SQLSET and CAPTURE_CURSOR_CACHE Procedure Parameters

Parameter	Description
basic_filter	Defines a filter to apply to the shared SQL area for each sample. If basic_filter is not set by the caller, then the subprogram captures only statements of type CREATE TABLE, INSERT, SELECT, UPDATE, DELETE, and MERGE.
sqlset_owner	Specifies the owner of the SQL tuning set or NULL for current schema owner
recursive_sql	Defines a filter that includes recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL) or excludes it (NO_RECURSIVE_SQL).

Examples

In this example capture takes place over a 30-second period, polling the cache once every five seconds. This captures all statements run during that period but not before or after. If the same statement appears a second time, the process replaces the stored statement with the new occurrence.

Note that in production systems the time limit and repeat interval would be set much higher. You should tune the time_limit and repeat_interval parameters based on the workload time and shared SQL area turnover properties of your system.

```
EXEC DBMS_SQLTUNE.CAPTURE_CURSOR_CACHE_SQLSET( -
                                        sqlset_name     => 'my_workload', -
                                        time_limit      =>  30, -
                                        repeat_interval =>  5);
```

In the following call you accumulate execution statistics as you go. This option produces an accurate picture of the cumulative activity of each cursor, even across age-outs, but it is more expensive than the previous example.

```
EXEC DBMS_SQLTUNE.CAPTURE_CURSOR_CACHE_SQLSET( -
                          sqlset_name     => 'my_workload', -
                          time_limit      => 30, -
                          repeat_interval => 5, -
                          capture_mode    =>
dbms_sqltune.MODE_ACCUMULATE_STATS);
```

This call performs a very inexpensive capture where you only insert new statements and do not update their statistics once they have been inserted into the SQL tuning set

```
EXEC DBMS_SQLTUNE.CAPTURE_CURSOR_CACHE_SQLSET( -
                                    sqlset_name     => 'my_workload', -
                                    time_limit      => 30, -
                                    repeat_interval => 5, -
                                    capture_option  => 'INSERT');
```

190.5.7 CREATE_SQL_PLAN_BASELINE Procedure

This procedure creates a SQL plan baseline for an execution plan. It can be used in the context of an Alternative Plan Finding made by SQL Tuning Advisor.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.CREATE_SQL_PLAN_BASELINE (
   task_name            IN VARCHAR2,
   object_id            IN NUMBER   := NULL,
   plan_hash_value      IN NUMBER,
   owner_name           IN VARCHAR2 := NULL,
   database_link_to     IN VARCHAR2 := NULL);
```

Parameters

Table 190-14 CREATE_SQL_PLAN_BASELINE Procedure Parameters

Parameter	Description
task_name	Name of the task for which to get a script
object_id	Object ID to which the SQL corresponds
plan_hash_value	Plan to create plan baseline
owner_name	Owner of the relevant tuning task. Defaults to the current schema owner.
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: `CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY password USING 'inst1';`

190.5.8 CREATE_SQLSET Procedure and Function

This procedure or function creates a SQL tuning set object in the database.

Syntax

```
DBMS_SQLTUNE.CREATE_SQLSET (
   sqlset_name  IN  VARCHAR2,
   description  IN  VARCHAR2 := NULL
   sqlset_owner IN  VARCHAR2 := NULL);


DBMS_SQLTUNE.CREATE_SQLSET (
   sqlset_name  IN  VARCHAR2 := NULL,
   description  IN  VARCHAR2 := NULL,
   sqlset_owner IN  VARCHAR2 := NULL)
 RETURN VARCHAR2;
```

Parameters

Table 190-15 CREATE_SQLSET Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the created SQL tuning set. The name is the name passed to the function. If no name is passed to the function, then the function generates an automatic name.
`description`	Provides an optional description of the SQL tuning set.
`sqlset_owner`	Specifies the owner of the SQL tuning set, or `NULL` for the current schema owner.

Examples

```
EXEC DBMS_SQLTUNE.CREATE_SQLSET(-
  sqlset_name => 'my_workload', -
  description => 'complete application workload');
```

190.5.9 CREATE_STGTAB_SQLPROF Procedure

This procedure creates the staging table used for copying SQL profiles from one system to another.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.CREATE_STGTAB_SQLPROF (
   table_name            IN VARCHAR2,
   schema_name           IN VARCHAR2 := NULL,
   tablespace_name       IN VARCHAR2 := NULL);
```

Parameters

Table 190-16 CREATE_STGTAB_SQLPROF Procedure Parameters

Parameter	Description
`table_name`	The name of the table to create (case-insensitive unless double quoted).
`schema_name`	The schema to create the table in, or `NULL` for the current schema (case-insensitive unless double quoted).
`tablespace_name`	The tablespace to store the staging table within, or `NULL` for the default tablespace of the current user (case-insensitive unless double quoted).

Usage Notes

- Call this procedure once before issuing a call to the PACK_STGTAB_SQLPROF Procedure.
- To put different SQL profiles in different staging tables, you can call this procedure multiple times.
- This is a DDL operation, so it does not occur within a transaction.

Examples

Create a staging table to store profile data that can be moved to another system.

```
EXEC DBMS_SQLTUNE.CREATE_STGTAB_SQLPROF (table_name  => 'PROFILE_STGTAB');
```

190.5.10 CREATE_STGTAB_SQLSET Procedure

This procedure creates a staging table through which SQL tuning sets are imported and exported.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.CREATE_STGTAB_SQLSET (
   table_name            IN VARCHAR2,
   schema_name           IN VARCHAR2 := NULL,
   tablespace_name       IN VARCHAR2 := NULL,
   db_version            IN NUMBER   := NULL);
```

Parameters

Table 190-17 CREATE_STGTAB_SQLSET and CREATE_STGTAB Procedure Parameters

Parameter	Description
table_name	Specifies the of the table to create. The name is case sensitive.
schema_name	Defines the schema in which to create the table, or NULL for the current schema. The name is case sensitive.
tablespace_name	Specifies the tablespace in which to store the staging table, or NULL for the default tablespace of the current user. The name is case sensitive.
db_version	Specifies the database version that determines the format of the staging table. You can also create an older database version staging table to export an STS to an older database version. Use one of the following values: • NULL (default) — Specifies the current database version. • STS_STGTAB_10_2_VERSION — Specifies the 10.2 database version. • STS_STGTAB_11_1_VERSION — Specifies the 11.1 database version. • STS_STGTAB_11_2_VERSION — Specifies the 11.2 database version. • STS_STGTAB_12_1_VERSION — Specifies the 12.1 database version. • STS_STGTAB_12_2_VERSION — Specifies the 12.2 database version.

Security Model

You must have CREATE TABLE permissions in the specified schema and tablespace.

Usage Notes

- Call this procedure once before packing the SQL set.
- To have different tuning sets in different staging tables, you can call this procedure multiple times.
- This is a DDL operation, so it does not occur within a transaction.
- The staging table contains nested table columns and indexes, so it should not be renamed.

Examples

Create a staging table for packing and eventually exporting a SQL tuning sets

```
EXEC DBMS_SQLTUNE.CREATE_STGTAB_SQLSET(table_name => 'STGTAB_SQLSET');
```

Create a staging table to pack a SQL tuning set in Oracle Database 11*g* Release 2 (11.2) format

```
BEGIN
  DBMS_SQLTUNE.CREATE_STGTAB_SQLSET(
      table_name => 'STGTAB_SQLSET'
  ,   db_version => DBMS_SQLTUNE.STS_STGTAB_11_2_VERSION );
END;
```

190.5.11 CREATE_TUNING_TASK Functions

This function creates a SQL Tuning Advisor task.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

You can use different forms of this function to:

- Create a tuning task for a single statement given its text.
- Create a tuning task for a single statement from the shared SQL area given its identifier.
- Create a tuning task for a single statement from the workload repository given a range of snapshot identifiers.
- Create a tuning task for a SQL tuning set.
- Create a tuning task for SQL Performance Analyzer.

In all cases, the function mainly creates a SQL Tuning Advisor task and sets its parameters.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

SQL text format:

```
DBMS_SQLTUNE.CREATE_TUNING_TASK (
  sql_text         IN CLOB,
  bind_list        IN sql_binds := NULL,
  user_name        IN VARCHAR2  := NULL,
  scope            IN VARCHAR2  := SCOPE_COMPREHENSIVE,
  time_limit       IN NUMBER    := TIME_LIMIT_DEFAULT,
  task_name        IN VARCHAR2  := NULL,
  description      IN VARCHAR2  := NULL,
  con_name         IN VARCHAR2  := NULL,
  database_link_to IN VARCHAR2  := NULL)
RETURN VARCHAR2;
```

SQL ID format:

```
DBMS_SQLTUNE.CREATE_TUNING_TASK (
  sql_id           IN VARCHAR2,
```

```
  plan_hash_value  IN NUMBER    := NULL,
  scope            IN VARCHAR2  := SCOPE_COMPREHENSIVE,
  time_limit       IN NUMBER    := TIME_LIMIT_DEFAULT,
  task_name        IN VARCHAR2  := NULL,
  description      IN VARCHAR2  := NULL,
  con_name         IN VARCHAR2  := NULL,
  database_link_to IN VARCHAR2  := NULL)
RETURN VARCHAR2;
```

AWR format:

```
DBMS_SQLTUNE.CREATE_TUNING_TASK (
  begin_snap       IN NUMBER,
  end_snap         IN NUMBER,
  sql_id           IN VARCHAR2,
  plan_hash_value  IN NUMBER    := NULL,
  scope            IN VARCHAR2  := SCOPE_COMPREHENSIVE,
  time_limit       IN NUMBER    := TIME_LIMIT_DEFAULT,
  task_name        IN VARCHAR2  := NULL,
  description      IN VARCHAR2  := NULL,
  con_name         IN VARCHAR2  := NULL,
  dbid             IN NUMBER    := NULL,
  database_link_to IN VARCHAR2  := NULL)
RETURN VARCHAR2;
```

SQL tuning set format:

```
DBMS_SQLTUNE.CREATE_TUNING_TASK (
  sqlset_name       IN VARCHAR2,
  basic_filter      IN VARCHAR2 :=  NULL,
  object_filter     IN VARCHAR2 :=  NULL,
  rank1             IN VARCHAR2 :=  NULL,
  rank2             IN VARCHAR2 :=  NULL,
  rank3             IN VARCHAR2 :=  NULL,
  result_percentage IN NUMBER   :=  NULL,
  result_limit      IN NUMBER   :=  NULL,
  scope             IN VARCHAR2 :=  SCOPE_COMPREHENSIVE,
  time_limit        IN NUMBER   :=  TIME_LIMIT_DEFAULT,
  task_name         IN VARCHAR2 :=  NULL,
  description       IN VARCHAR2 :=  NULL
  plan_filter       IN VARCHAR2 :=  'MAX_ELAPSED_TIME',
  sqlset_owner      IN VARCHAR2 :=  NULL,
  database_link_to IN VARCHAR2  :=  NULL)
RETURN VARCHAR2;
```

SQL Performance Analyzer format:

```
DBMS_SQLTUNE.CREATE_TUNING_TASK (
   spa_task_name     IN VARCHAR2,
   spa_task_owner    IN VARCHAR2 :=  NULL,
   spa_compare_exec  IN VARCHAR2 :=  NULL,
   basic_filter      IN VARCHAR2 :=  NULL,
   time_limit        IN NUMBER   :=  TIME_LIMIT_DEFAULT,
   task_name         IN VARCHAR2 :=  NULL,
```

```
   description       IN VARCHAR2 :=  NULL)
RETURN VARCHAR2;
```

Parameters

Table 190-18 *CREATE_TUNING_TASK Function Parameters*

Parameter	Description
sql_text	Specifies the text of a SQL statement.
begin_snap	Specifies the begin snapshot identifier.
end_snap	Specifies the end snapshot identifier.
sql_id	Specifies the identifier of a SQL statement.
bind_list	Defines an ordered list of bind values in ANYDATA type. **NOTE:** This parameter is not supported on a standby database.
plan_hash_value	Specifies the hash value of the SQL execution plan.
sqlset_name	Specifies the SQL tuning set name.
basic_filter	Specifies the predicate used to filter the SQL from the SQL tuning set.
object_filter	Specifies the object filter.
rank(i)	Specifies an ORDER BY clause on the selected SQL statement.
result_percentage	Specifies the percentage on the sum of a ranking measure.
result_limit	Specifies the top L(imit) SQL from the filtered or ranked SQL.
user_name	Specifies the user name for whom the statement is to be tuned.
scope	Specifies the tuning scope: • LIMITED: SQL Tuning Advisor produces recommendations based on statistical checks, access path analysis, and SQL structure analysis. SQL profile recommendations are not generated. • COMPREHENSIVE: SQL Tuning Advisor carries out all the analysis it performs under limited scope plus SQL profiling.
time_limit	Specifies the maximum duration in seconds for the tuning session.
task_name	Specifies an optional tuning task name.
description	Provides a description of the SQL tuning session, up to a maximum of 256 characters.

Table 190-18 (Cont.) *CREATE_TUNING_TASK Function Parameters*

Parameter	Description
plan_filter	Specifies the plan filter. It is applicable when multiple plans (plan_hash_value) are associated with the same statement. This filter allows for selecting one plan (plan_hash_value) only. Possible values are: • LAST_GENERATED: most recent timestamp • FIRST_GENERATED: earliest timestamp, the opposite to LAST_GENERATED • LAST_LOADED: most recent first_load_time statistics information • FIRST_LOADED: earliest first_load_time statistics information, the opposite to LAST_LOADED • MAX_ELAPSED_TIME: maximum elapsed time • MAX_BUFFER_GETS: maximum buffer gets • MAX_DISK_READS: maximum disk reads • MAX_DIRECT_WRITES: maximum direct writes • MAX_OPTIMIZER_COST: maximum optimizer cost
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for the current schema owner.
spa_task_name	Specifies the name of the SQL Performance Analyzer task whose regressions are to be tuned.
spa_task_owner	Specifies the owner of specified SQL Performance Analyzer task or NULL for current user.
spa_compare_exec	Specifies the execution name of the Compare Performance trial of SQL Performance Analyzer task. If NULL, then the advisor uses the most recent execution of the given SQL Performance Analyzer task, of type COMPARE PERFORMANCE.
dbid	Specifies the DBID for imported or PDB-level AWR data. If NULL, then the current database DBID is used.

Table 190-18 (Cont.) *CREATE_TUNING_TASK Function Parameters*

Parameter	Description
con_name	Specifies the container for the tuning task. The semantics depend on the function format: • For the SQL text format, this parameter specifies the container in which SQL Tuning Advisor tunes the SQL statement. If null (default), then SQL Tuning Advisor uses the current container. • For the SQL ID format, this parameter specifies the container from which the database fetches the SQL statement for tuning. SQL Tuning Advisor tunes the statement in this container. If null, then the database uses the current PDB for tuning, fetches the statement from the cursor cache of all valid containers executing the SQL statement, and tunes the most expensive statement in its container. • For the AWR format, this parameter specifies the container from whose AWR data the database fetches the SQL statement for tuning. SQL Tuning Advisor tunes the statement in this container. If null, then the database uses the current PDB for tuning, fetches the statement from the AWR of all valid containers that have this SQL statement, and tunes the most expensive statement in its container. The following statements are true of all function formats: • In a non-CDB, this parameter is ignored. • In a PDB, this parameter must be null or match the container name of the PDB. Otherwise, an error occurs. • In a CDB root, this parameter must be null or match the container name of a container in this CDB. Otherwise, an error occurs.
database_link_to	The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: `CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY password USING 'inst1';`

Return Values

A SQL tuning task name that is unique by user (two different users can give the same name to their advisor tasks).

Usage Notes

With regard to the form of this subprogram that takes a SQL tuning set, filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.

Example 190-1 Examples

The following examples assume the following variable definitions:

```
VARIABLE stmt_task     VARCHAR2(64);
VARIABLE sts_task      VARCHAR2(64);
VARIABLE spa_tune_task VARCHAR2(64);
```

Example 190-2 Create Tuning Task with SQL Text Format

```
EXEC :stmt_task := DBMS_SQLTUNE.CREATE_TUNING_TASK( -
  sql_text => 'SELECT quantity_sold FROM sales s, times t WHERE s.time_id =
t.time_id AND s.time_id = TO_DATE(''24-NOV-00'')');
```

Example 190-3 Create Tuning Task with SQL ID Format

```
EXEC :stmt_task := DBMS_SQLTUNE.CREATE_TUNING_TASK(sql_id => 'ay1m3ssvtrh24');

EXEC :stmt_task := DBMS_SQLTUNE.CREATE_TUNING_TASK(sql_id => 'ay1m3ssvtrh24',
-
   scope => 'LIMITED');

EXEC :stmt_task := DBMS_SQLTUNE.CREATE_TUNING_TASK(sql_id => 'ay1m3ssvtrh24',
-
   time_limit => 600);
```

Example 190-4 Create Tuning Task with AWR Snapshot Format

```
EXEC :stmt_task := DBMS_SQLTUNE.CREATE_TUNING_TASK(begin_snap => 1, -
   end_snap => 2, sql_id => 'ay1m3ssvtrh24');
```

Example 190-5 Create Tuning Task with SQL Tuning Set Format

This example creates a task that tunes SQL statements in order by buffer gets, and also sets a time limit of one hour. The default ranking measure is elapsed time.

```
EXEC :sts_task := DBMS_SQLTUNE.CREATE_TUNING_TASK( -
  sqlset_name  => 'my_workload', -
  rank1        => 'BUFFER_GETS', -
  time_limit   => 3600, -
  description  => 'tune my workload ordered by buffer gets');
```

Example 190-6 Create Tuning Task with SPA Task Format

This example tunes the SQL statement that were reported as having regressed from the compare performance execution of the SQL Performance Analyzer task named `task_123`.

```
EXEC :spa_tune_task := DBMS_SQLTUNE.CREATE_TUNING_TASK(
   spa_task_name     => 'task_123',
   spa_task_owner    => 'SCOTT',
   spa_compare_exec  => 'exec1');
```

Example 190-7 Creating SQL Tuning Task on Standby Database

This example creates a tuning task on the standby database. The `tune_stby_wkld` task uses the `lnk_to_primary` database link to write data to the primary database, which is open read/write.

```
VAR tname VARCHAR2(30);
VAR query VARCHAR2(500);
EXEC :tname := 'tune_stby_wkld';
EXEC :query := 'SELECT /*+ FULL(t)*/ col1 FROM table1 t WHERE col1=9000';
EXEC :tname := DBMS_SQLTUNE.CREATE_TUNING_TASK(sql_text => :query,-
                 task_name => :tname, database_link_to => 'lnk_to_primary');
```

190.5.12 DELETE_SQLSET Procedure

This procedure deletes a set of SQL statements from a SQL tuning set.

Syntax

```
DBMS_SQLTUNE.DELETE_SQLSET (
   sqlset_name   IN  VARCHAR2,
   basic_filter  IN  VARCHAR2 := NULL,
   sqlset_owner  IN  VARCHAR2 := NULL);
```

Parameters

Table 190-19 DELETE_SQLSET Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set.
`basic_filter`	Specifies the SQL predicate to filter the SQL from the SQL tuning set. This basic filter is used as a where clause on the SQL tuning set content to select a desired subset of SQL from the SQL tuning set.
`sqlset_owner`	Specifies the owner of the SQL tuning set, or `NULL` for current schema owner.

Examples

```
-- Delete all statements in a sql tuning set.
EXEC DBMS_SQLTUNE.DELETE_SQLSET(sqlset_name   => 'my_workload');

-- Delete all statements in a sql tuning set which ran for less than a second
EXEC DBMS_SQLTUNE.DELETE_SQLSET(sqlset_name   => 'my_workload', -
                                basic_filter  => 'elapsed_time < 1000000');
```

190.5.13 DROP_SQL_PROFILE Procedure

This procedure drops the named SQL profile from the database.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.DROP_SQL_PROFILE (
   name          IN  VARCHAR2,
   ignore        IN  BOOLEAN  := FALSE);
```

Parameters

Table 190-20 DROP_SQL_PROFILE Procedure Parameters

Parameter	Description
`name`	The (mandatory) name of SQL profile to be dropped. The name is case sensitive.
`ignore`	Ignores errors due to object not existing

Usage Notes

Requires the `DROP ANY SQL PROFILE` privilege.

Examples

```
-- Drop the profile:
EXEC DBMS_SQLTUNE.DROP_SQL_PROFILE(:pname);
```

190.5.14 DROP_SQLSET Procedure

This procedure drops a SQL tuning set if it is not active.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.DROP_SQLSET (
   sqlset_name   IN  VARCHAR2,
   sqlset_owner  IN  VARCHAR2 := NULL);
```

Parameters

Table 190-21 DROP_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the SQL tuning set.
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for current schema owner.

Usage Notes

You cannot drop a SQL tuning set when it is referenced by one or more clients.

Examples

```
-- Drop the sqlset.
EXEC DBMS_SQLTUNE.DROP_SQLSET ('my_workload');
```

190.5.15 DROP_TUNING_TASK Procedure

This procedure drops a SQL tuning task. The task and all its result data are deleted.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.DROP_TUNING_TASK (
 task_name          IN VARCHAR2);
```

Parameters

Table 190-22 DROP_TUNING_TASK Procedure Parameters

Parameter	Description
task_name	Specifies name of the tuning task to drop.

190.5.16 EXECUTE_TUNING_TASK Function and Procedure

This function and procedure executes a previously created tuning task. Both the function and the procedure run in the context of a new task execution. The difference is that the function version returns that new execution name.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.EXECUTE_TUNING_TASK (
   task_name         IN VARCHAR2,
   execution_name    IN VARCHAR2               := NULL,
   execution_params  IN dbms_advisor.argList   := NULL,
   execution_desc    IN VARCHAR2               := NULL,
   database_link_to  IN VARCHAR2               := NULL)
 RETURN VARCHAR2;


DBMS_SQLTUNE.EXECUTE_TUNING_TASK (
   task_name         IN VARCHAR2,
   execution_name    IN VARCHAR2               := NULL,
   execution_params  IN dbms_advisor.argList   := NULL,
   execution_desc    IN VARCHAR2               := NULL,
   database_link_to  IN VARCHAR2               := NULL);
```

Parameters

Table 190-23 EXECUTE_TUNING_TASK Function & Procedure Parameters

Parameter	Description
task_name	Name of the tuning task to execute.
execution_name	A name to qualify and identify an execution. If not specified, it is generated by the advisor and returned by function.
execution_params	List of parameters (name, value) for the specified execution. The execution parameters have effect only on the execution for which they are specified. They override the values for the parameters stored in the task (set through the SET_TUNING_TASK_PARAMETER Procedures).
execution_desc	A 256-length string describing the execution.

Table 190-23 (Cont.) EXECUTE_TUNING_TASK Function & Procedure Parameters

Parameter	Description
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY *password* USING 'inst1';

Usage Notes

A tuning task can be executed multiples times without having to reset it.

Examples

```
EXEC DBMS_SQLTUNE.EXECUTE_TUNING_TASK(:stmt_task);
```

190.5.17 IMPLEMENT_TUNING_TASK Procedure

This procedure implements a set of SQL profile recommendations made by SQL Tuning Advisor.

Executing IMPLEMENT_TUNING_TASK is equivalent to executing the SCRIPT_TUNING_TASK Function and then running the script.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.IMPLEMENT_TUNING_TASK(
    task_name          IN  VARCHAR2,
    rec_type           IN  VARCHAR2 := REC_TYPE_SQL_PROFILES,
    owner_name         IN  VARCHAR2 := NULL,
    execution_name     IN  VARCHAR2 := NULL,
    database_link_to   IN  VARCHAR2 := NULL);
```

Parameters

Table 190-24 IMPLEMENT_TUNING_TASK Procedure Parameters

Parameter	Description
task_name	Name of the tuning task for which to implement recommendations.
rec_type	Filter the types of recommendations to implement. Only 'PROFILES' is supported.
owner_name	Owner of the relevant tuning task or NULL for the current user.
execution_name	Name of the task execution to use. If NULL, then the procedure implements recommendations from the last task execution.
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: `CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY *password* USING 'inst1';`

190.5.18 INTERRUPT_TUNING_TASK Procedure

This procedure interrupts the currently executing tuning task. The task ends its operations as it would at normal exit so that the user can access the intermediate results.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.INTERRUPT_TUNING_TASK (
 task_name         IN VARCHAR2);
```

Parameters

Table 190-25 INTERRUPT_TUNING_TASK Procedure Parameters

Parameter	Description
task_name	Name of the tuning task to interrupt

Examples

```
EXEC DBMS_SQLTUNE.INTERRUPT_TUNING_TASK(:my_task);
```

190.5.19 LOAD_SQLSET Procedure

This procedure populates the SQL tuning set with a set of selected SQL statements. You can call the procedure multiple times to add new SQL statements or replace attributes of existing statements.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.LOAD_SQLSET (
   sqlset_name       IN  VARCHAR2,
   populate_cursor   IN  sqlset_cursor,
   load_option       IN VARCHAR2 := 'INSERT',
   update_option     IN VARCHAR2 := 'REPLACE',
   update_condition  IN VARCHAR2 :=  NULL,
   update_attributes IN VARCHAR2 :=  NULL,
   ignore_null       IN BOOLEAN  :=  TRUE,
   commit_rows       IN POSITIVE :=  NULL,
   sqlset_owner      IN VARCHAR2 := NULL);
```

Parameters

Table 190-26 LOAD_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of SQL tuning set to be loaded.
populate_cursor	Specifies the cursor reference to the SQL tuning set to be loaded.

Table 190-26 (Cont.) LOAD_SQLSET Procedure Parameters

Parameter	Description
load_option	Specifies which statements are loaded into the SQL tuning set. The possible values are: • INSERT (default) — Adds only new statements. • UPDATE — Updates existing the SQL statements and ignores any new statements. • MERGE — Inserts new statements and updates the information of the existing ones.
update_option	Specifies how existing SQL statements are updated. This parameter is considered only if load_option is specified with UPDATE or MERGE as an option. The possible values are: • REPLACE (default) — Updates the statement using the new statistics, bind list, object list, and so on. • ACCUMULATE — Combines attributes when possible (for example, statistics such as elapsed_time), and otherwise replaces the existing values (for example, module and action) with the provided values. The SQL statement attributes that can be accumulated are: elapsed_time, buffer_gets, direct_writes, disk_reads, row_processed, fetches, executions, end_of_fetch_count, stat_period and active_stat_period.
update_condition	Specifies when to perform the update. The procedure only performs the update when the specified condition is satisfied. The condition can refer to either the data source or destination. The condition must use the following prefixes to refer to attributes from the source or the destination: • OLD — Refers to statement attributes from the SQL tuning set (destination). • NEW — Refers to statement attributes from the input statements (source).
update_attributes	Specifies the list of SQL statement attributes to update during a merge or update. The possible values are: • NULL (default) — Specifies the content of the input cursor except the execution context. On other terms, it is equivalent to ALL without execution contexts such as module and action. • BASIC — Specifies statistics and binds only. • TYPICAL — Specifies BASIC with SQL plans (without row source statistics) and without an object reference list. • ALL — Specifies all attributes, including the execution context attributes such as module and action. • List of comma separated attribute names to update: – EXECUTION_CONTEXT – EXECUTION_STATISTICS – SQL_BINDS – SQL_PLAN – SQL_PLAN_STATISTICS (similar to SQL_PLAN with added row source statistics)
ignore_null	Specifies whether to update attributes when the new value is NULL. If TRUE, then the procedure does not update an attribute when the new value is NULL. That is, do not override with NULL values unless intentional.

Table 190-26 (Cont.) LOAD_SQLSET Procedure Parameters

Parameter	Description
commit_rows	Specifies whether to commit statements after DML. If a value is provided, then the load commits after each specified number of statements is inserted. If NULL is provided, then the load commits only once, at the end of the operation. Providing a value for this argument enables you to monitor the progress of a SQL tuning set load operation in the DBA_SQLSET views. The STATEMENT_COUNT value increases as new SQL statements are loaded.
sqlset_owner	Defines the owner of the SQL tuning set, or the current schema owner (or NULL for the current owner).

Exceptions

- This procedure returns an error when sqlset_name is invalid, or a corresponding SQL tuning set does not exist, or the populate_cursor is incorrect and cannot be executed.
- Exceptions are also raised when invalid filters are provided. Filters can be invalid either because they don't parse (for example, they refer to attributes not in sqlset_row), or because they violate the user's privileges.

Usage Notes

Rows in the input populate_cursor must be of type SQLSET_ROW.

Examples

In this example, you create and populate a SQL tuning set with all shared SQL area statements with an elapsed time of 5 seconds or more excluding statements that belong to SYS schema (to simulate an application user workload). You select all attributes of the SQL statements and load them in the tuning set using the default mode, which loads only new statements, since the SQL tuning set is empty.

```
-- create the tuning set
EXEC DBMS_SQLTUNE.CREATE_SQLSET('my_workload');
-- populate the tuning set from the shared SQL area
DECLARE
  cur DBMS_SQLTUNE.SQLSET_CURSOR;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
      FROM table(
        DBMS_SQLTUNE.SELECT_CURSOR_CACHE(
          'parsing_schema_name <> ''SYS'' AND elapsed_time > 5000000',
           NULL, NULL, NULL, NULL, 1, NULL,
          'ALL')) P;

  DBMS_SQLTUNE.LOAD_SQLSET(sqlset_name => 'my_workload', populate_cursor =>
cur);
END;
/
```

Suppose now you wish to augment this information with what is stored in the workload repository (AWR). You populate the tuning set with `'ACCUMULATE'` as your `update_option` because it is assumed the cursors currently in the cache had aged out since the snapshot was taken.

You omit the `elapsed_time` filter because it is assumed that any statement captured in AWR is important, but still you throw away the `SYS`-parsed cursors to avoid recursive SQL.

```
DECLARE
  cur DBMS_SQLTUNE.SQLSET_CURSOR;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
      FROM table(
        DBMS_SQLTUNE.SELECT_WORKLOAD_REPOSITORY(1,2,
                                                'parsing_schema_name <>
''SYS''',
                                                NULL, NULL,NULL,NULL,
                                                1,
                                                NULL,
                                                'ALL')) P;

  DBMS_SQLTUNE.LOAD_SQLSET(sqlset_name     => 'my_workload',
                           populate_cursor => cur,
                           Using DBMS_SQLTUNE
                           load_option => 'MERGE',
                           update_option => 'ACCUMULATE');
END;
```

The following example is a simple load that only inserts new statements from the workload repository, skipping existing ones (in the SQL tuning set). Note that `'INSERT'` is the default value for the `load_option` argument of the `LOAD_SQLSET` procedure.

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
  SELECT VALUE(P)
  FROM table(DBMS_SQLTUNE.SELECT_WORKLOAD_REPOSITORY(1,2)) P;
  DBMS_SQLTUNE.LOAD_SQLSET(sqlset_name => 'my_workload', populate_cursor =>
cur);
END;
/
```

The next example demonstrates a load with `UPDATE` option. This updates statements that already exist in the SQL tuning set but does not add new ones. By default, old statistics are replaced by their new values.

```
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE) P;
```

```
  DBMS_SQLTUNE.LOAD_SQLSET(sqlset_name     => 'my_workload',
                           populate_cursor => cur,
                           load_option     => 'UPDATE');
END;
/
```

190.5.20 PACK_STGTAB_SQLPROF Procedure

This procedure copies profile data from the `SYS.` schema into the staging table.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.PACK_STGTAB_SQLPROF (
   profile_name          IN VARCHAR2 := '%',
   profile_category      IN VARCHAR2 := 'DEFAULT',
   staging_table_name    IN VARCHAR2,
   staging_schema_owner  IN VARCHAR2 := NULL);
```

Parameters

Table 190-27 PACK_STGTAB_SQLPROF Procedure Parameters

Parameter	Description
`profile_name`	The name of the profile to pack (% wildcards acceptable, case-sensitive)
`profile_category`	The category to pack profiles from (% wildcards acceptable, case-sensitive)
`staging_table_name`	The name of the table to use (case-insensitive unless double quoted). Required.
`staging_schema_owner`	The schema where the table resides, or `NULL` for current schema (case-insensitive unless double quoted)

Security Model

This procedures requires `ADMINISTER SQL MANAGEMENT OBJECT` privilege and `INSERT` privilege on the staging table.

Usage Notes

This function issues a `COMMIT` after packing each SQL profile. If an error is raised mid-execution, then clear the staging table by deleting its rows.

Examples

Put only those profiles in the `DEFAULT` category into the staging table. This corresponds to all profiles used by default on this system.

```
EXEC DBMS_SQLTUNE.PACK_STGTAB_SQLPROF (staging_table_name =>
'PROFILE_STGTAB');
```

This is another example where you put all profiles into the staging table. Note this moves profiles that are not currently being used by default but are in other categories, such as for testing purposes.

```
EXEC DBMS_SQLTUNE.PACK_STGTAB_SQLPROF (profile_category => '%', -
                                       staging_table_name => 'PROFILE_STGTAB');
```

190.5.21 PACK_STGTAB_SQLSET Procedure

This procedure copies one or more SQL tuning sets from their location in the `SYS` schema to a staging table created by the `CREATE_STGTAB_SQLSET` procedure.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.PACK_STGTAB_SQLSET (
   sqlset_name          IN VARCHAR2,
   sqlset_owner         IN VARCHAR2 := NULL,
   staging_table_name   IN VARCHAR2,
   staging_schema_owner IN VARCHAR2 := NULL,
   db_version           IN NUMBER   := NULL);
```

Examples

Put all SQL tuning sets on the database in the staging table:

```
BEGIN
  DBMS_SQLTUNE.PACK_STGTAB_SQLSET(
      sqlset_name        => '%'
  ,   sqlset_owner       => '%'
  ,   staging_table_name => 'STGTAB_SQLSET');
END;
```

Put only those SQL tuning sets owned by the current user in the staging table:

```
BEGIN
  DBMS_SQLTUNE.PACK_STGTAB_SQLSET(
      sqlset_name        => '%'
```

```
  ,   staging_table_name => 'STGTAB_SQLSET');
END;
```

Pack a specific SQL tuning set:

```
BEGIN
  DBMS_SQLTUNE.PACK_STGTAB_SQLSET(
      sqlset_name         => 'my_workload'
  ,   staging_table_name  => 'STGTAB_SQLSET');
END;
```

Pack a second SQL tuning set:

```
BEGIN
  DBMS_SQLTUNE.PACK_STGTAB_SQLSET(
      sqlset_name         => 'workload_subset'
  ,   staging_table_name  => 'STGTAB_SQLSET');
END;
```

Pack the STS `my_workload_subset` into a staging table `stgtab_sqlset` created for Oracle Database 11*g* Release 1 (11.2):

```
BEGIN
  DBMS_SQLTUNE.PACK_STGTAB_SQLSET(
      sqlset_name          => 'workload_subset'
  ,   staging_table_name   => 'STGTAB_SQLSET'
  ,   db_version           => DBMS_SQLTUNE.STS_STGTAB_11_2_VERSION);
END;
```

190.5.22 REMAP_STGTAB_SQLPROF Procedure

This procedure changes the profile data values kept in the staging table prior to performing an unpack operation.

You can use this procedure to change the category of a profile. You can also use it to change the name of a profile if one already exists on the system with the same name.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REMAP_STGTAB_SQLPROF (
  old_profile_name      IN VARCHAR2,
  new_profile_name      IN VARCHAR2 := NULL,
  new_profile_category  IN VARCHAR2 := NULL,
  staging_table_name    IN VARCHAR2,
  staging_schema_owner  IN VARCHAR2 := NULL);
```

Parameters

Table 190-28 REMAP_STGTAB_SQLPROF Procedure Parameters

Parameter	Description
old_profile_name	The name of the profile to target for a remap operation (case-sensitive)
new_profile_name	The new name of the profile, or NULL to remain the same (case-sensitive)
new_profile_category	The new category for the profile, or NULL to remain the same (case-sensitive)
staging_table_name	The name of the table on which to perform the remap operation (case-sensitive). Required.
staging_schema_owner	The schema where the table resides, or NULL for current schema (case-sensitive)

Security Model

This procedure requires the UPDATE privilege on the staging table.

Examples

Change the name of a profile before we unpack, to avoid conflicts

```
BEGIN
  DBMS_SQLTUNE.REMAP_STGTAB_SQLPROF(
      old_profile_name    => :pname
  ,   new_profile_name    => 'IMP' || :pname
  ,   staging_table_name  => 'PROFILE_STGTAB');
END;
```

Change the SQL profile in the staging table to be 'TEST' category before we import it. This way users can test the profile on the new system before it is active.

```
BEGIN
  DBMS_SQLTUNE.REMAP_STGTAB_SQLPROF(
      old_profile_name     => :pname
  ,   new_profile_category => 'TEST'
  ,   staging_table_name   => 'PROFILE_STGTAB');
END;
```

190.5.23 REMAP_STGTAB_SQLSET Procedure

This procedure changes the tuning set names and owners in the staging table so that they can be unpacked with different values.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REMAP_STGTAB_SQLSET (
   old_sqlset_name        IN VARCHAR2,
   old_sqlset_owner       IN VARCHAR2 := NULL,
   new_sqlset_name        IN VARCHAR2 := NULL,
   new_sqlset_owner       IN VARCHAR2 := NULL,
   staging_table_name     IN VARCHAR2,
   staging_schema_owner   IN VARCHAR2 := NULL
   old_con_dbid           IN NUMBER   := NULL,
   new_con_dbid           IN NUMBER   := NULL);
);
```

Parameters

The parameters are identical for the DBMS_SQLTUNE.REMAP_STGTAB_SQLSET and DBMS_SQLSET.REMAP_SQLSET procedures.

Table 190-29 REMAP_STGTAB_SQLSET and REMAP_SQLSET Procedure Parameters

Parameter	Description
old_sqlset_name	Specifies the name of the tuning set to target for a remap operation. Wildcard characters (%) are not supported.
old_sqlset_owner	Specifies the new name of the tuning set owner to target for a remap operation. NULL for current schema owner
new_sqlset_name	Specifies the new name for the tuning set, or NULL to keep the same tuning set name.
new_sqlset_owner	Specifies the new owner for the tuning set, or NULL to keep the same owner name.
staging_table_name	Specifies the name of the table on which to perform the remap operation. The value is case sensitive.
staging_schema_owner	Specifies the name of staging table owner, or NULL for the current schema owner. The value is case sensitive.
old_con_dbid	Specifies the old container DBID to be remapped to a new container DBID. Specify NULL to use the same container DBID. You must provide both old_con_dbid and new_con_dbid for the remap to succeed.
new_con_dbid	Specifies the new container DBID to replace with the old container DBID. Specify NULL to use the same container DBID. You must provide both old_con_dbid and new_con_dbid for the remap to succeed.

Usage Notes

Call this procedure multiple times to remap more than one tuning set name or owner. This procedure only handles one tuning set per call.

Examples

```
-- Change the name of an STS in the staging table before unpacking it.
BEGIN
  DBMS_SQLTUNE.REMAP_STGTAB_SQLSET(
      old_sqlset_name    =>
```

```
'my_workload'
  ,   old_sqlset_owner   => 'SH'
  ,   new_sqlset_name    =>
'imp_workload'
  ,   staging_table_name => 'STGTAB_SQLSET');

-- Change the owner of an STS in the staging table before unpacking it.
  DBMS_SQLTUNE.REMAP_STGTAB_SQLSET(
      old_sqlset_name    => 'imp_workload'
  ,   old_sqlset_owner   => 'SH'
  ,   new_sqlset_owner   => 'SYS'
  ,   staging_table_name => 'STGTAB_SQLSET');
END;
```

190.5.24 REMOVE_SQLSET_REFERENCE Procedure

This procedure deactivates a SQL tuning set to indicate that it is no longer used by the client.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REMOVE_SQLSET_REFERENCE (
   sqlset_name    IN  VARCHAR2,
   reference_id   IN  NUMBER,
   sqlset_owner   IN  VARCHAR2 := NULL,
   force_remove   IN  NUMBER   := 0);
```

Parameters

The parameters are identical for the `DBMS_SQLTUNE.REMOVE_SQLSET_REFERENCE` and `DBMS_SQLSET.REMOVE_REFERENCE` procedures.

Table 190-30 REMOVE_SQLSET_REFERENCE and REMOVE_REFERENCE Procedure Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set.
`reference_id`	Specifies the identifier of the reference to remove.
`sqlset_owner`	Specifies the owner of the SQL tuning set (or `NULL` for the current schema owner).
`force_remove`	Specifies whether references can be removed for other users (`1`) or whether they cannot be removed (`0`). Setting this parameter to `1` only takes effect when the user has the `ADMINISTER ANY SQL TUNING SET` privilege. Otherwise, the database only removes references owned by the user.

Examples

You can remove references on a given SQL tuning set when you finish using it and want to make it writable again. The following example removes the reference to `my_workload`:

```
EXEC DBMS_SQLTUNE.REMOVE_SQLSET_REFERENCE( -
                                 sqlset_name   => 'my_workload', -
                                 reference_id  => :rid,
                                 sqlset_owner  => NULL,
                                 force_remove  => 0);
```

To find all references to a given SQL tuning set, query the `DBA_SQLSET_REFERENCES` view.

190.5.25 REPORT_AUTO_TUNING_TASK Function

This function displays a report from the automatic tuning task.

This function reports on a range of task executions, whereas the REPORT_TUNING_TASK Function reports on a single execution. Note that this function is deprecated with Oracle Database 11*g* Release 2 (11.2) in favor of `DBMS_AUTO_SQLTUNE.REPORT_AUTO_TUNING_TASK`.

> **See Also:**
>
> - DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group
> - REPORT_AUTO_TUNING_TASK Function

Syntax

```
DBMS_SQLTUNE.REPORT_AUTO_TUNING_TASK(
    begin_exec      IN VARCHAR2  := NULL,
    end_exec        IN VARCHAR2  := NULL,
    type            IN VARCHAR2  := TYPE_TEXT,
    level           IN VARCHAR2  := LEVEL_TYPICAL,
    section         IN VARCHAR2  := SECTION_ALL,
    object_id       IN NUMBER    := NULL,
    result_limit    IN NUMBER    := NULL)
  RETURN CLOB;
```

Parameters

Table 190-31 REPORT_AUTO_TUNING_TASK Function Parameters

Parameter	Description
`begin_exec`	Specifies the name of the execution from which to begin the report. `NULL` retrieves a report on the most recent run.
`end_exec`	Specifies the name of the execution at which to end the report. `NULL` retrieves a report on the most recent run.

Table 190-31 (Cont.) REPORT_AUTO_TUNING_TASK Function Parameters

Parameter	Description
type	Specifies the type of the report to produce. Possible values are TYPE_TEXT which produces a text report
level	Specifies the level of detail in the report: • LEVEL_BASIC: simple version of the report. Just show info about the actions taken by the advisor. • LEVEL_TYPICAL: show information about every statement analyzed, including requests not implemented. • LEVEL_ALL: highly detailed report level, also provides annotations about statements skipped over.
section	Limits the report to a single section (ALL for all sections): • SECTION_SUMMARY - summary information • SECTION_FINDINGS - tuning findings • SECTION_PLAN - explain plans • SECTION_INFORMATION - general information • SECTION_ERROR - statements with errors • SECTION_ALL - all statements
object_id	Specifies the advisor framework object ID that represents a single statement to restrict reporting to. Specify NULL for all statements. Only valid for reports that target a single execution.
result_limit	Specifies the maximum number of SQL statements to show in the report.

Return Values

A CLOB containing the desired report.

190.5.26 REPORT_SQL_DETAIL Function

This function builds a report for a specific SQLID. For each SQLID it gives various statistics and details as obtained from the V$ views and AWR.

> **See Also:**
>
> DBMS_SQLTUNE SQL Performance Reporting Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REPORT_SQL_DETAIL (
   sql_id                  IN  VARCHAR2   DEFAULT NULL,
   sql_plan_hash_value     IN  NUMBER     DEFAULT NULL,
   start_time              IN  DATE       DEFAULT NULL,
   duration                IN  NUMBER     DEFAULT NULL,
   inst_id                 IN  NUMBER     DEFAULT NULL,
   dbid                    IN  NUMBER     DEFAULT NULL,
   event_detail            IN  VARCHAR2   DEFAULT 'YES',
   bucket_max_count        IN  NUMBER     DEFAULT 128,
```

```
   bucket_interval          IN  NUMBER     DEFAULT NULL,
   top_n                    IN  NUMBER     DEFAULT 10,
   report_level             IN  VARCHAR2   DEFAULT 'TYPICAL',
   type                     IN  VARCHAR2   DEFAULT 'ACTIVE',
   data_source              IN  VARCHAR2   DEFAULT 'AUTO',
   end_time                 IN  DATE       DEFAULT NULL,
   duration_stats           IN  NUMBER     DEFAULT NULL,
   con_name                 IN  VARCHAR2   DEFAULT NULL)
RETURN CLOB;
```

Parameters

Table 190-32 REPORT_SQL_DETAIL Function Parameters

Parameter	Description
sql_id	SQLID for which monitoring information should be displayed. If NULL (the default), display statistics for the SQLID of the last SQL statement executed in the current session.
sql_plan_hash_value	Displays SQL statistics and details for a specific plan_hash_value. If NULL (default), displays statistics and details for all plans of the SQL_ID.
start_time	If specified, shows SQL activity (from GV$ACTIVE_SESSION_HISTORY) starting at this time. On Oracle RAC, the minimum start_time is the earliest sample_time of the in-memory ASH buffers across all instances. If NULL (default), one hour before the current time.
duration	Duration of activity in seconds for the report. If NULL (default) uses a value of 1 hour.
inst_id	Target instance to get SQL details from. If NULL, uses data from all instances. If 0 or -1, uses current instance.
dbid	DBID from which to get SQL details. If NULL, uses current DBID.
event_detail	When set to 'NO', the activity is aggregated by wait_class only. Use 'YES' (the default) to aggregate by (wait_class, event_name).
bucket_max_count	If specified, this should be the maximum number of histogram buckets created in the report. If not specified, a value of 128 is used.
bucket_interval	If specified, this represents the exact time interval in seconds, of all histogram buckets. If specified, bucket_max_count is ignored.
top_n	Controls the number of entries to display per dimension in the top dimensions section. If not specified, a default value of 10 is used.

Table 190-32 (Cont.) REPORT_SQL_DETAIL Function Parameters

Parameter	Description
report_level	Level of detail for the report, either 'BASIC', 'TYPICAL' or 'ALL'. Default assumes 'TYPICAL'. Their meanings are explained below. In addition, individual report sections can also be enabled or disabled by using a +/- *section_name*. Several sections are defined: • 'TOP'- Show top values for the ASH dimensions for a SQL statement; ON by default • 'SPM'- Show existing plan baselines for a SQL statement; OFF by default • 'MISMATCH'- Show reasons for creating new child cursors (sharing criteria violations); OFF by default. • 'STATS'- Show SQL execution statistics per plan from GV$SQLAREA_PLAN_HASH; ON by default • 'ACTIVITY' - Show top activity from ASH for each plan of a SQL statement; ON by default • 'ACTIVITY_ALL' - Show top activity from ASH for each line of the plan for a SQL statement; OFF by default • 'HISTOGRAM' - Show activity histogram for each plan of a SQL statement (plan time line histogram); ON by default • 'SESSIONS' - Show activity for top sessions for each plan of a SQL statement; OFF by default • 'MONITOR' - Show show one monitored SQL execution per execution plan; ON by default • 'XPLAN' - Show execution plans; ON by default • 'BINDS' - show captured bind data; ON by default In addition, SQL text can be specified at different levels: • -SQL_TEXT - No SQL text in report • +SQL_TEXT - OK with partial SQL text up to the first 2000 chars as stored in GV$SQL_MONITOR • -SQL_FULLTEXT - No full SQL text (+SQL_TEXT) • +SQL_FULLTEXT - Show full SQL text (default value) The meanings of the three top-level report levels are: • NONE - minimum possible • BASIC - SQL_TEXT+STATS+ACTIVITY+HISTOGRAM • TYPICAL - SQL_FULLTEXT+TOP+STATS+ACTIVITY+HISTOGRAM+XPLAN+MONITOR • ALL - everything Only one of these 4 levels can be specified and, if it is, it has to be at the start of the REPORT_LEVEL string
type	Report format: 'ACTIVE' by default. Can also be 'XML' (see Usage Notes).
data_source	Determines the data source of SQL data based on one of the following values: • MEMORY: The data source is GV$ view • DISK: The data source is DBA_HIST_* view • AUTO: Automatically determines the data source based on the time frame (default)
end_time	If specified, shows SQL activity from start_time to end_time. If NULL (default), shows SQL activity for systimestamp.

Table 190-32 (Cont.) REPORT_SQL_DETAIL Function Parameters

Parameter	Description
duration_stats	Duration of additional SQL execution statistics from AWR (in hours), for the report. If NULL (default), then the duration of 24 hours is considered.
con_name	Name of the multitenant container database (CDB).

Security Model

The invoker needs the EXECUTE privilege on the DBMS_XPLAN package.

Return Values

A CLOB containing the desired report.

Usage Notes

- ACTIVE reports have a rich, interactive user interface similar to Enterprise Manager while not requiring any EM installation. The report file built is in HTML format, so it can be interpreted by most modern browsers. The code powering the active report is downloaded transparently by the web browser when the report is first viewed, hence viewing it requires outside connectivity.
- The invoker needs the SELECT or READ privilege on the following views:
  - V$SESSION
  - DBA_ADVISOR_FINDINGS
  - V$DATABASE
  - GV$ASH_INFO
  - GV$ACTIVE_SESSION_HISTORY
  - GV$SQLAREA_PLAN_HASH
  - GV$SQL
  - DBA_HIST_SNAPSHOT
  - DBA_HIST_WR_CONTROL
  - DBA_HIST_ACTIVE_SESS_HISTORY
  - DBA_HIST_SQLSTAT
  - DBA_HIST_SQL_BIND_METADATA
  - DBA_HIST_SQLTEXT
  - DBA_SQL_PLAN_BASELINES
  - DBA_SQL_PROFILES
  - DBA_ADVISOR_TASKS
  - DBA_SERVICES
  - DBA_USERS
  - DBA_OBJECTS

- DBA_PROCEDURES

190.5.27 REPORT_SQL_HISTORY Function

This function generates a query history details report for a given execution in cases where the user has the required viewing privileges.

Syntax

```
DBMS_SQLTUNE.REPORT_SQL_HISTORY(
       sql_id                    IN varchar2  default  NULL,
       session_id                IN number    default  NULL,
       session_serial            IN number    default  NULL,
       sql_exec_start            IN date      default  NULL,
       sql_exec_id               IN number    default  NULL,
       inst_id                   IN number    default  NULL,
       instance_id_filter        IN number    default  NULL,
       base_path                 IN varchar2  default  NULL,
       report_level              IN varchar2  default  'TYPIC
       type                      IN varchar2  default  'TEXT'
       sql_plan_hash_value       IN number    default  NULL,
       con_name                  IN varchar2  default  NULL,
       report_id                 IN number    default  NULL)
  RETURN CLOB;
```

Parameters

Table 190-33 REPORT_SQL_HISTORY Function Parameters

Parameter	Description
sql_id	SQL ID for for execution.
session_id	ID of the session in which the execution ran.
session_serial	Serial number of the session in which execution ran.
sql_exec_start	SQL execution start time.
sql_exec_id	SQL execution ID.
inst_id	Instance ID on which the SQL ran.
base_path	Base path URL.
report_level	Report level: 'NONE', 'BASIC', 'TYPICAL' or 'ALL'.
type	Report type: 'TEXT', 'HTML', 'XML' or 'ACTIVE'.
sql_plan_hash_value	Plan hash value of the SQL execution.
con_name	PDB Container Name.
report_id	ID of the report in the auto-report repository.

190.5.28 REPORT_SQL_HISTORY_LIST Function

This function generates a query history list report for all executions in the given user session, or across sessions for a privileged user (who holds SYS user or DBA privileges).

Syntax

```
FUNCTION DBMS_SQLTUNE.REPORT_SQL_HISTORY_LIST(
    sql_id                  IN varchar2   default  NULL,
    session_id              IN number     default  NULL,
    session_serial          IN number     default  NULL,
    inst_id                 IN number     default  NULL,
    max_sqltext_length      IN number     default  NULL,
    top_n_count             IN number     default  NULL,
    top_n_rankby            IN varchar2   default  'last_active_time',
    report_level            IN varchar2   default  'TYPICAL',
    base_path               IN varchar2   default  NULL,
    type                    IN varchar2   default 'TEXT',
    con_name                IN varchar2   default  NULL)
  RETURN CLOB;
```

Parameters

Table 190-34 REPORT_SQL_HISTORY_LIST Function Parameters

Parameter	Description
sql_id	SQL ID to filter by.
session_id	Session ID to filter by.
session_serial	Session serial number to filter by.
inst_id	Database instance ID.
max_sqltext_length	Maximum SQL text length to display.
top_n_count	Top-N SQL statement history.
top_n_rankby	Rank by last active time.
report_level	Report level: 'NONE', 'BASIC', 'TYPICAL', or 'ALL'.
base_path	Base path URL.
type	Report type: 'TEXT', 'HTML', 'XML', or 'ACTIVE'.
con_name	PDB container name.

190.5.29 REPORT_SQL_MONITOR Function

This function builds a report (text, simple HTML, active HTML, XML) for the monitoring information collected on behalf of the targeted statement execution.

> **See Also:**
>
> Real-Time SQL Monitoring for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REPORT_SQL_MONITOR(
   sql_id                    IN VARCHAR2  DEFAULT  NULL,
   dbop_name                 IN VARCHAR2  DEFAULT  NULL,
   dbop_exec_id              IN NUMBER    DEFAULT  NULL,
   session_id                IN NUMBER    DEFAULT  NULL,
   session_serial            IN NUMBER    DEFAULT  NULL,
   sql_exec_start            IN DATE      DEFAULT  NULL,
   sql_exec_id               IN NUMBER    DEFAULT  NULL,
   inst_id                   IN NUMBER    DEFAULT  NULL,
   start_time_filter         IN DATE      DEFAULT  NULL,
   end_time_filter           IN DATE      DEFAULT  NULL,
   instance_id_filter        IN NUMBER    DEFAULT  NULL,
   parallel_filter           IN VARCHAR2  DEFAULT  NULL,
   plan_line_filter          IN NUMBER    DEFAULT  NULL,
   event_detail              IN VARCHAR2  DEFAULT  'YES',
   bucket_max_count          IN NUMBER    DEFAULT  128,
   bucket_interval           IN NUMBER    DEFAULT  NULL,
   base_path                 IN VARCHAR2  DEFAULT  NULL,
   last_refresh_time         IN DATE      DEFAULT  NULL,
   report_level              IN VARCHAR2  DEFAULT 'TYPICAL',
   type                      IN VARCHAR2  DEFAULT 'TEXT',
   sql_plan_hash_value       IN NUMBER    DEFAULT  NULL,
   con_name                  IN VARCHAR2  DEFAULT  NULL,
   report_id                 IN NUMBER    DEFAULT  NULL)
 RETURN CLOB;
```

Parameters

Table 190-35 REPORT_SQL_MONITOR Function Parameters

Parameter	Description
sql_id	SQL_ID for which monitoring information should be displayed. Use NULL (the default) to report on the last statement monitored by Oracle.
dbop_name	DBOP_NAME for which monitoring information of the composite database operation is displayed.
dbop_exec_id	Execution ID for the composite database operation for which monitoring information is displayed.
session_id	If not NULL, this parameters targets only the sub-set of statements executed by the specified session. Default is NULL. Use USERENV('SID') for current session.
session_serial	In addition to the session_id parameter, one can also specify its session serial to ensure that the desired session incarnation is targeted. This parameter is ignored when session_id is NULL.
sql_exec_start	This parameter, along with sql_exec_id, is only applicable when sql_id is also specified. Jointly, they can be used to display monitoring information associated to any execution of the statement identified by sql_id, assuming that this statement was monitored. When NULL (the default), the last monitored execution of SQL sql_id is shown.

Table 190-35 (Cont.) REPORT_SQL_MONITOR Function Parameters

Parameter	Description
sql_exec_id	This parameter, along with sql_exec_start, is only applicable when sql_id is also specified. Jointly, they can be used to display monitoring information associated to any execution of the statement identified by sql_id, assuming that this statement was monitored. When NULL (the default), the last monitored execution of SQL sql_id is shown.
inst_id	Only considers statements started on the specified instance. Use -1 to target the login instance. NULL (default) targets all instances.
start_time_filter	If not NULL, the report considers only the activity (from GV$ACTIVE_SESSION_HISTORY) recorded after the specified date. If NULL, the reported activity starts when the execution of the targeted SQL statement has started.
end_time_filter	If not NULL, the report shows only the activity (from GV$ACTIVE_SESSION_HISTORY) collected before the date end_time_filter. If NULL, the reported activity ends when the targeted SQL statement execution has ended or is the current time if the statement is still executing.
instance_id_filter	Only applies when the execution runs parallel across multiple Oracle Real Application Cluster (Oracle RAC) instances. This parameter allows to only report the activity of the specified instance. Use a NULL value (the default) to include the activity on all instances where the parallel query was executed.
parallel_filter	Applies only to parallel execution and allows reporting the activity of only a subset of the processes involved in the parallel execution (Query Coordinator and/or Parallel eXecution servers). The value of this parameter can be: • NULL to target all processes • [qc][servers(<svr_grp>[,] <svr_set>[,] <srv_num>)]: 'qc' stands for query coordinator and servers() stipulate which PX servers to consider. The following examples show how to target a subset of the parallel processes: • qc: targets only the query coordinator • servers(1): targets all parallel execution servers in group number 1. Note that statement running parallel have one main server group (group number 1) plus one additional group for each nested sub-query running parallel. • servers(,2): targets all parallel execution servers from any group but only running in set 1 of each group (each group has at most two set of parallel execution servers) • servers(1,1): consider only group 1, set 1 • servers(1,2,4): consider only group 1, set 2, server number 4. This reports for a single parallel server process • qc servers(1,2,4): same as above by also including the query coordinator
event_detail	When value is 'YES' (the default), reported activity from GV$ACTIVE_SESSION_HISTORY is aggregated by (wait_class, event_name). Use 'NO' to only aggregate by wait_class.
bucket_max_count	If specified, this should be the maximum number of histogram buckets created in the report

Table 190-35 (Cont.) REPORT_SQL_MONITOR Function Parameters

Parameter	Description
bucket_interval	If specified, this represents the exact time interval in seconds, of all histogram buckets. If specified, bucket_max_count is ignored.
base_path	URL path for flex HTML resources since flex HTML format is required to access external files (java scripts and the flash SWF file itself)
last_refresh_time	If not NULL (default is NULL), the time when the report was last retrieved (see SYSDATE attribute of the report tag). Use this option to display the report of a running query, and when the report is refreshed on a regular basis. This optimizes the size of the report since only the new or changed information is returned. In particular, the following are optimized: • SQL text is not returned when this option is specified • activity histogram starts at the bucket that intersect at that time. The entire content of the bucket is returned, even if last_refresh_time is after the start of that bucket
report_level	Level of detail for the report:'NONE', 'BASIC', 'TYPICAL' or 'ALL'. Default assumes 'TYPICAL'. In addition, individual report sections can also be enabled or disabled by using a +/- *section_name*. Several sections are defined: • 'XPLAN'- Show explain plan; ON by default • 'PLAN'- Show plan monitoring statistics; ON by default • 'SESSIONS'- Show session details. Applies only to parallel queries; ON by default • 'INSTANCE'- Show instance details. Applies only to parallel and cross instance; ON by default • 'PARALLEL'- An umbrella parameter for specifying sessions+instance details • 'ACTIVITY' - Show activity summary at global level, plan line level and session or instance level (if applicable); ON by default • 'BINDS' - Show bind information when available; ON by default • 'METRICS' - Show metric data (CPU, I/Os, ...) over time; ON by default • 'ACTIVITY_HISTOGRAM' - Show an histogram of the overall query activity; ON by default • 'PLAN_HISTOGRAM' - Show activity histogram at plan line level; OFF by default • 'OTHER' - Other info; ON by default In addition, SQL text can be specified at different levels: • SQL_TEXT - No SQL text in report • +SQL_TEXT - OK with partial SQL text up to the first 2000 chars as stored in GV$SQL_MONITOR • -SQL_FULLTEXT - No full SQL text (+SQL_TEXT) • +SQL_FULLTEXT - Show full SQL text (default value)

Table 190-35 (Cont.) REPORT_SQL_MONITOR Function Parameters

Parameter	Description
report_level (contd.)	The meanings of the three top-level report levels are: • NONE - minimum possible • +BASIC - SQL_TEXT-PLAN-XPLAN-SESSIONS-INSTANCE-ACTIVITY_HISTOGRAM-PLAN_HISTOGRAM-METRICS • TYPICAL - everything but PLAN_HISTOGRAM • ALL - everything Only one of these 4 levels can be specified and, if it is, it has to be at the start of the REPORT_LEVEL string
type	Report format, 'TEXT' by default. Can be 'TEXT', 'HTML', 'XML' or 'ACTIVE' (see Usage Notes).
sql_plan_hash_value	Target only those SQL executions with the specified plan_hash_value. Default is NULL.
con_name	Name of the multitenant container database (CDB).
report_id	ID of the report in auto-report repository. Report IDs can be found in DBA_HIST_REPORTS.

Return Values

A CLOB containing the desired report.

Usage Notes

- The target SQL statement for this report can be:
  - The most recent SQL statement monitored by Oracle Database. This is the default behavior, so there is no need to specify any parameter.
  - The most recent SQL statement executed by a specific session and monitored by Oracle. The session is identified by its session id and optionally it serial number. For example, use session_id => for the current session or session_id => 20, session_serial => 103 for session ID 20, serial number 103.
  - The most recent execution of a specific statement identified by its sql_id.
  - A specific execution of a SQL statement identified by its execution key (sql_id, sql_exec_start and sql_exec_id).
- This report produces performance data exposed by several fixed views, listed below. For this reason, the invoker of the report function must have privilege to select data from these fixed views (such as the SELECT_CATALOG role).
  - GV$SQL_MONITOR
  - GV$SQL_PLAN_MONITOR
  - GV$SQL_PLAN
  - GV$ACTIVE_SESSION_HISTORY
  - GV$SESSION_LONGOPS
  - GV$SQL
- The bucket_max_count and bucket_interval parameters control the activity histogram.

By default, the maximum number of buckets is set to 128. The database derives the `bucket_interval` value based on this count. The `bucket_interval` (value is in seconds) is computed such that it is the smallest possible power of 2 value (starting at 1 second) without exceeding the maximum number of buckets. For example, if the query has executed for 600 seconds, then the database selects a `bucket_interval` of 8 seconds (a power of two). The database chooses the value of 8 because 600/8 = 74, which is less than 128 buckets maximum. Smaller than 8 seconds would be 4 seconds, which would lead to more buckets than the 128 maximum. If `bucket_interval` is specified, then the database uses the specified value instead of deriving it from `bucket_max_count`.

- `ACTIVE` reports have a rich, interactive user interface similar to Enterprise Manager, while not requiring any EM installation.

  The report file is in HTML format. The code powering the active report is downloaded transparently by the web browser when the report is first viewed. Therefore, viewing the report requires outside connectivity.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* for more information about SQL real-time monitoring.

190.5.30 REPORT_SQL_MONITOR_LIST Function

This function builds a report for all or a subset of statements monitored by Oracle Database. For each statement, the subprogram gives key information and associated global statistics.

Use the REPORT_SQL_MONITOR Function to get detailed monitoring information for a single SQL statement.

> **See Also:**
>
> Real-Time SQL Monitoring for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REPORT_SQL_MONITOR_LIST(
   sql_id                    IN VARCHAR2  DEFAULT  NULL,
   session_id                IN NUMBER    DEFAULT  NULL,
   session_serial            IN NUMBER    DEFAULT  NULL,
   inst_id                   IN NUMBER    DEFAULT  NULL,
   active_since_date         IN DATE      DEFAULT  NULL,
   active_since_sec          IN NUMBER    DEFAULT  NULL,
   active_before_date        IN DATE      DEFAULT  NULL,
   last_refresh_time         IN DATE      DEFAULT  NULL,
   dbop_name                 IN VARCHAR2  DEFAULT  NULL,
   monitor_type              IN NUMBER    DEFAULT  MONITOR_TYPE_ALL,
   max_sqltext_length        IN NUMBER    DEFAULT  NULL,
   top_n_count               IN NUMBER    DEFAULT  NULL,
   top_n_rankby              IN VARCHAR2  DEFAULT  'LAST_ACTIVE_TIME',
   report_level              IN VARCHAR2  DEFAULT  'TYPICAL',
```

```
  auto_refresh              IN NUMBER    DEFAULT  NULL,
  base_path                 IN VARCHAR2  DEFAULT  NULL,
  type                      IN VARCHAR2  DEFAULT 'TEXT',
  con_name                  IN VARCHAR2  DEFAULT  NULL,
  top_n_detail_count        IN NUMBER    DEFAULT  NULL)
RETURN CLOB;
```

Parameters

Table 190-36 REPORT_SQL_MONITOR_LIST Function Parameters

Parameter	Description
sql_id	SQL_ID for which monitoring information should be displayed. Use NULL (the default) to report on the last statement monitored by Oracle.
session_id	If not NULL, then this parameter targets only the subset of statements executed by the specified session. Default is NULL. Use -1 or USERENV('SID') for current session.
session_serial	In addition to the session_id parameter, you can also specify its session serial to ensure that the desired session incarnation is targeted. This parameter is ignored when session_id is NULL.
inst_id	Only considers statements started on the specified instance. Use -1 to target the login instance. NULL (default) targets all instances.
active_since_date	If not NULL (default), returns only monitored statements active since the specified time. This includes all statements that are still executing along with all statements that have completed their execution after the specified date and time.
active_since_sec	Same as active_since_date but with the date specified relative to the current SYSDATE minus a specified number of seconds. For example, use 3600 to apply a limit of 1 hour.
active_before_date	If not NULL (default), returns only monitored statements that have been active before the specified date and time.
last_refresh_time	If not NULL (default), the date and time when the list report was last retrieved. This optimizes the case where an application shows the list and refreshes the report on a regular basis (such as once every 5 seconds). In this case, the report shows detail about the execution of monitored queries that active since the specified last_refresh_time. For other queries, the report returns the execution key (sql_id, sql_exec_start, sql_exec_id). For queries with a first refresh time after the specified date, the function returns only the SQL execution key and statistics.
dbop_name	DB operation name. Specify NULL to display all the monitored DB operations.
monitor_type	Type of the SQL Monitor operation. Specify one of the following values: • MONITOR_TYPE_SQL - Returns only SQL statements • MONITOR_TYPE_DBOP - Returns only database operations • MONITOR_TYPE_ALL - Returns SQL statements as well as database operations
max_sqltext_length	Maximum length of the SQL text. Default is NULL (no limit).
top_n_count	Limits the number of top-N SQL statements that need to be included in the report.

Table 190-36 (Cont.) REPORT_SQL_MONITOR_LIST Function Parameters

Parameter	Description
top_n_rankby	Specifies the attribute to rank the SQL statements. Specify this value when top_n_count value is not NULL. The ranking of an SQL statement is done based on one of the following values: • LAST_ACTIVE_TIME - Last active date and time (top N most recent) • DURATION - Total duration of execution • DB_TIME - DB time used • CPU_TIME - CPU time used • IO_REQUESTS - Number of I/O requests • IO_BYTES - Number of I/O bytes
report_level	Level of detail for the report. The level is one of the following: • BASIC - SQL text up to 200 characters • TYPICAL - include full SQL text assuming that cursor has not aged out, in which case the SQL text is included up to 2000 characters • ALL - currently the same as TYPICAL
auto_refresh	Currently non-operational, reserved for future use.
base_path	URL path for flex HTML resources because flex HTML format is required to access external files (java scripts and the flash SWF file itself).
type	Report format: TEXT (default), HTML, or XML.
con_name	Name of the multitenant container database (CDB)
top_n_detail_count	Limits the number of top-N SQL statements for which the SQL monitor details need to be included in the report.

Return Values

A report for the list of SQL statements that have been monitored. The report type is text, XML, or HTML.

Usage Notes

You must have the privilege to access the following fixed views: GV$SQL_MONITOR and GV$SQL.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* for more information about SQL real-time monitoring.

190.5.31 REPORT_TUNING_TASK Function

This function displays the results of a tuning task. By default the report is in text format.

> **See Also:**
>
> DBMS_SQLTUNE SQL Performance Reporting Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REPORT_TUNING_TASK(
   task_name        IN   VARCHAR2,
   type             IN   VARCHAR2   := 'TEXT',
   level            IN   VARCHAR2   := 'TYPICAL',
   section          IN   VARCHAR2   := ALL,
   object_id        IN   NUMBER     := NULL,
   result_limit     IN   NUMBER     := NULL,
   owner_name       IN   VARCHAR2   := NULL,
   execution_name   IN   VARCHAR2   := NULL,
   database_link_to IN   VARCHAR2   := NULL)
RETURN CLOB;
```

Parameters

Table 190-37 REPORT_TUNING_TASK Function Parameters

Parameter	Description
`task_name`	Name of the tuning task.
`type`	Type of the report to produce. Possible values are `TEXT` which produces a text report.
`level`	Level of detail in the report: • `BASIC`: simple version of the report. Just show info about the actions taken by the advisor. • `TYPICAL`: show information about every statement analyzed, including requests not implemented. • `ALL`: highly detailed report level, also provides annotations about statements skipped over.
`section`	Section of the report to include. You can limit the report to any of the following single sections (`ALL` for all sections): • `SUMMARY` - Summary information • `FINDINGS` - Tuning findings • `PLAN` - Explain plans • `INFORMATION` - General information • `ERROR` - Statements with errors • `ALL` - All statements

Table 190-37 (Cont.) REPORT_TUNING_TASK Function Parameters

Parameter	Description
object_id	Advisor framework object ID that represents a single statement to restrict reporting to. NULL for all statements. Only valid for reports that target a single execution.
result_limit	Maximum number of SQL statements to show in the report.
owner_name	Owner of the relevant tuning task. The default is the current schema owner.
execution_name	Name of the task execution to use. If NULL, then the function generates the report for the last task execution.
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY *password* USING 'inst1';

Return Values

A CLOB containing the desired report.

Examples

```
-- Display the report for a single statement.
SELECT DBMS_SQLTUNE.REPORT_TUNING_TASK(:stmt_task)
FROM   DUAL;

-- Display the summary for a SQL tuning set.
SELECT DBMS_SQLTUNE.REPORT_TUNING_TASK(:sts_task, 'TEXT', 'TYPICAL',
'SUMMARY')
FROM   DUAL;

-- Display the findings for a specific statement.
SELECT DBMS_SQLTUNE.REPORT_TUNING_TASK(:sts_task, 'TEXT',
'TYPICAL','FINDINGS', 5)
FROM   DUAL;
```

190.5.32 REPORT_TUNING_TASK_XML Function

This function displays an XML report of a tuning task.

> **See Also:**
>
> DBMS_SQLTUNE SQL Performance Reporting Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.REPORT_TUNING_TASK_LIST_XML(
   task_name         IN   VARCHAR2   := NULL,
   level             IN   VARCHAR2   := LEVEL_TYPICAL,
   section           IN   VARCHAR2   := SECTION_ALL,
   object_id         IN   NUMBER     := NULL,
   result_limit      IN   NUMBER     := 160,
   owner_name        IN   VARCHAR2   := NULL,
   execution_name    IN   VARCHAR2   := NULL,
   autotune_period   IN   NUMBER     := NULL,
   report_tag        IN   VARCHAR2   := NULL)
RETURN XMLTYPE;
```

Parameters

Table 190-38 REPORT_TUNING_TASK_XML Function Parameters

Parameter	Description
`task_name`	Name of the tuning task.
`level`	Level of detail in the report: • `BASIC`: simple version of the report. Just show info about the actions taken by the advisor. • `TYPICAL`: show information about every statement analyzed, including requests not implemented. • `ALL`: highly detailed report level, also provides annotations about statements skipped over.
`section`	Section of the report to include. You can limit the report to any of the following single sections (`ALL` for all sections): • `SUMMARY` - Summary information • `ALL` - All statements
`object_id`	Advisor framework object ID that represents a single statement to restrict reporting to. `NULL` for all statements. Only valid for reports that target a single execution.

Table 190-38 (Cont.) REPORT_TUNING_TASK_XML Function Parameters

Parameter	Description
result_limit	The number of statements in a SQL tuning set or snapshot range for which the report is generated. The default is 160 (20 statements * 8 categories). The categories are as follows: • Profile • Index • Restructure SQL • Alternate plan • Statistics • Errors • Information • No findings
owner_name	Owner of the relevant tuning task. The default is the current schema owner.
execution_name	Name of the task execution to use. If NULL, then the function generates the report for the most recent task execution.
autotune_period	The time period for the automatic SQL tuning. This setting applies only to the automatic SQL Tuning Advisor task. Possible values are as follows: • Null or negative value (default) — All or full. The result includes all task executions. • 0 — Result of the current or most recent task execution. • 1 — Result for the most recent 24-hour period. • 7 — Result for the most recent 7-day period. The procedure interprets any other value as the time of the most recent task execution minus the value of this argument.
report_tag	The name of the root XML tag. By default, the tag is the report reference generated by the reporting framework.

Return Values

A CLOB containing the desired report.

190.5.33 RESET_TUNING_TASK Procedure

This procedure is called on a tuning task that is not currently executing to prepare it for re-execution.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.RESET_TUNING_TASK(
 task_name         IN VARCHAR2);
```

Parameters

Table 190-39 RESET_TUNING_TASK Procedure Parameters

Parameter	Description
task_name	The name of the tuning task to reset

Examples

```
-- reset and re-execute a task
EXEC DBMS_SQLTUNE.RESET_TUNING_TASK(:sts_task);

-- re-execute the task
EXEC DBMS_SQLTUNE.EXECUTE_TUNING_TASK(:sts_task);
```

190.5.34 RESUME_TUNING_TASK Procedure

This procedure resumes a previously interrupted task that was created to process a SQL tuning set.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.RESUME_TUNING_TASK(
 task_name          IN VARCHAR2,
 basic_filter       IN VARCHAR2 := NULL);
```

Parameters

Table 190-40 RESUME_TUNING_TASK Procedure Parameters

Parameter	Description
task_name	The name of the tuning task to resume.
basic_filter	A SQL predicate to filter the SQL from the SQL tuning set. Note that this filter is applied in conjunction with the parameter basic_filteri when calling CREATE_TUNING_TASK Functions.

Usage Notes

Resuming a single SQL tuning task (a task that was created to tune a single SQL statement as compared to a SQL tuning set) is not supported.

Examples

```
-- Interrupt the task
EXEC DBMS_SQLTUNE.INTERRUPT_TUNING_TASK(:conc_task);

-- Once a task is interrupted, we can elect to reset it, resume it, or check
-- out its results and then decide.  For this example we will just resume.

EXEC DBMS_SQLTUNE.RESUME_TUNING_TASK(:conc_task);
```

190.5.35 SCHEDULE_TUNING_TASK Function

This function creates a tuning task for a single SQL statement and schedules a `DBMS_SCHEDULER` job to execute the tuning task. One form of the function finds the information about the statement to be tuned in the shared SQL area, whereas the other finds the information in AWR.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

Shared SQL Area Format:

```
DBMS_SQLTUNE.SCHEDULE_TUNING_TASK(
  sql_id          IN VARCHAR2,
  plan_hash_value IN NUMBER                    := NULL,
  start_date      IN TIMESTAMP WITH TIME ZONE  := NULL,
  scope           IN VARCHAR2                  := SCOPE_COMPREHENSIVE,
  time_limit      IN NUMBER                    := TIME_LIMIT_DEFAULT,
  task_name       IN VARCHAR2                  := NULL,
  description     IN VARCHAR2                  := NULL,
  con_name        IN VARCHAR2                  := NULL)
RETURN VARCHAR2;
```

AWR Format:

```
DBMS_SQLTUNE.SCHEDULE_TUNING_TASK(
  begin_snap      IN NUMBER,
  end_snap        IN NUMBER,
  sql_id          IN VARCHAR2,
  plan_hash_value IN NUMBER                    := NULL,
  start_date      IN TIMESTAMP WITH TIME ZONE  := NULL,
  scope           IN VARCHAR2                  := SCOPE_COMPREHENSIVE,
  time_limit      IN NUMBER                    := TIME_LIMIT_DEFAULT,
  task_name       IN VARCHAR2                  := NULL,
  description     IN VARCHAR2                  := NULL,
  con_name        IN VARCHAR2                  := NULL,
```

```
  dbid            IN NUMBER                    := NULL)
RETURN VARCHAR2;
```

Parameters

Table 190-41 SCHEDULE_TUNING_TASK Function Parameters

Parameter	Description
begin_snap	The beginning snapshot identifier. The range is exclusive, which means that SQL statements in this snapshot ID are not included.
end_snap	The end snapshot identifier. The range is inclusive, which means that SQL statements in this snapshot ID are included.
sql_id	The SQL ID of the statement to be tuned.
plan_hash_value	The plan hash value of the statement to be tuned. For example, the tuning job fetches captured binds for this SQL plan.
start_date	The date on which the schedule becomes valid. If null, then SQL Tuning Advisor immediately executes the task.
scope	The scope of the tuning job: limited, or comprehensive.
time_limit	The maximum duration in seconds for the SQL tuning session.
task_name	Optional SQL tuning task name.
description	Description of the SQL tuning session. The description can contain a maximum of 256 characters.
con_name	The container from which SQL Tuning Advisor accesses the SQL statement information.
dbid	DBID for imported or PDB-level AWR data. If NULL, then the current database DBID is used.

Security Model

The caller must possess the CREATE JOB privilege for the job.

Return Values

A SQL tuning task name that is unique for each user. Multiple users can assign the same name to their advisor tasks.

Usage Notes

- The task is scheduled only once.
- The name of the scheduler job is created as follows: sqltune_job_*taskid_orahash*(*systimestamp*).

190.5.36 SCRIPT_TUNING_TASK Function

This function creates a SQL*Plus script which can then be executed to implement a set of SQL Tuning Advisor recommendations.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Advisor Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.SCRIPT_TUNING_TASK(
  task_name          IN VARCHAR2,
  rec_type           IN VARCHAR2  := REC_TYPE_ALL,
  object_id          IN NUMBER    := NULL,
  result_limit       IN NUMBER    := NULL,
  owner_name         IN VARCHAR2  := NULL,
  execution_name     IN VARCHAR2  := NULL,
  database_link_to   IN VARCHAR2  := NULL)
 RETURN CLOB;
```

Parameters

Table 190-42 SCRIPT_TUNING_TASK Function Parameters

Parameter	Description
`task_name`	Name of the tuning task for which to apply a script.
`rec_type`	Filter the script by types of recommendations to include. You can use any subset of the following values, separated by commas: 'ALL: ''PROFILES''''STATISTICS''''INDEXES'. For example, a script with profiles and statistics would use the filter 'PROFILES,STATISTICS'.
`object_id`	Optionally filters by a single object ID.
`result_limit`	Optionally shows commands for only top *n* SQL (ordered by `object_id` and ignored if an `object_id` is also specified).
`owner_name`	Owner of the relevant tuning task. Defaults to the current schema owner.
`excution_name`	Name of the task execution to use. If `NULL`, the script is generated for the last task execution.

Table 190-42 (Cont.) SCRIPT_TUNING_TASK Function Parameters

Parameter	Description
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY *password* USING 'inst1';

Return Values

Returns a script in the form of a CLOB.

Usage Notes

- After the script is returned, check it before executing it.
- Wrap with a call to DBMS_ADVISOR.CREATE_FILE to put it into a file.

Examples

```
SET LINESIZE 140

-- Get a script for all actions recommended by the task.
SELECT DBMS_SQLTUNE.SCRIPT_TUNING_TASK(:stmt_task) FROM DUAL;

-- Get a script of only the sql profiles we should create.
SELECT DBMS_SQLTUNE.SCRIPT_TUNING_TASK(:stmt_task, 'PROFILES') FROM DUAL;

-- Get a script of only stale / missing stats
SELECT DBMS_SQLTUNE.SCRIPT_TUNING_TASK(:stmt_task, 'STATISTICS') FROM DUAL;

-- Get a script with recommendations about only one SQL statement when we have
-- tuned an entire STS.
SELECT DBMS_SQLTUNE.SCRIPT_TUNING_TASK(:sts_task, 'ALL', 5) FROM DUAL;
```

190.5.37 SELECT_CURSOR_CACHE Function

This function collects SQL statements from the shared SQL area.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.SELECT_CURSOR_CACHE (
  basic_filter         IN   VARCHAR2 := NULL,
  object_filter        IN   VARCHAR2 := NULL,
  ranking_measure1     IN   VARCHAR2 := NULL,
  ranking_measure2     IN   VARCHAR2 := NULL,
  ranking_measure3     IN   VARCHAR2 := NULL,
  result_percentage    IN   NUMBER   := 1,
  result_limit         IN   NUMBER   := NULL,
  attribute_list       IN   VARCHAR2 := NULL,
  recursive_sql        IN   VARCHAR2 := HAS_RECURSIVE_SQL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 190-43 SELECT_CURSOR_CACHE Function Parameters

Parameter	Description
`basic_filter`	Specifies the SQL predicate that filters the SQL from the shared SQL area defined on attributes of the `SQLSET_ROW`. If `basic_filter` is not set by the caller, then the subprogram captures only statements of the type `CREATE TABLE`, `INSERT`, `SELECT`, `UPDATE`, `DELETE`, and `MERGE`.
`object_filter`	Currently not supported.
`ranking_measure(n)`	Defines an `ORDER BY` clause on the selected SQL.
`result_percentage`	Specifies a filter that picks the top *n*% according to the supplied ranking measure. The value applies only if one ranking measure is supplied.
`result_limit`	Defines the top limit SQL from the filtered source ranked by the ranking measure.

Table 190-43 (Cont.) SELECT_CURSOR_CACHE Function Parameters

Parameter	Description
attribute_list	Specifies the list of SQL statement attributes to return in the result. Possible values are: • TYPICAL — Specifies BASIC plus SQL plan (without row source statistics) and without object reference list (default). • BASIC — Specifies all attributes (such as execution statistics and binds) except the plans. The execution context is always part of the result. • ALL — Specifies all attributes. • Comma-separated list of attribute names. This values returns only a subset of SQL attributes: – EXECUTION_STATISTICS – BIND_LIST – OBJECT_LIST – SQL_PLAN – SQL_PLAN_STATISTICS — Similar to SQL_PLAN plus row source statistics
recursive_sql	Specifies that the filter must include recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL, which is the default) or exclude it (NO_RECURSIVE_SQL).

Return Values

This function returns a one SQLSET_ROW per SQL_ID or PLAN_HASH_VALUE pair found in each data source.

Usage Notes

- Filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.
- Users need privileges on the shared SQL area views.

Examples

```
-- Get sql ids and sql text for statements with 500 buffer gets.
SELECT sql_id, sql_text
FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE('buffer_gets > 500'))
ORDER BY sql_id;

-- Get all the information we have about a particular statement.
SELECT *
FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE('sql_id = ''4rm4183czbs7j'''));

-- Notice that some statements can have multiple plans.  The output of the
-- SELECT_XXX table functions is unique by (sql_id, plan_hash_value).  This is
-- because a data source can store multiple plans per sql statement.
SELECT sql_id, plan_hash_value
FROM table(dbms_sqltune.select_cursor_cache('sql_id = ''ay1m3ssvtrh24'''))
ORDER BY sql_id, plan_hash_value;

-- PL/SQL examples: load_sqlset is called after opening a cursor, along the
```

```
-- lines given below

-- Select all statements in the shared SQL area.
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT value(P)
    FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE) P;

  -- Process each statement (or pass cursor to load_sqlset).

  CLOSE cur;
END;/


-- Look for statements not parsed by SYS.
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur for
    SELECT VALUE(P)
    FROM table(
     DBMS_SQLTUNE.SELECT_CURSOR_CACHE('parsing_schema_name <> ''SYS''')) P;

  -- Process each statement (or pass cursor to load_sqlset).

  CLOSE cur;
end;/


-- All statements from a particular module/action.
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(
      DBMS_SQLTUNE.SELECT_CURSOR_CACHE(
         'module = ''MY_APPLICATION'' and action = ''MY_ACTION''')) P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;/


-- all statements that ran for at least five seconds
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE('elapsed_time > 5000000')) P;

  -- Process each statement (or pass cursor to load_sqlset)
```

```
  CLOSE cur;
end;/


-- select all statements that pass a simple buffer_gets threshold and
-- are coming from an APPS user
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(
      DBMS_SQLTUNE.SELECT_CURSOR_CACHE(
        'buffer_gets > 100 and parsing_schema_name = ''APPS'''))P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
end;/


-- select all statements exceeding 5 seconds in elapsed time, but also
-- select the plans (by default we only select execution stats and binds
-- for performance reasons - in this case the SQL_PLAN attribute of sqlset_row
-- is NULL)
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(dbms_sqltune.select_cursor_cache(
      'elapsed_time > 5000000', NULL, NULL, NULL, NULL, 1, NULL,
      'EXECUTION_STATISTICS, SQL_BINDS, SQL_PLAN')) P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;/


-- Select the top 100 statements in the shared SQL area ordering by
elapsed_time.
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE(NULL,
                                                NULL,
                                                'ELAPSED_TIME', NULL, NULL,
                                                1,
                                                100)) P;

  -- Process each statement (or pass cursor to load_sqlset)
```

```
  CLOSE cur;
end;/


-- Select the set of statements which cumulatively account for 90% of the
-- buffer gets in the shared SQL area.  This means that the buffer gets of all
-- of these statements added up is approximately 90% of the sum of all
-- statements currently in the cache.
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE(P)
    FROM table(DBMS_SQLTUNE.SELECT_CURSOR_CACHE(NULL,
                                                NULL,
                                                'BUFFER_GETS', NULL, NULL,
                                                .9)) P;

  -- Process each statement (or pass cursor to load_sqlset).

  CLOSE cur;
END;
/
```

190.5.38 SELECT_SQL_TRACE Function

This table function reads the content of one or more trace files and returns the SQL statements it finds in the format of `sqlset_row`.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.SELECT_SQL_TRACE (
  directory              IN VARCHAR2,
  file_name              IN VARCHAR2 := NULL,
  mapping_table_name     IN VARCHAR2 := NULL,
  mapping_table_owner    IN VARCHAR2 := NULL,,
  select_mode            IN POSITIVE := SINGLE_EXECUTION,
  options                IN BINARY_INTEGER := LIMITED_COMMAND_TYPE,
  pattern_start          IN VARCHAR2 := NULL,
  pattern_end            IN VARCHAR2 := NULL,
  result_limit           IN POSITIVE := NULL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 190-44 SELECT_SQL_TRACE Function Parameters

Parameter	Description
directory	Defines the directory object containing the trace files. This field is mandatory.
file_name	Specifies all or part of the name of the trace files. If NULL, then the function uses the current or most recent file in the specified location or path. '%' wildcards are supported for matching trace file names.
mapping_table_name	Specifies the mapping table name. Note that the mapping table name is case insensitive. If the mapping table name is NULL, then the function uses the mappings in the current database.
mapping_table_owner	Specifies the mapping table owner. If it is NULL, then the function uses the current user.
select_mode	Specifies the mode for selecting SQL from the trace. Possible values are: • SINGLE_EXECUTION — Returns one execution of a SQL. This is the default. • ALL_EXECUTIONS — Returns all executions.
options	Specifies which types of SQL statements are returned. • LIMITED_COMMAND_TYPE — Returns the SQL statements with the command types CREATE, INSERT, SELECT, UPDATE, DELETE, and MERGE. This value is the default. • ALL_COMMAND_TYPE — Returns the SQL statements with all command types.
pattern_start	Specifies the delimiting pattern of the trace file sections to consider. CURRENTLY INOPERABLE.
pattern_end	Specifies the closing delimiting pattern of the trace file sections to process. CURRENTLY INOPERABLE.
result_limit	Specifies the top SQL from the filtered source. Default to MAXSB4 if NULL.

Return Values

This function returns a SQLSET_ROW object.

Usage Notes

The ability to create a directory object for the system directory creates a potential security issue. For example, in a CDB, all containers write trace files to the same directory. A local user with SELECT privileges on this directory can read the contents of trace files belonging to any container.

To prevent this type of unauthorized access, copy the files from the default SQL trace directory into a different directory, and then create a directory object. Use the PATH_PREFIX clause of the CREATE PLUGGABLE DATABASE statement to ensure that all directory object paths associated with the PDB are restricted to the specified directory or its subdirectories.

Examples

The following code shows how to enable SQL trace for a few SQL statements and load the results into a SQL tuning set:

```
-- turn on the SQL trace in the capture database
ALTER SESSION SET EVENTS '10046 TRACE NAME CONTEXT FOREVER, LEVEL 4'

-- run sql statements
SELECT 1 FROM DUAL;
SELECT COUNT(*) FROM dba_tables WHERE table_name = :mytab;

ALTER SESSION SET EVENTS '10046 TRACE NAME CONTEXT OFF';

-- create mapping table from the capture database
CREATE TABLE mapping AS
SELECT object_id id, owner, substr(object_name, 1, 30) name
   FROM dba_objects
   WHERE object_type NOT IN ('CONSUMER GROUP', 'EVALUATION CONTEXT',
                             'FUNCTION', 'INDEXTYPE', 'JAVA CLASS',
                             'JAVA DATA', 'JAVA RESOURCE', 'LIBRARY',
                             'LOB', 'OPERATOR', 'PACKAGE',
                             'PACKAGE BODY', 'PROCEDURE', 'QUEUE',
                             'RESOURCE PLAN', 'TRIGGER', 'TYPE',
                             'TYPE BODY')
UNION ALL
SELECT user_id id, username owner, NULL name
   FROM dba_users;

-- create the directory object where the SQL traces are stored
CREATE DIRECTORY SQL_TRACE_DIR as '/home/foo/trace';

-- create the STS
EXEC DBMS_SQLTUNE.CREATE_SQLSET('my_sts', 'test purpose');

-- load the SQL statements into STS from SQL TRACE
DECLARE
   cur sys_refcursor;
BEGIN
   OPEN cur FOR
   SELECT value(p)
     FROM TABLE(
        DBMS_SQLTUNE.SELECT_SQL_TRACE(
           directory=>'SQL_TRACE_DIR',
           file_name=>'%trc',
           mapping_table_name=>'mapping')) p;
   DBMS_SQLTUNE.LOAD_SQLSET('my_sts', cur);
   CLOSE cur;
END;
/
```

> **See Also:**
>
> *Oracle Database SQL Language Reference* to learn more about the `PATH_PREFIX` clause

190.5.39 SELECT_SQLPA_TASK Function

This function collects SQL statements from a SQL Performance Analyzer comparison task.

> **See Also:**
>
> - DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group
> - *Oracle Database Testing Guide* for a `SELECT_SQLPA_TASK` example

Syntax

```
DBMS_SQLTUNE.SELECT_SQLPA_TASK(
    task_name         IN VARCHAR2,
    task_owner        IN VARCHAR2 := NULL,
    execution_name    IN VARCHAR2 := NULL,
    level_filter      IN VARCHAR2 := 'REGRESSED',
    basic_filter      IN VARCHAR2 := NULL,
    object_filter     IN VARCHAR2 := NULL,
    attribute_list    IN VARCHAR2 := 'TYPICAL')
RETURN sys.sqlset PIPELINED;
```

Parameters

Table 190-45 SELECT_SQLPA_TASK Function Parameters

Parameter	Description
`task_name`	Specifies the name of the SQL Performance Analyzer task.
`task_owner`	Specifies the owner of the SQL Performance Analyzer task. If `NULL`, then assume the current user.
`execution_name`	Specifies the name of the SQL Performance Analyzer task execution (type `COMPARE PERFORMANCE`) from which the provided filters will be applied. If `NULL`, then assume the most recent `COMPARE PERFORMANCE` execution.

Table 190-45 (Cont.) SELECT_SQLPA_TASK Function Parameters

Parameter	Description
level_filter	Specifies which subset of SQL statements to include. Same format as DBMS_SQLPA.REPORT_ANALYSIS_TASK.LEVEL, with some possible strings removed. • IMPROVED includes only improved SQL. • REGRESSED includes only regressed SQL (default). • CHANGED includes only SQL with changed performance. • UNCHANGED includes only SQL with unchanged performance. • CHANGED_PLANS includes only SQL with plan changes. • UNCHANGED_PLANS includes only SQL with unchanged plans. • ERRORS includes only SQL with errors only. • MISSING_SQL includes only missing SQL statements (across STS). • NEW_SQL includes only new SQL statements (across STS).
basic filter	Specifies the SQL predicate to filter the SQL in addition to the level filters.
object_filter	Currently not supported.
attribute_list	Defines the SQL statement attributes to return in the result. Possible values are: • TYPICAL — Returns BASIC plus the SQL plan (without row source statistics) and without an object reference list. This is the default. • BASIC — Returns all attributes (such as execution statistics and binds) except the plans. The execution context is always part of the result. • ALL — Returns all attributes. • Comma-separated list of attribute names this allows to return only a subset of SQL attributes: EXECUTION_STATISTICS, SQL_BINDS, SQL_PLAN_STATISTICS (similar to SQL_PLAN + row source statistics).

Return Values

This function returns a SQL tuning set object.

Usage Notes

For example, you can use this function to create a SQL tuning set containing the subset of SQL statements that regressed during a SQL Performance Analyzer (SPA) experiment. You can also specify other arbitrary filters.

190.5.40 SELECT_SQLSET Function

This is a table function that reads the contents of a SQL tuning set.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.SELECT_SQLSET (
  sqlset_name          IN   VARCHAR2,
  basic_filter         IN   VARCHAR2 := NULL,
  object_filter        IN   VARCHAR2 := NULL,
  ranking_measure1     IN   VARCHAR2 := NULL,
  ranking_measure2     IN   VARCHAR2 := NULL,
  ranking_measure3     IN   VARCHAR2 := NULL,
  result_percentage    IN   NUMBER   := 1,
  result_limit         IN   NUMBER   := NULL)
  attribute_list       IN   VARCHAR2 := NULL,
  plan_filter          IN   VARCHAR2 := NULL,
  sqlset_owner         IN   VARCHAR2 := NULL,
  recursive_sql        IN   VARCHAR2 := HAS_RECURSIVE_SQL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 190-46 SELECT_SQLSET Function Parameters

Parameter	Description
`sqlset_name`	Specifies the name of the SQL tuning set to query.
`basic_filter`	Specifies the SQL predicate to filter the SQL from the SQL tuning set defined on attributes of the `SQLSET_ROW`.
`object_filter`	Currently not supported.
`ranking_measure(n)`	Specifies an `ORDER BY` clause on the selected SQL.
`result_percentage`	Specifies a filter that picks the top *n*% according to the supplied ranking measure. Note that this parameter applies only if one ranking measure is supplied.
`result_limit`	The top limit SQL from the filtered source, ranked by the ranking measure.
`attribute_list`	Defines the SQL statement attributes to return in the result. The possible values are: • `BASIC` — Returns all attributes (such as execution statistics and binds) except the plans. The execution context is included in the result. • `TYPICAL` — Returns `BASIC` plus the SQL plan, but without row source statistics and without the object reference list. This is the default. • `ALL` — Returns all attributes. • Comma-separated list of attribute names. This value enables the function to return only a subset of SQL attributes: – `EXECUTION_STATISTICS` – `SQL_BINDS` – `SQL_PLAN_STATISTICS` (similar to `SQL_PLAN` plus row source statistics)

Table 190-46 (Cont.) SELECT_SQLSET Function Parameters

Parameter	Description
plan_filter	Specifies the plan filter. This parameter enables you to select a single plan when a statement has multiple plans. Possible values are: • LAST_GENERATED — Returns the plan with the most recent timestamp. • FIRST_GENERATED — Returns the plan with the least recent timestamp. • LAST_LOADED — Returns the plan with the most recent FIRST_LOAD_TIME statistical information. • FIRST_LOADED — Returns the plan with the least recent FIRST_LOAD_TIME statistical information. • MAX_ELAPSED TIME — Returns the plan with the maximum elapsed time. • MAX_BUFFER_GETS — Returns the plan with the maximum buffer gets. • MAX_DISK_READS — Returns the plan with the maximum disk reads. • MAX_DIRECT_WRITES — Returns the plan with the maximum direct writes. • MAX_OPTIMIZER_COST — Returns the plan with the maximum optimizer cost value.
sqlset_owner	Specifies the owner of the SQL tuning set, or NULL for the current schema owner.
recursive_sql	Specifies that the filter must include recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL, which is the default) or exclude it (NO_RECURSIVE_SQL).

Return Values

This function returns one SQLSET_ROW per SQL_ID or PLAN_HASH_VALUE pair found in each data source.

Usage Notes

Filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.

Examples

```
-- select from a sql tuning set
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE (P)
    FROM table(dbms_sqltune.select_sqlset('my_workload')) P;

  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

190.5.41 SELECT_WORKLOAD_REPOSITORY Function

This function collects SQL statements from the workload repository.

The overloaded forms enable you to collect SQL statements from the following sources:

- Snapshots between `begin_snap` and `end_snap`
- A workload repository baseline

Syntax

```
DBMS_SQLTUNE.SELECT_WORKLOAD_REPOSITORY (
  begin_snap        IN NUMBER,
  end_snap          IN NUMBER,
  basic_filter      IN VARCHAR2 := NULL,
  object_filter     IN VARCHAR2 := NULL,
  ranking_measure1  IN VARCHAR2 := NULL,
  ranking_measure2  IN VARCHAR2 := NULL,
  ranking_measure3  IN VARCHAR2 := NULL,
  result_percentage IN NUMBER   := 1,
  result_limit      IN NUMBER   := NULL,
  attribute_list    IN VARCHAR2 := NULL,
  recursive_sql     IN VARCHAR2 := HAS_RECURSIVE_SQL,
  dbid              IN NUMBER   := NULL)
 RETURN sys.sqlset PIPELINED;

DBMS_SQLTUNE.SELECT_WORKLOAD REPOSITORY (
  baseline_name     IN VARCHAR2,
  basic_filter      IN VARCHAR2 := NULL,
  object_filter     IN VARCHAR2 := NULL,
  ranking_measure1  IN VARCHAR2 := NULL,
  ranking_measure2  IN VARCHAR2 := NULL,
  ranking_measure3  IN VARCHAR2 := NULL,
  result_percentage IN NUMBER   := 1,
  result_limit      IN NUMBER   := NULL,
  attribute_list    IN VARCHAR2 := NULL,
  recursive_sql     IN VARCHAR2 := HAS_RECURSIVE_SQL,
  dbid              IN NUMBER   := NULL)
 RETURN sys.sqlset PIPELINED;
```

Parameters

Table 190-47 SELECT_WORKLOAD_REPOSITORY Function Parameters

Parameter	Description
`begin_snap`	Defines the beginning AWR snapshot (non-inclusive).
`end_snap`	Defines the ending AWR snapshot (inclusive).
`baseline_name`	Specifies the name of the AWR baseline period.

Table 190-47 (Cont.) SELECT_WORKLOAD_REPOSITORY Function Parameters

Parameter	Description
basic_filter	Specifies the SQL predicate to filter the SQL from the workload repository. The filter is defined on attributes of the SQLSET_ROW. If basic_filter is not set by the caller, then the subprogram captures only statements of type CREATE TABLE, INSERT, SELECT, UPDATE, DELETE, and MERGE.
object_filter	Currently not supported.
ranking_measure(n)	Defines an ORDER BY clause on the selected SQL.
result_percentage	Specifies a filter that picks the top *n*% according to the supplied ranking measure. Note that this percentage applies only if one ranking measure is given.
result_limit	Specifies the top limit SQL from the source according to the supplied ranking measure.
attribute_list	Specifies the SQL statement attributes to return in the result. The possible values are: • TYPICAL — Returns BASIC plus SQL plan (without row source statistics) and without object reference list. This is the default. • BASIC — Returns all attributes (such as execution statistics and binds) are returned except the plans. The execution context is always part of the result. • ALL — Returns all attributes • Comma-separated list of attribute names this allows to return only a subset of SQL attributes: EXECUTION_STATISTICS, SQL_BINDS, SQL_PLAN_STATISTICS (similar to SQL_PLAN plus row source statistics).
recursive_sql	Specifies the filter that includes recursive SQL in the SQL tuning set (HAS_RECURSIVE_SQL) or excludes it (NO_RECURSIVE_SQL).
dbid	Specifies the DBID for imported or PDB-level AWR data. If NULL, then the function uses the current database DBID.

Return Values

This function returns one SQLSET_ROW per SQL_ID or PLAN_HASH_VALUE pair found in each data source.

Usage Notes

Filters provided to this function are evaluated as part of a SQL run by the current user. As such, they are executed with that user's security privileges and can contain any constructs and subqueries that user can access, but no more.

Examples

```
-- select statements from snapshots 1-2
DECLARE
  cur sys_refcursor;
BEGIN
  OPEN cur FOR
    SELECT VALUE (P)
    FROM table(dbms_sqltune.select_workload_repository(1,2)) P;
```

```
  -- Process each statement (or pass cursor to load_sqlset)

  CLOSE cur;
END;
/
```

190.5.42 SET_TUNING_TASK_PARAMETER Procedures

This procedure updates the value of a SQL tuning parameter of type `VARCHAR2` or `NUMBER`.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.SET_TUNING_TASK_PARAMETER (
   task_name        IN  VARCHAR2,
   parameter        IN  VARCHAR2,
   value            IN  VARCHAR2,
   database_link_to IN  VARCHAR2);

DBMS_SQLTUNE.SET_TUNING_TASK_PARAMETER (
   task_name        IN  VARCHAR2,
   parameter        IN  VARCHAR2,
   value            IN  NUMBER,
   database_link_to IN  VARCHAR2);
);
```

Parameters

Table 190-48 SET_TUNING_TASK_PARAMETER Procedure Parameters

Parameter	Description
`task_name`	Identifier of the task to execute

Table 190-48 (Cont.) SET_TUNING_TASK_PARAMETER Procedure Parameters

Parameter	Description
parameter	Name of the parameter to set. The possible tuning parameters that can be set by this procedure using the parameter in the form VARCHAR2: • APPLY_CAPTURED_COMPILENV: indicates whether the advisor could use the compilation environment captured with the SQL statements. The default is 0 (that is, NO). • BASIC_FILTER: basic filter for SQL tuning set • DAYS_TO_EXPIRE: number of days until the task is deleted • DEFAULT_EXECUTION_TYPE: the task defaults to this type of execution when none is specified by the EXECUTE_TUNING_TASK Function and Procedure • EXECUTION_DAYS_TO_EXPIRE: number of days until the tasks's executions is deleted (without deleting the task) • LOCAL_TIME_LIMIT: per-statement time out (seconds) • MODE: tuning scope (comprehensive, limited) • OBJECT_FILTER: object filter for SQL tuning set • PLAN_FILTER: plan filter for SQL tuning set (see SELECT_SQLSET for possible values) • RANK_MEASURE1: first ranking measure for SQL tuning set • RANK_MEASURE2: second possible ranking measure for SQL tuning set • RANK_MEASURE3: third possible ranking measure for SQL tuning set • RESUME_FILTER: a extra filter for SQL tuning sets besides BASIC_FILTER • SQL_LIMIT: maximum number of SQL statements to tune • SQL_PERCENTAGE: percentage filter of SQL tuning set statements • TEST_EXECUTE: FULL/AUTO/OFF. * FULL - test-execute for as much time as necessary, up to the local time limit for the SQL (or the global task time limit if no SQL time limit is set) * AUTO - test-execute for an automatically-chosen time proportional to the tuning time * OFF - do not test-execute • TIME_LIMIT: global time out (seconds) • USERNAME: username under which the statement is parsed
value	New value of the specified parameter

Table 190-48 (Cont.) SET_TUNING_TASK_PARAMETER Procedure Parameters

Parameter	Description
database_link_to	Name of a database link that exists on a standby database. The link specifies the connection to a primary database. By default, the value is null, which means that the SQL Tuning Advisor session is local. Use DBMS_SQLTUNE to tune high-load SQL statements running on a standby database in an Active Data Guard scenario. When you execute REPORT_TUNING_TASK locally on the standby database, the function uses the database link to obtain the data from the primary database, and then constructs it locally on the standby database. The database_link_to parameter must specify a private database link. This link must be owned by SYS and accessed by the default privileged user SYS$UMF. The following sample statement creates a link named lnk_to_pri: `CREATE DATABASE LINK lnk_to_pri CONNECT TO SYS$UMF IDENTIFIED BY password USING 'inst1';`

Usage Notes

When setting automatic tuning task parameters, use the SET_AUTO_TUNING_TASK_PARAMETER Procedures in the DBMS_AUTO_SQLTUNE package.

190.5.43 SQLTEXT_TO_SIGNATURE Function

This function returns a SQL text's signature. The signature can be used to identify SQL text in dba_sql_profiles.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.SQLTEXT_TO_SIGNATURE (
  sql_text     IN CLOB,
  force_match IN BOOLEAN  := FALSE)
RETURN NUMBER;
```

Parameters

Table 190-49 SQLTEXT_TO_SIGNATURE Function Parameters

Parameter	Description
sql_text	SQL text whose signature is required. Required.

Table 190-49 (Cont.) SQLTEXT_TO_SIGNATURE Function Parameters

Parameter	Description
force_match	If TRUE, this returns a signature that supports SQL matching with literal values transformed into bind variables. If FALSE, returns the signature based on the text with literals not transformed

Return Values

This function returns the signature of the specified SQL text.

190.5.44 UNPACK_STGTAB_SQLPROF Procedure

This procedure copies profile data stored in the staging table to create profiles on the system.

> **See Also:**
>
> DBMS_SQLTUNE SQL Profile Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.UNPACK_STGTAB_SQLPROF (
   profile_name          IN VARCHAR2 := '%',
   profile_category      IN VARCHAR2 := 'DEFAULT',
   replace               IN BOOLEAN,
   staging_table_name    IN VARCHAR2,
   staging_schema_owner  IN VARCHAR2 := NULL);
```

Parameters

Table 190-50 UNPACK_STGTAB_SQLPROF Procedure Parameters

Parameter	Description
profile_name	The name of the profile to unpack (% wildcards acceptable, case-sensitive)
profile_category	The category from which to unpack profiles (% wildcards acceptable, case-sensitive)
replace	The option to replace profiles if they already exist. Note that profiles cannot be replaced if one in the staging table has the same name as an active profile in a different SQL statement. If FALSE, this function raises errors if you try to create a profile that already exists
staging_table_name	The name of the table on which to perform the remap operation (case-insensitive unless double quoted). Required.
staging_schema_owner	The schema where the table resides, or NULL for current schema (case-insensitive unless double quoted)

Usage Notes

Using this procedure requires the `CREATE ANY SQL PROFILE` privilege and the `SELECT` privilege on staging table.

Examples

```
-- Unpack all profiles stored in a staging table.
BEGIN
  DBMS_SQLTUNE.UNPACK_STGTAB_SQLPROF(
      replace            => FALSE
  ,   staging_table_name => 'PROFILE_STGTAB');
END;

-- If there is a failure during the unpack operation, you can find the profile
-- that caused the error and perform a remap_stgtab_sqlprof operation
targeting it.
-- You can resume the unpack operation by setting replace to TRUE so that
-- the profiles that were already created are replaced.
BEGIN
  DBMS_SQLTUNE.UNPACK_STGTAB_SQLPROF(
      replace            => TRUE
  ,   staging_table_name => 'PROFILE_STGTAB');
END;
```

190.5.45 UNPACK_STGTAB_SQLSET Procedure

This procedure copies one or more SQL tuning sets from their location in the staging table into the SQL tuning sets schema, making them proper SQL tuning sets.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.UNPACK_STGTAB_SQLSET (
   sqlset_name          IN VARCHAR2 := '%',
   sqlset_owner         IN VARCHAR2 := NULL,
   replace              IN BOOLEAN,
   staging_table_name   IN VARCHAR2,
   staging_schema_owner IN VARCHAR2 := NULL);
```

Parameters

The parameters are identical for `DBMS_SQLTUNE.UNPACK_STGTAB_SQLSET` and `DBMS_SQLSET.UNPACK_STGTAB`.

Table 190-51 UNPACK_STGTAB_SQLSET and UNPACK_STGTAB Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the tuning set to unpack (not null). Wildcard characters (%) are supported to unpack multiple tuning sets in a single call. For example, specify % to unpack all tuning sets from the staging table.
sqlset_owner	Specifies the name of tuning set owner, or NULL for the current schema owner. Wildcard characters (%) are supported.
replace	Specifies whether to replace an existing SQL tuning set. If FALSE, then this procedure raises errors when you try to create a tuning set that already exists.
staging_table_name	Specifies the name of the staging table, moved after a call to the DBMS_SQLTUNE.PACK_STGTAB_SQLSET or DBMS_SQLSET.PACK_STGTAB procedure (case-sensitive).
staging_schema_owner	Specifies the name of staging table owner, or NULL for the current schema owner (case-sensitive).

Examples

```
 -- unpack all STS in the staging table
EXEC DBMS_SQLTUNE.UNPACK_STGTAB_SQLSET(sqlset_name         => '%', -
                                       sqlset_owner        => '%', -
                                       replace             => FALSE, -
                                       staging_table_name  =>
'STGTAB_SQLSET');

-- errors can arise during STS unpack when a STS in the staging table has the
-- same name/owner as STS on the system.  In this case, users should call
-- remap_stgtab_sqlset to patch the staging table and with which to call
unpack
-- Replace set to TRUE.
EXEC DBMS_SQLTUNE.UNPACK_STGTAB_SQLSET(sqlset_name         => '%', -
                                       sqlset_owner        => '%', -
                                       replace             => TRUE, -
                                       staging_table_name  =>
'STGTAB_SQLSET');
```

190.5.46 UPDATE_SQLSET Procedures

This overloaded procedure updates selected fields for SQL statements in a SQL tuning set.

> **See Also:**
>
> DBMS_SQLTUNE SQL Tuning Set Subprograms for other subprograms in this group

Syntax

```
DBMS_SQLTUNE.UPDATE_SQLSET (
   sqlset_name      IN  VARCHAR2,
   sql_id           IN  VARCHAR2,
   attribute_name   IN  VARCHAR2,
   attribute_value  IN  VARCHAR2 := NULL);

DBMS_SQLTUNE.UPDATE_SQLSET (
   sqlset_name      IN  VARCHAR2,
   sql_id           IN  VARCHAR2,
   attribute_name   IN  VARCHAR2,
   attribute_value IN NUMBER := NULL);
```

Parameters

Table 190-52 UPDATE_SQLSET Procedure Parameters

Parameter	Description
sqlset_name	Specifies the name of the SQL tuning set.
sql_id	Specifies the identifier of the SQL statement to be updated.
plan_hash value	Specifies the hash value of the execution plan for a SQL statement. Use this parameter when you want to update the attribute for a specific plan for a statement, but not all plans for the statement.
attribute_name	Specifies the name of the attribute to be modified. You can update the text field for MODULE, ACTION, PARSING_SCHEMA_NAME, and OTHER. The only numerical field that you can update is PRIORITY. If a statement has multiple plans, then the procedure changes the attribute value for all plans.
attribute_value	Specifies the new value of the attribute.

191
DBMS_STAT_FUNCS

The `DBMS_STAT_FUNCS` package provides statistical functions.

This chapter contains the following topic:

- Summary of DBMS_STAT_FUNCS Subprograms

191.1 Summary of DBMS_STAT_FUNCS Subprograms

This table lists and briefly describes the DBMS_STAT_FUNCS subprograms.

Table 191-1 DBMS_STAT_FUNCS Package Subprograms

Subprogram	Description
EXPONENTIAL_DIST_FIT Procedure	Tests how well a sample of values fits an exponential distribution
NORMAL_DIST_FIT Procedure	Tests how well a sample of values fits a normal distribution
POISSON_DIST_FIT Procedure	Tests how well a sample of values fits a Poisson distribution
SUMMARY Procedure	Summarizes a numerical column of a table
UNIFORM_DIST_FIT Procedure	Tests how well a sample of values fits a uniform distribution
WEIBULL_DIST_FIT Procedure	Tests how well a sample of values fits a Weibull distribution

191.1.1 EXPONENTIAL_DIST_FIT Procedure

This procedure tests how well a sample of values fits an exponential distribution.

Syntax

```
DBMS_STAT_FUNCS.EXPONENTIAL_DIST_FIT (
   ownername    IN    VARCHAR2,
   tablename    IN    VARCHAR2,
   columnname   IN    VARCHAR2,
   test_type    IN    VARCHAR2 DEFAULT 'KOLMOGOROV_SMIRNOV',
   lambda       IN    NUMBER,
   mu           IN    NUMBER,
   sig          OUT   NUMBER);
```

Parameters

Table 191-2 *EXPONENTIAL_DIST_FIT Procedure Parameters*

Parameter	Description
`ownername`	The schema where the table resides.
`tablename`	The table where the column resides.
`columnname`	The column of the table against which to run the test.

Table 191-2 (Cont.) *EXPONENTIAL_DIST_FIT Procedure Parameters*

Parameter	Description
test_type	The type of test to use: 'CHI_SQUARED', 'KOLMOGOROV_SMIRNOV' or 'ANDERSON_DARLING'.
lambda	The scale parameter.
mu	The location parameter.
sig	The goodness of fit value, based on test type. A small value indicates a significant difference between the sample and the exponential distribution. A number close to 1 indicates a close match.

191.1.2 NORMAL_DIST_FIT Procedure

This procedure tests how well a sample of values fits a normal distribution.

Syntax

```
DBMS_STAT_FUNCS.NORMAL_DIST_FIT (
   ownername    IN    VARCHAR2,
   tablename    IN    VARCHAR2,
   columnname   IN    VARCHAR2,
   test_type    IN    VARCHAR2 DEFAULT 'SHAPIRO_WILKS',
   mean         IN    NUMBER,
   stdev        IN    NUMBER,
   sig          OUT   NUMBER);
```

Parameters

Table 191-3 NORMAL_DIST_FIT Procedure Parameters

Parameter	Description
ownername	The schema where the table resides.
tablename	The table where the column resides.
columnname	The column of the table against which to run the test.
test_type	The type of test to use: 'CHI_SQUARED', 'KOLMOGOROV_SMIRNOV', 'ANDERSON_DARLING' or 'SHAPIRO_WILKS'.
mean	The mean of the distribution against which to compare.
stdev	The standard deviation of the distribution against which to compare.
sig	The goodness of fit value, based on test type. A small value indicates a significant difference between the sample and the normal distribution. A number close to 1 indicates a close match.

191.1.3 POISSON_DIST_FIT Procedure

This procedure tests how well a sample of values fits a Poisson distribution.

Syntax

```
DBMS_STAT_FUNCS.POISSON_DIST_FIT (
   ownername    IN    VARCHAR2,
```

```
   tablename    IN    VARCHAR2,
   columnname   IN    VARCHAR2,
   test_type    IN    VARCHAR2 DEFAULT 'KOLMOGOROV_SMIRNOV',
   lambda       IN    NUMBER,
   sig          OUT   NUMBER);
```

Parameters

Table 191-4 POISSON_DIST_FIT Procedure Parameters

Parameter	Description
ownername	The schema where the table resides.
tablename	The table where the column resides.
columnname	The column of the table against which to run the test.
test_type	The type of test to use: 'KOLMOGOROV_SMIRNOV' or 'ANDERSON_DARLING'.
lambda	The lambda parameter is the shape parameter.
sig	The goodness of fit value, based on test type. A small value indicates a significant difference between the sample and the Poisson distribution. A number close to 1 indicates a close match.

191.1.4 SUMMARY Procedure

This procedure summarizes the numerical column specified in the `columnname` of `tablename`.

The summary is returned as a Summary Type. Note that most of the output of `SUMMARY` can be obtained with currently available SQL.

Syntax

```
DBMS_STAT_FUNCS.SUMMARY (
   ownername    IN    VARCHAR2,
   tablename    IN    VARCHAR2,
   columnname   IN    VARCHAR2,
   sigma_value  IN    NUMBER DEFAULT 3,
   s           OUT    SummaryType);
```

Parameters

Table 191-5 SUMMARY Procedure Parameters

Parameter	Description
ownername	The schema where the table resides.
tablename	The table where the column resides.
columnname	The column of the table to be summarized.
sigma_value	The number of sigmas for the set of extreme values, defaults to 3.
s	The Record containing summary information about given column.

Definition of SummaryType

```
TYPE n_arr IS VARRAY(5) of NUMBER;
TYPE num_table IS TABLE of NUMBER;
```

```
TYPE summaryType IS RECORD (
   count             NUMBER,
   min               NUMBER,
   max               NUMBER,
   range             NUMBER,
   mean              NUMBER,
   cmode             num_table,
   variance          NUMBER,
   stddev            NUMBER,
   quantile_5        NUMBER,
   quantile_25       NUMBER,
   median            NUMBER,
   quantile_75       NUMBER,
   quantile_95       NUMBER,
   plus_x_sigma      NUMBER,
   minus_x_sigma     NUMBER,
   extreme_values    num_table,
   top_5_values      n_arr,
   bottom_5_values   n_arr);
```

191.1.5 UNIFORM_DIST_FIT Procedure

This procedure tests well a sample of values fits a uniform distribution.

Syntax

```
DBMS_STAT_FUNCS.UNIFORM_DIST_FIT (
   ownername    IN    VARCHAR2,
   tablename    IN    VARCHAR2,
   columnname   IN    VARCHAR2,
   var_type     IN    VARCHAR2 DEFAULT 'CONTINUOUS',
   test_type    IN    VARCHAR2 DEFAULT 'KOLMOGOROV_SMIRNOV',
   paramA       IN    NUMBER,
   paramB       IN    NUMBER,
   sig          OUT   NUMBER);
```

Parameters

Table 191-6 UNIFORM_DIST_FIT Procedure Parameters

Parameter	Description
ownername	The schema where the table resides.
tablename	The table where the column resides.
columnname	The column of the table against which to run the test.
var_type	The type of distribution: 'CONTINUOUS' (the default) or 'DISCRETE'
test_type	The type of test to use: 'CHI_SQUARED', 'KOLMOGOROV_SMIRNOV' or 'ANDERSON_DARLING'.
paramA	Parameter A estimated from the sample (the location parameter).
paramB	Parameter B estimated from the sample (the scale parameter).
sig	The goodness of fit value, based on test type. A small value indicates a significant difference between the sample and the uniform distribution. A number close to 1 indicates a close match.

191.1.6 WEIBULL_DIST_FIT Procedure

This procedure tests how well a sample of values fits a Weibull distribution.

Syntax

```
DBMS_STAT_FUNCS.WEIBULL_DIST_FIT (
   ownername    IN    VARCHAR2,
   tablename    IN    VARCHAR2,
   columnname   IN    VARCHAR2,
   test_type    IN    VARCHAR2 DEFAULT 'KOLMOGOROV_SMIRNOV',
   alpha        IN    NUMBER,
   mu           IN    NUMBER,
   beta         IN    NUMBER,
   sig          OUT   NUMBER);
```

Parameters

Table 191-7 WEIBULL_DIST_FIT Procedure Parameters

Parameter	Description
ownername	The schema where the table resides.
tablename	The table where the column resides.
columnname	The column of the table against which to run the test.
test_type	The type of test to use: 'CHI_SQUARED', 'KOLMOGOROV_SMIRNOV' or 'ANDERSON_DARLING'.
alpha	The scale parameter.
mu	The location parameter.
beta	The slope/shape parameter.
sig	The goodness of fit value, based on test type. A small value indicates a significant difference between the sample and the Weibull distribution. A number close to 1 indicates a close match.

192
DBMS_STATS

With the `DBMS_STATS` package you can view and modify optimizer statistics gathered for database objects.

Users can also collect statistics on Global Temporary Tables (GTTs) using the DBMS_STATS package. However, `DBMS_STATS` cannot collect statistics on Private Temporary Tables (PTTs).

This chapter contains the following topics:

- DBMS_STATS Overview
- DBMS_STATS Deprecated Subprograms
- DBMS_STATS Types
- DBMS_STATS Constants
- DBMS_STATS Operational Notes
- DBMS_STATS Data Structures
- Summary of DBMS_STATS Subprograms

> **See Also:**
>
> - *Oracle Database SQL Tuning Guide*
> - *Oracle Database Administrator's Guide*

192.1 DBMS_STATS Overview

To improve performance, the database enables you to collect optimizer statistics.

> **Note:**
>
> By default, the database collects statistics automatically, so this package is intended only for specialized cases.

Optimizer statistics can reside in the data dictionary or in a table created in the user's schema. You can also collect and manage user-defined statistics for tables and domain indexes using this package. For example, if you invoke the `DELETE_COLUMN_STATS` procedure on a column for which an association is defined, the database deletes both user-defined and standard statistics for this column.

Only optimizer statistics stored in the data dictionary have an effect on the cost-based optimizer. You can also use `DBMS_STATS` to gather statistics in parallel.

Optimizer Statistics Advisor inspects the statistics gathering process, automatically diagnoses problems in the existing practices for gathering statistics, and then generates a report of findings and recommendations. The advisor task runs automatically in the maintenance window. However, you can also run the job on demand.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to manage optimizer statistics

192.2 DBMS_STATS Deprecated Subprograms

Oracle recommends that you do not use deprecated subprograms. Support for deprecated features is for backward compatibility only.

Starting with Oracle Database 11*g* Release 2 (11.2), the following subprograms are obsolete:

- GET_PARAM Function

  Instead, use GET_PREFS Function
- SET_PARAM Procedure

  Instead, use SET_GLOBAL_PREFS Procedure
- RESET_PARAM_DEFAULTS Procedure

  Instead use RESET_GLOBAL_PLSQL_PREFS_DEF Procedure

The following subprogram is deprecated with Oracle Database 12*c* and later:

- `GENERATE_STATS`

  This procedure is replaced by the `GATHER_INDEX_STAT` procedure.

> **See Also:**
>
> "GATHER_INDEX_STATS Procedure"

192.3 DBMS_STATS Types

Histograms

Types for the minimum and maximum values and histogram endpoints include the following:

```
TYPE numarray  IS VARRAY(2050) OF NUMBER;
TYPE datearray IS VARRAY(2050) OF DATE;
TYPE chararray IS VARRAY(2050) OF VARCHAR2(4000);
TYPE rawarray  IS VARRAY(2050) OF RAW(2000);
TYPE fltarray  IS VARRAY(2050) OF BINARY_FLOAT;
TYPE dblarray  IS VARRAY(2050) OF BINARY_DOUBLE;
```

Stale Tables

Types for listing stale tables include the following:

```
TYPE ObjectElem IS RECORD (
  ownname     VARCHAR2(30),     -- owner
  objtype     VARCHAR2(6),      -- 'TABLE' or 'INDEX'
  objname     VARCHAR2(30),     -- table/index
  partname    VARCHAR2(30),     -- partition
  subpartname VARCHAR2(30));    -- subpartition
TYPE ObjectTab IS TABLE OF ObjectElem;
```

Statistics Difference Reports

Use the following type to displays a statistics difference report:

```
TYPE DiffRepElem IS RECORD (
  report      CLOB,            -- stats difference report
  maxdiffpct  NUMBER);         -- max stats difference (percentage)
TYPE DiffRepTab IS TABLE OF DiffRepElem;
```

Optimizer Statistics Advisor

The following type represents database objects for which you can gather statistics:

```
TYPE ObjectElem IS RECORD (
  ownname     dbms_quoted_id, -- owner
  objtype     VARCHAR2(6), -- 'TABLE' or 'INDEX'
  objname     dbms_quoted_id, -- table/index
  partname    dbms_quoted_id, -- partition
  subpartname dbms_quoted_id -- subpartition
);
TYPE ObjectTab IS TABLE OF ObjectElem;
```

> **Note:**
>
> Make sure to maintain `satisfy_obj_filter` when the `ObjectElem` type is changed

The following type represents an operation:

```
TYPE StatsAdvOpr IS RECORD (
  name  VARCHAR2(64), -- name of the operation
  param VARCHAR2(4000)
); -- XML containing parameters and their values
TYPE StatsAdvOprTab IS TABLE OF StatsAdvOpr;
```

The following type represents a filter list:

```
TYPE StatsAdvFilter IS RECORD (
  rulename VARCHAR2(64),   -- rule name
  objlist  ObjectTab,      -- object filter list
  oprlist  StatsAdvOprTab, -- operation filter list
  include  BOOLEAN);       -- include/exclude elements in the list
TYPE StatsAdvFilterTab IS TABLE OF StatsAdvFilter;
```

192.4 DBMS_STATS Constants

The DBMS_STATS package defines several constants to use specifying parameter values.

Table 192-1 DBMS_STATS Constants

Name	Type	Description
ADD_GLOBAL_PREFS	NUMBER	Copies global preferences
AUTO_CASCADE	BOOLEAN	Lets Oracle decide whether to collect statistics for indexes or not
AUTO_DEGREE	NUMBER	Lets Oracle select the degree of parallelism based on size of the object, number of CPUs and initialization parameters
AUTO_INVALIDATE	BOOLEAN	Lets Oracle decide when to invalidate dependent cursors
AUTO_SAMPLE_SIZE	NUMBER	Indicates that auto-sample size algorithms should be used
PURGE_ALL	TIMESTAMP WITH TIME ZONE	A flag that can be passed to the PURGE_STATS Procedure and unconditionally deletes all the history statistics. The deletion uses TRUNCATE statements on the various dictionary statistics tables holding the history of statistics.
RECLAIM_SYNOPSIS	TIMESTAMP WITH TIME ZONE	A constant used for reclaiming synopsis table space.

192.5 DBMS_STATS Operational Notes

Observe these operational notes when using the DBMS_STATS package.

The DBMS_STATS subprograms perform the following general operations:

- Gathering Optimizer Statistics
- Setting or Getting Statistics
- Deleting Statistics
- Transferring Statistics
- Locking or Unlocking Statistics
- Restoring and Purging Statistics History
- User-Defined Statistics
- Pending Statistics
- Comparing Statistics
- Extended Statistics
- Optimizer Statistics Advisor

Most of the DBMS_STATS procedures include the three parameters statown, stattab, and statid. These parameters enable you to store statistics in your own tables (outside of the dictionary), which does not affect the optimizer. Therefore, you can maintain and experiment with *sets* of statistics.

The stattab parameter specifies the name of a table in which to hold statistics, and it is assumed that it resides in the same schema as the object for which statistics are collected (unless the statown parameter is specified). You can create multiple tables with different stattab identifiers to hold separate sets of statistics.

Additionally, you can maintain different sets of statistics within a single stattab by using the statid parameter, which avoids cluttering the user's schema.

For the SET and GET procedures, if stattab is not provided (that is, NULL), then the operation works directly on the dictionary statistics; therefore, you do not need to create these statistics tables if they only plan to modify the dictionary directly. However, if stattab is not NULL, then the SET or GET operation works on the specified user statistics table, and not the dictionary.

You can change the default values of some of the parameters of DBMS_STATS procedures using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Most procedures in this package commit the current transaction, perform the operation, and then commit again.

Most of the procedures have a force parameter that enables you to override a lock on statistics. Whenever statistics in dictionary are modified, old versions of statistics are saved automatically for future restoring.

Gathering Optimizer Statistics

Use the following subprograms to gather certain classes of optimizer statistics, with possible performance improvements over the ANALYZE command:

- GATHER_DATABASE_STATS Procedures
- GATHER_DICTIONARY_STATS Procedure
- GATHER_FIXED_OBJECTS_STATS Procedure
- GATHER_INDEX_STATS Procedure
- GATHER_SCHEMA_STATS Procedures
- GATHER_SYSTEM_STATS Procedure
- GATHER_TABLE_STATS Procedure

The GATHER_* procedures also collect user-defined statistics for columns and domain indexes.

The statown, stattab, and statid parameters instruct the package to back up current statistics in the specified table before gathering new statistics.

Oracle also provides the following procedure for generating statistics for derived objects when you have sufficient statistics on related objects:

GENERATE_STATS Procedure

Setting or Getting Statistics

Use the following subprograms to store and retrieve individual column-related, index-related, and table-related statistics:

- PREPARE_COLUMN_VALUES Procedures
- PREPARE_COLUMN_VALUES_NVARCHAR Procedure

PREPARE_COLUMN_VALUES_ROWID Procedure

SEED_COL_USAGE Procedure
SET_INDEX_STATS Procedures
SET_SYSTEM_STATS Procedure
SET_TABLE_STATS Procedure

GET_COLUMN_STATS Procedures
GET_INDEX_STATS Procedures
GET_SYSTEM_STATS Procedure
GET_TABLE_STATS Procedure

In the special versions of the `SET_*_STATS` procedures for setting user-defined statistics, the following, if provided, are stored in the dictionary or user statistics table:

- User-defined statistics
- Owner of statistics type
- Name of statistics type

The user-defined statistics and the corresponding statistics type are inserted into the `USTATS$` dictionary table. You can specify user-defined statistics without specifying the statistics type name.

The special versions of the `GET_*_STATS` procedures return user-defined statistics and the statistics type owner and name as `OUT` arguments corresponding to the schema object specified. If user-defined statistics are not collected, `NULL` values are returned.

Deleting Statistics

The `DELETE_*` procedures delete both user-defined statistics and the standard statistics for the given schema object.

DELETE_COLUMN_STATS Procedure
DELETE_DATABASE_STATS Procedure
DELETE_DICTIONARY_STATS Procedure
DELETE_FIXED_OBJECTS_STATS Procedure
DELETE_INDEX_STATS Procedure
DELETE_SCHEMA_STATS Procedure
DELETE_SYSTEM_STATS Procedure
DELETE_TABLE_STATS Procedure
Note that `DELETE_TABLE_STATS`, `DELETE_DICTIONARY_STATS`, `DELETE_DATABASE_STATS` and `DELETE_SCHEMA_STATS` have a parameter `stat_category` which specifies which statistics to delete. The parameter accepts multiple values separated by comma. The supported values are `'OBJECT_STATS'` (table statistics, column statistics and index statistics) and `'SYNOPSES'` (auxiliary statistics created when statistics are incrementally maintained). The default is `'OBJECT_STATS, SYNOPSES'`.

Transferring Statistics

Use the following procedures for creating and dropping the user statistics table.

CREATE_STAT_TABLE Procedure
DROP_STAT_TABLE Procedure

Use the following procedures to transfer statistics

- from the dictionary to a user statistics table (`EXPORT_*`)

- from a user statistics table to the dictionary (`IMPORT_*`)

EXPORT_COLUMN_STATS Procedure
EXPORT_DATABASE_STATS Procedure
EXPORT_DICTIONARY_STATS Procedure
EXPORT_FIXED_OBJECTS_STATS Procedure
EXPORT_INDEX_STATS Procedure
EXPORT_SCHEMA_STATS Procedure
EXPORT_SYSTEM_STATS Procedure
EXPORT_TABLE_STATS Procedure

IMPORT_COLUMN_STATS Procedure
IMPORT_DATABASE_STATS Procedure
IMPORT_DICTIONARY_STATS Procedure
IMPORT_FIXED_OBJECTS_STATS Procedure
IMPORT_INDEX_STATS Procedure
IMPORT_SCHEMA_STATS Procedure
IMPORT_SYSTEM_STATS Procedure
IMPORT_TABLE_STATS Procedure

> **Note:**
>
> Oracle does not support export or import of statistics across databases of different character sets.

Locking or Unlocking Statistics

Use the following procedures to lock and unlock statistics on objects.

LOCK_PARTITION_STATS Procedure
LOCK_SCHEMA_STATS Procedure
LOCK_TABLE_STATS Procedure

UNLOCK_PARTITION_STATS Procedure
UNLOCK_SCHEMA_STATS Procedure
UNLOCK_TABLE_STATS Procedure

The `LOCK_*` procedures either freeze the current set of the statistics or to keep the statistics untouched.When statistics on a table are locked, all the statistics depending on the table, including table statistics, column statistics, histograms and statistics on all dependent indexes, are considered to be locked.

Restoring and Purging Statistics History

Use the following procedures to restore statistics as of a specified timestamp. This is useful in case newly collected statistics leads to some sub-optimal execution plans and the administrator wants to revert to the previous set of statistics.

RESET_GLOBAL_PLSQL_PREFS_DEF Procedure
RESTORE_DATABASE_STATS Procedure
RESTORE_DICTIONARY_STATS Procedure
RESTORE_FIXED_OBJECTS_STATS Procedure
RESTORE_SCHEMA_STATS Procedure
RESTORE_SYSTEM_STATS Procedure

RESTORE_TABLE_STATS Procedure

Whenever statistics in dictionary are modified, old versions of statistics are saved automatically for future restoring. The old statistics are purged automatically at regular intervals based on the statistics history retention setting and the time of recent statistics gathering performed in the system. Retention is configurable using the ALTER_STATS_HISTORY_RETENTION Procedure.

The other `DBMS_STATS` procedures related to restoring statistics are:

- PURGE_STATS Procedure: This procedure lets you manually purge old versions beyond a time stamp.
- GET_STATS_HISTORY_RETENTION Function: This function gets the current statistics history retention value.
- GET_STATS_HISTORY_AVAILABILITY Function: This function gets the oldest time stamp where statistics history is available. Users cannot restore statistics to a time stamp older than the oldest time stamp.

`RESTORE_*` operations are not supported for user defined statistics.

User-Defined Statistics

The `DBMS_STATS` package supports operations on user-defined statistics. When a domain index or column is associated with a statistics type (using the `associate` statement), operations on the index or column manipulate user-defined statistics. For example, gathering statistics for a domain index (for which an association with a statistics type exists) using the GET_INDEX_STATS Procedures invokes the user-defined statistics collection method of the associated statistics type. Similarly, delete, transfer, import, and export operations manipulate user-defined statistics.

`SET_*` and `GET_*` operations for user-defined statistics are also supported using a special version of the `SET` and `GET` interfaces for columns and indexes.

`EXPORT_*`, `IMPORT_*` and `RESTORE_*` operations are not supported for user defined statistics.

Pending Statistics

The package gathers statistics and stores it in the dictionary by default. User's can store these statistics in the system's private area instead of the dictionary by turning the `PUBLISH` option to `FALSE` using the `SET*PREFS` procedures. The default value for `PUBLISH` is `TRUE`.The statistics stored in private area are not used by Cost Based Optimizer unless parameter `optimizer_use_pending_statistics` is set to `TRUE`. The default value of this parameter is `FALSE` and this boolean parameter can be set at the session/system level. Users can verify the impact of the new statistics on query plans by using the pending statistics on a session.

Pending statistics provide a mechanism to verify the impact of the new statistics on query plans before making them available for general use. There are two scenarios to verify the query plans:

- Export the pending statistics (use the EXPORT_PENDING_STATS Procedure) to a test system, then run the query workload and check the performance or plans.
- Set `optimizer_use_pending_statistics` to `TRUE` in a session on the system where pending statistics have been gathered, run the workload, and check the performance or plans.

After the performance or query plans have been verified, you can publish the pending statistics using the PUBLISH_PENDING_STATS Procedure if the performance is acceptable, or delete the pending statistics using DELETE_PENDING_STATS Procedure if it is not.

Pending statistics can be published, exported, or deleted. The following procedures are provided to manage pending statistics:

- DELETE_PENDING_STATS Procedure
- EXPORT_PENDING_STATS Procedure
- PUBLISH_PENDING_STATS Procedure

Comparing Statistics

You can use the `DIFF_TABLE_STATS_*` statistics to compare statistics for a table from two different sources. The statistics can be from:

- Two different user statistics tables
- A single user statistics table containing two sets of statistics that can be identified using `statids`
- A user statistics table and dictionary history
- Pending statistics

The functions also compare the statistics of the dependent objects (indexes, columns, partitions). They display statistics of the objects from both sources if the difference between those statistics exceeds a certain threshold. The threshold can be specified as an argument to the function, with a default of 10%. The database uses the statistics corresponding to the first source (`stattab1` or `time1`) as a basis for computing the differential percentage.

Extended Statistics

This package enables you to collect statistics for column groups and expressions. The statistics collected for column groups and expressions are called "extended statistics".

Statistics on column groups are used by optimizer for accounting correlation between columns. For example, if a query has predicates c1=1 and c2=1 and if there are statistics on column group (c1, c2), the optimizer uses these statistics for estimating the combined cardinality of the predicates. The optimizer uses the expression statistics to estimate cardinality of predicates on those expressions. The extended statistics are similar to column statistics. The procedures that take columns names accept extended statistics names in place of column names.

Related subprograms:

- CREATE_EXTENDED_STATS Function
- DROP_EXTENDED_STATS Procedure
- SHOW_EXTENDED_STATS_NAME Function
- REPORT_COL_USAGE Function
- SEED_COL_USAGE Procedure
- RESET_COL_USAGE Procedure

Optimizer Statistics Advisor

Optimizer Statistics Advisor is built-in diagnostic software that helps use to use best practices to manage optimizer statistics. The advisor analyzes how you are currently gathering statistics (using manual scripts, explicitly setting parameters, and so on), the effectiveness of existing statistics gathering jobs, and the quality of the gathered statistics. The advisor generates findings for any issues it finds. Based on these findings, the advisor provides recommendations, which it stores in `DBA_ADVISOR_RECOMMENDATIONS`.

The advisor organizes rules into the following classes:

- System

  This class checks the preferences for statistics collection, status of the automated statistics gathering job, use of SQL plan directives, and so on. Rules in this class have the value `SYSTEM` in `V$STATS_ADVISOR_RULES.RULE_TYPE`.

- Operation

  This class checks whether statistics collection uses the defaults, test statistics are created using the `SET_*_STATS` procedures, and so on. Rules in this class have the value `OPERATION` in `V$STATS_ADVISOR_RULES.RULE_TYPE`.

- Object

  This class checks for the quality of the statistics, staleness of statistics, unnecessary collection of statistics, and so on. Rules in this class have the value `OBJECT` in `V$STATS_ADVISOR_RULES.RULE_TYPE`.

All Optimizer Statistics Advisor subprograms require the `ADVISOR` privilege. All procedures and functions execute using the invoker's privilege for the operation instead of the task owner's privilege. For example, if a user without the `ANALYZE ANY DICTIONARY` privilege creates a task `t1`, and if a DBA then executes this task, then the task execution checks for `SYS` objects. Another example is a task that is executed by `user1`, interrupted, and then resumed by `user2`. In this case, the checks of the resumed execution are based on the privilege of `user2` rather than `user1`.

You can use the following subprograms to manage Optimizer Statistics Advisor:

- CANCEL_ADVISOR_TASK Procedure
- CONFIGURE_ADVISOR_FILTER Function
- CONFIGURE_ADVISOR_OBJ_FILTER Function
- CONFIGURE_ADVISOR_OPR_FILTER Functions
- CONFIGURE_ADVISOR_RULE_FILTER Function
- CREATE_ADVISOR_TASK Function
- DROP_ADVISOR_TASK Procedure
- EXECUTE_ADVISOR_TASK Function
- GET_ADVISOR_OPR_FILTER Procedure
- GET_ADVISOR_RECS Function
- IMPLEMENT_ADVISOR_TASK Function
- INTERRUPT_ADVISOR_TASK Procedure
- REPORT_ADVISOR_TASK Function
- RESET_ADVISOR_TASK Procedure
- RESUME_ADVISOR_TASK Procedure
- SCRIPT_ADVISOR_TASK Function
- SET_ADVISOR_TASK_PARAMETER Procedure

> **See Also:**
>
> *Oracle Database SQL Tuning Guide*to learn how to analyze statistics using Optimizer Statistics Advisor

192.6 DBMS_STATS Data Structures

The `DBMS_STATS` package defines a `RECORD` type.

RECORD Types

- STAT_REC Record Type

192.6.1 DBMS_STATS STAT_REC Record Type

This record type is provided for users in case they want to set column statistics manually. Its fields allow specifying column min/max values, as well as a histogram for a column.

Syntax

```
TYPE STATREC IS RECORD (
  epc    NUMBER,
  minval RAW(2000),
  maxval RAW(2000),
  bkvals NUMARRAY,
  novals NUMARRAY,
  chvals CHARARRAY,
  eavals RAWARRAY,
  rpcnts NUMARRAY,
  eavs   NUMBER);
```

Fields of the Record type COMPARISON_TYPE (STAT_REC Attributes)

Table 192-2 STAT_REC Attributes

Field	Description
`epc`	Number of buckets in histogram
`minval`	Minimum value
`maxval`	Maximum value
`bkvals`	Array of bucket numbers
`novals`	Array of normalized end point values
`chvals`	Array of dumped end point values
`eavals`	Array of end point actual values
`rpcnts`	Array of end point value frequencies
`eavs`	A number indicating whether actual end point values are needed in the histogram. If using the PREPARE_COLUMN_VALUES Procedures, this field will be automatically filled.

192.7 Summary of DBMS_STATS Subprograms

This table lists the `DBMS_STATS` subprograms and briefly describes them.

Table 192-3 ***DBMS_STATS Package Subprograms***

Subprogram	Description
ALTER_STATS_HISTORY_RETENTION Procedure	Changes the statistics history retention value
CANCEL_ADVISOR_TASK Procedure	Cancels an Optimizer Statistics Advisor execution
CONFIGURE_ADVISOR_FILTER Function	Configures the filter list for an Optimizer Statistics Advisor task
CONFIGURE_ADVISOR_OBJ_FILTER Function	Configures an object filter for an Optimizer Statistics Advisor task
CONFIGURE_ADVISOR_OPR_FILTER Functions	Configures an operation filter for an Optimizer Statistics Advisor task
CONFIGURE_ADVISOR_RULE_FILTER Function	Configures a rule filter for an Optimizer Statistics Advisor task
CREATE_ADVISOR_TASK Function	Creates an advisor task for the Optimizer Statistics Advisor
CONVERT_RAW_VALUE Procedures	Converts the internal representation of a minimum value, maximum value, or histogram endpoint actual value into a datatype-specific value
CONVERT_RAW_VALUE_NVARCHAR Procedure	Converts the internal representation of a minimum value, maximum value, or histogram endpoint actual value into a datatype-specific value
CONVERT_RAW_VALUE_ROWID Procedure	Converts the internal representation of a minimum value, maximum value, or histogram endpoint actual value into a datatype-specific value
COPY_TABLE_STATS Procedure	Copies the statistics of the source [sub] partition to the destination [sub] partition after scaling
CREATE_EXTENDED_STATS Function	Creates a virtual column for a user specified column group or an expression in a table
CREATE_STAT_TABLE Procedure	Creates a table with name `stattab` in `ownname`'s schema which is capable of holding statistics
DELETE_COLUMN_STATS Procedure	Deletes column-related statistics
DELETE_DATABASE_PREFS Procedure	Deletes the statistics preferences of all the tables
DELETE_DATABASE_STATS Procedure	Deletes statistics for the entire database
DELETE_DICTIONARY_STATS Procedure	Deletes statistics for all dictionary schemas ('`SYS`', '`SYSTEM`' and database component schemas)
DELETE_FIXED_OBJECTS_STATS Procedure	Deletes statistics of all fixed tables
DELETE_INDEX_STATS Procedure	Deletes index-related statistics
DELETE_PENDING_STATS Procedure	Deletes the private statistics that have been collected but have not been published
DELETE_PLSQL_PREFS Procedure	Deletes the preferences for a given function. The values then default to the global preferences.

Table 192-3 (Cont.) *DBMS_STATS Package Subprograms*

Subprogram	Description
DELETE_PROCESSING_RATE Procedure	Deletes the processing rate of a given statistics source. If the source is not specified, it deletes the statistics of all the sources
DELETE_SCHEMA_PREFS Procedure	Deletes the statistics preferences of all the tables owned by the specified owner name
DELETE_SCHEMA_STATS Procedure	Deletes schema-related statistics
DELETE_SYSTEM_STATS Procedure	Deletes system statistics
DELETE_TABLE_PREFS Procedure	Deletes statistics preferences of the specified table in the specified schema
DELETE_TABLE_STATS Procedure	Deletes table-related statistics
DIFF_TABLE_STATS_IN_HISTORY Function	Compares statistics for a table from two timestamps in past and compare the statistics as of that timestamps
DIFF_TABLE_STATS_IN_PENDING Function	Compares pending statistics and statistics as of a timestamp or statistics from dictionary
DIFF_TABLE_STATS_IN_STATTAB Function	Compares statistics for a table from two different sources
DROP_ADVISOR_TASK Procedure	Drops the specified Optimizer Statistics Advisor task
DROP_EXTENDED_STATS Procedure	Drops the statistics entry that is created for the user specified extension
DROP_STAT_TABLE Procedure	Drops a user statistics table created by `CREATE_STAT_TABLE`
EXECUTE_ADVISOR_TASK Function	Executes a previously created Optimizer Statistics Advisor task
EXPORT_COLUMN_STATS Procedure	Retrieves statistics for a particular column and stores them in the user statistics table identified by `stattab`
EXPORT_DATABASE_PREFS Procedure	Exports the statistics preferences of all the tables
EXPORT_DATABASE_STATS Procedure	Retrieves statistics for all objects in the database and stores them in the user statistics table identified by `statown.stattab`
EXPORT_DICTIONARY_STATS Procedure	Retrieves statistics for all dictionary schemas ('`SYS`', '`SYSTEM`' and RDBMS component schemas) and stores them in the user statistics table identified by `stattab`
EXPORT_FIXED_OBJECTS_STATS Procedure	Retrieves statistics for fixed tables and stores them in the user statistics table identified by `stattab`
EXPORT_INDEX_STATS Procedure	Retrieves statistics for a particular index and stores them in the user statistics table identified by `stattab`
EXPORT_PENDING_STATS Procedure	Exports the statistics gathered and stored as pending
EXPORT_SCHEMA_PREFS Procedure	Exports the statistics preferences of all the tables owned by the specified owner name
EXPORT_SCHEMA_STATS Procedure	Retrieves statistics for all objects in the schema identified by `ownname` and stores them in the user statistics table identified by `stattab`

Table 192-3 (Cont.) *DBMS_STATS Package Subprograms*

Subprogram	Description
EXPORT_SYSTEM_STATS Procedure	Retrieves system statistics and stores them in the user statistics table
EXPORT_TABLE_PREFS Procedure	Exports statistics preferences of the specified table in the specified schema into the specified statistics table
EXPORT_TABLE_STATS Procedure	Retrieves statistics for a particular table and stores them in the user statistics table
FLUSH_DATABASE_MONITORING_INFO Procedure	Flushes in-memory monitoring information for all the tables to the dictionary
GATHER_DATABASE_STATS Procedures	Gathers statistics for all objects in the database
GATHER_DICTIONARY_STATS Procedure	Gathers statistics for dictionary schemas '`SYS`', '`SYSTEM`' and schemas of RDBMS components
GATHER_FIXED_OBJECTS_STATS Procedure	Gathers statistics of fixed objects
GATHER_INDEX_STATS Procedure	Gathers index statistics
GATHER_PROCESSING_RATE Procedure	Starts the job of gathering the processing rates which end after `interval` defined in minutes
GATHER_SCHEMA_STATS Procedures	Gathers statistics for all objects in a schema
GATHER_SYSTEM_STATS Procedure	Gathers system statistics
GATHER_TABLE_STATS Procedure	Gathers table and column (and index) statistics
GENERATE_STATS Procedure	Generates object statistics from previously collected statistics of related objects
GET_ADVISOR_OPR_FILTER Procedure	Creates an operation filter for an Optimizer Statistics Advisor operation
GET_ADVISOR_RECS Function	Generates a recommendation report for the specified item
GET_COLUMN_STATS Procedures	Gets all column-related information
GET_INDEX_STATS Procedures	Gets all index-related information
GET_PARAM Function	Gets the default value of parameters of `DBMS_STATS` procedures (see DBMS_STATS Deprecated Subprograms)
GET_PLSQL_PREFS Procedure	Get the preferences for the given function.
GET_PREFS Function	Gets the default value of the specified preference
GET_STATS_HISTORY_AVAILABILITY Function	Gets the oldest timestamp where statistics history is available
GET_STATS_HISTORY_RETENTION Function	Returns the current statistics history retention value
GET_SYSTEM_STATS Procedure	Gets system statistics from `stattab`, or from the dictionary if `stattab` is `NULL`

Table 192-3 (Cont.) *DBMS_STATS Package Subprograms*

Subprogram	Description
GET_TABLE_STATS Procedure	Gets all table-related information
IMPLEMENT_ADVISOR_TASK Function	Implements the recommendations made by Optimizer Statistics Advisor
IMPORT_COLUMN_STATS Procedure	Retrieves statistics for a particular column from the user statistics table identified by `stattab` and stores them in the dictionary
IMPORT_DATABASE_PREFS Procedure	Imports the statistics preferences of all the tables
IMPORT_DATABASE_STATS Procedure	Retrieves statistics for all objects in the database from the user statistics table and stores them in the dictionary
IMPORT_DICTIONARY_STATS Procedure	Retrieves statistics for all dictionary schemas ('`SYS`', '`SYSTEM`' and RDBMS component schemas) from the user statistics table and stores them in the dictionary
IMPORT_FIXED_OBJECTS_STATS Procedure	Retrieves statistics for fixed tables from the user statistics table identified by `stattab` and stores them in the dictionary
IMPORT_INDEX_STATS Procedure	Retrieves statistics for a particular index from the user statistics table identified by `stattab` and stores them in the dictionary
IMPORT_SCHEMA_PREFS Procedure	Imports the statistics preferences of all the tables owned by the specified owner name
IMPORT_SCHEMA_STATS Procedure	Retrieves statistics for all objects in the schema identified by `ownname` from the user statistics table and stores them in the dictionary
IMPORT_SYSTEM_STATS Procedure	Retrieves system statistics from the user statistics table and stores them in the dictionary
IMPORT_TABLE_PREFS Procedure	Sets the statistics preferences of the specified table in the specified schema
IMPORT_TABLE_STATS Procedure	Retrieves statistics for a particular table from the user statistics table identified by `stattab` and stores them in the dictionary
INTERRUPT_ADVISOR_TASK Procedure	Interrupts a currently executing Optimizer Statistics Advisor task.
LOCK_PARTITION_STATS Procedure	Locks statistics for a partition
LOCK_SCHEMA_STATS Procedure	Locks the statistics of all tables of a schema
LOCK_TABLE_STATS Procedure	Locks the statistics on the table
MERGE_COL_USAGE Procedure	Merges column usage information from a source database, by means of a `dblink`, into the local database
PREPARE_COLUMN_VALUES Procedures	Converts user-specified minimum, maximum, and histogram endpoint datatype-specific values into Oracle's internal representation for future storage using the SEED_COL_USAGE Procedure
PREPARE_COLUMN_VALUES_NVARCHAR Procedure	Converts user-specified minimum, maximum, and histogram endpoint datatype-specific values into Oracle's internal representation for future storage using the SEED_COL_USAGE Procedure
PREPARE_COLUMN_VALUES_ROWID Procedure	Converts user-specified minimum, maximum, and histogram endpoint datatype-specific values into Oracle's internal representation for future storage using the SEED_COL_USAGE Procedure

Table 192-3 (Cont.) *DBMS_STATS Package Subprograms*

Subprogram	Description
PUBLISH_PENDING_STATS Procedure	Publishes the statistics gathered and stored as pending
PURGE_STATS Procedure	Purges old versions of statistics saved in the dictionary
REMAP_STAT_TABLE Procedure	Remaps the names of objects in the user statistics table
REPORT_ADVISOR_TASK Function	Reports the results of an Optimizer Advisor Task.
REPORT_COL_USAGE Function	Reports the recorded column (group) usage information
REPORT_GATHER_AUTO_STATS Function	Runs the auto statistics gathering job in reporting mode
REPORT_GATHER_DATABASE_STATS Functions	Runs the GATHER_DATABASE_STATS Procedures in reporting mode.
REPORT_GATHER_DICTIONARY_STATS Functions	Runs the GATHER_DICTIONARY_STATS Procedure in reporting mode
REPORT_GATHER_FIXED_OBJ_STATS Function	Runs the GATHER_FIXED_OBJECTS_STATS Procedure in reporting mode
REPORT_GATHER_SCHEMA_STATS Functions	Runs the GATHER_SCHEMA_STATS Procedures in reporting mode
REPORT_GATHER_TABLE_STATS Function	Runs the GATHER_TABLE_STATS Procedure in reporting mode
REPORT_STATS_OPERATIONS Function	Generates a report of all statistics operations that take place between two timestamps which may or may not have been provided
RESET_ADVISOR_TASK Procedure	Resets an Optimizer Statistics Advisor task execution to its initial state. Only reset a task that is not currently executing
RESET_COL_USAGE Procedure	Resets the recorded column (group) usage information
RESET_GLOBAL_PLSQL_PREFS_DEF Procedure	Reset all global PLSQL preferences for dynamic sampling to the defaults.
RESET_PARAM_DEFAULTS Procedure	Resets global preferences to default values (see DBMS_STATS Deprecated Subprograms)
RESTORE_DICTIONARY_STATS Procedure	Restores statistics of all dictionary tables (tables of 'SYS', 'SYSTEM' and RDBMS component schemas) as of a specified timestamp
RESTORE_FIXED_OBJECTS_STATS Procedure	Restores statistics of all fixed tables as of a specified timestamp
RESTORE_SCHEMA_STATS Procedure	Restores statistics of all tables of a schema as of a specified timestamp
RESTORE_SYSTEM_STATS Procedure	Restores statistics of all tables of a schema as of a specified timestamp
RESTORE_TABLE_STATS Procedure	Restores statistics of a table as of a specified timestamp (as_of_timestamp), as well as statistics of associated indexes and columns
RESUME_ADVISOR_TASK Procedure	Resumes an interrupted task. It only resumes the execution that was most recently interrupted
SCRIPT_ADVISOR_TASK Function	Retrieves the script that implements the recommended actions for the problems found by Optimizer Statistics Advisor

Table 192-3 (Cont.) *DBMS_STATS Package Subprograms*

Subprogram	Description
SEED_COL_USAGE Procedure	Iterates over the SQL statements in the specified SQL tuning set, compiles them and seeds column usage information for the columns that appear in these statements
SET_ADVISOR_TASK_PARAMETER Procedure	Updates the value of an Optimizer Statistics Advisor task parameter
SET_COLUMN_STATS Procedures	Sets column-related information
SET_DATABASE_PREFS Procedure	Sets the statistics preferences of all the tables
SET_GLOBAL_PREFS Procedure	Sets the global statistics preferences
SET_INDEX_STATS Procedures	Sets index-related information
SET_PARAM Procedure	Sets default values for parameters of `DBMS_STATS` procedures (see DBMS_STATS Deprecated Subprograms)
SET_PLSQL_PREFS Procedure	Sets preferences for a specific function – ownername, package_name, function_name, value.
SET_GLOBAL_PLSQL_PREFS Procedure	Sets the global preference for performing dynamic sampling for PL/SQL functions - pname (dynamic_stats by default), pvalue (`ON`, `OFF`, `CHOOSE`, `DEFAULT`).
SET_PROCESSING_RATE Procedure	Sets the value of rate of processing for a given operation
SET_SCHEMA_PREFS Procedure	Sets the statistics preferences of all the tables owned by the specified owner name
SET_SYSTEM_STATS Procedure	Sets system statistics
SET_TABLE_PREFS Procedure	Sets the statistics preferences of the specified table in the specified schema
SET_TABLE_STATS Procedure	Sets table-related information
SHOW_EXTENDED_STATS_NAME Function	Returns the name of the virtual column that is created for the user-specified extension
TRANSFER_STATS Procedure	Transfers statistics for specified table(s) from a remote database specified by `dblink` to the local database
UNLOCK_PARTITION_STATS Procedure	Unlocks the statistics for a partition
UNLOCK_SCHEMA_STATS Procedure	Unlocks the statistics on all the tables in schema
UNLOCK_TABLE_STATS Procedure	Unlocks the statistics on the table
UPGRADE_STAT_TABLE Procedure	Upgrades user statistics on an older table

192.7.1 ALTER_STATS_HISTORY_RETENTION Procedure

This procedure changes the statistics history retention value.

Statistics history retention is used by both the automatic purge and PURGE_STATS Procedure.

Syntax

```
DBMS_STATS.ALTER_STATS_HISTORY_RETENTION (
   retention       IN     NUMBER);
```

Parameters

Table 192-4 ALTER_STATS_HISTORY_RETENTION Procedure Parameters

Parameter	Description
retention	The retention time in days. The statistics history will be retained for at least these many number of days.The valid range is [1,365000]. Also you can use the following values for special purposes: • -1: Statistics history is never purged by automatic purge • 0: Old statistics are never saved. The automatic purge will delete all statistics history • NULL: Change statistics history retention to default value

Usage Notes

To run this procedure, you must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege.

Exceptions

ORA-20000: Insufficient privileges

192.7.2 CANCEL_ADVISOR_TASK Procedure

This procedure cancels an Optimizer Statistics Advisor execution. The advisor removes all intermediate results of the current execution from the task.

Syntax

```
DBMS_STATS.CANCEL_ADVISOR_TASK (
  task_name IN VARCHAR2);
```

Parameters

Table 192-5 CANCEL_ADVISOR_TASK Procedure Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.

Security Model

Note the following:

- To execute this subprogram, you must have the `ADVISOR` privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Consider a case in which a task is executed by one user, interrupted, and then resumed by a different user. In this case, Optimizer Statistics Advisor bases its checks of the resumed execution on the privilege of the user who resumed the task.

Exceptions

- `ORA-20000`: Insufficient privileges
- `ORA-20001`: Invalid input values
- `ORA-20012`: Optimizer Statistics Advisor errors

Usage Notes

To be canceled or interrupted, the specified task must be currently executing.

Example 192-1 Canceling an Optimizer Statistics Advisor

In this example, you start a SQL*Plus session, and then create and execute an advisor task named `my_task`:

```
DECLARE
  v_tname   VARCHAR2(128) := 'my_task';
  v_ename   VARCHAR2(128) := NULL;
BEGIN
  -- create a task
  v_tname := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);

  -- execute the task
  v_ename := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);
END;
/
```

In a separate terminal, you start a second SQL*Plus session, and then execute the following program:

```
EXEC DBMS_STATS.CANCEL_ADVISOR_TASK('my_task');
```

The first session returns an `ORA-13632` to indicate the cancelation of the task:

```
ORA-13632: The user cancelled the current operation.
```

192.7.3 CONFIGURE_ADVISOR_FILTER Function

This function configures the filter list for an Optimizer Statistics Advisor task. Filters are useful for excluding irrelevant findings from a report.

Syntax

```
DBMS_STATS.CONFIGURE_ADVISOR_FILTER (
  task_name           IN   VARCHAR2,
  stats_adv_opr_type  IN   VARCHAR2,
  configuration_type  IN   VARCHAR2,
  filter              IN   StatsAdvFilterTab   := NULL)
RETURN CLOB;
```

Parameters

Table 192-6 CONFIGURE_ADVISOR_FILTER Function Parameters

Parameter	Description
`task_name`	The name of the Optimizer Statistics Advisor task.
`stats_adv_opr_type`	The type of operation to configure. Possible values are `EXECUTE`, `REPORT`, `SCRIPT`, and `IMPLEMENT`. The function permits you to specify a combination of operation types by using the plus (+) operator, for example, `EXECUTE +REPORT`. If this parameter is null, then the filter applies to all types of advisor operations.
`configuration_type`	The type of configuration. Possible values are as follows: • `SET`: Sets the specified filter list values. The submitted filter overrides existing filter values. • `CLEAR`: Clears the existing values for the specified filter. • `SHOW`: Shows the current values of the specified filter.
`filter`	The list of filter items for the script.

Security Model

Note the following:

- To execute this subprogram, you must have the `ADVISOR` privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Return Values

This function returns a CLOB that contains the configuration of the provided filter in XML format.

Exceptions

- `ORA-20000`: Insufficient privileges
- `ORA-20001`: Invalid input values
- `ORA-20012`: Optimizer Statistics Advisor errors

Usage Notes

To provide fine-grained control and a unified interface across all procedures, `DBMS_STATS` provides the `StatsAdvFilter` type. You can use this data type to instantiate and construct a table of filters. You can then pass a parameter of type `StatsAdvFilter` to `CONFIGURE_ADVISOR_FILTER` along with a Boolean variable that specifies either of the following:

- Inclusion list

  Only include these objects in the check.

- Exclusion list

  Do not include these objects in the check.

You can also pass in a parameter specifying whether to replace the existing list. This list only filters object-level and operation-level items. The advisor always checks system-level rules.

You can create the following types of filters:

- Rule filter

  This filter takes a rule name as input. Obtain rule names from the `V$STATS_ADVISOR_RULES` view.

- Operation filter

  This filter is an exact match filter that takes in the name of the operation and an XML string representation of all the parameter values in the call. To obtain the XML, see the notes section of the `DBA_OPTSTAT_OPERATIONS` view. To obtain the filter for an operation, use `DBMS_STATS.GET_ADVISOR_OPR_FILTER`.

- Object filter

  This filter accept an owner name and an object name. Wildcards (`%`) are supported in the owner name and object name. When an object name is null or `%`, it means a filter for all the objects in the specified schema. If the owner name is also null or `%`, it means a default filter for all objects in the system.

If none of the filters is specified, then the function recognizes the filter as setting the global default value of filtering (include or exclude). During the check, if no filter has been specified for a rule, operation, or object, then the function uses the default value to determine whether to include or exclude it.

Example 192-2 Enabling and Disabling Rules

You may want to turn off checks for all rules except for a specific rule. In this example, you want to check whether SQL plan directives have been disabled.

```
DECLARE
  v_task_name    VARCHAR2(128)                := 'my_task';
  v_ret          VARCHAR2(128);
  filter         DBMS_STATS.StatsAdvFilter    := null;
  filterTab      DBMS_STATS.StatsAdvFilterTab := null;
  v_counter      NUMBER                       := 0;
  v_filterReport CLOB;
BEGIN
  -- Create the advisor task
  v_ret := DBMS_STATS.CREATE_ADVISOR_TASK(v_task_name);

  -- Initialize the filter table
```

```
  filterTab := DBMS_STATS.StatsAdvFilterTab();

  -- First filter: set filters to be FALSE by default
  filter.include := FALSE;

  -- Add this filter to the filter table
  v_counter := v_counter + 1;
  filterTab.extend;
  filterTab(v_counter) := filter;

  -- Second filter: turn on filter for one rule
  filter.include := TRUE;
  filter.rulename := 'TurnOnSQLPlanDirective';

  -- Add the SQL plan directive filter to the filter table
  v_counter := v_counter + 1;
  filterTab.extend;
  filterTab(v_counter) := filter;

  v_filterReport := DBMS_STATS.CONFIGURE_ADVISOR_FILTER(
                      task_name          => v_task_name,
                      stats_adv_opr_type => NULL,
                      configuration_type => 'SET',
                      filter             => filterTab);
  -- Drop the task
  DBMS_STATS.DROP_ADVISOR_TASK(v_task_name);
END;
```

Example 192-3 Configuring an Operations Filter

In this example, your shop uses customized scripts to gather statistics for a table. If you do not want to see a specific statistics operation in the report, then you can specify an operations filter.

```
DECLARE
  v_task_name    VARCHAR2(128)                := 'my_task';
  v_ret          VARCHAR2(128);
  filter         DBMS_STATS.StatsAdvFilter    := null;
  filterTab      DBMS_STATS.StatsAdvFilterTab := null;
  opr            DBMS_STATS.StatsAdvOpr;
  oprTab         DBMS_STATS.StatsAdvOprTab;
  v_oprCnt       NUMBER                       := 0;
  TYPE numTab IS TABLE OF NUMBER;
  opr_tab        numTab;
  v_filterReport CLOB;

BEGIN
  -- Create the advisor task
  v_ret := DBMS_STATS.CREATE_ADVISOR_TASK(v_task_name);

  -- Initialize filter table
  filterTab := DBMS_STATS.StatsAdvFilterTab();

  -- Initialize operations filter
  oprTab := DBMS_STATS.StatsAdvOprTab();
```

```
  SELECT ID
    BULK COLLECT INTO opr_tab
  FROM   WRI$_OPTSTAT_OPR
  WHERE  OPERATION = 'set_table_stats'
  AND    TARGET = 'HR.EMPLOYEES';

  -- Populate the operations table
  FOR i IN 1..opr_tab.count LOOP

    -- Use the procedure GET_ADVISOR_OPR_FILTER to construct
    -- an operation filter
    DBMS_STATS.GET_ADVISOR_OPR_FILTER(opr_tab(i), opr);

    v_oprCnt := v_oprCnt + 1;
    oprTab.extend;
    oprTab(v_oprCnt) := opr;

  END LOOP;

  filter.include := FALSE;
  filter.oprlist := oprTab;

  -- Add to filter table
  filterTab.extend;
  filterTab(1) := filter;

  v_filterReport := DBMS_STATS.CONFIGURE_ADVISOR_FILTER(
                      task_name          => v_task_name,
                      stats_adv_opr_type => NULL,
                      configuration_type => 'SET',
                      filter             => filterTab);

  -- Drop the task
  DBMS_STATS.DROP_ADVISOR_TASK(v_task_name);

END;
```

Example 192-4 Reporting on a Specific Schema

In this example, you want to generate a report only for the `sh` schema. Also, you want to skip the `sh.products` table. You create an object filter as follows:

```
DECLARE
  v_task_name     VARCHAR2(128)                := 'my_task';
  v_ret           VARCHAR2(128);
  filter          DBMS_STATS.StatsAdvFilter    := null;
  filterTab       DBMS_STATS.StatsAdvFilterTab := null;
  v_filterReport  CLOB;
  v_counter       NUMBER                       := 0;
  obj             DBMS_STATS.ObjectElem;
  objTab          DBMS_STATS.ObjectTab;
  v_objCnt        NUMBER                       := 0;
BEGIN
  -- Create the advisor task
  v_ret := DBMS_STATS.CREATE_ADVISOR_TASK(v_task_name);
```

```
-- Initialize filter table
filterTab := DBMS_STATS.StatsAdvFilterTab();

-- Set object filter to be off by default
filter.include := FALSE;

objTab := DBMS_STATS.ObjectTab();

obj.ownname := NULL;
obj.objname := NULL;

-- Add to the object table

v_objCnt := v_objCnt + 1;
objTab.extend;
objTab(v_objCnt) := obj;

filter.objlist := objTab;

-- Add the object filter to the filter table
v_counter := v_counter + 1;
filterTab.extend;
filterTab(v_counter) := filter;

-- In filter 1, turn on the check only for schema SH

filter.include := TRUE;

objTab := DBMS_STATS.ObjectTab();
v_objCnt := 0;

obj.ownname := 'SH';
obj.objname := NULL;

-- add to the object table
v_objCnt := v_objCnt + 1;
objTab.extend;
objTab(v_objCnt) := obj;

filter.objlist := objTab;

-- Add the object filter to the filter table
v_counter := v_counter + 1;
filterTab.extend;
filterTab(v_counter) := filter;

-- In filter 2, exclude the check for object sh.products

filter.include := FALSE;


objTab := dbms_stats.ObjectTab();
v_objCnt := 0;

-- Specify another object filter for sh.products
obj.ownname := 'SH';
```

```
  obj.objname := 'PRODUCTS';

  -- Add to the object table
  v_objCnt := v_objCnt + 1;
  objTab.extend;
  objTab(v_objCnt) := obj;

  filter.objlist := objTab;


  -- Add the object filter to the filter table
  v_counter := v_counter + 1;
  filterTab.extend;
  filterTab(v_counter) := filter;

  v_filterReport :=
  DBMS_STATS.CONFIGURE_ADVISOR_FILTER(
    task_name          => v_task_name,
    stats_adv_opr_type => NULL,
    configuration_type => 'SET',
    filter             => filterTab);

  -- Drop the task
  DBMS_STATS.DROP_ADVISOR_TASK(v_task_name);
END;
```

192.7.4 CONFIGURE_ADVISOR_OBJ_FILTER Function

This function configures an object filter for an Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.CONFIGURE_ADVISOR_OBJ_FILTER (
  task_name             IN VARCHAR2,
  stats_adv_opr_type    IN VARCHAR2,
  rule_name             IN VARCHAR2,
  ownname               IN VARCHAR2,
  tabname               IN VARCHAR2,
  action                IN VARCHAR2)
RETURN CLOB;
```

Parameters

Table 192-7 CONFIGURE_ADVISOR_OBJ_FILTER Function Parameters

Parameter	Description
`task_name`	The name of the Optimizer Statistics Advisor task.
`stats_adv_opr_type`	The type of operation to configure. Possible values are `EXECUTE`, `REPORT`, `SCRIPT`, and `IMPLEMENT`. See"CONFIGURE_ADVISOR_RULE_FILTER Function".
`rule_name`	The name of the rule to configure. If null, the function applies the filter to all operation-level rules.

Table 192-7 (Cont.) CONFIGURE_ADVISOR_OBJ_FILTER Function Parameters

Parameter	Description
ownname	The owner name of the operation target. If null, the function applies the filter to all owner names.
tabname	The table name of the operation target.
action	The configuration action to take for the specified rule. See"CONFIGURE_ADVISOR_RULE_FILTER Function".

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Return Values

This function returns an XML CLOB that contains the updated values of the filter.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

192.7.5 CONFIGURE_ADVISOR_OPR_FILTER Functions

This overloaded function configures an operation filter for an Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.CONFIGURE_ADVISOR_OPR_FILTER (
  task_name             IN VARCHAR2,
  stats_adv_opr_type    IN VARCHAR2,
  rule_name             IN VARCHAR2,
  operation_name        IN VARCHAR2,
  action                IN VARCHAR2)
RETURN CLOB;


DBMS_STATS.CONFIGURE_ADVISOR_OPR_FILTER (
  task_name             IN VARCHAR2,
  stats_adv_opr_type    IN VARCHAR2,
  rule_name             IN VARCHAR2,
  operation_name        IN VARCHAR2,
  ownname               IN VARCHAR2,
  tabname               IN VARCHAR2,
```

```
  action                IN VARCHAR2)
RETURN CLOB;


DBMS_STATS.CONFIGURE_ADVISOR_OPR_FILTER (
  task_name             IN VARCHAR2,
  stats_adv_opr_type    IN VARCHAR2,
  rule_name             IN VARCHAR2,
  operation_id          IN NUMBER,
  action                IN VARCHAR2)
RETURN CLOB;
```

Parameters

Table 192-8 CONFIGURE_ADVISOR_OPR_FILTER Function Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.
stats_adv_opr_type	The type of operation to configure. Possible values are EXECUTE, REPORT, SCRIPT, and IMPLEMENT. See"CONFIGURE_ADVISOR_RULE_FILTER Function".
rule_name	The name of the rule to configure. If null, the function applies the filter to all operation-level rules.
operation_name	The name of the operation. For example, an operation name could be gather_table_stats. This value cannot be null.
operation_id	The ID of the operation to configure. The filter applies to any operation with the same signature as the specified operation ID. If two operations have the same signature, then they have the same value for every parameter. View the operation ID in DBA_OPSTAT_OPERATIONS.ID. This value cannot be null.
ownname	The owner name of the operation target. This value cannot be null.
tabname	The table name of the operation target.
action	The configuration action to take for the specified rule. See"CONFIGURE_ADVISOR_RULE_FILTER Function".

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Return Values

This function returns an XML CLOB that contains the updated values of the filter.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

Example 192-5 Excluding Operations for Gathering Table Statistics

In this example, your goal is to exclude operations that gather table statistics in the `hr` schema. User account `stats` has been granted the `DBA` role, `ADVISOR` privilege, and `SELECT ON DBA_OPTSTAT_OPERATIONS` privilege. You perform the following steps:

1. Log in to the database as `stats`.
2. Drop any existing task named `opt_adv_task1`.

```
DECLARE
  v_tname VARCHAR2(32767);
BEGIN
  v_tname := 'opt_adv_task1';
  DBMS_STATS.DROP_ADVISOR_TASK(v_tname);
END;
/
```

3. Create a procedure named `opr_filter` that configures a task to advise on all operations *except* those that gather statistics for tables in the `hr` schema.

```
CREATE OR REPLACE PROCEDURE opr_filter(p_tname IN VARCHAR2) IS
   v_retc CLOB;
BEGIN
   -- For all rules, prevent the advisor from operating
   -- on the operations selected in the following query
   FOR rec IN
     (SELECT ID FROM DBA_OPTSTAT_OPERATIONS WHERE OPERATION =
'gather_table_stats' AND TARGET LIKE 'HR.%')
   LOOP
     v_retc := DBMS_STATS.CONFIGURE_ADVISOR_OPR_FILTER(
                 task_name          => p_tname
               , stats_adv_opr_type => NULL
               , rule_name          => NULL
               , operation_id       => rec.id
               , action             => 'DISABLE');
   END LOOP;
END;
/
SHOW ERRORS
```

4. Create a task named `opt_adv_task1`, and then execute the `opr_filter` procedure for this task.

```
DECLARE
  v_tname VARCHAR2(32767);
  v_ret VARCHAR2(32767);
BEGIN
  v_tname := 'opt_adv_task1';
  v_ret   := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);
  opr_filter(v_tname);
END;
/
```

5. Execute the task `opt_adv_task1`.

```
DECLARE
  v_tname VARCHAR2(32767);
  v_ret   VARCHAR2(32767);
begin
  v_tname := 'opt_adv_task1';
  v_ret   := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);
END;
/
```

6. Print the report.

```
SPOOL /tmp/rep.txt
SET LONG 1000000
COLUMN report FORMAT A200
SET LINESIZE 250
SET PAGESIZE 1000

SELECT DBMS_STATS.REPORT_ADVISOR_TASK(
         task_name      => 'opt_adv_task1'
       , execution_name => NULL
       , type           => 'TEXT'
       , section        => 'ALL'
       ) AS report
FROM   DUAL;
SPOOL OFF
```

See Also:

- *Oracle Database Reference* to learn more about `DBA_OPTSTAT_OPERATIONS`
- *Oracle Database SQL Tuning Guide* to learn how to manage Optimizer Statistics Advisor

192.7.6 CONFIGURE_ADVISOR_RULE_FILTER Function

This function configures a rule filter for an Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.CONFIGURE_ADVISOR_RULE_FILTER (
  task_name            IN   VARCHAR2,
  stats_adv_opr_Type   IN   VARCHAR2,
  rule_name            IN   VARCHAR2,
  action               IN   VARCHAR2)
RETURN CLOB;
```

Parameters

Table 192-9 SCRIPT_ADVISOR_TASK Function Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.
stats_adv_opr_type	The type of operation to configure. Possible values are EXECUTE, REPORT, SCRIPT, and IMPLEMENT. You can specify a combination of operation types, for example, EXECUTE +REPORT. If this parameter is null, then the filter applies to all types of Optimizer Statistics Advisor operations.
rule_name	The name of the rule to configure. If null, the function applies the filter to all rules.
action	The configuration action to take for the specified rule. Possible values are: • ENABLE: Enables the filter • DISABLE: Disables the filter • DELETE: Deletes the filter • SHOW: Shows the current filter value

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Return Values

This function returns an XML CLOB that contains the updated values of the filter.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

192.7.7 CONVERT_RAW_VALUE Procedures

This procedure converts the internal representation of a minimum value, maximum value, or histogram endpoint actual value into a datatype-specific value.

The minval, maxval, and eavals fields of the StatRec structure as filled in by GET_COLUMN_STATS or PREPARE_COLUMN_VALUES are appropriate values for input.

Syntax

```
DBMS_STATS.CONVERT_RAW_VALUE (
   rawval     RAW,
   resval OUT BINARY_FLOAT);
```

```
DBMS_STATS.CONVERT_RAW_VALUE (
   rawval     RAW,
   resval OUT BINARY_DOUBLE);

DBMS_STATS.CONVERT_RAW_VALUE (
   rawval     RAW,
   resval OUT DATE);

DBMS_STATS.CONVERT_RAW_VALUE (
   rawval     RAW,
   resval OUT NUMBER);

DBMS_STATS.CONVERT_RAW_VALUE (
   rawval     RAW,
   resval OUT VARCHAR2);
```

Parameters

Table 192-10 CONVERT_RAW_VALUE Procedure Parameters

Parameter	Description
rawval	Raw representation of a column minimum, maximum, histogram end point actual value
resval	Converted, type-specific value

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.8 CONVERT_RAW_VALUE_NVARCHAR Procedure

This procedure converts the internal representation of a a minimum value, maximum value, or histogram end point actual value.

The `minval`, `maxval` and `eavals` fields of the `StatRec` structure as filled in by `GET_COLUMN_STATS` or `PREPARE_COLUMN_VALUES` are appropriate values for input.

Syntax

```
DBMS_STATS.CONVERT_RAW_VALUE_NVARCHAR (
   rawval     RAW,
   resval OUT NVARCHAR2);
```

Parameters

Table 192-11 CONVERT_RAW_VALUE_NVARCHAR Procedure Parameters

Parameter	Description
rawval	The raw representation of a column minimum or maximum datatype-specific output parameters
resval	The converted, type-specific value

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.9 CONVERT_RAW_VALUE_ROWID Procedure

This procedure converts the internal representation of a a minimum value, maximum value, or histogram end point actual value.

The `minval`, `maxval` and `eavals` fields of the `StatRec` structure as filled in by `GET_COLUMN_STATS` or `PREPARE_COLUMN_VALUES` are appropriate values for input.

Syntax

```
DBMS_STATS.CONVERT_RAW_VALUE_ROWID (
   rawval     RAW,
   resval OUT ROWID);
```

Pragmas

```
pragma restrict_references(convert_raw_value_rowid, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 192-12 CONVERT_RAW_VALUE_ROWID Procedure Parameters

Parameter	Description
`rawval`	The raw representation of a column minimum or maximum datatype-specific output parameters
`resval`	The converted, type-specific value

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.10 COPY_TABLE_STATS Procedure

This procedure copies statistics of all dependent object such as columns and local indexes. If the statistics for source are not available then nothing is copied. It can optionally scale the statistics (such as the number of `blks`, or number of rows) based on the given `scale_factor`.

Syntax

```
DBMS_STATS.COPY_TABLE_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   srcpartname      VARCHAR2,
   dstpartname      VARCHAR2,
   scale_factor     VARCHAR2 DEFAULT 1,
   flags            NUMBER DEFAULT 0,
   force            BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-13 COPY_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	Schema of the table of source and destination [sub] partitions
tabname	Table name of source and destination [sub] partitions
srcpartname	Source [sub] partition
dtspartname	Destination [sub] partition
scale_factor	Scale factor to scale nblks, nrows etc. in dstpartname
flags	For internal Oracle use.
force	When value of this argument is TRUE copy statistics even if the destination [sub]partition is locked

Security Model

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Invalid [sub]partition name

ORA-20001: Bad input value

Usage Notes

This procedure updates the minimum and maximum values of destination partition for the first partitioning column as follows:

- If the partitioning type is HASH, then the minimum and maximum values of the destination partition are same as that of the source partition.
- If the partitioning type is LIST, then the behavior depends on the setting of the destination partition:
  - If the destination partition is a NOT DEFAULT partition, then the following statements are true:
    - The minimum value of the destination partition is set to the minimum value of the value list that describes the destination partition.
    - The maximum value of the destination partition is set to the maximum value of the value list that describes the destination partition.
  - Alternatively, if the destination partition is a DEFAULT partition, then the following statements are true:
    - The minimum value of the destination partition is set to the minimum value of the source partition.
    - The maximum value of the destination partition is set to the maximum value of the source partition.
- If the partitioning type is RANGE, then the following statements are true:

  - The minimum value of the destination partition is set to the high bound of previous partition unless the destination partition is the first partition. For the first partition, the minimum value is set to the high bound of the destination partition.
  - The maximum value of the destination partition is set to the high bound of the destination partition unless the high bound of the destination partition is `MAXVALUE`, in which case the maximum value of the destination partition is set to the high bound of the previous partition.
  - If the source partition column's minimum value is equal to its maximum value, and if both are equal to the source partition's lower bound, and if it has a single distinct value, then the destination partition column's minimum and maximum values are both set to the destination partition's lower bound. This is done for all partitioning columns.

    If the above condition does not apply, second and subsequent partitioning columns are updated as follows. The destination partition column's maximum value is set to the greater of the destination partition upper bound and the source partition column's maximum value, with one exception. If the destination partition is `D` and its preceding partition is `D-1` and the key column to be adjusted is `Cn`, the maximum value for `Cn` is set to the upper bound of `D` (ignoring the maximum value of the source partition column) provided that the upper bounds of the previous key column `Cn-1` are the same in partitions `D` and `D-1`.

- If the minimum and maximum values are different for a column after modifications, and if the number of distinct values is less than 1, then the number of distinct values is updated as 2.
- If the source or destination is a partition of a composite partitioned table, then this procedure does not copy statistics of the underlying subpartitions.

192.7.11 CREATE_ADVISOR_TASK Function

This function creates an advisor task for the Optimizer Statistics Advisor.

Syntax

```
DBMS_STATS.CREATE_ADVISOR_TASK (
   task_name    IN   VARCHAR2   := NULL)
 RETURN VARCHAR2;
```

Parameters

Table 192-14 CREATE_ADVISOR_TASK Function Parameters

Parameter	Description
`task_name`	Name of the task. If the task name is already specified, then the function uses the specified task name. Otherwise, the function generates a new task name automatically.

Security Model

Note the following:

- To execute this subprogram, you must have the `ADVISOR` privilege.
- This subprogram executes using invoker's rights.

Return Values

This function returns the unique name of the Optimizer Statistics Advisor task.

Exceptions

`ORA-20000`: Insufficient privileges / creating extension is not supported

`ORA-20001`: Error when processing extension

`ORA-20012`: Optimizer Statistics Advisor errors

Example 192-6 Creating and Executing a Task

This example creates an Optimizer Statistics Advisor task named `my_task`, and then executes it.

```
DECLARE
  v_tname   VARCHAR2(128) := 'my_task';
BEGIN
  -- create a task
  v_tname := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);
  -- execute the task
  v_tname := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);
END;
```

> **Note:**
>
> *Oracle Database SQL Tuning Guide* to learn how manage Optimizer Statistics Advisor

192.7.12 CREATE_EXTENDED_STATS Function

This function creates a column statistics entry in the system for a user-specified column group or an expression in a table.

The database gathers statistics for this extension when a user-initiated or automatic statistics gathering job gathers statistics for the table. Statistics for such an extension are called **extended statistics**. This function returns the name of this newly created entry for the extension.
This second form of this function creates statistics extensions based on the column group usage recorded by the SEED_COL_USAGE Procedure. This function returns a report of extensions created.

Syntax

```
DBMS_STATS.CREATE_EXTENDED_STATS (
   ownname    VARCHAR2,
   tabname    VARCHAR2,
   extension  VARCHAR2)
 RETURN VARCHAR2;

DBMS_STATS.CREATE_EXTENDED_STATS (
   ownname    VARCHAR2,
```

```
   tabname    VARCHAR2)
 RETURN CLOB;
```

Parameters

Table 192-15 CREATE_EXTENDED_STATS Function Parameters

Parameter	Description
ownname	Owner name of a table
tabname	Name of the table
extension	Can be either a column group or an expression. Suppose the specified table has columns c1, c2. An example column group is "(c1, c2)". An example expression is "(c1 + c2)".

Return Values

This function returns the name of this newly created entry for the extension.

Exceptions

ORA-20000: Insufficient privileges / creating extension is not supported

ORA-20001: Error when processing extension

ORA-20007: Extension already exists

ORA-20008: Reached the upper limit on number of extensions

Usage Notes

To invoke this procedure you must be owner of the table, or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table, or have either the ANALYZE ANY DICTIONARY or SYSDBA privilege.

The extension has the following restrictions:

- The extension cannot contain a virtual column.
- Extensions cannot be created on tables owned by SYS.
- Extensions cannot be created on cluster tables, index organized tables, temporary tables, or external tables.
- The total number of extensions in a table cannot be greater than a maximum of (20, 10% of number of non-virtual columns in the table).
- The number of columns in a column group must be in the range [2, 32].
- A column cannot appear more than once in a column group.
- The extension can contain an expression only if a corresponding virtual column has been created.
- An expression must contain at least one column.
- An expression cannot contain a subquery.
- The COMPATIBLE parameter must be 11.0.0.0.0 or greater.

192.7.13 CREATE_STAT_TABLE Procedure

This procedure creates a table with name `stattab` in `ownname`'s schema which is capable of holding statistics. The columns and types that compose this table are not relevant as it should be accessed solely through the procedures in this package.

Syntax

```
DBMS_STATS.CREATE_STAT_TABLE (
   ownname             VARCHAR2,
   stattab             VARCHAR2,
   tblspace            VARCHAR2 DEFAULT NULL,
   global_temporary    BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-16 CREATE_STAT_TABLE Procedure Parameters

Parameter	Description
`ownname`	Name of the schema
`stattab`	Name of the table to create. This value should be passed as the `stattab` parameter to other procedures when the user does not want to modify the dictionary statistics directly.
`tblspace`	Tablespace in which to create the statistics tables. If none is specified, then they are created in the user's default tablespace.
`global_temporary`	Whether or not the table should be created as a global temporary table

Security Model

To invoke this procedure you need whichever privileges are required for creating a table in the specified schema.

Exceptions

`ORA-20000`: Table already exists or insufficient privileges

`ORA-20001`: Tablespace does not exist

192.7.14 DELETE_COLUMN_STATS Procedure

This procedure deletes column-related statistics.

Syntax

```
DBMS_STATS.DELETE_COLUMN_STATS (
   ownname        VARCHAR2,
   tabname        VARCHAR2,
   colname        VARCHAR2,
   partname       VARCHAR2 DEFAULT NULL,
   stattab        VARCHAR2 DEFAULT NULL,
   statid         VARCHAR2 DEFAULT NULL,
   cascade_parts  BOOLEAN  DEFAULT TRUE,
   statown        VARCHAR2 DEFAULT NULL,
```

```
   no_invalidate  BOOLEAN  DEFAULT to_no_invalidate_type (
                                    get_param('NO_INVALIDATE')),
   force          BOOLEAN  DEFAULT FALSE,
   col_stat_type  VARCHAR2 DEFAULT 'ALL');
```

Parameters

Table 192-17 DELETE_COLUMN_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
tabname	Name of the table to which this column belongs
colname	Name of the column or extension
partname	Name of the table partition for which to delete the statistics. If the table is partitioned and if partname is NULL, then global column statistics are deleted.
stattab	User statistics table identifier describing from where to delete the statistics. If stattab is NULL, then the statistics are deleted directly from the dictionary.
statid	Identifier (optional) to associate with these statistics within stattab (Only pertinent if stattab is not NULL).
cascade_parts	If the table is partitioned and if partname is NULL, then setting this to true causes the deletion of statistics for this column for all underlying partitions as well.
statown	Schema containing stattab (if different than ownname)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	When value of this argument is TRUE, deletes column statistics even if locked
col_stat_type	Type of column statistics to be deleted. This argument takes the following values: • HISTOGRAM - delete column histogram only • ALL - delete base column statistics and histogram

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20005: Object statistics are locked

Usage Notes

To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

192.7.15 DELETE_DATABASE_PREFS Procedure

This procedure deletes the statistics preferences set for all non-system tables. You can include system tables by passing `TRUE` for the `add_sys` parameter.

Syntax

```
DBMS_STATS.DELETE_DATABASE_PREFS (
    pname            IN   VARCHAR2,
    add_sys          IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-18 DELETE_DATABASE_PREFS Procedure Parameters

Parameter	Description
pname	Preference name. The existing value for following preferences can be deleted and default preference values will be used: • APPROXIMATE_NDV_ALGORITHM • AUTO_STAT_EXTENSIONS • CASCADE • DEGREE • ESTIMATE_PERCENT • GLOBAL_TEMP_TABLE_STATS • GRANULARITY • INCREMENTAL • INCREMENTAL_LEVEL • INCREMENTAL_STALENESS • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • STAT_CATEGORY • TABLE_CACHED_BLOCKS
add_sys	Determines whether `SYS` tables will be included.

Table 192-19 Statistics Preferences

Preference	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. Consider a synopsis as an internal management structure that samples distinct values. You can specify the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
CASCADE	Determines whether index statistics are collected as part of gathering table statistics.
DEGREE	Determines degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	Determines the percentage of rows to estimate.
GLOBAL_TEMP_TABLE_STATS	Controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics.
GRANULARITY	Determines the granularity of statistics to collect. This value is only relevant for partitioned tables.
INCREMENTAL	Determines whether the global statistics of a partitioned table will be maintained without doing a full table scan.
INCREMENTAL_LEVEL	Controls which synopses to collect when INCREMENTAL preference is set to TRUE.

Table 192-19 (Cont.) Statistics Preferences

Preference	Description

INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'.
	The parameter accepts the following values:
	• USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference.
	For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics.
	Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics.
	• USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes.
	For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics.
	• ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist.
	• NULL—A partition or subpartition is considered stale when it has any DML changes.
	For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.
	Note that the following two executions are different:
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	The first execution uses single quotes to set the preference to the value NULL, whereas the second sets the preference to the default, which is ALLOW_MIXED_FORMAT.

METHOD_OPT	Controls column statistics collection and histogram creation. When setting preference on global, schema, database or dictionary level, only'FOR ALL' syntax is allowed.

Table 192-19 (Cont.) Statistics Preferences

Preference	Description
NO_INVALIDATE	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
OPTIONS	Determines the options parameter used in the GATHER_TABLE_STATS Procedure.
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • TRUE — Ignores input parameter values, and uses the value of the corresponding preference. • FALSE — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether the database publishes newly gathered statistics after the gathering job completes. You can gather statistics without publishing them immediately. This technique enables you to test new statistics before publishing them.
STALE_PERCENT	Determines the percentage of rows in a table that have to change before the statistics on that table are deemed stale and should be regathered.
STAT_CATEGORY	Specifies which statistics to import or export, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics (default) • SYNOPSES — information to support incremental statistics • REALTIME_STATS — Specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. The value 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.

Security Model

To run this procedure, you must have the `SYSDBA` role or both `ANALYZE ANY DICTIONARY` and `ANALYZE ANY` system privileges.

Exceptions

`ORA-20000`: Insufficient privileges

`ORA-20001`: Invalid or Illegal input values

Usage Notes

All `pname` arguments are of type `VARCHAR2` and values are enclosed in quotes, even when they represent numbers.

Example 192-7 Examples

```
DBMS_STATS.DELETE_DATABASE_PREFS('CASCADE', FALSE);
DBMS_STATS.DELETE_DATABASE_PREFS('ESTIMATE_PERCENT', TRUE);
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to manage optimizer statistics preferences

192.7.16 DELETE_DATABASE_STATS Procedure

This procedure deletes statistics for all the tables in a database.

Syntax

```
DBMS_STATS.DELETE_DATABASE_STATS (
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   force            BOOLEAN  DEFAULT FALSE,
   stat_category    VARCHAR2 DEFAULT DEFAULT_DEL_STAT_CATEGORY);
```

Parameters

Table 192-20 DELETE_DATABASE_STATS Procedure Parameters

Parameter	Description
`stattab`	User statistics table identifier describing from where to delete the statistics. If `stattab` is `NULL`, then the statistics are deleted directly in the dictionary.
`statid`	Identifier (optional) to associate with these statistics within `stattab` (Only pertinent if `stattab` is not `NULL`)
`statown`	Schema containing `stattab` (if different from current schema)

Table 192-20 (Cont.) DELETE_DATABASE_STATS Procedure Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	When the value of this argument is TRUE, deletes statistics of tables in a database even if they are locked
stat_category	Statistics to delete. It accepts multiple values separated by comma: • OBJECT_STATS — table statistics, column statistics and index statistics • SYNOPSES — information to support incremental statistics • REALTIME_STATS — specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. The default is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

To run this procedure, you need to have the SYSDBA role or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.

192.7.17 DELETE_DICTIONARY_STATS Procedure

This procedure deletes statistics for all dictionary schemas ('SYS', 'SYSTEM' and RDBMS component schemas).

Syntax

```
DBMS_STATS.DELETE_DICTIONARY_STATS (
   stattab        VARCHAR2 DEFAULT NULL,
   statid         VARCHAR2 DEFAULT NULL,
   statown        VARCHAR2 DEFAULT NULL,
   no_invalidate  BOOLEAN  DEFAULT to_no_invalidate_type (
                                 get_param('NO_INVALIDATE')),
   stattype       VARCHAR2 DEFAULT 'ALL',
   force          BOOLEAN  DEFAULT FALSE,
   stat_category  VARCHAR2 DEFAULT DEFAULT_DEL_STAT_CATEGORY);
```

Parameters

Table 192-21 DELETE_DICTIONARY_STATS Procedure Parameters

Parameter	Description
stattab	User statistics table identifier describing from where to delete the statistics. If stattab is NULL, then the statistics are deleted directly in the dictionary.
statid	Identifier (optional) to associate with these statistics within stattab (Only pertinent if stattab is not NULL)
statown	Schema containing stattab (if different from current schema)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattype	Statistics type
force	When the value of this argument is TRUE, deletes statistics of tables in a database even if they are locked
stat_category	Statistics to delete. It accepts multiple values separated by comma: • OBJECT_STATS — table statistics, column statistics and index statistics • SYNOPSES — information to support incremental statistics • REALTIME_STATS — Specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. The default is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'

Usage Notes

You must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege to execute this procedure.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20002: Bad user statistics table, may need to upgrade it

192.7.18 DELETE_FIXED_OBJECTS_STATS Procedure

This procedure deletes statistics of all fixed tables.

Syntax

```
DBMS_STATS.DELETE_FIXED_OBJECTS_STATS (
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   force            BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 192-22 DELETE_FIXED_OBJECTS_STATS Procedure Parameters

Parameter	Description
stattab	The user statistics table identifier describing from where to delete the current statistics. If stattab is NULL, the statistics will be deleted directly in the dictionary.
statid	The (optional) identifier to associate with these statistics within stattab. This only applies if stattab is not NULL.
statown	Schema containing stattab (if different from current schema)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Ignores the statistics lock on objects and deletes the statistics if set to TRUE

Usage Notes

You must have the SYSDBA or ANALYZE ANY DICTIONARY system privilege to execute this procedure.

Exceptions

ORA-20000: Insufficient privileges

ORA-20002: Bad user statistics table, may need to upgrade it

192.7.19 DELETE_INDEX_STATS Procedure

This procedure deletes index-related statistics.

Syntax

```
DBMS_STATS.DELETE_INDEX_STATS (
   ownname          VARCHAR2,
   indname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   cascade_parts    BOOLEAN  DEFAULT TRUE,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   stattype         VARCHAR2 DEFAULT 'ALL',
   force            BOOLEAN  DEFAULT FALSE);
   stat_category    VARCHAR2 DEFAULT DEFAULT_DEL_STAT_CATEGORY);
```

Parameters

Table 192-23 DELETE_INDEX_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
indname	Name of the index
partname	Name of the index partition for which to delete the statistics. If the index is partitioned and if partname is NULL, then index statistics are deleted at the global level.
stattab	User statistics table identifier describing from where to delete the statistics. If stattab is NULL, then the statistics are deleted directly from the dictionary.
statid	Identifier (optional) to associate with these statistics within stattab (Only pertinent if stattab is not NULL)
cascade_parts	If the index is partitioned and if partname is NULL, then setting this to TRUE causes the deletion of statistics for this index for all underlying partitions as well
statown	Schema containing stattab (if different than ownname)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattype	Statistics type
force	When value of this argument is TRUE, deletes index statistics even if locked

Table 192-23 (Cont.) DELETE_INDEX_STATS Procedure Parameters

Parameter	Description
stat_category	Statistics to delete. It accepts multiple values separated by comma: • OBJECT_STATS — table statistics, column statistics and index statistics • SYNOPSES — information to support incremental statistics • REALTIME_STATS — Specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. The default is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20005: Object statistics are locked

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.20 DELETE_PENDING_STATS Procedure

This procedure is used to delete the pending statistics that have been collected but have not been published.

Syntax

```
DBMS_STATS.DELETE_PENDING_STATS (
    ownname     IN  VARCHAR2  DEFAULT USER,
    tabname     IN  VARCHAR2);
```

Parameters

Table 192-24 DELETE_PENDING_STATS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name

Security Model

To run this procedure, you need to have the same privilege for gathering statistics on the tables that will be affected by this procedure. The default owner is the user who runs the procedure.

Exceptions

ORA-20000: Insufficient privileges

Usage Notes

If the parameter `tabname` is `NULL` delete applies to all tables of the specified schema.

Examples

```
DBMS_STATS.DELETE_PENDING_STATS('SH', 'SALES');
```

192.7.21 DELETE_PLSQL_PREFS Procedure

Deletes the named preference from the named PL/SQL function or from all functions in the same package.

Syntax

```
DBMS_STATS.DELETE_PLSQL_PREFS(
    user           IN VARCHAR2,
    package_name   IN VARCHAR2,
    function_name  IN VARCHAR2,
    pname          IN VARCHAR2
);
```

Parameters

Table 192-25 DELETE_PLSQL_PREFS Procedure Parameters

Parameter	Description
`ownname`	Owner of the function. If NULL, the current schema is used.
`package_name`	The PL/SQL package name. If NULL, this specifies a stand-alone, top-level function.
`function_name`	Name of the PL/SQL function. The preference will be deleted on this function. If NULL, the preference is deleted on all functions in the package.
`pname`	The preference that you want to delete. Currently only `'dynamic_stats'` is accepted.

Examples

In this example the owner is the current schema and so the `user` parameter is set to `NULL`. The dynamic statistics preference for the named function is deleted and the global preference is used for this function.

```
EXEC DELETE_PLSQL_PREFS(NULL, '<PACKAGE_NAME>', '<FUNCTION_NAME>',
'dynamic_stats')
```

In this example, the `function_name` parameter is set to NULL. Therefore the dynamic statistics preference is deleted on all functions within the named PL/SQL package.

```
EXEC DELETE_PLSQL_PREFS('<USER>', '<PACKAGE_NAME>', NULL,'dynamic_stats')
```

192.7.22 DELETE_PROCESSING_RATE Procedure

This procedure deletes the processing rate of a given statistics source. If the source is not specified, it deletes the statistics of all the sources.

Syntax

```
DBMS_STATS.DELETE_PROCESSING_RATE (
   stat_source      IN    VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-26 DELETE_PROCESSING_RATE Procedure Parameters

Parameter	Description
stat_source	Source of processing rates: • 'MANUAL': values set by the user manually using the SET_PROCESSING_RATE Procedure • 'CALIBRATION' : values collected by the calibration GATHER_PROCESSING_RATE Procedure run explicitly by the user • 'FEEDBACK': values obtained by time feedback

Usage Notes

You require the OPTIMIZER_PROCESSING_RATE role to run this procedure since AUTO DOP uses processing rates to determine the optimal degree of parallelism for a SQL statement.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or illegal input value

192.7.23 DELETE_SCHEMA_PREFS Procedure

This procedure is used to delete the statistics preferences of all the tables owned by the specified owner name.

Syntax

```
DBMS_STATS.DELETE_SCHEMA_PREFS (
    ownname    IN   VARCHAR2,
    pname      IN   VARCHAR2);
```

Parameters

Table 192-27 DELETE_SCHEMA_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name

Table 192-27 (Cont.) DELETE_SCHEMA_PREFS Procedure Parameters

Parameter	Description
pname	Preference name. The existing value for following preferences can be deleted and default preference values will be used: • APPROXIMATE_NDV_ALGORITHM • AUTO_STAT_EXTENSIONS • CASCADE • DEGREE • ESTIMATE_PERCENT • GLOBAL_TEMP_TABLE_STATS • GRANULARITY • INCREMENTAL • INCREMENTAL_LEVEL • INCREMENTAL_STALENESS • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • TABLE_CACHED_BLOCKS

Table 192-28 Statistics Preferences

Preference	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. You can consider a synopsis as an internal management structure that samples distinct values. You can set the following values: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses.

Table 192-28 (Cont.) Statistics Preferences

Preference	Description
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
CASCADE	Determines whether index statistics are collected as part of gathering table statistics.
COORDINATOR_TRIGGER_SHARD	User of each shard uses this preference to determine whether to allow shard coordinator to interact with the statistics gathering in each shards. While gathering the statistics in shard coordinator, if the statistics in one of the shards are not up to date, the shard coordinator will try to trigger the statistics gathering in that shard. By using this preference, user can execute or ignore that command from the shard coordinator. You can set the following values: • TRUE—Allows the shard coordinator trigger the statistics gathering on sharded table in local shard if the statistics on local shard are stale. • FALSE—Ignores the statistics gathering command triggered from the shard coordinator. The default value is FALSE.
DEGREE	Determines the degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	The value determines the percentage of rows to estimate.
METHOD_OPT	Controls column statistics collection and histogram creation. When setting preferences at the global, schema, database, or dictionary level, onlyFOR ALL syntax is allowed.
NO_INVALIDATE	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
GRANULARITY	The value determines granularity of statistics to collect (only pertinent if the table is partitioned)

Table 192-28 (Cont.) Statistics Preferences

Preference	Description
PUBLISH	This value determines whether or not newly gathered statistics will be published once the gather job has completed.
INCREMENTAL	This value determines whether or not the global statistics of a partitioned table will be maintained without doing a full table scan.
INCREMENTAL_LEVEL	This value controls what synopses to collect when INCREMENTAL preference is set to TRUE.

Table 192-28 (Cont.) Statistics Preferences

Preference	Description

INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'.
	
	The parameter accepts the following values:
	• USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference.
	For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics.
	Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics.
	• USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes.
	For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics.
	• ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist.
	• NULL—A partition or subpartition is considered stale when it has any DML changes.
	For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.
	Note that the following two executions are different:
	
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	
	The first execution uses single quotes to set the preference to the value NULL, whereas the second sets the preference to the default, which is ALLOW_MIXED_FORMAT.

Table 192-28 (Cont.) Statistics Preferences

Preference	Description
ROOT_TRIGGER_PDB	The application PDB user, uses this preference to determine whether to allow the application root to interact with the statics gathering in PDB. During the statistics gathering of a metadata linked table in the application root, if the statistics in a PDB are in stale state, the application root triggers the statistics gathering for the particular PDB. Using this preference, the user can either execute or ignore the command from the application root. You can set the following values: • TRUE—Allows the application root trigger the statistics gathering on metadata linked table in application PDB if the statistics on PDB are stale. • FALSE—Ignores the statistics gathering command triggered from application root. The default value is FALSE. **Note:** CDB root, different from application root, never triggers statistics gathering on the PDBs and it is not controlled by this preference.
STALE_PERCENT	This value determines the percentage of rows in a table that have to change before the statistics on that table are deemed stale and should be regathered.
GLOBAL_TEMP_TABLE_STATS	This controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.
OPTIONS	Determines the options parameter used in the GATHER_TABLE_STATS Procedure.

Security Model

To run this procedure, you must be the object owner, or have the SYSDBA privilege, or have the ANALYZE ANY system privilege.

Exceptions

ORA-20000: Insufficient privileges / Schema "<schema>" does not exist

ORA-20001: Invalid or Illegal input values

Usage Notes

All arguments are of type `VARCHAR2` and values are enclosed in quotes, even when they represent numbers.

Examples

```
DBMS_STATS.DELETE_SCHEMA_PREFS('SH', 'CASCADE');
DBMS_STATS.DELETE_SCHEMA_PREFS('SH', 'ESTIMATE_PERCENT');
DBMS_STATS.DELETE_SCHEMA_PREFS('SH', 'DEGREE');
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to manage optimizer statistics preferences

192.7.24 DELETE_SCHEMA_STATS Procedure

This procedure deletes statistics for an entire schema.

Syntax

```
DBMS_STATS.DELETE_SCHEMA_STATS (
   ownname          VARCHAR2,
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN DEFAULT to_no_invalidate_type (
                                    get_param('NO_INVALIDATE')),
   force            BOOLEAN DEFAULT FALSE,
   stat_category   VARCHAR2 DEFAULT DEFAULT_DEL_STAT_CATEGORY);
```

Parameters

Table 192-29 DELETE_SCHEMA_STATS Procedure Parameters

Parameter	Description
`ownname`	Specifies the name of the schema.
`stattab`	Identifies the table where statistics are stored. If `stattab` is `NULL`, then the procedure deletes statistics directly from the data dictionary.
`statid`	Specifies the identifier (optional) associated with these statistics within `stattab`. This parameter is only relevant if `stattab` is not `NULL`.
`statown`	Specifies the schema containing `stattab` (if different than `ownname`).

Table 192-29 (Cont.) DELETE_SCHEMA_STATS Procedure Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Indicates whether to force the deletion for locked statistics. When the value is TRUE, this procedure deletes table statistics even if locked.
stat_category	Specifies which statistics to process. The following values are supported: • OBJECT_STATS — Specifies table statistics, column statistics, and index statistics. • SYNOPSES — Specifies metadata for incremental statistics. • REALTIME_STATS — Specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. You can specify a list of comma-delimited values. For example, 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

192.7.25 DELETE_SYSTEM_STATS Procedure

This procedure deletes workload statistics (collected using the 'INTERVAL' or 'START' and 'STOP' options) and resets the default to noworkload statistics (collected using 'NOWORKLOAD' option), if stattab is not specified. If stattab is specified, the subprogram deletes all system statistics with the associated statid from the stattab.

Syntax

```
DBMS_STATS.DELETE_SYSTEM_STATS (
   stattab       VARCHAR2 DEFAULT NULL,
   statid        VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-30 DELETE_SYSTEM_STATS Procedure Parameters

Parameter	Description
stattab	Identifier of the user statistics table where the statistics will be saved
statid	Optional identifier associated with the statistics saved in the stattab
statown	Schema containing stattab (if different from current schema)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20002: Bad user statistics table; may need to be upgraded

Usage Notes

To run this procedure, you need the GATHER_SYSTEM_STATISTICS role.

192.7.26 DELETE_TABLE_PREFS Procedure

This procedure deletes the optimizer statistics preferences of the specified table in the specified schema.

Syntax

```
DBMS_STATS.DELETE_TABLE_PREFS (
    ownname     IN  VARCHAR2,
    tabname     IN  VARCHAR2,
    pname       IN  VARCHAR2);
```

Parameters

Table 192-31 DELETE_TABLE_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name

Table 192-31 (Cont.) DELETE_TABLE_PREFS Procedure Parameters

Parameter	Description
pname	Preference name. The existing value for following preferences can be deleted and default preference values will be used: • APPROXIMATE_NDV_ALGORITHM • AUTO_STAT_EXTENSIONS • CASCADE • DEGREE • ESTIMATE_PERCENT • GRANULARITY • GLOBAL_TEMP_STATS • INCREMENTAL • INCREMENTAL_LEVEL • INCREMENTAL_STALENESS • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • STAT_CATEGORY • TABLE_CACHED_BLOCKS

Table 192-32 Statistics Preferences

Preference	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. Consider a synopsis as an internal management structure that samples distinct values. You can specify the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
CASCADE	Determines whether index statistics are collected as part of gathering table statistics.
DEGREE	Determines the degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	Determines the percentage of rows to estimate.
GRANULARITY	Determines granularity of statistics to collect. This value is only relevant for partitioned tables.
GLOBAL_TEMP_TABLE_STATS	Controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics.

Table 192-32 (Cont.) Statistics Preferences

Preference	Description
INCREMENTAL	This value determines whether or not the global statistics of a partitioned table will be maintained without doing a full table scan.
INCREMENTAL_LEVEL	This value controls what synopses to collect when INCREMENTAL preference is set to TRUE.

Table 192-32 (Cont.) Statistics Preferences

Preference	Description

INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'.
	The parameter accepts the following values:
	• USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference.
	For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics.
	Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics.
	• USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes.
	For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics.
	• ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist.
	• NULL—A partition or subpartition is considered stale when it has any DML changes.
	For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.
	Note that the following two executions are different:
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	The first execution uses single quotes to set the preference to the value NULL, whereas the second sets the preference to the default, which is ALLOW_MIXED_FORMAT.

Table 192-32 (Cont.) Statistics Preferences

Preference	Description
METHOD_OPT	Controls column statistics collection and histogram creation. When setting preference at the global, schema, database, or dictionary level, only 'FOR ALL' syntax is allowed.
NO_INVALIDATE	The value controls the invalidation of dependent cursors of the tables for which statistics are being gathered.
OPTIONS	Determines the options parameter used in the GATHER_TABLE_STATS Procedure.
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • TRUE — Ignores input parameter values, and uses the value of the corresponding preference. • FALSE — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether newly gathered statistics will be published after the statistics gathering job has completed.
STALE_PERCENT	Determines the percentage of rows in a table that have to change before the statistics on that table are deemed stale and should be regathered.
STAT_CATEGORY	Specifies which statistics to import or export, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics (default) • SYNOPSES — information to support incremental statistics • REALTIME_STATS — specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. The value 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Invalid or Illegal input values

Usage Notes

- To run this procedure, you need to connect as owner of the table, be granted ANALYZE privilege on the table, or ANALYZE ANY system privilege.

- All arguments are of type `VARCHAR2` and values are enclosed in quotes, even when they represent numbers.

Examples

```
DBMS_STATS.DELETE_TABLE_PREFS('SH', 'SALES', 'CASCADE');
DBMS_STATS.DELETE_TABLE_PREFS('SH', 'SALES', 'DEGREE');
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to manage optimizer statistics preferences

192.7.27 DELETE_TABLE_STATS Procedure

This procedure deletes table-related statistics.

Syntax

```
DBMS_STATS.DELETE_TABLE_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   cascade_parts    BOOLEAN  DEFAULT TRUE,
   cascade_columns  BOOLEAN  DEFAULT TRUE,
   cascade_indexes  BOOLEAN  DEFAULT TRUE,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   force            BOOLEAN  DEFAULT FALSE,
   stat_category    VARCHAR2 DEFAULT DEFAULT_DEL_STAT_CATEGORY);
```

Parameters

Table 192-33 DELETE_TABLE_STATS Procedure Parameters

Parameter	Description
`ownname`	Specifies the name of the schema.
`tabname`	Specifies the name of the table to which this column belongs.
`partname`	Specifies the name of the table partition or subpartition from which to get the statistics. If the table is partitioned and if `partname` is `NULL`, then the statistics are retrieved from the global table level.
`stattab`	Identifies the user statistics table where statistics will be retrieved. If `stattab` is `NULL`, then the procedure retrieves statistics directly from the dictionary.
`statid`	Specifies the identifier (optional) associated with these statistics within `stattab`. This parameter is only relevant if `stattab` is not `NULL`.

Table 192-33 (Cont.) DELETE_TABLE_STATS Procedure Parameters

Parameter	Description
cascade_parts	Specifies whether the procedure should operate on underlying partitions. If the table is partitioned, and if partname is NULL, then specifying TRUE deletes statistics for underlying partitions.
cascade_columns	Indicates whether to invoke the DELETE_COLUMN_STATS procedure. If TRUE, then this procedure calls DELETE_COLUMN_STATS for all underlying columns.
cascade_indexes	Indicates whether to invoke the DELETE_INDEX_STATS procedure. If TRUE, then this procedure calls DELETE_INDEX_STATS for all underlying columns.
statown	Specifies the schema containing stattab (if different than ownname).
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Indicates whether to force the deletion for locked statistics. When the value is TRUE, this procedure deletes table statistics even if locked.
stat_category	Specifies which statistics to process. The following values are supported: • OBJECT_STATS — specifies table statistics, column statistics, and index statistics. • SYNOPSES — specifies metadata for incremental statistics. • REALTIME_STATS — specifies only real-time statistics. • MODELS — supports import, export, and deletion for regression models in real-time stats. You can specify a list of comma-delimited values. For example, 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS, MODELS'.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20002: Bad user statistics table, may need to upgrade it

ORA-20005: Object statistics are locked

192.7.28 DIFF_TABLE_STATS_IN_HISTORY Function

This function compares statistics for a table as of two specified timestamps.

Syntax

```
DBMS_STATS.DIFF_TABLE_STATS_IN_HISTORY(
   ownname        IN  VARCHAR2,
   tabname        IN  VARCHAR2,
   time1          IN  TIMESTAMP WITH TIME ZONE,
   time2          IN  TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   pctthreshold   IN  NUMBER                   DEFAULT 10)
RETURN DiffRepTab pipelined;
```

Parameters

Table 192-34 DIFF_TABLE_STATS_IN_HISTORY Function Parameters

Parameter	Description
ownname	Specifies the owner of the table. Specify NULL for current schema.
tabname	Specifies the table for which statistics are to be compared.
time1	Specifies the first timestamp for comparison.
time2	Specifies the second timestamp for comparison.
pctthreshold	Specifies the threshold limit. The function reports differences in statistics only if the change percentage exceeds this limit. The default value is 10.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Usage Notes

If the second timestamp is NULL, then the function compares the current statistics in the data dictionary with the statistics as of the first timestamp.

192.7.29 DIFF_TABLE_STATS_IN_PENDING Function

This function compares pending statistics to either the current statistics in the data dictionary, or user-specified historical statistics.

Syntax

```
DBMS_STATS.DIFF_TABLE_STATS_IN_PENDING(
      ownname        IN  VARCHAR2,
      tabname        IN  VARCHAR2,
      timestamp      IN  TIMESTAMP WITH TIME ZONE,
      pctthreshold   IN  NUMBER  DEFAULT 10)
RETURN DiffRepTab pipelined;
```

Parameters

Table 192-35 DIFF_TABLE_STATS_IN_PENDING Function Parameters

Parameter	Description
ownname	Owner of the table. Specify NULL for the current schema.
tabname	Table for which statistics are to be compared.
timestamp	Timestamp in the statistics history that corresponds to the desired statistics. If the timestamp is NULL, then this function compares the current statistics in the dictionary with the pending statistics (default).
pctthreshold	Limit for reporting. The function reports difference in statistics only if it exceeds the specified limit. The default value is 10.

Security Model

To invoke this procedure you must be owner of the table, or you must have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table, or you must have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.30 DIFF_TABLE_STATS_IN_STATTAB Function

This function compares table statistics from two sources.

The function can obtain statistics from the following sources:

- Two user statistics tables
- A single user statistics table containing two sets of statistics that can be identified using statids
- A user statistics table and dictionary

The function also compares the statistics of the dependent objects: indexes, columns, and partitions. It displays statistics of the objects from both sources when the difference between those statistics exceeds a certain threshold (%). You can specify this threshold as an argument to the function. The function uses the statistics corresponding to the first source (stattab1 or time1) as the basis for computing the difference percentage.

Syntax

```
DBMS_STATS.DIFF_TABLE_STATS_IN_STATTAB(
   ownname         IN  VARCHAR2,
   tabname         IN  VARCHAR2,
   stattab1        IN  VARCHAR2,
   stattab2        IN  VARCHAR2 DEFAULT NULL,
   pctthreshold    IN  NUMBER   DEFAULT 10,
   statid1         IN  VARCHAR2 DEFAULT NULL,
   statid2         IN  VARCHAR2 DEFAULT NULL,
   stattab1own     IN  VARCHAR2 DEFAULT NULL,
   stattab2own     IN  VARCHAR2 DEFAULT NULL)
RETURN DiffRepTab pipelined;
```

Parameters

Table 192-36 DIFF_TABLE_STATS_IN_STATTAB Function Parameters

Parameter	Description
ownname	Specifies the owner of the table. Specify NULL for current schema.
tabname	Specifies the table for which statistics are to be compared.
stattab1	Specifies the user statistics table 1.
stattab2	Specifies the user statistics table 2. If NULL, the function compares statistics in stattab1 with current statistics in the data dictionary. This is the default. To compare two sets within the statistics table, specify the same table as stattab1 (see statid below).
pctthreshold	Specifies the percent thresholds for comparison. The function reports difference in statistics only if it exceeds this limit. The default value is 10.
stadid1	(optional) Identifies statistics set within stattab1.
stadid2	(optional) Identifies statistics set within stattab2.
stattab1own	Specifies the schema containing stattab1 (if other than ownname).
stattab2own	Specifies the schema containing stattab2 (if other than ownname).

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.31 DROP_ADVISOR_TASK Procedure

This procedure drops the specified Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.DROP_ADVISOR_TASK (
  task_name IN VARCHAR2);
```

Parameters

Table 192-37 DROP_ADVISOR_TASK Procedure Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

Example 192-8 Dropping an Optimizer Statistics Advisor Task

This example drops the Optimizer Statistics Advisor task named my_task:

```
EXEC DBMS_STATS.DROP_ADVISOR_TASK('my_task');
```

192.7.32 DROP_EXTENDED_STATS Procedure

This function drops the statistics entry that is created for the user specified extension.

This cancels the effects of the CREATE_EXTENDED_STATS Function.

Syntax

```
DBMS_STATS.DROP_EXTENDED_STATS (
   ownname    VARCHAR2,
   tabname    VARCHAR2,
   extension  VARCHAR2);
```

Parameters

Table 192-38 DROP_EXTENDED_STATS Procedure Parameters

Parameter	Description
ownname	Owner name of a table
tabname	Name of the table
extension	Can be either a column group or an expression. Suppose the specified table has two column c1, c2. An example column group can be "(c1, c2)" and an example expression can be "(c1 + c2)".

Exceptions

- ORA-20000: Insufficient privileges or extension does not exist
- ORA-20001: Error when processing extension

Usage Notes

- To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.
- If no extended statistics set is created for the extension, this function throws an error.

192.7.33 DROP_STAT_TABLE Procedure

This procedure drops a user statistics table.

Syntax

```
DBMS_STATS.DROP_STAT_TABLE (
   ownname VARCHAR2,
   stattab VARCHAR2);
```

Parameters

Table 192-39 DROP_STAT_TABLE Procedure Parameters

Parameter	Description
ownname	Name of the schema
stattab	User statistics table identifier

Exceptions

ORA-20000: Table does not exists or insufficient privileges.

Usage Notes

To invoke this procedure you need the privileges for dropping the specified table.

192.7.34 EXECUTE_ADVISOR_TASK Function

This function executes a previously created Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.EXECUTE_ADVISOR_TASK (
   task_name        IN   VARCHAR2,
   execution_name   IN   VARCHAR2 := NULL)
 RETURN VARCHAR2;
```

Parameters

Table 192-40 EXECUTE_ADVISOR_TASK Parameters

Parameter	Description
task_name	Name of the Optimizer Statistics Advisor task.
execution_name	A name that qualifies and identifies an advisor execution. If not specified, then the advisor automatically generates it. If the specified execution conflicts with the name of an existing execution, then the function returns an error.

Security Model

Note the following:

- To execute this subprogram, you must have the `ADVISOR` privilege.
- You must be the owner of the task.
- You can execute this subprogram for `AUTO_STATS_ADVISOR_TASK`, which is predefined.
- This subprogram executes using invoker's rights.

The results of performing this task depend on the privileges of the executing user:

- `SYSTEM` level

  Only users with both the `ANALYZE ANY` and `ANALYZE ANY DICTIONARY` privileges can perform this task on system-level rules.
- Operation level

  The results depend on the following privileges:

  – Users with both the `ANALYZE ANY` and `ANALYZE ANY DICTIONARY` privileges can perform this task for all statistics operations.

  – Users with the `ANALYZE ANY` privilege but *not* the `ANALYZE ANY DICTIONARY` privilege can perform this task for statistics operations related to any schema except `SYS`.

  – Users with the `ANALYZE ANY DICTIONARY` privilege but *not* the `ANALYZE ANY` privilege can perform this task for statistics operations related to their own schema and the `SYS` schema.

  – Users with neither the `ANALYZE ANY` nor the `ANALYZE ANY DICTIONARY` privilege can only perform this operation for statistics operations relating to their own schema.
- Object level

  Users can perform this task for any object for which they have statistics collection privileges.

Exceptions

- `ORA-20000`: Insufficient privileges
- `ORA-20001`: Invalid input values
- `ORA-20012`: Optimizer Statistics Advisor errors

Returns

This function returns the name of the new execution.

Usage Notes

The results of the execution depend on user privileges and the type of rules:

- System

  To perform the operation on system-level rules, you must have both the `ANALYZE ANY` and `ANALYZE ANY DICTIONARY` privileges.
- Operation

If you have the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges, then you can execute this function for all operations. If you have only the ANALYZE ANY privilege, then you can execute this function for operations related to any schemas except SYS. If you have only the ANALYZE ANY DICTIONARY privilege, then you can execute this function for operations related to any schemas, including SYS. If you have neither the ANALYZE ANY nor the ANALYZE ANY DICTIONARY privilege, then you can execute this function only for operations in your own schema.

- Object

  If you have the privilege to collect statistics for an object, then you can execute this function for the object.

Example 192-9 Creating and Executing a Task

This example creates an Optimizer Statistics Advisor task named my_task, and then executes it.

```
DECLARE
  v_tname   VARCHAR2(128) := 'my_task';
BEGIN
  -- create a task
  v_tname := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);
  -- execute the task
  v_tname := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);
END;
```

192.7.35 EXPORT_COLUMN_STATS Procedure

This procedure exports statistics for a specified column and stores them in the user statistics table identified by stattab.

Syntax

```
DBMS_STATS.EXPORT_COLUMN_STATS (
   ownname  VARCHAR2,
   tabname  VARCHAR2,
   colname  VARCHAR2,
   partname VARCHAR2 DEFAULT NULL,
   stattab  VARCHAR2,
   statid   VARCHAR2 DEFAULT NULL,
   statown  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-41 EXPORT_COLUMN_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
tabname	Name of the table to which this column belongs
colname	Name of the column or extension

Table 192-41 (Cont.) EXPORT_COLUMN_STATS Procedure Parameters

Parameter	Description
partname	Name of the table partition. If the table is partitioned and if partname is NULL, then global and partition column statistics are exported.
stattab	User statistics table identifier describing where to store the statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.36 EXPORT_DATABASE_PREFS Procedure

This procedure is used to export the statistics preferences of all the tables, excluding the tables owned by Oracle. These tables can by included by passing TRUE for the add_sys parameter.

Syntax

```
DBMS_STATS.EXPORT_DATABASE_PREFS (
    stattab    IN  VARCHAR2,
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT NULL
    add_sys    IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-42 EXPORT_DATABASE_PREFS Procedure Parameters

Parameter	Description
stattab	Statistics table name to where statistics should be exported
statid	(Optional) Identifier to associate with these statistics within stattab
statown	Schema containing stattab (if other than ownname)
add_sys	Value TRUE will include the Oracle-owned tables

Exceptions

ORA-20000: Insufficient privileges

Usage Notes

- To run this procedure, you need to have the SYSDBA role, or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.
- All arguments are of type VARCHAR2 and values are enclosed in quotes.
- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.EXPORT_DATABASE_PREFS('STATTAB', statown=>'SH');
```

192.7.37 EXPORT_DATABASE_STATS Procedure

This procedure exports statistics for all objects in the database and stores them in the user statistics tables identified by statown.stattab.

Syntax

```
DBMS_STATS.EXPORT_DATABASE_STATS (
   stattab         VARCHAR2,
   statid          VARCHAR2 DEFAULT NULL,
   statown         VARCHAR2 DEFAULT NULL,
   stat_category   VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-43 EXPORT_DATABASE_STATS Procedure Parameters

Parameter	Description
stattab	User statistics table identifier describing where to store the statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different from current schema)
stat_category	Specifies what statistics to import, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics (default). • SYNOPSES — information to support incremental statistics. • REALTIME_STATS — specifies only real-time statistics. If 'OBJECT_STATS, SYNOPSES' is specified, table statistics, column statistics, index statistics and synopses are deleted.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.38 EXPORT_DICTIONARY_STATS Procedure

This procedure exports statistics for all data dictionary schemas (`SYS`, `SYSTEM`, and RDBMS component schemas) and stores them in the user statistics table identified by `stattab`.

Syntax

```
DBMS_STATS.EXPORT_DICTIONARY_STATS (
   stattab         VARCHAR2,
   statid          VARCHAR2 DEFAULT NULL,
   statown         VARCHAR2 DEFAULT NULL,
   stat_category   VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-44 EXPORT_DICTIONARY_STATS Procedure Parameters

Parameter	Description
`stattab`	User statistics table identifier describing where to store the statistics
`statid`	Identifier (optional) to associate with these statistics within `stattab`
`statown`	Schema containing `stattab` (if different from current schema)
`stat_category`	Specifies what statistics to import, accepting multiple values separated by a comma. Values supported: • `OBJECT_STATS` — table statistics, column statistics and index statistics (default) • `SYNOPSES` — information to support incremental statistics • `REALTIME_STATS` — specifies only real-time statistics. If `'OBJECT_STATS, SYNOPSES'` is specified, table statistics, column statistics, index statistics and synopses are deleted.

Security Model

To invoke this procedure you must be owner of the table or have the `ANALYZE ANY` privilege. For objects owned by `SYS`, you must be either the owner of the table or have either the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20002`: Bad user statistics table, may need to upgrade it

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.39 EXPORT_FIXED_OBJECTS_STATS Procedure

This procedure exports statistics for fixed tables and stores them in the user statistics table identified by `stattab`.

Syntax

```
DBMS_STATS.EXPORT_FIXED_OBJECTS_STATS (
   stattab  VARCHAR2,
   statid   VARCHAR2 DEFAULT NULL,
   statown  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-45 EXPORT_FIXED_OBJECTS_STATS Procedure Parameters

Parameter	Description
`stattab`	User statistics table identifier describing where to store the statistics
`statid`	Identifier (optional) to associate with these statistics within `stattab`
`statown`	Schema containing `stattab` (if different from current schema)

Security Model

To invoke this procedure you must be owner of the table or have the `ANALYZE ANY` privilege. For objects owned by `SYS`, you must be either the owner of the table or have either the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20002`: Bad user statistics table, may need to upgrade it

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.40 EXPORT_INDEX_STATS Procedure

This procedure retrieves statistics for a particular index and stores them in the user statistics table identified by `stattab`.

Syntax

```
DBMS_STATS.EXPORT_INDEX_STATS (
   ownname  VARCHAR2,
   indname  VARCHAR2,
   partname VARCHAR2 DEFAULT NULL,
   stattab  VARCHAR2,
   statid   VARCHAR2 DEFAULT NULL,
   statown  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-46 EXPORT_INDEX_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
indname	Name of the index
partname	Name of the index partition. If the index is partitioned and if partname is NULL, then global and partition index statistics are exported.
stattab	User statistics table identifier describing where to store the statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

- To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.
- Oracle does not support export or import of statistics across databases of different character sets.

192.7.41 EXPORT_PENDING_STATS Procedure

This procedure is used to export the statistics gathered and stored as pending.

Syntax

```
DBMS_STATS.EXPORT_PENDING_STATS (
    ownname    IN  VARCHAR2  DEFAULT USER,
    tabname    IN  VARCHAR2,
    stattab    IN  VARCHAR2,
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT USER);
```

Parameters

Table 192-47 EXPORT_PENDING_STATS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name
stattab	Statistics table name to where to export the statistics
statid	(Optional) Identifier to associate with these statistics within stattab
statown	Schema containing stattab (if other than ownname)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

- If the parameter tabname is NULL then export applies to all tables of the specified schema.
- The default owner/schema is the user who runs the procedure.
- To run this procedure, you need to have the same privilege for gathering statistics on the tables that will be touched by this procedure.
- All arguments are of type VARCHAR2 and values are enclosed in quotes.
- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.EXPORT_PENDING_STATS(NULL, NULL, 'MY_STAT_TABLE');
```

192.7.42 EXPORT_SCHEMA_PREFS Procedure

This procedure is used to export the statistics preferences of all the tables owned by the specified owner name.

Syntax

```
DBMS_STATS.EXPORT_SCHEMA_PREFS (
    ownname    IN  VARCHAR2,
    stattab    IN  VARCHAR2,
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-48 EXPORT_SCHEMA_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
stattab	Statistics table name to where to export the statistics
statid	(Optional) Identifier to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

- To run this procedure, you need to connect as owner, or have the SYSDBA privilege, or have the ANALYZE ANY system privilege.
- All arguments are of type VARCHAR2 and values are enclosed in quotes.

- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.EXPORT_SCHEMA_PREFS('SH', 'STAT');
```

192.7.43 EXPORT_SCHEMA_STATS Procedure

This procedure exports statistics for all objects in the schema identified by `ownname` and stores them in the user statistics tables identified by `stattab`.

Syntax

```
DBMS_STATS.EXPORT_SCHEMA_STATS (
   ownname         VARCHAR2,
   stattab         VARCHAR2,
   statid          VARCHAR2 DEFAULT NULL,
   statown         VARCHAR2 DEFAULT NULL,
   stat_category   VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-49 EXPORT_SCHEMA_STATS Procedure Parameters

Parameter	Description
`ownname`	Specifies the name of the schema.
`stattab`	Identifies the user statistics table in which to store the exported statistics.
`statid`	Specifies the identifier (optional) associated with these statistics within `stattab`.
`statown`	Specifies the schema containing `stattab` (if different than `ownname`).
`stat_category`	Specifies which statistics to process. The following values are supported: • `OBJECT_STATS` — specifies table statistics, column statistics, and index statistics • `SYNOPSES` — specifies metadata for incremental statistics • `REALTIME_STATS` — specifies only real-time statistics You can specify a list of comma-delimited values. For example, '`OBJECT_STATS, SYNOPSES`' specifies table statistics, column statistics, index statistics, and synopses. The default value is '`OBJECT_STATS, SYNOPSES, REALTIME_STATS`'.

Security Model

To invoke this procedure you must be owner of the table or have the `ANALYZE ANY` privilege. For objects owned by `SYS`, you must be either the owner of the table or have either the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.44 EXPORT_SYSTEM_STATS Procedure

This procedure retrieves system statistics and stores them in the user statistics table, identified by `stattab`.

Syntax

```
DBMS_STATS.EXPORT_SYSTEM_STATS (
   stattab       VARCHAR2,
   statid        VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-50 EXPORT_SYSTEM_STATS Procedure Parameters

Parameter	Description
`stattab`	Identifier of the user statistics table that describes where the statistics will be stored
`statid`	Optional identifier associated with the statistics stored from the `stattab`
`statown`	Schema containing `stattab` (if different from current schema)

Security Model

To run this procedure, you must have the `GATHER_SYSTEM_STATISTICS` role.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20002`: Bad user statistics table; may need to be upgraded

`ORA-20003`: Unable to export system statistics

Usage Notes

Oracle Database does not support the export or import of statistics across databases of different character sets.

192.7.45 EXPORT_TABLE_PREFS Procedure

This procedure is used to export the statistics preferences of the specified table in the specified schema into the specified statistics table.

Syntax

```
DBMS_STATS.EXPORT_TABLE_PREFS (
    ownname    IN  VARCHAR2,
    tabname    IN  VARCHAR2,
    stattab    IN  VARCHAR2,
```

```
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-51 EXPORT_TABLE_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name
stattab	Statistics table name where to export the statistics
statid	Optional identifier to associate with these statistics within stattab
statown	Schema containing stattab (if other than ownname)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

- To run this procedure, you need to connect as owner of the table, or have the ANALYZE ANY system privilege.
- All arguments are of type VARCHAR2 and values are enclosed in quotes.
- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.EXPORT_TABLE_PREFS('SH', 'SALES', 'STAT');
```

192.7.46 EXPORT_TABLE_STATS Procedure

This procedure exports statistics for a specified table (including associated index statistics) and stores them in the user statistics table identified by stattab.

Syntax

```
DBMS_STATS.EXPORT_TABLE_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   stattab          VARCHAR2,
   statid           VARCHAR2 DEFAULT NULL,
   cascade          BOOLEAN  DEFAULT TRUE,
   statown          VARCHAR2 DEFAULT NULL,
   stat_category    VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-52 EXPORT_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	Specifies the name of the schema.
tabname	Specifies the name of the table.
partname	Specifies the name of the table partition. If the table is partitioned, and if partname is NULL, then the procedure exports global and partition statistics.
stattab	Specifies the identifier (optional) associated with these statistics within stattab.
statid	Specifies the identifier (optional) associated with these statistics within stattab.
cascade	Indicates whether to export column and index statistics. If TRUE, then the procedure exports column and index statistics for the specified table. This is the default.
statown	Specifies the schema containing stattab (if different than ownname).
stat_category	Specifies which statistics to process. The following values are supported: • OBJECT_STATS — specifies table statistics, column statistics, and index statistics • SYNOPSES — specifies metadata for incremental statistics • REALTIME_STATS — specifies only real-time statistics You can specify a list of comma-delimited values. For example, 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS'.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.47 FLUSH_DATABASE_MONITORING_INFO Procedure

This procedure saves monitoring information for all tables in the dictionary. The database immediately updates corresponding entries in the `*_TAB_MODIFICATIONS`, `*_TAB_STATISTICS` and `*_IND_STATISTICS` views.

Syntax

```
DBMS_STATS.FLUSH_DATABASE_MONITORING_INFO;
```

Security Model

The `ANALYZE_ANY` system privilege is required to run this procedure.

Exceptions

`ORA-20000`: Insufficient privileges

Usage Notes

Starting in Oracle Database 12c Release 2 (12.2), you do not need to call `FLUSH_DATABASE_MONITORING_INFO` to view the latest information in `*_TAB_STATISTICS` and `*_IND_STATISTICS` because these views show statistics cached in the SGA and stored on disk. Because the `GATHER_*_STATS` procedures internally save monitoring information to disk, it is not necessary to run this procedure before gathering statistics.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set optimizer statistics preferences

192.7.48 GATHER_DATABASE_STATS Procedures

This procedure gathers statistics for all objects in the database.

Syntax

```
DBMS_STATS.GATHER_DATABASE_STATS (
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                        (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   options          VARCHAR2 DEFAULT 'GATHER',
   objlist          OUT      ObjectTab,
   statown          VARCHAR2 DEFAULT NULL,
   gather_sys       BOOLEAN  DEFAULT TRUE,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   obj_filter_list ObjectTab DEFAULT NULL);
```

```
DBMS_STATS.GATHER_DATABASE_STATS (
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                        (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   options          VARCHAR2 DEFAULT 'GATHER',
   statown          VARCHAR2 DEFAULT NULL,
   gather_sys       BOOLEAN  DEFAULT TRUE,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   obj_filter_list ObjectTab DEFAULT NULL);
```

Parameters

Table 192-53 GATHER_DATABASE_STATS Procedure Parameters

Parameter	Description
estimate_percent	Determines the percentage of rows to sample. The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for optimal statistics. This is the default. You can change the default value using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Determines whether the database uses random block sampling (TRUE) or random row sampling (FALSE). The default is FALSE. Random block sampling is more efficient, but if the data is not randomly distributed on disk, then sample values may be somewhat correlated. This parameter is only relevant when estimating statistics.
method_opt	When setting preference on global, schema, database or dictionary level, only'FOR ALL' syntax is allowed.: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms. – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO.The value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Table 192-53 (Cont.) GATHER_DATABASE_STATS Procedure Parameters

Parameter	Description
degree	Determines the degree of parallelism used for gathering statistics. The default for degree is NULL. NULL means to use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Change the default using the SET_DATABASE_PREFS, SET_GLOBAL_PREFS, SET_SCHEMA_PREFS, and SET_TABLE_PREFS procedures. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. The degree is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters), according to the size of the object. When using DEGREE=>NULL, DEGREE=>*n*, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution when the size of the object does not warrant parallel execution.
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics
cascade	Determines whether to collect index statistics as part of gathering table statistics. Specifying this option is equivalent to running the GATHER_INDEX_STATS procedure on each index of the table. Use the constant DBMS_STATS.AUTO_CASCADE to enable the database to determine whether index statistics need to be collected. This is the default. You can change the default using the SET_DATABASE_PREFS, SET_GLOBAL_PREFS, SET_SCHEMA_PREFS, and SET_TABLE_PREFS procedures.
stattab	User statistics table identifier describing where to save the current statistics. The statistics table is assumed to reside in the same schema as the object being analyzed, so there must be one such table in each schema to use this option.
statid	Identifier (optional) to associate with these statistics within stattab.

Table 192-53 (Cont.) GATHER_DATABASE_STATS Procedure Parameters

Parameter	Description
options	Specifies which objects require statistics to be gathered. Valid values are as follows: • GATHER — Gathers statistics on all objects in the database. This is the default. • GATHER AUTO — Gathers all necessary statistics automatically. The database implicitly determines which objects need new statistics and determines how to gather those statistics. When GATHER AUTO is specified, the only additional valid parameters are comp_id, no_invalidate, stattab, statid, and statown; all other parameter settings will be ignored. Also, the database returns a list of objects processed. • GATHER STALE — Gathers statistics on stale objects by querying the *_TAB_MODIFICATIONS views. Also, the database returns a list of objects found to be stale. • GATHER EMPTY — Gathers statistics on objects that currently have no statistics. Also, the database returns a list of objects found to have no statistics. • LIST AUTO — Returns a list of objects to be processed with GATHER AUTO. • LIST STALE — Returns list of stale objects as determined by looking at the *_TAB_MODIFICATIONS views. • LIST EMPTY — Returns a list of objects that currently have no statistics.
objlist	List of objects found to be stale or empty
statown	Schema containing stattab (if different from current schema)
gather_sys	Gathers statistics on the objects owned by the SYS user.
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. If set to TRUE, then the database not invalidate dependent cursors. If set to FALSE, then the procedure invalidates dependent cursors immediately.
obj_filter_list	A list of object filters. When provided, GATHER_DATABASE_STATS will gather statistics only on objects which satisfy at least one object filter in the list as needed. In a single object filter, we can specify the constraints on the object attributes. The attribute values specified in the object filter are case- insensitive unless double-quoted. Wildcard is allowed in the attribute values. Suppose non-NULL values s1, s2, ... are specified for attributes a1, a2, ... in one object filter. An object o is said to satisfy this object filter if (o.a1 like s1) or (o.a2 like s2) or ... is true.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Bad input value

Usage Notes

To run this procedure, you need to have the SYSDBA role or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.

If the GATHER AUTO option is used then frequency histograms will be created using a sample rather than a full scan. The following scenario shows where GATHER AUTO is used:

- A table is created like this: CREATE TABLE NEWTAB as SELECT * FROM .....

  This will create statistics on NEWTAB but no histograms.
- Next, the DBA creates the histograms using GATHER AUTO on gather_table_stats.
- The FREQUENCY histograms on NEWTAB will be created using a sample rather than a full table scan.

192.7.49 GATHER_DICTIONARY_STATS Procedure

This procedure gathers statistics for dictionary schemas SYS, SYSTEM and schemas of RDBMS components.

Syntax

```
DBMS_STATS.GATHER_DICTIONARY_STATS (
   comp_id          VARCHAR2 DEFAULT NULL,
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                        (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   options          VARCHAR2 DEFAULT 'GATHER AUTO',
   objlist    OUT   ObjectTab,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   obj_filter_list ObjectTab DEFAULT NULL);

DBMS_STATS.GATHER_DICTIONARY_STATS (
   comp_id           VARCHAR2 DEFAULT NULL,
   estimate_percent  NUMBER DEFAULT
                        to_estimate_percent_type(GET_PARAM('ESTIMATE_PERCENT')),
   block_sample      BOOLEAN DEFAULT FALSE,
   method_opt        VARCHAR2 DEFAULT GET_PARAM('METHOD_OPT'),
   degree            NUMBER DEFAULT to_degree_type(GET_PARAM('DEGREE')),
   granularity       VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade           BOOLEAN DEFAULT to_cascade_type(GET_PARAM('CASCADE')),
   stattab           VARCHAR2 DEFAULT NULL,
   statid            VARCHAR2 DEFAULT NULL,
   options           VARCHAR2 DEFAULT 'GATHER AUTO',
   statown           VARCHAR2 DEFAULT NULL,
   no_invalidate     BOOLEAN DEFAULT
                             to_no_invalidate_type(get_param('NO_INVALIDATE')),
   obj_filter_list   ObjectTab DEFAULT NULL);
```

Parameters

Table 192-54 GATHER_DICTIONARY_STATS Procedure Parameters

Parameter	Description
comp_id	Component id of the schema to analyze. NULL results in the analysis of schemas for all RDBMS components. Refer to the COMP_ID column of DBA_REGISTRY view. The procedure always gather statistics on SYS and SYSTEM schemas regardless of this argument.
estimate_percent	Percentage of rows to estimate (NULL means compute). The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Determines whether or not to use random block sampling instead of random row sampling. Random block sampling is more efficient, but if the data is not randomly distributed on disk then the sample values may be somewhat correlated. Only pertinent when performing estimate statistics.
method_opt	The method options. This parameter accepts the following values: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms. – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.

Table 192-54 (Cont.) GATHER_DICTIONARY_STATS Procedure Parameters

Parameter	Description
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - gathers subpartition-level statistics
cascade	Gathers statistics on indexes also. Index statistics gathering will not be parallelized. Using this option is equivalent to running the GATHER_INDEX_STATS Procedure on each of the indexes in the schema in addition to gathering table and column statistics. Use the constant DBMS_STATS.AUTO_CASCADE to have Oracle determine whether index statistics to be collected or not. This is the default.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics
statid	The (optional) identifier to associate with these statistics within stattab

Table 192-54 (Cont.) GATHER_DICTIONARY_STATS Procedure Parameters

Parameter	Description
options	Specifies which objects require statistics to be gathered. Valid values are as follows: • GATHER AUTO — Gathers all necessary statistics automatically. This is the default. The database implicitly determines which objects need new statistics and determines how to gather those statistics. When GATHER AUTO is specified, the only additional valid parameters are comp_id, no_invalidate, stattab, statid, and statown; all other parameter settings will be ignored. Also, the database returns a list of objects processed. • GATHER — Gathers statistics on all objects in the relevant schema. • GATHER STALE — Gathers statistics on stale objects by querying the *_TAB_MODIFICATIONS views. Also, the database returns a list of objects found to be stale. • GATHER EMPTY — Gathers statistics on objects that currently have no statistics. Also, the database returns a list of objects found to have no statistics. • LIST AUTO — Returns a list of objects to be processed with GATHER AUTO. • LIST STALE — Returns list of stale objects as determined by looking at the *_TAB_MODIFICATIONS views. • LIST EMPTY — Returns a list of objects that currently have no statistics.
objlist	The list of objects found to be stale or empty.
statown	Schema containing stattab, if different from the current schema.
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
obj_filter_list	A list of object filters. When provided, this will gather statistics only on objects which satisfy at least one object filter in the list as needed. In a single object filter, we can specify the constraints on the object attributes. The attribute values specified in the object filter are case-insensitive unless double-quoted. Wildcard is allowed in the attribute values. Suppose non-NULL values s1, s2, ... are specified for attributes a1, a2, ... in one object filter. An object o is said to satisfy this object filter if (o.a1 like s1) and (o.a2 like s2) and ... is true. See Applying an Object Filter List.

Usage Notes

You must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege to execute this procedure.

If the GATHER AUTO option is used then frequency histograms will be created using a sample rather than a full scan. The following scenario shows where GATHER AUTO is used:

- A table is created like this: CREATE TABLE NEWTAB as SELECT * FROM .....

  This will create statistics on NEWTAB but no histograms.
- Next, the DBA creates the histograms using GATHER AUTO on gather_table_stats.
- The FREQUENCY histograms on NEWTAB will be created using a sample rather than a full table scan.

Exceptions

ORA-20000: Index does not exist or insufficient privileges

ORA-20001: Bad input value

ORA-20002: Bad user statistics table, may need to upgrade it

192.7.50 GATHER_FIXED_OBJECTS_STATS Procedure

This procedure gathers statistics for all fixed objects (dynamic performance tables).

Syntax

```
DBMS_STATS.GATHER_FIXED_OBJECTS_STATS (
   stattab        VARCHAR2 DEFAULT NULL,
   statid         VARCHAR2 DEFAULT NULL,
   statown        VARCHAR2 DEFAULT NULL,
   no_invalidate  BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')));
```

Parameters

Table 192-55 GATHER_FIXED_OBJECTS_STATS Procedure Parameters

Parameter	Description
stattab	User statistics table identifier describing where to save the current statistics
statid	Identifier to associate with these statistics within stattab (optional)
statown	Schema containing stattab (if different from current schema)
no_invalidate	Does not invalidate the dependent cursors if set to TRUE. The procedure invalidates the dependent cursors immediately if set to FALSE. Use DBMS_STATS.AUTO_INVALIDATE. to have Oracle decide when to invalidate dependent cursors. This is the default. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Usage Notes

You must have the SYSDBA or ANALYZE ANY DICTIONARY system privilege to execute this procedure.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Bad input value

ORA-20002: Bad user statistics table, may need to upgrade it

192.7.51 GATHER_INDEX_STATS Procedure

This procedure gathers index statistics. It attempts to parallelize as much of the work as possible.

Restrictions are described in the individual parameters. This operation will not parallelize with certain types of indexes, including cluster indexes, domain indexes, and bitmap join indexes. The granularity and no_invalidate arguments are not relevant to these types of indexes.

Syntax

```
DBMS_STATS.GATHER_INDEX_STATS (
   ownname          VARCHAR2,
   indname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                        (GET_PARAM('ESTIMATE_PERCENT')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type
                                       (GET_PARAM('NO_INVALIDATE')),
   force            BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-56 GATHER_INDEX_STATS Procedure Parameters

Parameter	Description
ownname	Schema of index to analyze
indname	Name of index
partname	Name of partition
estimate_percent	Percentage of rows to estimate (NULL means compute). The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)

Table 192-56 (Cont.) GATHER_INDEX_STATS Procedure Parameters

Parameter	Description
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics.
no_invalidate	Does not invalidate the dependent cursors if set to TRUE. The procedure invalidates the dependent cursors immediately if set to FALSE. Use DBMS_STATS.AUTO_INVALIDATE. to have Oracle decide when to invalidate dependent cursors. This is the default. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Gather statistics on object even if it is locked

Exceptions

ORA-20000: Index does not exist or insufficient privileges

ORA-20001: Bad input value

Usage Notes

To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

192.7.52 GATHER_PROCESSING_RATE Procedure

This procedure starts the job of gathering the processing rates which end after an `interval` defined in minutes.

Syntax

```
DBMS_STATS.GATHER_PROCESSING_RATE (
   gathering_mode      IN    VARCHAR2  DEFAULT 'START',
   interval            IN    NUMBER    DEFAULT  NULL);
```

Parameters

Table 192-57 GATHER_PROCESSING_RATE Procedure Parameters

Parameter	Description
`gathering_mode`	Mode: `'START'` or `'END'`. The mode is based on the Active Session History (ASH) data when invoked with `'START'` option. It stops gathering when invoked with `'END'` option. When invoked with `'START'`, 'interval' option can be specified optionally. If `interval` is not specified, its default value is set to 60 minutes.
`interval`	Time interval (number of minutes) for which the processing must be gathered

Usage Notes

- You require the `OPTIMIZER_PROCESSING_RATE` role to run this procedure.
- `AUTO DOP` uses processing rates to determine the optimal degree of parallelism for a SQL statement.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20001`: Invalid or illegal input value

192.7.53 GATHER_SCHEMA_STATS Procedures

This procedure gathers statistics for all objects in a schema.

Syntax

```
DBMS_STATS.GATHER_SCHEMA_STATS (
   ownname          VARCHAR2,
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                          (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
```

```
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   options          VARCHAR2 DEFAULT 'GATHER',
   objlist          OUT      ObjectTab,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
  force             BOOLEAN DEFAULT FALSE,
  obj_filter_list  ObjectTab DEFAULT NULL);

DBMS_STATS.GATHER_SCHEMA_STATS (
   ownname          VARCHAR2,
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                          (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   options          VARCHAR2 DEFAULT 'GATHER',
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type (
                                     get_param('NO_INVALIDATE')),
   force            BOOLEAN DEFAULT FALSE,
   obj_filter_list  ObjectTab DEFAULT NULL);
```

Parameters

Table 192-58 GATHER_SCHEMA_STATS Procedure Parameters

Parameter	Description
ownname	Schema to analyze (NULL means current schema)
estimate_percent	Determines the percentage of rows to sample. The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for optimal statistics. This is the default. You can change the default value using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Whether or not to use random block sampling instead of random row sampling. Random block sampling is more efficient, but if the data is not randomly distributed on disk, then the sample values may be somewhat correlated. Only pertinent when doing an estimate statistics.

Table 192-58 (Cont.) GATHER_SCHEMA_STATS Procedure Parameters

Parameter	Description
method_opt	Accepts: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics.

Table 192-58 (Cont.) GATHER_SCHEMA_STATS Procedure Parameters

Parameter	Description
cascade	Gather statistics on the indexes as well. Using this option is equivalent to running the GATHER_INDEX_STATS Procedure on each of the indexes in the schema in addition to gathering table and column statistics. Use the constant DBMS_STATS.AUTO_CASCADE to have Oracle determine whether index statistics to be collected or not. This is the default. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics
statid	Identifier (optional) to associate with these statistics within stattab
options	Specifies which objects require statistics to be gathered. Valid values are as follows: • GATHER — Gathers statistics on all objects in the schema. This is the default. • GATHER AUTO — Gathers all necessary statistics automatically. The database implicitly determines which objects need new statistics and determines how to gather those statistics. When GATHER AUTO is specified, the only additional valid parameters are comp_id, no_invalidate, stattab, statid, and statown; all other parameter settings will be ignored. Also, the database returns a list of objects processed. • GATHER STALE — Gathers statistics on stale objects by querying the *_TAB_MODIFICATIONS views. Also, the database returns a list of objects found to be stale. • GATHER EMPTY — Gathers statistics on objects that currently have no statistics. Also, the database returns a list of objects found to have no statistics. • LIST AUTO — Returns a list of objects to be processed with GATHER AUTO. • LIST STALE — Returns list of stale objects as determined by looking at the *_TAB_MODIFICATIONS views. • LIST EMPTY — Returns a list of objects that currently have no statistics.
objlist	List of objects found to be stale or empty
statown	Schema containing stattab (if different than ownname)
no_invalidate	Does not invalidate the dependent cursors if set to TRUE. The procedure invalidates the dependent cursors immediately if set to FALSE. Use DBMS_STATS.AUTO_INVALIDATE. to have Oracle decide when to invalidate dependent cursors. This is the default. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Gather statistics on objects even if they are locked

Table 192-58 (Cont.) GATHER_SCHEMA_STATS Procedure Parameters

Parameter	Description
obj_filter_list	A list of object filters. When provided, GATHER_SCHEMA_STATS will gather statistics only on objects which satisfy at least one object filter in the list as needed. In a single object filter, we can specify the constraints on the object attributes. The attribute values specified in the object filter are case- insensitive unless double-quoted. Wildcard is allowed in the attribute values. Suppose non-NULL values s1, s2, ... are specified for attributes a1, a2, ... in one object filter. An object o is said to satisfy this object filter if (o.a1 like s1) and (o.a2 like s2) and ... is true.

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

If the GATHER AUTO option is used then frequency histograms will be created using a sample rather than a full scan. The following scenario shows where GATHER AUTO is used:

- A table is created like this: CREATE TABLE NEWTAB as SELECT * FROM .....

  This will create statistics on NEWTAB but no histograms.
- Next, the DBA creates the histograms using GATHER AUTO on gather_table_stats.
- The FREQUENCY histograms on NEWTAB will be created using a sample rather than a full table scan.

Exceptions

ORA-20000: Schema does not exist or insufficient privileges

ORA-20001: Bad input value

Examples

Applying an Object Filter List

The following example specifies that the tables SH.SALES and SH.COSTS, if stale, will have statistics gathered upon them.

```
DECLARE
   filter_lst  DBMS_STATS.OBJECTTAB := DBMS_STATS.OBJECTTAB();
BEGIN
   filter_lst.extend(2);
   filter_lst(1).ownname := 'SH';
   filter_lst(1).objname := 'SALES';
   filter_lst(2).ownname := 'SH';
   filter_lst(2).objname := 'COSTS';
   DBMS_STATS.GATHER_SCHEMA_STATS(ownname=>'SH',obj_filter_list=>filter_lst);
 END;
```

192.7.54 GATHER_SYSTEM_STATS Procedure

This procedure gathers system statistics.

Syntax

```
DBMS_STATS.GATHER_SYSTEM_STATS (
   gathering_mode   VARCHAR2 DEFAULT 'NOWORKLOAD',
   interval         INTEGER  DEFAULT NULL,
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-59 GATHER_SYSTEM_STATS Procedure Parameters

Parameter	Description
gathering_mode	Specifies the mode in which the database gathers system statistics. Possible values are: • NOWORKLOAD The database captures performance characteristics of the I/O system. Gathering may take a few minutes and depends on the size of the database. During this period, the database estimates the average read seek time and transfer speed for the I/O system. This mode is suitable for the all workloads. To fine tune system statistics for the workload, use START and STOP or the INTERVAL option. If you gather both NOWORKLOAD and workload-specific statistics (statistics collected using INTERVAL or START and STOP), the optimizer uses the workload statistics. Collected components include cpuspeednw, ioseektim, and iotfrspeed. • INTERVAL The database captures system activity during a specified interval in minutes. This parameter works in combination with the interval parameter. The database creates or updates system statistics in the dictionary or stattab. You can use GATHER_SYSTEM_STATS (gathering_mode=>'STOP') to stop gathering earlier than scheduled. Collected components include maxthr, slavethr, cpuspeed, sreadtim, mreadtim, and mbrc. • START | STOP The database captures system activity during specified start and stop times and refreshes the dictionary or stattab with statistics for the elapsed period. The database ignores the INTERVAL value. Collected components include maxthr, slavethr, cpuspeed, sreadtim, mreadtim, and mbrc. • EXADATA In this mode, gathered system statistics take into account the unique capabilities of Oracle Exadata, such as large I/O size and high I/O throughput. The database sets multiblock read count and I/O throughput statistics along with CPU speed.
interval	Specifies the number of minutes in which to gather system statistics. This parameter applies only when gathering_mode='INTERVAL'.

Table 192-59 (Cont.) GATHER_SYSTEM_STATS Procedure Parameters

Parameter	Description
stattab	Specifies the table in which the database stores the statistics.
statid	Specifies an optional identifier associated with the statistics saved in stattab.
statown	Specifies the schema containing stattab, if different from the current schema.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid input value

ORA-20002: Bad user statistics table; may need to be upgraded

ORA-20003: Unable to gather system statistics

ORA-20004: Error in the INTERVAL mode: system parameter job_queue_processes must be >0

Usage Notes

To run this procedure, you must have the GATHER_SYSTEM_STATISTICS role.

Examples

Assume that you want to perform database application processing OLTP transactions during the day and run reports at night.

To collect daytime system statistics, gather statistics for 720 minutes. Store the statistics in the MYSTATS table.

```
BEGIN
   DBMS_STATS.GATHER_SYSTEM_STATS (
      interval => 720,
      stattab  => 'mystats',
      statid   => 'OLTP');
END;
```

To collect nighttime system statistics, gather statistics for 720 minutes. Store the statistics in the MYSTATS table.

```
BEGIN
   DBMS_STATS.GATHER_SYSTEM_STATS (
      interval => 720,
      stattab  => 'mystats',
      statid   => 'OLAP');
END;
```

Update the dictionary with the gathered statistics.

```
VARIABLE   jobno   number;
BEGIN
```

```
   DBMS_JOB.SUBMIT (:jobno, 'DBMS_STATS.IMPORT_SYSTEM_STATS
   (''mystats'',''OLTP'');'
   sysdate, 'sysdate + 1');
   COMMIT;
END;

BEGIN
   DBMS_JOB.SUBMIT (:jobno, 'DBMS_STATS.IMPORT_SYSTEM_STATS
   (''mystats'',''OLAP'');'
   sysdate + 0.5, 'sysdate + 1');
   COMMIT;
END;
```

192.7.55 GATHER_TABLE_STATS Procedure

This procedure gathers table, column, and index statistics. It attempts to parallelize as much work as possible, but there are some restrictions, which are described in the individual parameters.

Syntax

```
DBMS_STATS.GATHER_TABLE_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                          (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT  to_no_invalidate_type (
                                            get_param('NO_INVALIDATE')),
   stattype         VARCHAR2 DEFAULT 'DATA',
   force            BOOLEAN  DEFAULT FALSE,
   context          DBMS_STATS.CCONTEXT DEFAULT NULL, -- non operative
   options          VARCHAR2 DEFAULT get_param('OPTIONS'));
```

Parameters

Table 192-60 GATHER_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	Schema containing the table.
tabname	Name of the table.
partname	Name of the partition.

Table 192-60 (Cont.) GATHER_TABLE_STATS Procedure Parameters

Parameter	Description
estimate_percent	Determines the percentage of rows to sample. The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for optimal statistics. This is the default. You can change the default value using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Determines whether the database uses random block sampling (TRUE) or random row sampling (FALSE). The default is FALSE. Random block sampling is more efficient, but if the data is not randomly distributed on disk, then sample values may be somewhat correlated. This parameter is only relevant when estimating statistics.
method_opt	METHOD_OPT - When setting preference on global, schema, database or dictionary level, only 'FOR ALL' syntax is allowed. Other than that, method_opt accepts either of the following options, or both in combination: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] • FOR COLUMNS [column_clause] [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} column_clause is defined as column_clause := column_name | extension name | extension – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. - column_name : Name of a column - extension : can be either a column group in the format of (column_name, Colume_name [, ...]) or an expression The default is FOR ALL COLUMNS SIZE AUTO. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Table 192-60 (Cont.) GATHER_TABLE_STATS Procedure Parameters

Parameter	Description
degree	Determines the degree of parallelism used for gathering statistics. The default for degree is NULL. NULL means to use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Change the default using the SET_DATABASE_PREFS, SET_GLOBAL_PREFS, SET_SCHEMA_PREFS, and SET_TABLE_PREFS procedures. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. The degree is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters), according to the size of the object. When using DEGREE=>NULL, DEGREE=>*n*, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution when the size of the object does not warrant parallel execution.
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'APPROX_GLOBAL AND PARTITION' - similar to 'GLOBAL AND PARTITION' but in this case the global statistics are aggregated from partition level statistics. This option will aggregate all statistics except the number of distinct values for columns and number of distinct keys of indexes. The existing histograms of the columns at the table level are also aggregated.The aggregation will use only partitions with statistics, so to get accurate global statistics, users should make sure to have statistics for all partitions. Global statistics are gathered if partname is NULL or if the aggregation cannot be performed (for example, if statistics for one of the partitions is missing). 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics.

Table 192-60 (Cont.) GATHER_TABLE_STATS Procedure Parameters

Parameter	Description
cascade	Determines whether to collect index statistics as part of gathering table statistics. Specifying this option is equivalent to running the GATHER_INDEX_STATS procedure on each index of the table. Use the constant DBMS_STATS.AUTO_CASCADE to enable the database to determine whether index statistics need to be collected. This is the default. You can change the default using the SET_DATABASE_PREFS, SET_GLOBAL_PREFS, SET_SCHEMA_PREFS, and SET_TABLE_PREFS procedures.
stattab	User statistics table identifier describing where to save the current statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. If set to TRUE, then the database not invalidate dependent cursors. If set to FALSE, then the procedure invalidates dependent cursors immediately.
stattype	Statistics type. The only value allowed is DATA.
force	Gather statistics of table even if it is locked
context	Not used.
options	Determines the options parameter used in the GATHER_TABLE_STATS procedure. The preference takes the following values: • GATHER — Gathers statistics for all objects in the table. This is the default. • GATHER AUTO — Oracle recommends using GATHER AUTO to gather necessary statistics, such as histograms, after a table has been bulk loaded and acquired online statistics. This is applicable only to tables that are not using INCREMENTAL statistics. For partitioned tables using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO will gather statistics for a table if it is marked stale or has no statistics. In addition, statistics will be gathered for partitions and sub-partitions that are marked stale or have no statistics. For tables not using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO may gather statistics for partitions and sub-partitions, even if they are not marked stale.

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Index statistics collection can be parellelized except for cluster, domain and join indexes.

If the `GATHER AUTO` option is used then frequency histograms will be created using a sample rather than a full scan. The following scenario shows where `GATHER AUTO` is used:

- A table is created like this: `CREATE TABLE NEWTAB as SELECT * FROM .....`

  This will create statistics on `NEWTAB` but no histograms.
- Next, the DBA creates the histograms using `GATHER AUTO` on `gather_table_stats`.
- The `FREQUENCY` histograms on `NEWTAB` will be created using a sample rather than a full table scan.

Exceptions

`ORA-20000`: Table does not exist or insufficient privileges

`ORA-20001`: Bad input value

Examples

An extension can be either a column group (see Example 1) or an expression (see Example 2).

Example 1

```
DBMS_STATS.GATHER_TABLE_STATS(
    'SH', 'SALES', method_opt => 'FOR COLUMNS (empno, deptno)');
```

Example 2

```
DBMS_STATS.GATHER_TABLE_STATS(
    'SH', 'SALES', method_opt => 'FOR COLUMNS (sal+comm)');
```

192.7.56 GENERATE_STATS Procedure

This **deprecated procedure** generates object statistics from previously collected statistics of related objects. The currently supported objects are b-tree and bitmap indexes.

> **Note:**
>
> This subprogram has been deprecated and replaced by improved technology. It is maintained only for purposes of backward compatibility. As an alternative, use the `GATHER_INDEX_STAT` procedure. See "GATHER_INDEX_STATS Procedure".

Syntax

```
DBMS_STATS.GENERATE_STATS (
   ownname    VARCHAR2,
   objname    VARCHAR2,
   organized  NUMBER DEFAULT 7,
   force      BOOLEAN default FALSE);
```

Parameters

Table 192-61 GENERATE_STATS Procedure Parameters

Parameter	Description
ownname	Schema of object
objname	Name of object
organized	Amount of ordering associated between the index and its underlying table. A heavily organized index would have consecutive index keys referring to consecutive rows on disk for the table (the same block). A heavily disorganized index would have consecutive keys referencing different table blocks on disk. This parameter is only used for b-tree indexes. The number can be in the range of 0-10, with 0 representing a completely organized index and 10 a completely disorganized one.
force	If TRUE, generates statistics for the target object even if it is locked

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

For fully populated schemas, the gather procedures should be used instead when more accurate statistics are desired.

Exceptions

ORA-20000: Unsupported object type of object does not exist

ORA-20001: Invalid option or invalid statistics

192.7.57 GET_ADVISOR_OPR_FILTER Procedure

This procedure creates an operation filter for an Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.GET_ADVISOR_OPR_FILTER (
  opr_id        IN             NUMBER,
  opr_filter    IN OUT NOCOPY  StatsAdvOpr);
```

Parameters

Table 192-62 GET_ADVISOR_OPR_FILTER Procedure Parameters

Parameter	Description
opr_id	The ID of the statistics operation stored in the DBA_OPTSTAT_OPERATIONS view.
opr_filter	The Optimizer Statistics Advisor filter that is generated based on the specified statistics operation.

Security Model

Note the following:

- To execute this subprogram, you must have the `ADVISOR` privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Exceptions

- `ORA-20000`: Insufficient privileges
- `ORA-20001`: Invalid input values
- `ORA-20012`: Optimizer Statistics Advisor errors

Usage Notes

You can specify the filter using either the operation ID or the filter ID, but not both at the same time.

192.7.58 GET_ADVISOR_RECS Function

This function generates a recommendation report for the specified item.

Syntax

```
DBMS_STATS.GET_ADVISOR_RECS(
  ownname     IN   VARCHAR2,
  tabname     IN   VARCHAR2,
  rec         IN   VARCHAR2,
  type        IN   VARCHAR2    := 'TEXT')
RETURN CLOB;
```

Parameters

Table 192-63 GET_ADVISOR_RECS Function Parameters

Parameter	Description
`ownname`	The owner of the table.
`tabname`	The name of the table.

Table 192-63 (Cont.) GET_ADVISOR_RECS Function Parameters

Parameter	Description
rec	The Optimizer Statistics Advisor recommendation. • INCREMENTAL When only a small number of range partitions are modified, this option improves the performance of statistics gathering dramatically. However, it requires additional space to store synopses for maintaining incremental statistics. The report analyzes this trade-off. • CONCURRENT The report recommends either setting the CONCURRENT preference, or specifying AUTO_DEGREE for individual tables. If the system resources and usage satisfies the conditions, the advisor always recommends setting CONCURRENT first. The advisor only recommends AUTO_DEGREE when statistics gathering on an individual table take a long time and the CONCURRENT preference is already set.
type	Type of the report: TEXT, HTML, or XML.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must have the privileges to gather statistics for the objects for which recommendations are generated.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Usage Notes

The advisor does not make recommendations for manual statistics gathering. The database only make recommendations for automatic statistics gathering jobs, with the main goal of finishing the job within the maintenance window. As long as the automatic job finishes, the database does not make further recommendations.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

192.7.59 GET_COLUMN_STATS Procedures

These overloaded procedures get column-related statistics. In the user-defined statistics version, the procedure returns the type of statistics stored.

Syntax

```
DBMS_STATS.GET_COLUMN_STATS (
   ownname        VARCHAR2,
```

```
   tabname        VARCHAR2,
   colname        VARCHAR2,
   partname       VARCHAR2 DEFAULT NULL,
   stattab        VARCHAR2 DEFAULT NULL,
   statid         VARCHAR2 DEFAULT NULL,
   distcnt OUT    NUMBER,
   density OUT    NUMBER,
   nullcnt OUT    NUMBER,
   srec    OUT    StatRec,
   avgclen OUT    NUMBER,
   statown        VARCHAR2 DEFAULT NULL,
   realtime_stats BOOLEAN iDEFAULT TRUE);
```

Use the following for user-defined statistics:

```
DBMS_STATS.GET_COLUMN_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   colname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   ext_stats   OUT  RAW,
   stattypown  OUT  VARCHAR2 DEFAULT NULL,
   stattypname OUT  VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-64 GET_COLUMN_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
tabname	Specifies the name of the table to which this column belongs.
colname	Specifies the name of the column or extension.
partname	Specifies the name of the table partition from which to get the statistics. If the table is partitioned, and if partname is NULL, then the procedure retrieves statistics at the global table level.
stattab	Specifies the statistics table ID describing where to retrieve the statistics. If stattab is NULL, then the procedure retrieves statistics directly from the data dictionary.
statid	Specifies an optional identifier associated with these statistics within stattab. This parameter is only relevant when stattab is not NULL.
ext_stats	Specifies the user-defined statistics.
stattypown	Specifies the schema of the statistics type.
stattypname	Specifies the name of the statistics type.
distcnt	Specifies the number of distinct values.
density	Specifies the column density.
nullcnt	Specifies the number of NULL values.

Table 192-64 (Cont.) GET_COLUMN_STATS Procedure Parameters

Parameter	Description
srec	Specifies the structure holding the internal representation of the column minimum, maximum, and histogram values.
avgclen	Specifies the average length of the column (in bytes).
statown	Specifies the schema containing stattab (if different than ownname).
realtime_stats	Specifies whether to include real-time statistics. The default value is TRUE. When realtime_stats is FALSE, the database only includes optimizer statistics that were gathered by the GATHER_*_STATS procedures.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be the owner of the table, or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges or no statistics have been stored for requested object

Usage Notes

Before invoking this procedure, ensure that the table exists.

192.7.60 GET_INDEX_STATS Procedures

This overloaded procedure gets all index-related statistics. In the form of this procedure that deals with user-defined statistics, the statistics type returned is the type stored, in addition to the user-defined statistics.

Syntax

```
DBMS_STATS.GET_INDEX_STATS (
   ownname         VARCHAR2,
   indname         VARCHAR2,
   partname        VARCHAR2 DEFAULT NULL,
   stattab         VARCHAR2 DEFAULT NULL,
   statid          VARCHAR2 DEFAULT NULL,
   numrows     OUT NUMBER,
   numlblks    OUT NUMBER,
   numdist     OUT NUMBER,
   avglblk     OUT NUMBER,
   avgdblk     OUT NUMBER,
   clstfct     OUT NUMBER,
   indlevel    OUT NUMBER,
   statown         VARCHAR2 DEFAULT NULL);

DBMS_STATS.GET_INDEX_STATS (
   ownname         VARCHAR2,
   indname         VARCHAR2,
   partname        VARCHAR2 DEFAULT NULL,
```

```
   stattab        VARCHAR2 DEFAULT NULL,
   statid         VARCHAR2 DEFAULT NULL,
   numrows    OUT NUMBER,
   numlblks   OUT NUMBER,
   numdist    OUT NUMBER,
   avglblk    OUT NUMBER,
   avgdblk    OUT NUMBER,
   clstfct    OUT NUMBER,
   indlevel   OUT NUMBER,
   statown        VARCHAR2 DEFAULT NULL,
   guessq     OUT NUMBER);
```

Use the following form of the procedure for user-defined statistics:

```
DBMS_STATS.GET_INDEX_STATS (
   ownname          VARCHAR2,
   indname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   ext_stats   OUT  RAW,
   stattypown  OUT  VARCHAR2 DEFAULT NULL,
   stattypname OUT  VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-65 GET_INDEX_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
indname	Name of the index
partname	Name of the index partition for which to get the statistics. If the index is partitioned and if partname is NULL, then the statistics are retrieved for the global index level.
stattab	User statistics table identifier describing from where to retrieve the statistics. If stattab is NULL, then the statistics are retrieved directly from the dictionary.
statid	Identifier (optional) to associate with these statistics within stattab (Only pertinent if stattab is not NULL)
ext_stats	User-defined statistics
stattypown	Schema of the statistics type
stattypname	Name of the statistics type
numrows	Number of rows in the index (partition)
numlblks	Number of leaf blocks in the index (partition)
numdist	Number of distinct keys in the index (partition)
avglblk	Average integral number of leaf blocks in which each distinct key appears for this index (partition)
avgdblk	Average integral number of data blocks in the table pointed to by a distinct key for this index (partition)

Table 192-65 (Cont.) GET_INDEX_STATS Procedure Parameters

Parameter	Description
clstfct	Clustering factor for the index (partition)
indlevel	Height of the index (partition)
statown	Schema containing stattab (if different than ownname)
guessq	Guess quality for the index (partition)

Security Model

Before invoking this procedure, ensure that the table exists. To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges or no statistics have been stored for requested object

Usage Notes

- The optimizer uses the cached data to estimate number of cached blocks for index or statistics table access. The database calculates the total cost of the operation by combining the I/O cost of reading not cached blocks from disk, the CPU cost of getting cached blocks from the buffer cache, and the CPU cost of processing the data.
- The database maintains cachedblk and cachehit at all times. However, the database uses the corresponding caching statistics for optimization as part of the table and index statistics only when the user calls the DBMS_STATS.GATHER_[TABLE/INDEX/SCHEMA/DATABASE]_STATS procedure for automatic mode or DBMS_STATS.GATHER_SYSTEM_STATS for manual mode. To prevent the user from utilizing inaccurate and unreliable data, the optimizer computes a “confidence factor” for each cachehit and a cachedblk for each object. If the confidence factor for the value meets confidence criteria, then the database uses this value; otherwise, the database uses defaults.
- The automatic maintenance algorithm for object caching statistics assumes that only one major database workload exists. The algorithm adjusts statistics to this workload, ignoring other "minor" workloads. If this assumption is false, then you must use manual mode for maintaining object caching statistics.
- The object caching statistics maintenance algorithm for automatic mode prevents you from using statistics in the following situations:
  - When not enough data has been analyzed, such as when an object has been recently created
  - When the system does not have one major workload resulting in averages not corresponding to real values

See Also:

Oracle Database SQL Tuning Guide to learn how to manage optimizer statistics

192.7.61 GET_PARAM Function

This function returns the default value of parameters of `DBMS_STATS` procedures.

> **Note:**
>
> This subprogram has been replaced by improved technology and is maintained only for purposes of backward compatibility. In this case, use the GET_PREFS Function.
>
> See also DBMS_STATS Deprecated Subprograms.

Syntax

```
DBMS_STATS.GET_PARAM (
   pname     IN   VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 192-66 GET_PARAM Function Parameters

Parameter	Description
`pname`	Parameter name

Exceptions

`ORA-20001`: Invalid input values

192.7.62 GET_PLSQL_PREFS Procedure

Returns the preference used for the named function. If a function has no preference set, the procedure returns the global preference.

Syntax

```
DBMS_STATS.GET_PLSQL_PREFS(
      pname           IN VARCHAR2,
      user            IN VARCHAR2,
      package_name    IN VARCHAR2,
      function_name   IN VARCHAR2)
RETURN VARCHAR2;
```

You can get the global preference by passing in NULL for `user`, `package_name`, and `function_name`.

Parameters

Table 192-67 GET_PLSQL_PREFS Parameters

Parameter	Description
pname	The preference that you want to get. Currently only 'dynamic_stats' is accepted.
user	The owner of the function. If NULL, the current user is used.
package_name	The name of the PL/SQL package. If NULL, this specifies a standalone, top-level function.
function_name	The name of the PL/SQL function.

Example

Get the preference value that is set for an individual PL/SQL function within a package. The parameter `user` is NULL in order to specify the current user:

```
SELECT DBMS_STATS.GET_PLSQL_PREFS('dynamic_stats',NULL,'<PL/SQL_PACKAGE
NAME>','<PL/SQL FUNCTION
      NAME>') pref_value FROM dual;
```

192.7.63 GET_PREFS Function

This function returns the default value of the specified preference.

Syntax

```
DBMS_STATS.GET_PREFS (
   pname     IN   VARCHAR2,
   ownname   IN   VARCHAR2 DEFAULT NULL,
   tabname   IN   VARCHAR2 DEFAULT NULL)
 RETURN VARCHAR2;
```

Parameters

Table 192-68 GET_PREFS Function Parameters

Parameter	Description
pname	Preference name. The possible values are: • APPROXIMATE_NDV_ALGORITHM • AUTO_STAT_EXTENSIONS • AUTO_TASK_STATUS • AUTO_TASK_MAX_RUN_TIME • AUTO_TASK_INTERVAL • CASCADE • CONCURRENT • DEGREE • ESTIMATE_PERCENT • GLOBAL_TEMP_TABLE_STATS • GRANULARITY • INCREMENTAL • INCREMENTAL_STALENESS • INCREMENTAL_LEVEL • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • STAT_CATEGORY • TABLE_CACHED_BLOCKS • WAIT_TIME_TO_UPDATE_STATUS
ownname	Owner name
tabname	Table name

Table 192-69 Preference Descriptions

Preference Name	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. Consider a synopsis as an internal management structure that samples distinct values. You can specify the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
AUTO_TASK_STATUS	Enables or disables the high-frequency automatic optimizer statistics collection. Values are: • ON — Enables high-frequency automatic optimizer statistics collection. • OFF — Disables high-frequency automatic optimizer statistics collection. This is the default.

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description
AUTO_TASK_MAX_RUN_TIME	Configures the maximum run time in seconds of an execution of high-frequency automatic optimizer statistics collection. The maximum value is 3600 (equal to 1 hour), which is the default.
AUTO_TASK_INTERVAL	Specifies the interval in seconds between executions of high-frequency automatic optimizer statistics collection. The default value is 3600 but Oracle sets it differently for some platforms and Oracle Database releases.
CASCADE	Determines whether index statistics are collected as part of gathering table statistics.
CONCURRENT	Determines whether statistics are gathered concurrently on multiple objects, or serially, one object at a time. Valid values are: • MANUAL — Concurrency is enabled only for manual statistics gathering. • AUTOMATIC — Concurrency is enabled only for the automatic statistics gathering. • ALL — Concurrency is enabled for both manual and automatic statistics gathering. • OFF — Concurrency is disabled for both manual and automatic statistics.
DEGREE	Determines degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	Determines the percentage of rows to sample. The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for optimal statistics. This is the default.
GLOBAL_TEMP_TABLE_STATS	Controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics. This preference takes two values: • SHARED — All sessions see the same set of statistics • SESSION — Statistics gathered by the GATHER_TABLE_STATS procedure on a global temporary table are session-specific. Thus, the database only uses them for queries issued in the same session as the statistics gathering process. The database deletes session-specific statistics when a session terminates.

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description
GRANULARITY	Determines the granularity of statistics to collect. This preference is only relevant for partitioned tables. The following values are valid: • ALL — Gathers all statistics: subpartition, partition, and global. • AUTO — Determines the granularity based on the partitioning type. This is the default value. • DEFAULT — Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. Use GLOBAL AND PARTITION for this functionality. • GLOBAL — Gathers global statistics. • GLOBAL AND PARTITION — Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. • PARTITION — Gathers partition-level statistics. • SUBPARTITION — Gathers subpartition-level statistics.
INCREMENTAL	Determines whether the global statistics for a partitioned table are maintained without performing a full table scan. When a table is partitioned, an application typically loads data into a new partition. As new partitions are added and data is loaded, global table statistics must be kept up to date. If the following conditions are met, then the database updates the global table statistics by scanning only the changed partitions instead of the entire table: • The INCREMENTAL value for the partitioned table is set to TRUE. • The PUBLISH value for the partitioned table is set to TRUE. • The user specifies AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT and AUTO for GRANULARITY when gathering statistics on the table. If the INCREMENTAL value for the partitioned table was set to FALSE (default value), then the database uses a full table scan to maintain the global statistics. This technique is a much more resource-intensive and time-consuming operation for large tables.

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description
INCREMENTAL_LEVEL	Controls which synopses to collect when INCREMENTAL preference is set to TRUE. It takes the following values: • PARTITION — Gathers partition-level synopses. This is the default value. If PARTITION is set on a nonpartitioned table, then the database does not gather synopses. • TABLE — Gathers table-level synopses. Specify this value when you want to exchange this table with a partition. Before the exchange, you can run GATHER_TABLE_STATS on this table with INCREMENTAL set to TRUE and INCREMENTAL_LEVEL to TABLE. The result is that the database gathers table-level synopses on this table. After the exchange, the partition has synopses that come from the table-level synopses of the table before the exchange. You can only use preference value in the SET_TABLE_PREFS procedure: this value is not allowed in the other SET_*_PREFS procedures.

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description

INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'.
	The parameter accepts the following values:
	• USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference.
	For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics.
	Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics.
	• USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes.
	For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics.
	• ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist.
	• NULL—A partition or subpartition is considered stale when it has any DML changes.
	For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.
	Note that the following two executions are different:
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	The first execution uses single quotes to set the preference to the value NULL, whereas the second sets

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description
	the preference to the default, which is `ALLOW_MIXED_FORMAT`.
`METHOD_OPT`	Controls column statistics collection and histogram creation. When setting preferences at the global, schema, database, or dictionary level, only `FOR ALL` syntax is allowed: • `FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause]` The `size_clause` is defined as `size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY}` `integer` — Specifies the number of histogram buckets. The number must be between 1 and 2048. `REPEAT` — Collects histograms only on the columns that already have histograms. `AUTO` — Determines the columns on which to collect histograms based on data distribution and the workload of the columns. `SKEWONLY` — Determines the columns on which to collect histograms based on the data distribution of the columns. The default is `FOR ALL COLUMNS SIZE AUTO`. You can change the value using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
`NO_INVALIDATE`	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • `TRUE`: Dependent cursors are not invalidated. • `FALSE`: Dependent cursors are marked for immediate invalidation. • `AUTO`: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. If set to `TRUE`, then the database not invalidate dependent cursors. If set to `FALSE`, then the procedure invalidates dependent cursors immediately.

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description
OPTIONS	Determines the `options` parameter used in the `GATHER_TABLE_STATS` procedure. The preference takes the following values: • `GATHER` — Gathers statistics for all objects in the table. This is the default. • `GATHER AUTO` — Oracle recommends using `GATHER AUTO` to gather necessary statistics, such as histograms, after a table has been bulk loaded and acquired online statistics. This is applicable only to tables that are not using `INCREMENTAL` statistics. For partitioned tables using `INCREMENTAL` statistics, `GATHER_TABLE_STATS` with `GATHER AUTO` will gather statistics for a table if it is marked stale or has no statistics. In addition, statistics will be gathered for partitions and sub-partitions that are marked stale or have no statistics. For tables not using `INCREMENTAL` statistics, `GATHER_TABLE_STATS` with `GATHER AUTO` may gather statistics for partitions and sub-partitions, even if they are not marked stale.
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • `TRUE` — Ignores input parameter values, and uses the value of the corresponding preference. • `FALSE` — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether the database publishes newly gathered statistics after the gathering job completes. You can gather statistics without publishing them immediately. This technique enables you to test new statistics before publishing them.
STALE_PERCENT	Determines the percentage of rows in a table that must change before the statistics on that table are stale and need to be regathered. The valid domain for `stale_percent` is non-negative numbers. The default value is `10`, which means that a table having more than 10% of changes is considered stale.

Table 192-69 (Cont.) Preference Descriptions

Preference Name	Description
STAT_CATEGORY	Specifies which statistics to import or export, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS - table statistics, column statistics and index statistics (default) • SYNOPSES - information to support incremental statistics • REALTIME_STATS — specifies only real-time statistics. The value 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.
WAIT_TIME_TO_UPDATE_STATS	Specifies the number of minutes before timing out for locks and pins required for updating statistics. It accepts values in the range 0 to 65535. The default value is 15 minutes. The value 0 gets the locks and pins in no-wait mode.

Security Model

No special privilege or role is needed to invoke this procedure. To gather statistics concurrently, however, you must either have the DBA role, or have the following privileges in addition to privileges that are required for gathering statistics: CREATE JOB, MANAGE SCHEDULER, and MANAGE ANY QUEUE.

Exceptions

- ORA-20000: Unable to gather statistics concurrently: Resource Manager is not enabled.
- ORA-20001: Invalid input values

Usage Notes

Note the following guidelines:

- The CONCURRENT preference determines whether statistics are gathered concurrently when the user issues GATHER_*_STATS procedures. DBMS_STATS can collect statistics for a single object in parallel based on the value of the DEGREE parameter. However, parallelism is limited to one object. The CONCURRENT preference extends the scope of parallelism to multiple database objects. This approach is primarily intended for multi-CPU systems, and may not be suitable for small databases on single-CPU computers.

  To gather statistics concurrently, Resource Manager must be enabled, and the setting for the JOB_QUEUE_PROCESSES initialization parameter must be at least 4.

- If the ownname and tabname are provided, and if a preference has been entered for the table, then the function returns the preference as specified for the table. In all other cases,

it returns the global preference if it has been specified, otherwise it returns the default value.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to get optimizer statistics preferences

192.7.64 GET_STATS_HISTORY_AVAILABILITY Function

This function returns oldest timestamp where statistics history is available. Users cannot restore statistics to a timestamp older than this one.

Syntax

```
DBMS_STATS.GET_STATS_HISTORY_AVAILABILITY
 RETURN TIMESTAMP WITH TIMEZONE;
```

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.65 GET_STATS_HISTORY_RETENTION Function

This function returns the current statistics history retention value.

Syntax

```
DBMS_STATS.GET_STATS_HISTORY_RETENTION
 RETURN NUMBER;
```

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.66 GET_SYSTEM_STATS Procedure

This procedure gets system statistics from `stattab`, or from the dictionary if `stattab` is `NULL`.

Syntax

```
DBMS_STATS.GET_SYSTEM_STATS (
   status    OUT  VARCHAR2,
   dstart    OUT  DATE,
   dstop     OUT  DATE,
   pname     IN   VARCHAR2,
   pvalue    OUT  NUMBER,
   stattab   IN   VARCHAR2 DEFAULT NULL,
   statid    IN   VARCHAR2 DEFAULT NULL,
   statown   IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-70 GET_SYSTEM_STATS Procedure Parameters

Parameter	Description
status	Output is one of the following: • COMPLETED: • AUTOGATHERING: • MANUALGATHERING: • BADSTATS:
dstart	Date when statistics gathering started. If status = MANUALGATHERING, the start date is returned.
dstop	Date when statistics gathering stopped. • If status = COMPLETE, the finish date is returned. • If status = AUTOGATHERING, the future finish date is returned. • If status = BADSTATS, the must-finished-by date is returned.
pname	The parameter name to get, which can have one of the following values: • iotfrspeed - I/O transfer speed in bytes for each millisecond • ioseektim - seek time + latency time + operating system overhead time, in milliseconds • sreadtim - average time to read single block (random read), in milliseconds • mreadtim - average time to read an mbrc block at once (sequential read), in milliseconds • cpuspeed - average number of CPU cycles for each second, in millions, captured for the workload (statistics collected using 'INTERVAL' or 'START' and 'STOP' options) • cpuspeednw - average number of CPU cycles for each second, in millions, captured for the no-workload (statistics collected using 'NOWORKLOAD' option. • mbrc - average multiblock read count for sequential read, in blocks • maxthr - maximum I/O system throughput, in bytes/second • slavethr - average slave I/O throughput, in bytes/second
pvalue	Parameter value to get
stattab	Identifier of the user statistics table where the statistics will be obtained. If stattab is NULL, the statistics will be obtained from the dictionary.
statid	Optional identifier associated with the statistics saved in the stattab
statown	Schema containing stattab (if different from current schema)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20002: Bad user statistics table; may need to be upgraded

ORA-20003: Unable to gather system statistics

ORA-20004: Parameter does not exist

Usage Notes

To run this procedure, you need the GATHER_SYSTEM_STATISTICS role.

192.7.67 GET_TABLE_STATS Procedure

This overloaded procedure gets all table-related statistics.

Syntax

```
DBMS_STATS.GET_TABLE_STATS (
   ownname            VARCHAR2,
   tabname            VARCHAR2,
   partname           VARCHAR2 DEFAULT NULL,
   stattab            VARCHAR2 DEFAULT NULL,
   statid             VARCHAR2 DEFAULT NULL,
   numrows        OUT NUMBER,
   numblks        OUT NUMBER,
   avgrlen        OUT NUMBER,
   statown            VARCHAR2 DEFAULT NULL,
   realtime_stats     BOOLEAN DEFAULT TRUE);


DBMS_STATS.GET_TABLE_STATS (
   ownname            VARCHAR2,
   tabname            VARCHAR2,
   partname           VARCHAR2 DEFAULT NULL,
   stattab            VARCHAR2 DEFAULT NULL,
   statid             VARCHAR2 DEFAULT NULL,
   numrows        OUT NUMBER,
   numblks        OUT NUMBER,
   avgrlen        OUT NUMBER,
   statown            VARCHAR2 DEFAULT NULL,
   im_imcu_count  OUT NUMBER,
   im_block_count OUT NUMBER,
   scanrate       OUT NUMBER,
   realtime_stats     BOOLEAN DEFAULT TRUE);


DBMS_STATS.GET_TABLE_STATS (
   ownname           VARCHAR2,
   tabname           VARCHAR2,
   partname          VARCHAR2 DEFAULT NULL,
   stattab           VARCHAR2 DEFAULT NULL,
   statid            VARCHAR2 DEFAULT NULL,
   numrows       OUT NUMBER,
   numblks       OUT NUMBER,
   avgrlen       OUT NUMBER,
   cachedblk     OUT NUMBER,
   cachehit      OUT NUMBER,
   realtime_stats     BOOLEAN DEFAULT TRUE);
```

Parameters

Table 192-71 GET_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	Specifies the name of the schema.
tabname	Specifies the name of the table to which this column belongs.
partname	Specifies the name of the table partition from which to get the statistics. If the table is partitioned and if partname is NULL, then the statistics are retrieved from the global table level.
stattab	Specifies the user statistics table ID. This ID describes where to retrieve the statistics. If stattab is NULL, then the procedure gathers statistics directly from the data dictionary.
statid	Specifies the optional ID associates with these statistics within stattab. This ID is only relevant when stattab is not NULL.
numrows	Specifies the number of rows in the table or partition.
numblks	Specifies the number of blocks in the table or partition.
avgrlen	Specifies the average row length for the table or partition.
statown	Specifies the schema containing stattab (if different from ownname).
im_imcu_count	Specifies the number of In-Memory Compression Units (IMCUs) in the table or partition.
im_block_count	Specifies the number of In-Memory blocks in the table or partition. An In-Memory block corresponds to a specific data block on disk. If the table is fully populated in the IM column store, then the number of In-Memory blocks equals the number of data blocks.
scanrate	Specifies the rate, in MB/s, at which the database scans external tables. This parameter is relevant only for external tables.
realtime_stats	Specifies whether to include real-time statistics. The default value is TRUE. When realtime_stats is FALSE, the database only includes optimizer statistics that were gathered by the GATHER_*_STATS procedures.
cachedblk	For internal use only.
cachehit	For internal use only.

Security Model

Before invoking this procedure, ensure that the table exists. To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table, or have the ANALYZE ANY DICTIONARY or SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges or no statistics have been stored for requested object

ORA-20002: Bad user statistics table; may need to upgrade it

Usage Notes

- The optimizer uses the cached data to estimate number of cached blocks for index or statistics table access. The database calculates the total cost of the operation by combining the I/O cost of reading not cached blocks from disk, the CPU cost of getting cached blocks from the buffer cache, and the CPU cost of processing the data.
- The database maintains `cachedblk` and `cachehit` at all times. However, the database uses the corresponding caching statistics for optimization as part of the table and index statistics only when the user calls the `DBMS_STATS.GATHER_*_STATS` procedure for automatic mode or `DBMS_STATS.GATHER_SYSTEM_STATS` for manual mode. To prevent the user from utilizing inaccurate and unreliable data, the optimizer computes a “confidence factor” for each `cachehit` and a `cachedblk` for each object. If the confidence factor for the value meets confidence criteria, then the database uses this value; otherwise, the database uses defaults.
- The automatic maintenance algorithm for object caching statistics assumes that only one major database workload exists. The algorithm adjusts statistics to this workload, ignoring other "minor" workloads. If this assumption is false, then you must use manual mode for maintaining object caching statistics.
- The object caching statistics maintenance algorithm for automatic mode prevents you from using statistics in the following situations
  - When not enough data has been analyzed, such as when an object has been recently created
  - When the system does not have one major workload resulting in averages not corresponding to real values
- The database does not support export or import of statistics across databases of different character sets.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to manage optimizer statistics preferences

192.7.68 IMPLEMENT_ADVISOR_TASK Function

This function implements the recommendations made by Optimizer Statistics Advisor.

Syntax

```
DBMS_STATS.IMPLEMENT_ADVISOR_TASK (
  task_name          IN   VARCHAR2,
  execution_name     IN   VARCHAR2    := NULL,
  level              IN   VARCHAR2    := 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 192-72 IMPLEMENT_ADVISOR_TASK Function Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.
execution_name	A name that qualifies and identifies an advisor execution. If not specified, then the advisor automatically generates it. If the specified execution conflicts with the name of an existing execution, then the function returns an error.
level	The level of the implementation. Possible values are • ALL: Ignores the filters and implements all recommendations. • TYPICAL: Implements the recommendations according to the filters in place.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- You can execute this subprogram for AUTO_STATS_ADVISOR_TASK, which is predefined.
- This subprogram executes using invoker's rights.

The results of performing this task depend on the privileges of the executing user:

- SYSTEM level

  Only users with both the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges can perform this task on system-level rules.
- Operation level

  The results depend on the following privileges:

  – Users with both the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges can perform this task for all statistics operations.

  – Users with the ANALYZE ANY privilege but *not* the ANALYZE ANY DICTIONARY privilege can perform this task for statistics operations related to any schema except SYS.

  – Users with the ANALYZE ANY DICTIONARY privilege but *not* the ANALYZE ANY privilege can perform this task for statistics operations related to their own schema and the SYS schema.

  – Users with neither the ANALYZE ANY nor the ANALYZE ANY DICTIONARY privilege can only perform this operation for statistics operations relating to their own schema.
- Object level

  Users can perform this task for any object for which they have statistics collection privileges.

Return Values

This function returns an XML CLOB that indicates which recommendations were successfully implemented.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

Example 192-10 Implementing Optimizer Statistics Advisor Recommendations

This script illustrates a basic Optimizer Statistics Advisor session. It creates a task, executes it, generates a report, and then implements the recommendations.

```
DECLARE
  v_tname      VARCHAR2(128) := 'my_task';
  v_ename      VARCHAR2(128) := NULL;
  v_report     CLOB := null;
  v_script     CLOB := null;
  v_imp_result CLOB;
BEGIN
  -- create a task
  v_tname := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);

  -- execute the task
  v_ename := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);

  -- view the task report
  v_report := DBMS_STATS.REPORT_ADVISOR_TASK(v_tname);
  DBMS_OUTPUT.PUT_LINE(v_report);

  -- implement all recommendations
  v_imp_result := DBMS_STATS.IMPLEMENT_ADVISOR_TASK(v_tname);
END;
```

192.7.69 IMPORT_COLUMN_STATS Procedure

This procedure retrieves statistics for a particular column from the user statistics table identified by `stattab` and stores them in the dictionary.

Syntax

```
DBMS_STATS.IMPORT_COLUMN_STATS (
   ownname       VARCHAR2,
   tabname       VARCHAR2,
   colname       VARCHAR2,
   partname      VARCHAR2 DEFAULT NULL,
   stattab       VARCHAR2,
   statid        VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL,
   no_invalidate BOOLEAN DEFAULT to_no_invalidate_type (
                                  get_param('NO_INVALIDATE')),
   force         BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-73 IMPORT_COLUMN_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
tabname	Name of the table to which this column belongs
colname	Name of the column or extension
partname	Name of the table partition. If the table is partitioned and if partname is NULL, then global and partition column statistics are imported.
stattab	User statistics table identifier describing from where to retrieve the statistics
statid	Identifier to associate with these statistics within stattab (optional)
statown	Schema containing stattab (if different than ownname)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	If set to TRUE, imports statistics even if statistics are locked

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

ORA-20005: Object statistics are locked

Usage Notes

Oracle does not support export or import of statistics across databases of different character sets.

192.7.70 IMPORT_DATABASE_PREFS Procedure

This procedure is used to import the statistics preferences of all the tables, excluding the tables owned by Oracle. These tables can by included by passing `TRUE` for the `add_sys` parameter.

Syntax

```
DBMS_STATS.IMPORT_DATABASE_PREFS (
    stattab    IN  VARCHAR2,
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT NULL
    add_sys    IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-74 IMPORT_DATABASE_PREFS Procedure Parameters

Parameter	Description
`stattab`	Statistics table name where to import the statistics
`statid`	Optional identifier to associate with these statistics within `stattab`
`statown`	Schema containing `stattab` (if different than `ownname`)
`add_sys`	Value `TRUE` will include the Oracle-owned tables

Exceptions

`ORA-20000`: Insufficient privileges.

Usage Notes

- To run this procedure, you need to have the SYSDBA role, or both `ANALYZE ANY DICTIONARY` and `ANALYZE ANY` system privileges.
- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.IMPORT_DATABASE_PREFS('STATTAB', statown=>'SH');
```

192.7.71 IMPORT_DATABASE_STATS Procedure

This procedure imports statistics for all objects in the database from the user statistics table and stores them in the data dictionary.

Syntax

```
DBMS_STATS.IMPORT_DATABASE_STATS (
   stattab         VARCHAR2,
   statid          VARCHAR2 DEFAULT NULL,
   statown         VARCHAR2 DEFAULT NULL,
   no_invalidate   BOOLEAN DEFAULT to_no_invalidate_type(
                                   get_param('NO_INVALIDATE')),
```

```
force           BOOLEAN DEFAULT FALSE,
stat_category   VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-75 IMPORT_DATABASE_STATS Procedure Parameters

Parameter	Description
stattab	Specifies the statistics table that contains the statistics to be imported.
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different from current schema)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Overrides statistics locked at the object (table) level: • TRUE - Ignores the statistics lock and imports the statistics • FALSE - The statistics will be imported only if they are not locked
stat_category	Specifies what statistics to import, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics • SYNOPSES — information to support incremental statistics • REALTIME_STATS — specifies only real-time statistics. The default value is 'OBJECT_STATS, SYNOPSES, MODELS'.

Security Model

You must have either the SYSDBA privilege or both the ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.72 IMPORT_DICTIONARY_STATS Procedure

This procedure imports statistics for all data dictionary schemas (SYS, SYSTEM, and RDBMS component schemas) from the user statistics table and stores them in the dictionary.

Syntax

```
DBMS_STATS.IMPORT_DICTIONARY_STATS (
   stattab         VARCHAR2,
   statid          VARCHAR2 DEFAULT NULL,
   statown         VARCHAR2 DEFAULT NULL,
   no_invalidate   BOOLEAN DEFAULT to_no_invalidate_type(
                                 get_param('NO_INVALIDATE')),
   force           BOOLEAN DEFAULT FALSE,
   stat_category   VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-76 IMPORT_DICTIONARY_STATS Procedure Parameters

Parameter	Description
stattab	User statistics table identifier describing from where to retrieve the statistics
statid	The (optional) identifier to associate with these statistics within stattab
statown	Schema containing stattab (if different from current schema)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Overrides statistics lock at the object (table) level: • TRUE - Ignores the statistics lock and imports the statistics. • FALSE - The statistics will be imported only if there is no lock.
stat_category	Specifies what statistics to import, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics • SYNOPSES — information to support incremental statistics • REALTIME_STATS — Specifies only real-time statistics. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS'.

Security Model

You must have either the SYSDBA privilege or both the ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

ORA-20002: Bad user statistics table, may need to upgrade it

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.73 IMPORT_FIXED_OBJECTS_STATS Procedure

This procedure imports statistics for fixed tables from the user statistics table and stores them in the data dictionary.

Syntax

```
DBMS_STATS.IMPORT_FIXED_OBJECTS_STATS (
   stattab       VARCHAR2,
   statid        VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL,
   no_invalidate BOOLEAN  DEFAULT to_no_invalidate_type(
                                    get_param('NO_INVALIDATE')),
   force         BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 192-77 IMPORT_FIXED_OBJECTS_STATS Procedure Parameters

Parameter	Description
stattab	User statistics table identifier describing from where to retrieve the statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different from current schema)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Overrides statistics lock: • TRUE - Ignores the statistics lock and imports the statistics • FALSE - The statistics will be imported only if there is no lock

Security Model

You must have the SYSDBA or ANALYZE ANY DICTIONARY system privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

ORA-20002: Bad user statistics table, may need to upgrade it

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.74 IMPORT_INDEX_STATS Procedure

This procedure retrieves statistics for a particular index from the user statistics table identified by stattab and stores them in the dictionary.

Syntax

```
DBMS_STATS.IMPORT_INDEX_STATS (
   ownname       VARCHAR2,
   indname       VARCHAR2,
   partname      VARCHAR2 DEFAULT NULL,
   stattab       VARCHAR2,
   statid        VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL,
   no_invalidate BOOLEAN DEFAULT to_no_invalidate_type(
                                 get_param('NO_INVALIDATE')),
   force         BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-78 IMPORT_INDEX_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
indname	Name of the index
partname	Name of the index partition. If the index is partitioned and if partname is NULL, then global and partition index statistics are imported.
stattab	User statistics table identifier describing from where to retrieve the statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)

Table 192-78 (Cont.) IMPORT_INDEX_STATS Procedure Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Imports statistics even if index statistics are locked

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

ORA-20005: Object statistics are locked

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Oracle does not support export or import of statistics across databases of different character sets.

192.7.75 IMPORT_SCHEMA_PREFS Procedure

This procedure is used to import the statistics preferences of all the tables owned by the specified owner name.

Syntax

```
DBMS_STATS.IMPORT_SCHEMA_PREFS (
    ownname    IN  VARCHAR2,
    stattab    IN  VARCHAR2,
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-79 IMPORT_SCHEMA_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
stattab	Statistics table name from where to import the statistics

Table 192-79 (Cont.) IMPORT_SCHEMA_PREFS Procedure Parameters

Parameter	Description
statid	(Optional) Identifier to associate with these statistics within stattab
statown	Schema containing stattab (if other than ownname)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

- To run this procedure, you need to connect as owner, or have the SYSDBA privilege, or have the ANALYZE ANY system privilege.
- All arguments are of type VARCHAR2 and values are enclosed in quotes.
- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.IMPORT_SCHEMA_PREFS('SH', 'STAT');
```

192.7.76 IMPORT_SCHEMA_STATS Procedure

This procedure imports statistics for all objects in the schema identified by ownname from the user statistics table and stores them in the data dictionary.

Syntax

```
DBMS_STATS.IMPORT_SCHEMA_STATS (
   ownname          VARCHAR2,
   stattab          VARCHAR2,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type(
                                       get_param('NO_INVALIDATE')),
   force            BOOLEAN DEFAULT FALSE,
   stat_category    VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-80 IMPORT_SCHEMA_STATS Procedure Parameters

Parameter	Description
ownname	Specifies the name of the schema.
stattab	Identifies the user table that stores the statistics to be imported.
statid	Specifies the ID associated with these statistics within stattab.
statown	Specifies the schema containing stattab (if different than ownname).

Table 192-80 (Cont.) IMPORT_SCHEMA_STATS Procedure Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Specifies whether to override statistics locked at the object level. The possible values are as follows: • TRUE — Ignores the statistics lock and imports the statistics. • FALSE — Imports the statistics only if there is no lock. This is the default.
stat_category	Specifies which statistics to process. The following values are supported: • OBJECT_STATS — specifies table statistics, column statistics, and index statistics • SYNOPSES — specifies metadata for incremental statistics • REALTIME_STATS — specifies only real-time statistics You can specify a list of comma-delimited values. For example, 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS '.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.77 IMPORT_SYSTEM_STATS Procedure

This procedure retrieves system statistics from the user statistics table, identified by stattab, and stores the statistics in the dictionary.

Syntax

```
DBMS_STATS.IMPORT_SYSTEM_STATS (
   stattab       VARCHAR2,
```

```
   statid        VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-81 IMPORT_SYSTEM_STATS Procedure Parameters

Parameter	Description
stattab	Identifier of the user statistics table where the statistics will be retrieved
statid	Optional identifier associated with the statistics retrieved from the stattab
statown	Schema containing stattab (if different from current schema)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

ORA-20002: Bad user statistics table; may need to be upgraded

ORA-20003: Unable to import system statistics

Usage Notes

To run this procedure, you need the GATHER_SYSTEM_STATISTICS role.

Oracle does not support export or import of statistics across databases of different character sets.

192.7.78 IMPORT_TABLE_PREFS Procedure

This procedure is used to set the statistics preferences of the specified table in the specified schema.

Syntax

```
DBMS_STATS.IMPORT_TABLE_PREFS (
    ownname    IN  VARCHAR2,
    tabname    IN  VARCHAR2,
    stattab    IN  VARCHAR2,
    statid     IN  VARCHAR2 DEFAULT NULL,
    statown    IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-82 IMPORT_TABLE_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name
stattab	Statistics table name from where to import the statistics
statid	(Optional) Identifier to associate with these statistics within stattab
statown	Schema containing stattab (if other than ownname)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

Usage Notes

- To run this procedure, you need to connect as owner of the table, or have the ANALYZE ANY system privilege.
- All arguments are of type VARCHAR2 and values are enclosed in quotes.
- Oracle does not support export or import of statistics across databases of different character sets.

Examples

```
DBMS_STATS.IMPORT_TABLE_PREFS('SH', 'SALES', 'STAT');
```

192.7.79 IMPORT_TABLE_STATS Procedure

This procedure import statistics for a specified table from the user statistics table identified by stattab and stores them in the data dictionary.

Syntax

```
DBMS_STATS.IMPORT_TABLE_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   stattab          VARCHAR2,
   statid           VARCHAR2 DEFAULT NULL,
   cascade          BOOLEAN  DEFAULT TRUE,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT to_no_invalidate_type(
                                     get_param('NO_INVALIDATE')),
   force            BOOLEAN  DEFAULT FALSE,
   stat_category    VARCHAR2 DEFAULT DEFAULT_STAT_CATEGORY);
```

Parameters

Table 192-83 IMPORT_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	Specifies the name of the schema.
tabname	Specifies the name of the table.
partname	Name of the table partition. If the table is partitioned and if partname is NULL, then global and partition table statistics are imported.
stattab	Identifies the user statistics table that describes where to retrieve the statistics.
statid	Specifies the ID associated with these statistics within stattab.
cascade	Indicates whether to import column and index statistics for this table. The default is TRUE.
statown	Specifies the schema containing stattab (if different than ownname).

Table 192-83 (Cont.) IMPORT_TABLE_STATS Procedure Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure
force	Specifies whether to override statistics locked at the object level. The possible values are as follows: • TRUE — Ignores the statistics lock and imports the statistics. • FALSE — Imports the statistics only if there is no lock. This is the default.
stat_category	Specifies which statistics to process. The following values are supported: • OBJECT_STATS — specifies table statistics, column statistics, and index statistics • SYNOPSES — specifies metadata for incremental statistics • REALTIME_STATS — specifies only real-time statistics You can specify a list of comma-delimited values. For example, 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses. The default value is 'OBJECT_STATS, SYNOPSES, REALTIME_STATS'.

Security Model

To invoke this procedure you must be owner of the table or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table or have either the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values in the user statistics table

Usage Notes

Oracle Database does not support export or import of statistics across databases of different character sets.

192.7.80 INTERRUPT_ADVISOR_TASK Procedure

This procedure interrupts a currently executing Optimizer Statistics Advisor task.

The task ends its operations as it does when at a normal exit, at which point you can access intermediate results. You can also resume the task using the "RESUME_ADVISOR_TASK Procedure".

Syntax

```
DBMS_STATS.INTERRUPT_ADVISOR_TASK (
  task_name IN VARCHAR2);
```

Parameters

Table 192-84 INTERRUPT_ADVISOR_TASK Procedure Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Consider a case in which a task is executed by one user, interrupted, and then resumed by a different user. In this case, Optimizer Statistics Advisor bases its checks of the resumed execution on the privilege of the user who resumed the task.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

Example 192-11 Interrupting an Optimizer Statistics Advisor Task

In this example, you start a SQL*Plus session, and then create and execute an advisor task named my_task:

```
DECLARE
  v_tname   VARCHAR2(128) := 'my_task';
  v_ename   VARCHAR2(128) := NULL;
BEGIN
  -- create a task
  v_tname := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);

  -- execute the task
  v_ename := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);
END;
/
```

In a separate terminal, you start a second SQL*Plus session, and then execute the following program:

```
XEC DBMS_STATS.INTERRUPT_ADVISOR_TASK('my_task');
```

The first session returns an ORA-13632 to indicate the cancelation of the task:

```
ORA-13638: The user interrupted the current operation.
```

192.7.81 LOCK_PARTITION_STATS Procedure

This procedure enables the user to lock statistics for a partition.

Syntax

```
DBMS_STATS.LOCK_PARTITION_STATS (
    ownname    VARCHAR2,
    tabname    VARCHAR2,
    partname   VARCHAR2);
```

Parameters

Table 192-85 LOCK_PARTITION_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema to lock
tabname	Name of the table
partname	[Sub]Partition name

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.82 LOCK_SCHEMA_STATS Procedure

This procedure locks the statistics of all tables of a schema.

Syntax

```
DBMS_STATS.LOCK_SCHEMA_STATS (
   ownname    VARCHAR2);
```

Parameters

Table 192-86 LOCK_SCHEMA_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema to lock

Usage Notes

- To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

- When statistics on a table are locked, all the statistics depending on the table, including table statistics, column statistics, histograms and statistics on all dependent indexes, are considered to be locked.
- The `SET_*, DELETE_*, IMPORT_*, GATHER_*` procedures that modify statistics in the dictionary of an individual table, index or column will raise an error if statistics of the object is locked.
- Procedures that operates on multiple objects (such as `GATHER_SCHEMA_STATS`) will skip modifying the statistics of an object if it is locked. Many procedures have force argument to override the lock.
- This procedure either freezes the current set of the statistics or keeps the statistics empty (uncollected) to use dynamic statistics.
- The locked or unlocked state is not exported along with the table statistics when using `EXPORT_*_STATS` procedures.
- Neither the UNLOCK_SCHEMA_STATS Procedure nor the UNLOCK_TABLE_STATS Procedure is designed to unlock statistics of corresponding partitions. When you invoke the LOCK_TABLE_STATS Procedure, it sets the statistics lock bit at the table level. In that case, you cannot gather statistics on dependent objects such as partitions and indexes. By the same token, if table statistics are locked, the dependents are locked and you do not need to explicitly invoke the LOCK_PARTITION_STATS Procedure.

192.7.83 LOCK_TABLE_STATS Procedure

This procedure locks the statistics on the table.

Syntax

```
DBMS_STATS.LOCK_TABLE_STATS (
   ownname    VARCHAR2,
   tabname    VARCHAR2);
```

Parameters

Table 192-87 LOCK_TABLE_STATS Procedure Parameters

Parameter	Description
`ownname`	Name of the schema
`tabname`	Name of the table

Usage Notes

- To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.
- When statistics on a table are locked, all the statistics depending on the table, including table statistics, column statistics, histograms and statistics on all dependent indexes, are considered to be locked.
- The `SET_*, DELETE_*, IMPORT_*, GATHER_*` procedures that modify statistics in the dictionary of an individual table, index or column will raise an error if statistics of the object is locked.

- Procedures that operates on multiple objects (such as `GATHER_SCHEMA_STATS`) will skip modifying the statistics of an object if it is locked. Many procedures have force argument to override the lock.
- This procedure either freezes the current set of the statistics or keeps the statistics empty (uncollected) to use dynamic statistics.
- The locked or unlocked state is not exported along with the table statistics when using `EXPORT_*_STATS` procedures.
- Neither the UNLOCK_SCHEMA_STATS Procedure nor the UNLOCK_TABLE_STATS Procedure is designed to unlock statistics of corresponding partitions. When you invoke the LOCK_TABLE_STATS Procedure, it sets the statistics lock bit at the table level. In that case, you cannot gather statistics on dependent objects such as partitions and indexes. By the same token, if table statistics are locked, the dependents are locked and you do not need to explicitly invoke the LOCK_PARTITION_STATS Procedure.

192.7.84 MERGE_COL_USAGE Procedure

This procedure merges column usage information from a source database by means of a `dblink` into the local database.

If column usage information already exists for a given table or column `MERGE_COL_USAGE` will combine both the local and the remote information.

Syntax

```
DBMS_STATS.MERGE_COL_USAGE (
   dblink    IN    VARCHAR2);
```

Parameters

Table 192-88 MERGE_COL_USAGE Procedure Parameters

Parameter	Description
`dblink`	Name of `dblink`

Usage Notes

User must be `SYS` to execute this procedure. addition the user specified during the creation of the `dblink` is expected to have privileges to select from tables in the `SYS` schema.

Exceptions

`ORA-20000`: Insufficient privileges

`ORA-20001`: Parameter `dblink` cannot be `NULL`

`ORA-20002`: Unable to create a `TEMP` table

192.7.85 PREPARE_COLUMN_VALUES Procedures

These procedures convert user-specified minimum, maximum, and histogram endpoint actual values into Oracle's internal representation for future storage using `SET_COLUMN_STATS`.

Syntax

```
DBMS_STATS.PREPARE_COLUMN_VALUES (
   srec     IN OUT StatRec,
   charvals        CHARARRAY);

DBMS_STATS.PREPARE_COLUMN_VALUES (
   srec      IN OUT StatRec,
   datevals         DATEARRAY);

DBMS_STATS.PREPARE_COLUMN_VALUES (
   srec     IN OUT StatRec,
   dblvals         DBLARRAY);

DBMS_STATS.PREPARE_COLUMN_VALUES (
   srec     IN OUT StatRec,
   fltvals         FLTARRAY);

DBMS_STATS.PREPARE_COLUMN_VALUES (
   srec     IN OUT StatRec,
   numvals         NUMARRAY);

DBMS_STATS.PREPARE_COLUMN_VALUES (
   srec     IN OUT StatRec,
   rawvals         RAWARRAY);
```

Parameters

Table 192-89 PREPARE_COLUMN_VALUES Procedure Parameters

Parameter	Description
`srec.epc`	Number of values specified in `charvals`, `datevals`, `dblvals`, `fltvals`, `numvals`, or `rawvals`. This value must be between 2 and 2050, inclusive, and it should be set to `2` for procedures which do not allow histogram information (`nvarchar` and `rowid`). The first corresponding array entry should hold the minimum value for the column, and the last entry should hold the maximum. If there are more than two entries, then all the others hold the remaining height-balanced or frequency histogram endpoint values (with in-between values ordered from next-smallest to next-largest). This value may be adjusted to account for compression, so the returned value should be left as is for a call to `SET_COLUMN_STATS`.
`srec.bkvals`	If you want a frequency or hybrid histogram, this array contains the number of occurrences of each distinct value specified in `charvals`, `datevals`, `dblvals`, `fltvals`, `numvals`, or `rawvals`. Otherwise, it is merely an output parameter, and it must be set to `NULL` when this procedure is called.

Table 192-89 (Cont.) PREPARE_COLUMN_VALUES Procedure Parameters

Parameter	Description
srec.rpcnts	If you want a hybrid histogram, this array contains the total frequency of values that are less than or equal to each distinct value specified in charvals, datevals, numvals, or rawvals. Otherwise, it is merely an output argument and must be set to NULL when this procedure is called. As an example, for a given array numvals with numvals(i)=4, rpcnts(i)=13 means that there are 13 rows in the column which are less than or equal to 4. Note: • Whenever srec.rpcnts is populated, srec.bkvals must be populated as described above. • Whenever bkvals and/or rpcnts are populated, there should not be any duplicates in charvals, datevals, numvals, or rawvals.

Datatype-specific input parameters (use one) are shown in Table 192-90.

Table 192-90 Datatype-Specific Input Parameters

Type	Description
charvals	The array of values when the column type is character-based. Up to the first 64 bytes of each string should be provided. Arrays must have between 2 and 2050 entries, inclusive. If the datatype is fixed CHAR, the strings must be space-padded to 15 characters for correct normalization.
datevals	Array of values when the column type is date-based
dblvals	Array of values when the column type is double-based
fltvals	Array of values when the column type is float-based
numvals	Array of values when the column type is numeric-based
rawvals	Array of values when the column type is RAW. Up to the first 64 bytes of each value should be provided.
nvmin, nvmax	Minimum and maximum values when the column type is national character set based. No histogram information can be provided for a column of this type. If the datatype is fixed CHAR, the strings must be space-padded to 15 characters for correct normalization.
rwmin, rwmax	Minimum and maximum values when the column type is rowid. No histogram information is provided for a column of this type.

Output Parameters

Table 192-91 PREPARE_COLUMN_VALUES Procedure Output Parameters

Parameter	Description
srec.minval	Internal representation of the minimum suitable for use in a call to SET_COLUMN_STATS
srec.maxval	Internal representation of the maximum suitable for use in a call to SET_COLUMN_STATS
srec.bkvals	Array suitable for use in a call to SET_COLUMN_STATS

Table 192-91 (Cont.) PREPARE_COLUMN_VALUES Procedure Output Parameters

Parameter	Description
srec.novals	Array suitable for use in a call to SET_COLUMN_STATS
srec.eavals	Array suitable for use in a call to SET_COLUMN_STATS
srec.rpcnts	Array suitable for use in a call to SET_COLUMN_STATS

Exceptions

ORA-20001: Invalid or inconsistent input values

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.86 PREPARE_COLUMN_VALUES_NVARCHAR Procedure

This procedure converts user-specified minimum, maximum, and histogram endpoint actual values into Oracle's internal representation for future storage using the SET_COLUMN_STATS Procedures.

Syntax

```
DBMS_STATS.PREPARE_COLUMN_VALUES_NVARCHAR (
   srec     IN OUT StatRec,
   nvmin           NVARCHAR2,
   nvmax           NVARCHAR2);
```

Parameters

Table 192-92 PREPARE_COLUMN_VALUES_NVARCHAR Procedure Parameters

Parameter	Description
srec.epc	Number of values specified in charvals, datevals, dblvals, fltvals, numvals, or rawvals. This value must be between 2 and 2050, inclusive, and it should be set to 2 for procedures which do not allow histogram information (nvarchar and rowid). The first corresponding array entry should hold the minimum value for the column, and the last entry should hold the maximum. If there are more than two entries, then all the others hold the remaining height-balanced or frequency histogram endpoint values (with in-between values ordered from next-smallest to next-largest). This value may be adjusted to account for compression, so the returned value should be left as is for a call to SET_COLUMN_STATS.
srec.bkvals	If you want a frequency or hybrid histogram, then this array contains the number of occurrences of each distinct value specified in charvals, datevals, dblvals, fltvals, numvals, or rawvals. Otherwise, it is merely an output parameter, and it must be set to NULL when this procedure is called.

Table 192-92 (Cont.) PREPARE_COLUMN_VALUES_NVARCHAR Procedure Parameters

Parameter	Description
srec.rpcnts	If you want a hybrid histogram, this array contains the total frequency of values that are less than or equal to each distinct value specified in charvals, datevals, numvals, or rawvals. Otherwise, it is merely an output argument and must be set to NULL when this procedure is called. As an example, for a given array numvals with numvals(i)=4, rpcnts(i)=13 means that there are 13 rows in the column which are less than or equal to 4. Note: • Whenever srec.rpcnts is populated, srec.bkvals must be populated as described above. • Whenever bkvals and/or rpcnts are populated, there should not be any duplicates in charvals, datevals, numvals, or rawvals.

Datatype-specific input parameters (use one) are shown in Table 192-90.

Table 192-93 PREPARE_COLUMN_VALUES_NVARCHAR Datatype-Specific Input Parameters

Type	Description
nvmin, nvmax	The minimum and maximum values when the column type is national character set based. No histogram information can be provided for a column of this type. If the datatype is fixed CHAR, the strings must be space-padded to 15 characters for correct normalization.

Output Parameters

Table 192-94 PREPARE_COLUMN_VALUES_NVARCHAR Procedure Output Parameters

Parameter	Description
srec.minval	Internal representation of the minimum suitable for use in a call to SET_COLUMN_STATS
srec.maxval	Internal representation of the maximum suitable for use in a call to SET_COLUMN_STATS
srec.bkvals	Array suitable for use in a call to SET_COLUMN_STATS.
srec.novals	Array suitable for use in a call to SET_COLUMN_STATS
srec.eavals	Array suitable for use in a call to SET_COLUMN_STATS
srec.rpcnts	Array suitable for use in a call to SET_COLUMN_STATS

Exceptions

ORA-20001: Invalid or inconsistent input values

Usage Notes

No special privilege or role is needed to invoke this procedure.

Related Topics

- SET_COLUMN_STATS Procedures
  This procedure sets column-related information.

192.7.87 PREPARE_COLUMN_VALUES_ROWID Procedure

This procedure converts user-specified minimum, maximum, and histogram endpoint datatype-specific values into Oracle's internal representation for future storage using `SET_COLUMN_STATS`.

Syntax

```
DBMS_STATS.PREPARE_COLUMN_VALUES_ROWID (
   srec  IN OUT StatRec,
   rwmin        ROWID,
   rwmax        ROWID);
```

Pragmas

```
pragma restrict_references(prepare_column_values_rowid, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 192-95 PREPARE_COLUMN_VALUES_ROWID Procedure Parameters

Parameter	Description
srec	Values (`IN`): • `epc` • `bkvals` • `rpcnts` Values (`OUT`): • `minval` • `maxval` • `bkvals` • `novals` • `eavals` • `rpcnts`
rwmin	Minimum value when the column type is `rowid`. No histogram information is provided for a column of this type.
rwmax	Maximum value when the column type is `rowid`. No histogram information is provided for a column of this type.

Table 192-96 StatRec Record Type Fields

Field	Description
epc (IN)	Number of values specified in charvals, datevals, dblvals, fltvals, numvals, or rawvals. This value must be between 2 and 2050, inclusive, and it should be set to 2 for procedures which do not allow histogram information (nvarchar and rowid). The first corresponding array entry should hold the minimum value for the column, and the last entry should hold the maximum. If there are more than two entries, then all the others hold the remaining height-balanced or frequency histogram endpoint values (with in-between values ordered from next-smallest to next-largest). This value may be adjusted to account for compression, so the returned value should be left as is for a call to SET_COLUMN_STATS.
bkvals (IN)	If you want a frequency or hybrid histogram, this array contains the number of occurrences of each distinct value specified in charvals, datevals, dblvals, fltvals, numvals, or rawvals. Otherwise, it is merely an output parameter, and it must be set to NULL when this procedure is called.
rpcnts (IN)	If you want a hybrid histogram, this array contains the total frequency of values that are less than or equal to each distinct value specified in charvals, datevals, numvals, or rawvals. Otherwise, it is merely an output argument and must be set to NULL when this procedure is called. As an example, for a given array numvals with numvals(i)=4, rpcnts(i)=13 means that there are 13 rows in the column which are less than or equal to 4. Note: • Whenever srec.rpcnts is populated, srec.bkvals must be populated as described above. • Whenever bkvals and/or rpcnts are populated, there should not be any duplicates in charvals, datevals, numvals, or rawvals.
minval (OUT)	Internal representation of the minimum suitable for use in a call to SET_COLUMN_STATS.
maxval (OUT)	Internal representation of the maximum suitable for use in a call to SET_COLUMN_STATS.
bkvals (OUT)	Array suitable for use in a call to SET_COLUMN_STATS.
novals (OUT)	Array suitable for use in a call to SET_COLUMN_STATS.
eavals (OUT)	Array suitable for use in a call to SET_COLUMN_STATS
rpcnts (OUT)	Array suitable for use in a call to SET_COLUMN_STATS

Usage Notes

No special privilege or role is needed to invoke this procedure.

192.7.88 PUBLISH_PENDING_STATS Procedure

This procedure is used to publish the statistics gathered and stored as pending.

Syntax

```
DBMS_STATS.PUBLISH_PENDING_STATS (
    ownname          IN  VARCHAR2 DEFAULT USER,
    tabname          IN  VARCHAR2,
    no_invalidate        BOOLEAN DEFAULT
       TO_NO_INVALIDATE_TYPE(GET_PARAM('NO_INVALIDATE')),
    force            IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-97 PUBLISH_PENDING_STATS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	If TRUE, will override the lock

Security Model

To run this procedure, you must have the same privilege for gathering statistics on the tables that will be touched by this procedure.

Exceptions

ORA-20000: Insufficient privileges

Usage Notes

- If the parameter tabname is NULL then publish applies to all tables of the specified schema.
- The default owner/schema is the user who runs the procedure.

Examples

```
DBMS_STATS.PUBLISH_PENDING_STATS ('SH', null);
```

192.7.89 PURGE_STATS Procedure

This procedure purges old versions of statistics saved in the dictionary.

To run this procedure, you must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege.

Syntax

```
DBMS_STATS.PURGE_STATS(
    before_timestamp       TIMESTAMP WITH TIME ZONE);
```

Parameters

Table 192-98 PURGE_STATS Procedure Parameters

Parameter	Description
before_timestamp	Versions of statistics saved before this timestamp are purged. If NULL, it uses the purging policy used by automatic purge. The automatic purge deletes all history older than the older of (current time - statistics history retention) and (time of recent analyze in the system - 1). The statistics history retention value can be changed using ALTER_STATS_HISTORY_RETENTION Procedure.The default is 31 days.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values

Usage Notes

To invoke this procedure you need the ANALYZE ANY privilege and the ANALYZE ANY DICTIONARY privilege.

192.7.90 REMAP_STAT_TABLE Procedure

This procedure remaps the names of objects in the user statistics table. It allows you to import the statistics to objects with same definition but with different names.

Syntax

```
DBMS_STATS.REMAP_STAT_TABLE (
   ownname    IN    VARCHAR2,
   stattab    IN    VARCHAR2,
   src_own    IN    VARCHAR2,
   src_tab    IN    VARCHAR2,
   tgt_own    IN    VARCHAR2,
   tgt_tab    IN    VARCHAR2);
```

Parameters

Table 192-99 REMAP_STAT_TABLE Procedure Parameters

Parameter	Description
ownname	Owner of the statistics table. NULL means the current schema.
stattab	User statistics table identifier
src_own	Owner of the table to be renamed. This argument cannot be NULL.
src_tab	Name of the table to be renamed. If NULL, all tables are owned by src_own.
tgt_own	New name of the owner of the table. The owner name is also updated for the dependent objects such as columns and indexes. Note that an index of src_tab not owned by src_own is not renamed.This argument cannot be NULL.
tgt_tab	New name of the table. This argument is valid only if src_tab is not NULL.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Invalid input

Examples

The following statement remaps all objects of sh to shsave in user statistics table sh.ustat:

```
DBMS_STATS.REMAP_STAT_TABLE ('sh', 'ustat', 'sh', NULL, 'shsave', NULL);
```

The following statement can be used to import statistics into objects of shsave once the preceding remap procedure is completed:

```
DBMS_STATS.IMPORT_SCHEMA_STATS ('shsave', 'ustat', statown => 'sh');
```

The following statement remaps sh.customers to shsave.customers_sav:

```
DBMS_STATS.REMAP_STAT_TABLE ('sh', 'ustat', 'sh', 'customers','shsave', 'customers_sav');
```

192.7.91 REPORT_ADVISOR_TASK Function

This function reports the results of an Optimizer Statistics Advisor task.

Syntax

```
DBMS_STATS.REPORT_ADVISOR_TASK(
  task_name          IN   VARCHAR2,
  execution_name     IN   VARCHAR2   := NULL,
  type               IN   VARCHAR2   := 'TEXT',
  section            IN   VARCHAR2   := 'ALL',
  level              IN   VARCHAR2   := 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 192-100 REPORT_ADVISOR_TASK Function Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.
execution_name	A name that qualifies and identifies an advisor task execution. If not specified, the function uses the latest execution of the specified task.
type	The type of the Optimizer Statistics Advisor report. Possible values are TEXT, HTML, and XML.
section	A section in the report. Possible values are SUMMARY, FINDINGS, ERRORS, and ALL. You can specify combinations of different values can be using the plus (+) and minus (-) operator, as in 'SUMMARY +FINDINGS +ERRORS', and 'ALL -ERRORS'.
level	The format of the report. Possible values are BASIC, TYPICAL, ALL, and SHOW_HIDDEN. You can specify SHOW_HIDDEN together with the other three input values, as in 'BASIC +SHOW_HIDDEN' and 'TYPICAL +SHOW_HIDDEN'.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- You can execute this subprogram for AUTO_STATS_ADVISOR_TASK, which is predefined.
- This subprogram executes using invoker's rights.

The results of performing this task depend on the privileges of the executing user:

- SYSTEM level

  Only users with both the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges can perform this task on system-level rules.
- Operation level

  The results depend on the following privileges:

  – Users with both the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges can perform this task for all statistics operations.

  – Users with the ANALYZE ANY privilege but *not* the ANALYZE ANY DICTIONARY privilege can perform this task for statistics operations related to any schema except SYS.

  – Users with the ANALYZE ANY DICTIONARY privilege but *not* the ANALYZE ANY privilege can perform this task for statistics operations related to their own schema and the SYS schema.

  – Users with neither the ANALYZE ANY nor the ANALYZE ANY DICTIONARY privilege can only perform this operation for statistics operations relating to their own schema.
- Object level

  Users can perform this task for any object for which they have statistics collection privileges.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: User input errors
- ORA-20012: Optimizer Statistics Advisor errors

Returns

This function returns a CLOB that contains the report.

Examples

(Optional) List and briefly describe the examples for using the API or subprogram here.

Example 192-12 Generating an HTML Report

This example creates a procedure named myrep, and then calls this procedure to generate an HTML report.

```
SET ECHO ON
SET FEEDBACK ON
SET SERVEROUTPUT ON
SET TRIMS ON
SET LINESIZE 300

EXECUTE DBMS_OUTPUT.ENABLE (buffer_size => 10000000);

CREATE OR REPLACE PROCEDURE myrep(p_tname VARCHAR2, p_ftype VARCHAR2, which
VARCHAR2)
IS
  v_report CLOB          := null;
  v_script CLOB          := null;
  v_ftype  VARCHAR2(400) := p_ftype;
  v_tname  VARCHAR2(400) := p_tname;
  v_len    NUMBER(10);
  v_ps     NUMBER(10) := 10000;
  v_pn     NUMBER(10) := 1;
  v_ret    VARCHAR2(32767);
BEGIN
   IF which = 'REPORT'
   THEN
      -- generate a report
      v_report := DBMS_STATS.REPORT_ADVISOR_TASK(
                    task_name => v_tname,
                    type      => v_ftype,
                    section   => 'ALL',
                    level     => 'ALL');
      v_len := DBMS_LOB.getlength(v_report);
      WHILE (v_pn < v_len)
      LOOP
        DBMS_OUTPUT.PUT_LINE(DBMS_LOB.SUBSTR(v_report, v_ps, v_pn));
        v_pn := v_pn + v_ps;
      END LOOP;
   ELSE
      -- generate a script
```

```
      v_script := DBMS_STATS.SCRIPT_ADVISOR_TASK(v_tname);
      v_len := DBMS_LOB.getlength(v_script);
      WHILE (v_pn < v_len)
      LOOP
        DBMS_OUTPUT.PUT_LINE(DBMS_LOB.SUBSTR( v_script, v_ps, v_pn));
        v_pn := v_pn + v_ps;
      END LOOP;
   END IF;
END;
/

SHOW ERRORS

SPOOL report.txt
EXECUTE myrep('my_task','HTML','REPORT');
SPOOL OFF
```

Example 192-13 Generating a Textual Report for AUTO_STATS_ADVISOR_TASK

The following example invokes the `myrep` procedure created in the preceding example for `AUTO_STATS_ADVISOR_TASK`:

```
EXEC myrep('AUTO_STATS_ADVISOR_TASK','TEXT','REPORT');
```

The following sample output shows part of the report:

```
GENERAL INFORMATION
-------------------------------------------------------------------------------
-

 Task Name       : AUTO_STATS_ADVISOR_TASK
 Execution Name  : EXEC_97
 Created         : 07-08-16 10:18:10
 Last Modified   : 07-11-16 03:02:30

-------------------------------------------------------------------------------
-
SUMMARY
-------------------------------------------------------------------------------
-
 For execution EXEC_97 of task AUTO_STATS_ADVISOR_TASK, the Statistics
Advisor
 has 10 finding(s). The findings are related to the following rules:
 COMPLETEAUTOJOB, MAINTAINSTATSHISTORY, USEDEFAULTPREFERENCE,
 AVOIDSETPROCEDURES, USEDEFAULTPARAMS, USEGATHERSCHEMASTATS, AVOIDSTALESTATS,
 UNLOCKNONVOLATILETABLE, USEINCREMENTAL, AVOIDANALYZETABLE. Please refer to
the
 finding section for detailed information.
-------------------------------------------------------------------------------
-
FINDINGS
-------------------------------------------------------------------------------
-
 Rule Name:         MaintainStatsHistory
 Rule Description:  Maintain Statistics History
```

```
 Finding:  Statistics history tables are too big.

 Recommendation:  Check the other findings of this rule, as well as
the
                  findings for the rules
AvoidFrequentStatsCollection,
                  UseDefaultPreference, UseDefaultParams for possible
causes
                  and
recommendations.
 Rationale:  The size of the statistics history table could be big because
of
             violations of other rules.

-------------------------------------------------
 Rule Name:         UseDefaultPreference
 Rule Description:  Use Default Preference for Stats Collection
 Finding:  Global preference SYS_FLAGS is set to a non-default value '1'.


Recommendation:  Set the value of preference SYS_FLAGS to '0'.
 Example:
 -- Setting preference cascade to default value:
 dbms_stats.set_global_prefs('CASCADE', NULL);

 Rationale:  Preference SYS_FLAGS is for Oracle internal use only, setting
it
             to nondefault value '1' could cause unforeseen consequences.
.
.
.
```

192.7.92 REPORT_COL_USAGE Function

This function reports the recorded column (group) usage information.

Syntax

```
DBMS_STATS.REPORT_COL_USAGE (
   ownname    IN    VARCHAR2,
   tabname    IN    VARCHAR2)
RETURN CLOB;
```

Parameters

Table 192-101 REPORT_COL_USAGE Function Parameters

Parameter	Description
ownname	Owner name. If NULL it reports column usage information for tables in all schemas in the database.
tabname	Table name. If NULL it reports column usage information for all tables of ownname.

Usage Notes

To run this procedure, you need to have the `SYSDBA` administrative privilege or both `ANALYZE ANY DICTIONARY` and `ANALYZE ANY` system privileges.

192.7.93 REPORT_GATHER_AUTO_STATS Function

This function runs the auto statistics gathering job in reporting mode. That is, statistics are not actually collected, but all the objects that will be affected when auto statistics gathering is invoked are reported.

Syntax

```
DBMS_STATS.REPORT_GATHER_AUTO_STATS (
   detail_level      VARCHAR2  DEFAULT 'TYPICAL',
   format            VARCHAR2  DEFAULT 'TEXT')
 RETURN CLOB;
```

Parameters

Table 192-102 REPORT_GATHER_AUTO_STATS Function Parameters

Parameter	Description
detail_level	Detail level for the content of the report • BASIC: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • TYPICAL: In addition to the information provided at level BASIC, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed-out objects (applies to only auto statistics gathering) With regard to target objects: - owner and name of each target object - target object type (such as: table, index) - start time - end time - completion status • ALL: In addition to the information provided at level TYPICAL, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name - session ID - parameter values - error message if the operation failed With regard to target objects: - job name - batching details - estimated cost - rank in the target list - columns for which histograms were collected - list of collected extended statistics (if any) - reason for including the object in the target list - additional error details if the task has failed. Note that several fields (such as job name, estimated task cost) in the report are populated only when an operation is executed concurrently (CONCURRENT preference is turned on).

Table 192-102 (Cont.) REPORT_GATHER_AUTO_STATS Function Parameters

Parameter	Description
format	Report format: • XML • HTML • TEXT (Default)

Usage Notes

Only user SYS can run the REPORT_GATHER_AUTO_STATS function.

192.7.94 REPORT_GATHER_DATABASE_STATS Functions

This function runs the GATHER_DATABASE_STATS function in reporting mode.

The database does not collect statistics, but reports all objects that would be affected when invoking GATHER_DATABASE_STATS. The input set of parameters is the same as in GATHER_DATABASE_STATS, with two extra parameters.

Syntax

```
DBMS_STATS.REPORT_GATHER_DATABASE_STATS (
   estimate_percent     IN    NUMBER     DEFAULT to_estimate_percent_type (
                                             GET_PARAM('ESTIMATE_PERCENT')),
   block_sample         IN    BOOLEAN    DEFAULT FALSE,
   method_opt           IN    VARCHAR2   DEFAULT GET_PARAM('METHOD_OPT'),
   degree               IN    NUMBER     DEFAULT TO_DEGREE_TYPE(
                                             GET_PARAM('DEGREE')),
   granularity          IN    VARCHAR2   DEFAULT GET_PARAM('GRANULARITY'),
   cascade              IN    BOOLEAN    DEFAULT to_cascade_type (
                                             GET_PARAM('CASCADE')),
   stattab              IN    VARCHAR2   DEFAULT NULL,
   statid               IN    VARCHAR2   DEFAULT NULL,
   options              IN    VARCHAR2   DEFAULT 'GATHER',
   statown              IN    VARCHAR2   DEFAULT NULL,
   gather_sys           IN    BOOLEAN    DEFAULT TRUE,
   no_invalidate        IN    BOOLEAN    DEFAULT TO_NO_INVALIDATE_TYPE (
                                             GET_PARAM('NO_INVALIDATE')),
   gather_temp          IN    BOOLEAN    DEFAULT FALSE,
   gather_fixed         IN    BOOLEAN    DEFAULT FALSE,
   stattype             IN    VARCHAR2   DEFAULT DATA,
   obj_filter_list      IN    ObjectTab  DEFAULT NULL,
   detail_level         IN    VARCHAR2   DEFAULT 'TYPICAL',
   format               IN    VARCHAR2   DEFAULT 'TEXT')
 RETURN CLOB;
```

Parameters

Table 192-103 REPORT_GATHER_DATABASE_STATS Function Parameters

Parameter	Description
estimate_percent	The percentage of rows to use for the sample size. The valid range is between 0.000001 and 100. The null value means to compute. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for good statistics. This is the default. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Whether or not to use random block sampling instead of random row sampling. Random block sampling is more efficient, but if the data is not randomly distributed on disk, then the sample values may be somewhat correlated. This parameter is only relevant when estimating statistics.
method_opt	Method options. This parameter accepts the following values: FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause]. size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} - integer : Number of histogram buckets. Must be in the range [1,2048]. - REPEAT : Collects histograms only on the columns that already have histograms. - AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. - SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO.The value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.

Table 192-103 (Cont.) REPORT_GATHER_DATABASE_STATS Function Parameters

Parameter	Description
granularity	Determines the granularity of statistics to collect. This preference is only relevant for partitioned tables. The following values are valid: • ALL — Gathers all statistics: subpartition, partition, and global. • AUTO — Determines the granularity based on the partitioning type. This is the default value. • DEFAULT — Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. Use GLOBAL AND PARTITION for this functionality. • GLOBAL — Gathers global statistics. • GLOBAL AND PARTITION — Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. • PARTITION — Gathers partition-level statistics. • SUBPARTITION — Gathers subpartition-level statistics.
cascade	Gather statistics on the indexes as well. Using this option is equivalent to running the GATHER_INDEX_STATS Procedure on each of the indexes in the database in addition to gathering table and column statistics. Use the constant DBMS_STATS.AUTO_CASCADE to have Oracle determine whether index statistics to be collected or not. This is the default. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics. The database assumes that the statistics table resides in the same schema as the object being analyzed. Thus, to use this option, one such table must exist in each schema.
statid	Identifier (optional) to associate with these statistics within stattab.
options	Further specification of which objects to gather statistics for: GATHER: Gathers statistics on all objects in the schema. GATHER AUTO: Gathers all necessary statistics automatically. Oracle implicitly determines which objects need new statistics, and determines how to gather those statistics. When GATHER AUTO is specified, the only additional valid parameters are stattab, statid, objlist and statown; all other parameter settings are ignored. Returns a list of processed objects. GATHER STALE: Gathers statistics on stale objects as determined by looking at the *_tab_modifications views. Also, return a list of objects found to be stale. GATHER EMPTY: Gathers statistics on objects which currently have no statistics. Return a list of objects found to have no statistics. LIST AUTO: Returns a list of objects to be processed with GATHER AUTO LIST STALE: Returns a list of stale objects as determined by looking at the *_tab_modifications views LIST EMPTY: Returns a list of objects which currently have no statistics
statown	Schema containing stattab, if different from current schema.

Table 192-103 (Cont.) REPORT_GATHER_DATABASE_STATS Function Parameters

Parameter	Description
gather_sys	Gathers statistics on the objects owned by the SYS user.
no_invalidate	Does not invalidate the dependent cursors if set to TRUE. The procedure invalidates the dependent cursors immediately if set to FALSE. Use DBMS_STATS.AUTO_INVALIDATE. to have Oracle decide when to invalidate dependent cursors. This is the default. The default can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
gather_temp	Gathers statistics on global temporary tables when TRUE. The default is FALSE. The temporary table must be created with the ON COMMIT PRESERVE ROWS clause. Also, the statistics collected are based on the data in the session in which this procedure is run, but they are shared across all sessions.
gather_fixed	Gather statistics on fixed tables when TRUE. The default is FALSE. Only user SYS can collect statistics for fixed tables. The ownname must be SYS or null. When gathering statistics for fixed tables, the database ignores specified values for the following arguments: • estimate_percent • block_sample • stattab • statid • statown The database does not invalidate the dependent cursors on fixed tables on which stats is collected. This option is meant for internal use only.
stattype	The type of statistics: • DATA — Data statistics only • CACHE — Cache statistics only • ALL — All statistics
obj_filter_list	A list of object filters. The attribute values specified in the object filter are case-insensitive unless double-quoted. Wildcards are allowed in the attribute values. When specified, GATHER_DATABASE_STATS gathers statistics only on objects that satisfy at least one object filter in the list as needed. In a single-object filter, the database can specify the constraints on the object attributes. For example, non-NULL values s1 and s2 are specified for attributes a1 and a2 in one object filter. An object o is said to satisfy this object filter when (o.a1 LIKE s1) AND (o.a2 LIKE s2) is true.

Table 192-103 (Cont.) REPORT_GATHER_DATABASE_STATS Function Parameters

Parameter	Description
detail_level	The level of detail for the content of the report. Valid values are as follows: • BASIC: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • TYPICAL: In addition to the information provided at level BASIC, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed-out objects (applies to only auto statistics gathering) With regard to target objects: - owner and name of each target object - target object type (such as: table, index) - start time - end time - completion status • ALL: In addition to the information provided at level TYPICAL, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name - session ID - parameter values - error message if the operation failed With regard to target objects: - job name - batching details - estimated cost - rank in the target list - columns for which histograms were collected - list of collected extended statistics (if any) - additional error details if the task has failed. Note that several fields (such as job name, estimated task cost) in the report are populated only when an operation is executed concurrently (CONCURRENT preference is turned on).
format	The format of the report. Valid values are: • XML • HTML • TEXT (Default)

Return Values

A CLOB object that contains the report

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Bad input value

Usage Notes

To run this procedure, you need to have the SYSDBA role or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.

192.7.95 REPORT_GATHER_DICTIONARY_STATS Functions

This function runs the GATHER_DICTIONARY_STATS procedure in reporting mode.

The database does not collect statistics, but reports all objects affected when invoking GATHER_DICTIONARY_STATS. The detail level for the report is defined by the detail_level input parameter. See the descriptions of detail_level and format in REPORT_GATHER_DICTIONARY_STATS Functions. For all other input parameters, see GATHER_DICTIONARY_STATS Procedure.

Syntax

```
DBMS_STATS.REPORT_GATHER_DICTIONARY_STATS (
   comp_id               IN    VARCHAR2   DEFAULT NULL,
   estimate_percent      IN    NUMBER     DEFAULT TO_ESTIMATE_PERCENT_TYPE

(GET_PARAM('ESTIMATE_PERCENT')),
   block_sample          IN    BOOLEAN    DEFAULT FALSE,
   method_opt            IN    VARCHAR2   DEFAULT GET_PARAM('METHOD_OPT'),
   degree                IN    NUMBER     DEFAULT TO_DEGREE_TYPE
                                              (GET_PARAM('DEGREE')),
   granularity           IN    VARCHAR2   DEFAULT GET_PARAM('GRANULARITY'),
   cascade               IN    BOOLEAN    DEFAULT TO_CASCADE_TYPE
                                              (GET_PARAM('CASCADE')),
   stattab               IN    VARCHAR2   DEFAULT NULL,
   options               IN    VARCHAR2   DEFAULT 'GATHER AUTO',
   no_invalidate         IN    BOOLEAN    DEFAULT TO_NO_INVALIDATE_TYPE
                                              (GET_PARAM('NO_INVALIDATE')),
   stattype              IN    VARCHAR2   DEFAULT 'DATA',
   obj_filter_list       IN    ObjectTab  DEFAULT NULL,
   detail_level          IN    VARCHAR2   DEFAULT 'TYPICAL',
   format                IN    VARCHAR2   DEFAULT 'TEXT')
 RETURN CLOB;
```

Parameters

Table 192-104 REPORT_GATHER_DICTIONARY_STATS Function Parameters

Parameter	Description
comp_id	Component ID of the schema to analyze. NULL results in analyzing schemas of all RDBMS components. Refer refer to the COMP_ID column of the DBA_REGISTRY view. The procedure always gather statistics on SYS and SYSTEM schemas regardless of this argument.
estimate_percent	Percentage of rows to sample (NULL means compute). The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to let the database determine the appropriate sample size for good statistics. This is the default. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Determines whether to use random block sampling instead of random row sampling. Random block sampling is more efficient, but if the data is not randomly distributed on disk then the sample values may be somewhat correlated. Only pertinent when performing estimate statistics.
method_opt	The method options. This parameter accepts the following values: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms. – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.

Table 192-104 (Cont.) REPORT_GATHER_DICTIONARY_STATS Function Parameters

Parameter	Description
granularity	Determines the granularity of statistics to collect. This preference is only relevant for partitioned tables. The following values are valid: • ALL — Gathers all statistics: subpartition, partition, and global. • AUTO — Determines the granularity based on the partitioning type. This is the default value. • DEFAULT — Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. Use GLOBAL AND PARTITION for this functionality. • GLOBAL — Gathers global statistics. • GLOBAL AND PARTITION — Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. • PARTITION — Gathers partition-level statistics. • SUBPARTITION — Gathers subpartition-level statistics.
cascade	Gathers statistics on indexes also. Index statistics gathering will not be parallelized. Using this option is equivalent to running the GATHER_INDEX_STATS Procedure on each of the indexes in the schema in addition to gathering table and column statistics. Use the constant DBMS_STATS.AUTO_CASCADE to have Oracle determine whether index statistics to be collected or not. This is the default.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics.
options	Further specification of objects for which to gather statistics: • 'GATHER' - Gathers statistics on all objects in the schema • 'GATHER AUTO' - Gathers all necessary statistics automatically. Oracle implicitly determines which objects need new statistics and determines how to gather those statistics. When 'GATHER AUTO' is specified, the only additional valid parameters are comp_id, stattab, statid and statown; all other parameter settings will be ignored. Also, returns a list of objects processed. • 'GATHER STALE' - Gathers statistics on stale objects as determined by looking at the *_tab_modifications views. Also, returns a list of objects found to be stale. • 'GATHER EMPTY' - Gathers statistics on objects which currently have no statistics. Also, returns a list of objects found to have no statistics. • 'LIST AUTO' - Returns list of objects to be processed with 'GATHER AUTO' • 'LIST STALE' - Returns list of stale objects as determined by looking at the *_tab_modifications views • 'LIST EMPTY' - Returns list of objects which currently have no statistics

Table 192-104 (Cont.) REPORT_GATHER_DICTIONARY_STATS Function Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattype	The type of statistics: • DATA — Data statistics only • CACHE — Cache statistics only • ALL — All statistics
obj_filter_list	A list of object filters. When provided, this will gather statistics only on objects which satisfy at least one object filter in the list as needed. In a single object filter, we can specify the constraints on the object attributes. The attribute values specified in the object filter are case-insensitive unless double-quoted. Wildcard is allowed in the attribute values. Suppose non-NULL values s1, s2, ... are specified for attributes a1, a2, ... in one object filter. An object o is said to satisfy this object filter if (o.a1 like s1) and (o.a2 like s2) and ... is true.
detail_level	See the description in REPORT_GATHER_DICTIONARY_STATS Functions.
format	See the description in REPORT_GATHER_DICTIONARY_STATS Functions.

Return Values

A CLOB object that contains the report

Usage Notes

You must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege to execute this procedure.

Exceptions

ORA-20000: Index does not exist or insufficient privileges

ORA-20001: Bad input value

ORA-20002: Bad user statistics table, may need to upgrade it

192.7.96 REPORT_GATHER_FIXED_OBJ_STATS Function

This function runs the GATHER_FIXED_OBJECTS_STATS Procedure in reporting mode.

That is, statistics are not actually collected, but all the objects that will be affected when GATHER_FIXED_OBJ_STATS is invoked are reported. The input set of parameters are exactly the same as in GATHER_FIXED_OBJ_STATS with two extra parameters.

Syntax

```
DBMS_STATS.REPORT_GATHER_FIXED_OBJ_STATS (
   stattab          IN  VARCHAR2 DEFAULT NULL,
   statid           IN  VARCHAR2 DEFAULT NULL,
   statown          IN  VARCHAR2 DEFAULT NULL,
   no_invalidate    IN  BOOLEAN  DEFAULT TO_NO_INVALIDATE_TYPE (
                                    GET_PARAM('NO_INVALIDATE')),
   detail_level     IN  VARCHAR2   DEFAULT 'TYPICAL',
   format           IN  VARCHAR2   DEFAULT 'TEXT')
 RETURN CLOB;
```

Parameters

Table 192-105 REPORT_GATHER_FIXED_OBJ_STATS Procedure Parameters

Parameter	Description
`stattab`	User statistics table identifier describing where to save the current statistics
`statid`	Identifier to associate with these statistics within `stattab` (optional)
`statown`	Schema containing `stattab` (if different from current schema)
`no_invalidate`	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • `TRUE`: Dependent cursors are not invalidated. • `FALSE`: Dependent cursors are marked for immediate invalidation. • `AUTO`: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Table 192-105 (Cont.) REPORT_GATHER_FIXED_OBJ_STATS Procedure Parameters

Parameter	Description
`detail_level`	Detail level for the content of the report • `BASIC`: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • `TYPICAL`: In addition to the information provided at level `BASIC`, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed-out objects (applies to only auto statistics gathering) With regard to target objects: - owner and name of each target object - target object type (such as: table, index) - start time - end time - completion status • `ALL`: In addition to the information provided at level `TYPICAL`, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name - session ID - parameter values - error message if the operation failed With regard to target objects: - job name - batching details - estimated cost - rank in the target list - columns for which histograms were collected - list of collected extended statistics (if any) - additional error details if the task has failed. Note that several fields (such as job name, estimated task cost) in the report are populated only when an operation is executed concurrently (`CONCURRENT` preference is turned on).
`format`	Report format: • XML • HTML • TEXT (Default)

Return Values

A CLOB object that contains the report

Usage Notes

You must have the SYSDBA or ANALYZE ANY DICTIONARY system privilege to execute this procedure.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Bad input value

ORA-20002: Bad user statistics table, may need to upgrade it

Related Topics

- GATHER_FIXED_OBJECTS_STATS Procedure
  This procedure gathers statistics for all fixed objects (dynamic performance tables).

192.7.97 REPORT_GATHER_SCHEMA_STATS Functions

This function runs the GATHER_SCHEMA_STATS procedure in reporting mode.

The database does not actually gather statistics, but reports all objects that would be affected when invoking GATHER_SCHEMA_STATS. The input set of parameters is exactly the same as in GATHER_SCHEMA_STATS, with two extra parameters.

Syntax

```
DBMS_STATS.REPORT_GATHER_SCHEMA_STATS (
   ownname           IN   VARCHAR2,
   estimate_percent  IN   NUMBER      DEFAULT TO_ESTIMATE_PERCENT_TYPE (
                                        GET_PARAM ('ESTIMATE_PERCENT')),
   block_sample      IN   BOOLEAN     DEFAULT FALSE,
   method_opt        IN   VARCHAR2    DEFAULT GET_PARAM ('METHOD_OPT'),
   degree            IN   NUMBER      DEFAULT TO_DEGREE_TYPE (
                                        GET_PARAM('DEGREE')),
   granularity       IN   VARCHAR2    DEFAULT GET_PARAM('GRANULARITY'),
   cascade           IN   BOOLEAN     DEFAULT TO_CASCADE_TYPE (
                                        GET_PARAM ('CASCADE')),
   stattab           IN   VARCHAR2    DEFAULT NULL,
   statid            IN   VARCHAR2    DEFAULT NULL,
   options           IN   VARCHAR2    DEFAULT 'GATHER',
   statown           IN   VARCHAR2    DEFAULT NULL,
   no_invalidate     IN   BOOLEAN     DEFAULT TO_NO_INVALIDATE_TYPE (
                                        GET_PARAM ('NO_INVALIDATE')),
   force             IN   BOOLEAN     DEFAULT FALSE,
   obj_filter_list   IN   ObjectTab   DEFAULT NULL,
   detail_level      IN   VARCHAR2    DEFAULT 'TYPICAL',
   format            IN   VARCHAR2    DEFAULT 'TEXT')
 RETURN CLOB;
```

Parameters

Table 192-106 REPORT_GATHER_SCHEMA_STATS Function Parameters

Parameter	Description
ownname	Schema to analyze (NULL means current schema)
estimate_percent	Percentage of rows to estimate (NULL means compute): The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Whether or not to use random block sampling instead of random row sampling. Random block sampling is more efficient, but if the data is not randomly distributed on disk, then the sample values may be somewhat correlated. Only pertinent when doing an estimate statistics.
method_opt	Accepts: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.

Table 192-106 (Cont.) REPORT_GATHER_SCHEMA_STATS Function Parameters

Parameter	Description
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics.
cascade	Gather statistics on the indexes as well. Using this option is equivalent to running the GATHER_INDEX_STATS Procedure on each of the indexes in the schema in addition to gathering table and column statistics. Use the constant DBMS_STATS.AUTO_CASCADE to have Oracle determine whether index statistics to be collected or not. This is the default. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics
statid	Identifier (optional) to associate with these statistics within stattab
options	Further specification of which objects to gather statistics for: GATHER: Gathers statistics on all objects in the schema. GATHER AUTO: Gathers all necessary statistics automatically. Oracle implicitly determines which objects need new statistics, and determines how to gather those statistics. When GATHER AUTO is specified, the only additional valid parameters are ownname, stattab, statid, objlist and statown; all other parameter settings are ignored. Returns a list of processed objects. GATHER STALE: Gathers statistics on stale objects as determined by looking at the *_tab_modifications views. Also, return a list of objects found to be stale. GATHER EMPTY: Gathers statistics on objects which currently have no statistics. also, return a list of objects found to have no statistics. LIST AUTO: Returns a list of objects to be processed with GATHER AUTO. LIST STALE: Returns list of stale objects as determined by looking at the *_tab_modifications views. LIST EMPTY: Returns list of objects which currently have no statistics.
objlist	List of objects found to be stale or empty
statown	Schema containing stattab (if different than ownname)

Table 192-106 (Cont.) REPORT_GATHER_SCHEMA_STATS Function Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Gather statistics on objects even if they are locked
obj_filter_list	A list of object filters. When provided, GATHER_SCHEMA_STATS will gather statistics only on objects which satisfy at least one object filter in the list as needed. In a single object filter, we can specify the constraints on the object attributes. The attribute values specified in the object filter are case- insensitive unless double-quoted. Wildcard is allowed in the attribute values. Suppose non-NULL values s1, s2, ... are specified for attributes a1, a2, ... in one object filter. An object o is said to satisfy this object filter if (o.a1 like s1) and (o.a2 like s2) and ... is true. See Applying an Object Filter List.

Table 192-106 (Cont.) REPORT_GATHER_SCHEMA_STATS Function Parameters

Parameter	Description
`detail_level`	Detail level for the content of the report • `BASIC`: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • `TYPICAL`: In addition to the information provided at level `BASIC`, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed-out objects (applies to only auto statistics gathering) With regard to target objects: - owner and name of each target object - target object type (such as: table, index) - start time - end time - completion status • `ALL`: In addition to the information provided at level `TYPICAL`, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name - session ID - parameter values - error message if the operation failed With regard to target objects: - job name - batching details - estimated cost - rank in the target list - columns for which histograms were collected - list of collected extended statistics (if any) - additional error details if the task has failed. Note that several fields (such as job name, estimated task cost) in the report are populated only when an operation is executed concurrently (`CONCURRENT` preference is turned on).
`format`	Report format: • XML • HTML • TEXT (Default)

Usage Notes

To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

Exceptions

`ORA-20000`: Schema does not exist or insufficient privileges

`ORA-20001`: Bad input value

Examples

Applying an Object Filter List

The following example specifies that any table with a "`T`" prefix in the `SAMPLE` schema or any table in the `HR` schema, if stale, will have statistics gathered upon it.

```
DECLARE
  filter_lst  DBMS_STATS.OBJECTTAB := DBMS_STATS.OBJECTTAB();
BEGIN
  filter_lst.extend(2);
  filter_lst(1).ownname := 'SAMPLE';
  filter_lst(1).objname := 'T%';
  filter_lst(2).ownname := 'HR';
  DBMS_STATS.GATHER_SCHEMA_STATS(NULL, obj_filter_list => filter_lst,
                                 options => 'GATHER STALE');
END;
```

192.7.98 REPORT_GATHER_TABLE_STATS Function

This procedure runs the GATHER_TABLE_STATS Procedure in reporting mode.

That is, statistics are not actually collected, but all the objects that will be affected when `GATHER_TABLE_STATS` is invoked are reported.

Syntax

```
DBMS_STATS.REPORT_GATHER_TABLE_STATS (
   ownname          VARCHAR2,
   tabname          VARCHAR2,
   partname         VARCHAR2 DEFAULT NULL,
   estimate_percent NUMBER   DEFAULT to_estimate_percent_type
                                        (get_param('ESTIMATE_PERCENT')),
   block_sample     BOOLEAN  DEFAULT FALSE,
   method_opt       VARCHAR2 DEFAULT get_param('METHOD_OPT'),
   degree           NUMBER   DEFAULT to_degree_type(get_param('DEGREE')),
   granularity      VARCHAR2 DEFAULT GET_PARAM('GRANULARITY'),
   cascade          BOOLEAN  DEFAULT to_cascade_type(get_param('CASCADE')),
   stattab          VARCHAR2 DEFAULT NULL,
   statid           VARCHAR2 DEFAULT NULL,
   statown          VARCHAR2 DEFAULT NULL,
   no_invalidate    BOOLEAN  DEFAULT  to_no_invalidate_type (
                                 get_param('NO_INVALIDATE')),
   stattype         VARCHAR2 DEFAULT 'DATA',
   force            BOOLEAN  DEFAULT FALSE)
   detail_level     VARCHAR2 DEFAULT 'TYPICAL',   format          VARCHAR2 DEFAULT
'TEXT')
 RETURN CLOB;
```

Parameters

Table 192-107 REPORT_GATHER_TABLE_STATS Function Parameters

Parameter	Description
ownname	Schema of table to analyze
tabname	Name of table
partname	Name of partition
estimate_percent	Percentage of rows to estimate (NULL means compute) The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
block_sample	Whether or not to use random block sampling instead of random row sampling. Random block sampling is more efficient, but if the data is not randomly distributed on disk, then the sample values may be somewhat correlated. Only pertinent when doing an estimate statistics.
method_opt	Accepts either of the following options, or both in combination: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] • FOR COLUMNS [column_clause] [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} column_clause is defined as column_clause := column_name | extension name | extension – integer : Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns. – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns. - column_name : Name of a column - extension : can be either a column group in the format of (column_name, Colume_name [, ...]) or an expression The default is FOR ALL COLUMNS SIZE AUTO. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Table 192-107 (Cont.) REPORT_GATHER_TABLE_STATS Function Parameters

Parameter	Description
degree	Degree of parallelism. The default for degree is NULL. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure. NULL means use the table default value specified by the DEGREE clause in the CREATE TABLE or ALTER TABLE statement. Use the constant DBMS_STATS.DEFAULT_DEGREE to specify the default value based on the initialization parameters. The AUTO_DEGREE value determines the degree of parallelism automatically. This is between 1 (serial execution) and DEFAULT_DEGREE (the system default value based on number of CPUs and initialization parameters) according to the size of the object. When using DEGREE=>NULL, DEGREE=>n, or DEGREE=>DBMS_STATS.DEFAULT_DEGREE, the current implementation of DBMS_STATS may use serial execution if the size of the object does not warrant parallel execution.
granularity	Granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'APPROX_GLOBAL AND PARTITION' - similar to 'GLOBAL AND PARTITION' but in this case the global statistics are aggregated from partition level statistics. This option will aggregate all statistics except the number of distinct values for columns and number of distinct keys of indexes. The existing histograms of the columns at the table level are also aggregated.The aggregation will use only partitions with statistics, so to get accurate global statistics, users should make sure to have statistics for all partitions. Global statistics are gathered if partname is NULL or if the aggregation cannot be performed (for example, if statistics for one of the partitions is missing). 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION '- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics.
cascade	Gathers statistics on the indexes for this table. Using this option is equivalent to running the GATHER_INDEX_STATS Procedure on each of the table's indexes. Use the constant DBMS_STATS.AUTO_CASCADE to have Oracle determine whether index statistics are to be collected or not. This is the default. The default value can be changed using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattab	User statistics table identifier describing where to save the current statistics
statid	Identifier (optional) to associate with these statistics within stattab
statown	Schema containing stattab (if different than ownname)

Table 192-107 (Cont.) REPORT_GATHER_TABLE_STATS Function Parameters

Parameter	Description
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
stattype	Statistics type. The only value allowed is DATA.
force	Gather statistics of table even if it is locked

Table 192-107 (Cont.) REPORT_GATHER_TABLE_STATS Function Parameters

Parameter	Description
`detail_level`	Detail level for the content of the report • `BASIC`: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • `TYPICAL`: In addition to the information provided at level `BASIC`, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed-out objects (applies to only auto statistics gathering) With regard to target objects: - owner and name of each target object - target object type (such as: table, index) - start time - end time - completion status • `ALL`: In addition to the information provided at level `TYPICAL`, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name - session ID - parameter values - error message if the operation failed With regard to target objects: - job name - batching details - estimated cost - rank in the target list - columns for which histograms were collected - list of collected extended statistics (if any) - additional error details if the task has failed. Note that several fields (such as job name, estimated task cost) in the report are populated only when an operation is executed concurrently (`CONCURRENT` preference is turned on).
`format`	Report format: • XML • HTML • TEXT (Default)

Return Values

A `CLOB` object that contains the report

Usage Notes

To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

Related Topics

- GATHER_TABLE_STATS Procedure
  This procedure gathers table, column, and index statistics. It attempts to parallelize as much work as possible, but there are some restrictions, which are described in the individual parameters.

192.7.99 REPORT_SINGLE_STATS_OPERATION Function

This function generates a report for the provided operation optionally in a particular pluggable database (PDB) in a multitenant environment.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

Syntax

```
DBMS_STATS.REPORT_SINGLE_STATS_OPERATIONS (
   opid              NUMBER,
   detail_level      VARCHAR2  DEFAULT 'TYPICAL',
   format            VARCHAR2  DEFAULT 'TEXT'
   container_id      NUMBER    DEFAULT NULL)
 RETURN CLOB;
```

Parameters

Table 192-108 REPORT_SINGLE_STATS_OPERATION Function Parameters

Parameter	Description
opid	Operation ID

Table 192-108 (Cont.) REPORT_SINGLE_STATS_OPERATION Function Parameters

Parameter	Description
detail_level	Detail level for the content of the report • BASIC: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • TYPICAL: In addition to the information provided at level BASIC, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed-out objects (applies to only auto statistics gathering) With regard to target objects: - owner and name of each target object - target object type (such as: table, index) - start time - end time - completion status • ALL: In addition to the information provided at level TYPICAL, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name - session ID - parameter values - error message if the operation failed With regard to target objects: - job name - batching details - estimated cost - rank in the target list - columns for which histograms were collected - list of collected extended statistics (if any) - reason for including the object in the target list (applies to only automatic statistics gathering operation tasks) - additional error details if the task has failed. Note that several fields (such as job name, estimated task cost) in the report are populated only when an operation is executed concurrently (CONCURRENT preference is turned on).

Table 192-108 (Cont.) REPORT_SINGLE_STATS_OPERATION Function Parameters

Parameter	Description
format	Report format: • XML • HTML • TEXT (Default)
container_id	ID of the pluggable database (PDB) on which this operation was performed. Note that in a multitenant environment, operation ID does not uniquely identify an operation. That is, different operations from distinct PDBs may have the same operation ID. Hence, in a multitenant environment, if a PDB ID is not provided, then the report may contain multiple operations. In a typical (non-CDB) database environment, operation ID is unique to each operation.

Usage Notes

To invoke this procedure you need the ANALYZE ANY privilege and the ANALYZE ANY DICTIONARY privilege.

192.7.100 REPORT_STATS_OPERATIONS Function

This function generates a report of all statistics operations that take place between two timestamps which may or may not have been provided.

It allows the scope of the report to be narrowed down so that report will include only auto statistics gathering runs. Furthermore, in a multitenant environment, users may optionally provide a set of pluggable database (PDB) IDs so that only statistics operations from the specified pluggable databases will be reported.

Syntax

```
DBMS_STATS.REPORT_STATS_OPERATIONS (
   detail_level       VARCHAR2 DEFAULT 'TYPICAL',
   format             VARCHAR2 DEFAULT 'TEXT',
   latestN            NUMBER DEFAULT NULL,
   since              TIMESTAMP WITH TIME ZONE DEFAULT NULL,
   until              TIMESTAMP WITH TIME ZONE DEFAULT NULl,
   auto_only          BOOLEAN DEFAULT FALSE,
   container_ids      DBMS_UTILITY.NUMBER_ARRAY DEFAULT DBMS_STATS.NULL_NUMTAB)
 RETURN CLOB;
```

Parameters

Table 192-109 REPORT_STATS_OPERATIONS Function Parameters

Parameter	Description
detail_level	Detail level for the content of the report • BASIC: The report includes - operation ID - operation name - operation target object - start time - end time - completion status (such as: succeeded, failed) • TYPICAL: In addition to the information provided at level BASIC, the report includes individual target objects for which statistics are gathered in this operation. Specifically, with regard to operation related details: - total number of target objects - total number of successfully completed objects - total number of failed objects - total number of timed -out objects (applies to only auto statistics gathering) • ALL: In addition to the information provided at level TYPICAL, the report includes further information on each target object. Specifically, with regard to operation-related details: - job name (if the operation was run in a job) - session ID - parameter values - additional error details if the operation has failed
format	Report format: • XML • HTML • TEXT (Default)
latestN	Restricts the report to contain only the latest N operations that took place between the provided time points (since and until). The default value is NULL, meaning that all qualifying operations will be reported.
since	The report will include only statistics operations that started after this timestamp.
until	The report will include only statistics operations that before after this timestamp.
auto_only	When TRUE, the report will contain only auto statistics gathering job runs.
container_ids	A multitenant environment contains one or more pluggable databases (PDBs). container_ids represents a set of PDB IDs so that only statistics operations from the specified PDBs are reported (applies to only multitenant environments).

Usage Notes

To invoke this procedure you need the ANALYZE ANY privilege and the ANALYZE ANY DICTIONARY privilege.

Examples

Note that the type for `container_ids` input parameter is `DBMS_UTILITY.NUMBER_ARRAY` which is an associative PL/SQL array collection. Although associative array type allows for more flexible `harvals` table-like organization of entries, this function treats `container_ids` as a regular table collection with the first ID located at index 1 and the last id located at index `container_ids.count` without any empty array slot left between any two IDs. An example for 3 container ids is provided.

```
DECLARE
    conid_tab  DBMS_UTILITY.NUMBER_ARRAY;
    report clob;
BEGIN
    conid_tab(1) := 124;
    conid_tab(2) := 63;
    conid_tab(3) := 98;
    report := DBMS_STATS.REPORT_STATS_OPERATIONS (container_ids => conid_tab);
END;
```

192.7.101 RESET_ADVISOR_TASK Procedure

This procedure resets an Optimizer Statistics Advisor task execution to its initial state. Only reset a task that is not currently executing.

Syntax

```
DBMS_STATS.RESET_ADVISOR_TASK (
  task_name IN VARCHAR2);
```

Parameters

Table 192-110 RESET_ADVISOR_TASK Procedure Parameters

Parameter	Description
`task_name`	The name of the Optimizer Statistics Advisor task.

Security Model

Note the following:

- To execute this subprogram, you must have the `ADVISOR` privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Exceptions

- `ORA-20000`: Insufficient privileges
- `ORA-20001`: Invalid input values
- `ORA-20012`: Optimizer Statistics Advisor errors

192.7.102 RESET_COL_USAGE Procedure

This procedure deletes the recorded column (group) usage information.

This procedure should only be used in very rare cases when the seed column usage needs to be initialized.

Syntax

```
DBMS_STATS.RESET_COL_USAGE (
   ownname     IN    VARCHAR2,
   tabname     IN    VARCHAR2);
```

Parameters

Table 192-111 RESET_COL_USAGE Procedure Parameters

Parameter	Description
ownname	Owner name. If NULL it deletes column usage information for tables in all schemas in the database.
tabname	Table name. If NULL it deletes column usage information for all tables of ownname. If both the owner and tabname is NULL, the seed column usage is stopped if applicable. See :SEED_COL_USAGE Procedure for more information.

Usage Notes

To run this procedure, you need to have the SYSDBA administrative privilege, or both the ANALYZE ANY DICTIONARY and the ANALYZE ANY system privileges.

192.7.103 RESET_GLOBAL_PLSQL_PREFS_DEF Procedure

Reset all global PL/SQL preferences to their default value.

Syntax

```
EXEC DBMS_STATS.RESET_GLOBAL_PLSQL_PREF_DEFAULTS
```

Parameters

None.

192.7.104 RESET_PARAM_DEFAULTS Procedure

This deprecated procedure resets the default values of all parameters to Oracle recommended values.

> **Note:**
>
> This subprogram has been replaced by improved technology and is maintained only for purposes of backward compatibility.
>
> See also . DBMS_STATS Deprecated Subprograms

Syntax

```
DBMS_STATS.RESET_PARAM_DEFAULTS;
```

192.7.105 RESTORE_DATABASE_STATS Procedure

This procedure restores statistics of all tables of the database as of a specified timestamp (`as_of_timestamp`).

Syntax

```
DBMS_STATS.RESTORE_DATABSE_STATS(
   as_of_timestamp        TIMESTAMP WITH TIME ZONE,
   force                  BOOLEAN DEFAULT FALSE,
   no_invalidate          BOOLEAN DEFAULT to_no_invalidate_type
                                              (GET_PARAM('NO_INVALIDATE')));
```

Parameters

Table 192-112 RESTORE_DATABASE_STATS Procedure Parameters

Parameter	Description
`as_of_timestamp`	The timestamp to which to restore statistics
`force`	Restores statistics even if their statistics are locked
`no_invalidate`	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • `TRUE`: Dependent cursors are not invalidated. • `FALSE`: Dependent cursors are marked for immediate invalidation. • `AUTO`: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values

ORA-20006: Unable to restore statistics, statistics history not available

192.7.106 RESTORE_DICTIONARY_STATS Procedure

This procedure restores statistics of all dictionary tables (tables of 'SYS', 'SYSTEM' and RDBMS component schemas) as of a specified timestamp (as_of_timestamp).

Syntax

```
DBMS_STATS.RESTORE_DICTIONARY_STATS(
   as_of_timestamp        TIMESTAMP WITH TIME ZONE,
   force                  BOOLEAN DEFAULT FALSE,
   no_invalidate          BOOLEAN DEFAULT to_no_invalidate_type
                                              (GET_PARAM('NO_INVALIDATE')));
```

Parameters

Table 192-113 RESTORE_DICTIONARY_STATS Procedure Parameters

Parameter	Description
as_of_timestamp	Timestamp to which to restore statistics
force	Restores statistics even if their statistics are locked
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Usage Notes

To run this procedure, you must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values

ORA-20006: Unable to restore statistics, statistics history not available

192.7.107 RESTORE_FIXED_OBJECTS_STATS Procedure

This procedure restores statistics of all fixed tables as of a specified timestamp (`as_of_timestamp`).

Syntax

```
DBMS_STATS.RESTORE_FIXED_OBJECTS_STATS(
   as_of_timestamp        TIMESTAMP WITH TIME ZONE,
   force                  BOOLEAN DEFAULT FALSE,
   no_invalidate          BOOLEAN DEFAULT to_no_invalidate_type
                                            (GET_PARAM('NO_INVALIDATE')));
```

Parameters

Table 192-114 RESTORE_FIXED_OBJECTS_STATS Procedure Parameters

Parameter	Description
`as_of_timestamp`	The timestamp to which to restore statistics
`force`	Restores statistics even if their statistics are locked
`no_invalidate`	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • `TRUE`: Dependent cursors are not invalidated. • `FALSE`: Dependent cursors are marked for immediate invalidation. • `AUTO`: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Usage Notes

To run this procedure, you must have the `SYSDBA` or `ANALYZE ANY DICTIONARY` system privilege.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20001`: Invalid or inconsistent values

`ORA-20006`: Unable to restore statistics, statistics history not available

192.7.108 RESTORE_SCHEMA_STATS Procedure

This procedure restores statistics of all tables of a schema as of a specified timestamp (`as_of_timestamp`).

Syntax

```
DBMS_STATS.RESTORE_SCHEMA_STATS(
   ownname                VARCHAR2,
   as_of_timestamp        TIMESTAMP WITH TIME ZONE,
   force                  BOOLEAN DEFAULT FALSE,
```

```
   no_invalidate          BOOLEAN DEFAULT to_no_invalidate_type
                                                 (GET_PARAM('NO_INVALIDATE')));
```

Parameters

Table 192-115 RESTORE_SCHEMA_STATS Procedure Parameters

Parameter	Description
ownname	Schema of the tables for which the statistics are to be restored
as_of_timestamp	The timestamp to which to restore statistics
force	Restores statistics even if their statistics are locked
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values

ORA-20006: Unable to restore statistics, statistics history not available

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.109 RESTORE_SYSTEM_STATS Procedure

This procedure restores system statistics as of a specified timestamp (as_of_timestamp).

Syntax

```
DBMS_STATS.RESTORE_SCHEMA_STATS(
   as_of_timestamp        TIMESTAMP WITH TIME ZONE);
```

Parameters

Table 192-116 RESTORE_SYSTEM_STATS Procedure Parameters

Parameter	Description
as_of_timestamp	The timestamp to which to restore statistics

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values

ORA-20006: Unable to restore statistics, statistics history not available

Usage Notes

To run this procedure, you need the GATHER_SYSTEM_STATISTICS role.

192.7.110 RESTORE_TABLE_STATS Procedure

This procedure restores statistics of a table as of a specified timestamp (as_of_timestamp). It also restores statistics of associated indexes and columns.

If the table statistics were locked at the specified timestamp the procedure will lock the statistics. The procedure will not restore user defined statistics.

Syntax

```
DBMS_STATS.RESTORE_TABLE_STATS (
   ownname                  VARCHAR2,
   tabname                  VARCHAR2,
   as_of_timestamp          TIMESTAMP WITH TIME ZONE,
   restore_cluster_index    BOOLEAN DEFAULT FALSE,
   force                    BOOLEAN DEFAULT FALSE,
   no_invalidate            BOOLEAN DEFAULT to_no_invalidate_type
                                          (GET_PARAM('NO_INVALIDATE')));
```

Parameters

Table 192-117 RESTORE_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	The schema of the table for which the statistics are to be restored
tabname	The table name
as_of_timestamp	The timestamp to which to restore statistics
restore_cluster_index	If the table is part of a cluster, restore statistics of the cluster index if set to TRUE
force	Restores statistics even if the table statistics are locked. If the table statistics were not locked at the specified timestamp, it unlocks the statistics.
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent values

ORA-20006: Unable to restore statistics, statistics history not available

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.111 RESUME_ADVISOR_TASK Procedure

This procedure resumes an interrupted task. It only resumes the execution that was most recently interrupted.

Syntax

```
DBMS_STATS.RESUME_ADVISOR_TASK (
  task_name IN VARCHAR2);
```

Parameters

Table 192-118 RESUME_ADVISOR_TASK Procedure Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Consider a case in which a task is executed by one user, interrupted, and then resumed by a different user. In this case, Optimizer Statistics Advisor bases its checks of the resumed execution on the privilege of the user who resumed the task.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

Example 192-14 Resuming an Interrupted Task

In this example, you start a SQL*Plus session, and then create and execute an advisor task named `my_task`:

```
DECLARE
  v_tname   VARCHAR2(128) := 'my_task';
  v_ename   VARCHAR2(128) := NULL;
BEGIN
  -- create a task
  v_tname := DBMS_STATS.CREATE_ADVISOR_TASK(v_tname);

  -- execute the task
  v_ename := DBMS_STATS.EXECUTE_ADVISOR_TASK(v_tname);
END;
/
```

In a separate terminal, you start a second SQL*Plus session, and then execute the following program:

```
XEC DBMS_STATS.INTERRUPT_ADVISOR_TASK('my_task');
```

The first session returns an `ORA-13632` to indicate the cancelation of the task:

```
ORA-13638: The user interrupted the current operation.
```

In the second SQL*Plus session, you resume the task execution as follows:

```
XEC DBMS_STATS.RESUME_ADVISOR_TASK('my_task');
```

192.7.112 SCRIPT_ADVISOR_TASK Function

Retrieves the script that implements the recommended actions for the problems found by Optimizer Statistics Advisor.

The generated script contains PL/SQL statements that you can choose to execute. Preceding the commands for each action are comments that list the potential side effects. You can review the comments, and choose to execute only the desired sections.

Syntax

```
DBMS_STATS.SCRIPT_ADVISOR_TASK (
  task_name          IN   VARCHAR2,
  execution_name     IN   VARCHAR2    := NULL,
  dir_name           IN   VARCHAR2    := NULL,
  level              IN   VARCHAR2    := 'TYPICAL')
RETURN CLOB;
```

Parameters

Table 192-119 SCRIPT_ADVISOR_TASK Function Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.
execution_name	A name that qualifies and identifies an advisor execution. If not specified, then the advisor automatically generates it. If the specified execution conflicts with the name of an existing execution, then the function returns an error.
dir_name	Directory name to which to write the generated script. If the name is not specified (NULL), then the function includes the script in the returned CLOB. If the name is specified, then the function returns the script as a CLOB and as a new file in the specified directory.
level	The level of the script to generate. Possible values are • ALL: Ignores the filter and generates a script for all findings • TYPICAL: Generates a script according to the filters in place

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- You can execute this subprogram for AUTO_STATS_ADVISOR_TASK, which is predefined.
- This subprogram executes using invoker's rights.

The results of performing this task depend on the privileges of the executing user:

- SYSTEM level

  Only users with both the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges can perform this task on system-level rules.

- Operation level

  The results depend on the following privileges:

  – Users with both the ANALYZE ANY and ANALYZE ANY DICTIONARY privileges can perform this task for all statistics operations.

  – Users with the ANALYZE ANY privilege but *not* the ANALYZE ANY DICTIONARY privilege can perform this task for statistics operations related to any schema except SYS.

  – Users with the ANALYZE ANY DICTIONARY privilege but *not* the ANALYZE ANY privilege can perform this task for statistics operations related to their own schema and the SYS schema.

  – Users with neither the ANALYZE ANY nor the ANALYZE ANY DICTIONARY privilege can only perform this operation for statistics operations relating to their own schema.

- Object level

  Users can perform this task for any object for which they have statistics collection privileges.

Return Values

This function returns a CLOB that contains the script.

Exceptions

- ORA-20000: Insufficient privileges
- ORA-20001: Invalid input values
- ORA-20012: Optimizer Statistics Advisor errors

Example 192-15 Creating an Optimizer Statistics Advisor Script

This example creates a procedure named myrep, and then calls this procedure to print the script the implements the recommendations.

```
SET ECHO ON
SET FEEDBACK ON
SET SERVEROUTPUT ON
SET TRIMS ON
SET LINESIZE 300

EXECUTE DBMS_OUTPUT.ENABLE (buffer_size => 10000000);

CREATE OR REPLACE PROCEDURE myrep(p_tname VARCHAR2, p_ftype VARCHAR2, which
VARCHAR2)
IS
  v_report CLOB          := null;
  v_script CLOB          := null;
  v_ftype  VARCHAR2(400) := p_ftype;
  v_tname  VARCHAR2(400) := p_tname;
  v_len    NUMBER(10);
  v_ps     NUMBER(10)    := 10000;
  v_pn     NUMBER(10)    := 1;
  v_ret    VARCHAR2(32767);
BEGIN
   IF which = 'REPORT'
   THEN
      -- generate a report
      v_report := DBMS_STATS.REPORT_ADVISOR_TASK(
                    task_name => v_tname,
                    type      => v_ftype,
                    section   => 'ALL',
                    level     => 'ALL');
      v_len := DBMS_LOB.getlength(v_report);
      WHILE (v_pn < v_len)
      LOOP
        DBMS_OUTPUT.PUT_LINE(DBMS_LOB.SUBSTR(v_report, v_ps, v_pn));
        v_pn := v_pn + v_ps;
      END LOOP;
   ELSE
      -- generate a script
      v_script := DBMS_STATS.SCRIPT_ADVISOR_TASK(v_tname);
      v_len := DBMS_LOB.getlength(v_script);
      WHILE (v_pn < v_len)
      LOOP
```

```
        DBMS_OUTPUT.PUT_LINE(DBMS_LOB.SUBSTR(v_script, v_ps, v_pn));
        v_pn := v_pn + v_ps;
      END LOOP;
   END IF;
END;
/

SHOW ERRORS

SPOOL report.txt
EXECUTE myrep('my_task','-','SCRIPT');
SPOOL OFF
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide*

192.7.113 SEED_COL_USAGE Procedure

This procedure seeds column usage information from a statements in the specified SQL tuning set, or in the database.

The procedure iterates over the SQL statements, compiles them, and then seeds column usage information for the columns that appear in these statements. You can monitor the workload on the system for given amount of time and seed the and seed the column usage information based on the columns that appear in statements executed during the monitoring window.

Syntax

```
DBMS_STATS.SEED_COL_USAGE (
   sqlset_name     IN    VARCHAR2,
   owner_name      IN    VARCHAR2,
   time_limit      IN    POSITIVE DEFAULT NULL);
```

Parameters

Table 192-120 SEED_COL_USAGE Procedure Parameters

Parameter	Description
`sqlset_name`	Name of the SQL tuning set that contains the statements to be monitored. If this parameter and `owner_name` are both null, then the procedure monitors all statements in the database for the specified time limit.
`owner_name`	Owner of the SQL tuning set that contains the statements to be monitored. If this parameter and `sqlset_name` are both null, then the procedure monitors all statements in the database for the specified time limit.
`time_limit`	Time limit (in seconds).

Security Model

To invoke this procedure you must have the `ANALYZE ANY` privilege and the `ANALYZE ANY DICTIONARY` privilege.

Exceptions

ORA-20000: Insufficient privileges

Usage Notes

This procedure also records group of columns. You can create extensions for the recorded group of columns using the CREATE_EXTENDED_STATS Function procedure. If sqlset_name and owner_name are NULL, then the procedure records the column (group) usage information for the statements executed in the system in next time_limit seconds.

This monitoring procedure records different information from the traditional column usage information that is visible in SYS.COL_USAGE$. The procedure stores information in SYS.COL_GROUP_USAGE$.

Examples

The following example turns on monitoring for 5 minutes or 300 seconds.

```
BEGIN
   DBMS_STATS.SEED_COL_USAGE (null,null,300);
END;
```

192.7.114 SET_ADVISOR_TASK_PARAMETER Procedure

This procedure updates the value of an Optimizer Statistics Advisor task parameter.

Syntax

```
DBMS_STATS.SET_ADVISOR_TASK_PARAMETER (
  task_name     IN   VARCHAR2,
  parameter     IN   VARCHAR2,
  value         IN   VARCHAR2);
```

Parameters

Table 192-121 SET_ADVISOR_TASK_PARAMETER Procedure Parameters

Parameter	Description
task_name	The name of the Optimizer Statistics Advisor task.
parameter	The name of the parameter to set. The function returns an error if the specified parameter does not exist.
value	The new value of the parameter.

Security Model

Note the following:

- To execute this subprogram, you must have the ADVISOR privilege.
- You must be the owner of the task.
- This subprogram executes using invoker's rights.

Exceptions

- ORA-20000: Insufficient privileges

- `ORA-20001`: Invalid input values
- `ORA-20012`: Optimizer Statistics Advisor errors

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to manage Optimizer Statistics Advisor

192.7.115 SET_COLUMN_STATS Procedures

This procedure sets column-related information.

In the version of this procedure that deals with user-defined statistics, the statistics type specified is the type to store in the dictionary, in addition to the actual user-defined statistics. If this statistics type is `NULL`, the statistics type associated with the index or column is stored.

Syntax

```
DBMS_STATS.SET_COLUMN_STATS (
   ownname       VARCHAR2,
   tabname       VARCHAR2,
   colname       VARCHAR2,
   partname      VARCHAR2 DEFAULT NULL,
   stattab       VARCHAR2 DEFAULT NULL,
   statid        VARCHAR2 DEFAULT NULL,
   distcnt       NUMBER DEFAULT NULL,
   density       NUMBER DEFAULT NULL,
   nullcnt       NUMBER DEFAULT NULL,
   srec          StatRec DEFAULT NULL,
   avgclen       NUMBER DEFAULT NULL,
   flags         NUMBER DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL,
   no_invalidate BOOLEAN DEFAULT to_no_invalidate_type(
                                 get_param('NO_INVALIDATE')),
   force         BOOLEAN DEFAULT FALSE);
```

Use the following for user-defined statistics:

```
DBMS_STATS.SET_COLUMN_STATS (
   ownname       VARCHAR2,
   tabname       VARCHAR2,
   colname       VARCHAR2,
   partname      VARCHAR2 DEFAULT NULL,
   stattab       VARCHAR2 DEFAULT NULL,
   statid        VARCHAR2 DEFAULT NULL,
   ext_stats     RAW,
   stattypown    VARCHAR2 DEFAULT NULL,
   stattypname   VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL,
   no_invalidate BOOLEAN DEFAULT to_no_invalidate_type(
                                 get_param('NO_INVALIDATE')),
   force         BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-122 SET_COLUMN_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema.
tabname	Name of the table to which this column belongs.
colname	Name of the column or extension
partname	Name of the table partition in which to store the statistics. If the table is partitioned and partname is NULL, then the statistics are stored at the global table level.
stattab	User statistics table identifier describing where to store the statistics. If stattab is NULL, then the statistics are stored directly in the dictionary.
statid	Identifier (optional) to associate with these statistics within stattab (Only pertinent if stattab is not NULL)
ext_stats	User-defined statistics
stattypown	Schema of the statistics type
stattypname	Name of the statistics type
distcnt	Number of distinct values
density	Column density. If this value is NULL and if distcnt is not NULL, then density is derived from distcnt.
nullcnt	Number of NULLs
srec	StatRec structure filled in by a call to PREPARE_COLUMN_VALUES or GET_COLUMN_STATS
avgclen	Average length for the column (in bytes)
flags	For internal Oracle use (should be left as NULL)
statown	Schema containing stattab (if different than ownname)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
force	Sets the values even if statistics of the column are locked

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or inconsistent input values

ORA-20005: Object statistics are locked

Usage Notes

To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

192.7.116 SET_DATABASE_PREFS Procedure

This procedure sets the statistics preferences of all the tables, excluding the tables owned by the database. These tables can by included by passing `TRUE` for the `add_sys` parameter.

Syntax

```
DBMS_STATS.SET_DATABASE_PREFS (
    pname           IN   VARCHAR2,
    pvalue          IN   VARCHAR2,
    add_sys         IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 192-123 SET_DATABASE_PREFS Procedure Parameters

Parameter	Description
`pname`	Preference name. The existing value for following preferences can be set and default preference values will be used: • `APPROXIMATE_NDV_ALGORITHM` • `AUTO_STAT_EXTENSIONS` • `CASCADE` • `DEGREE` • `ESTIMATE_PERCENT` • `GLOBAL_TEMP_TABLE_STATS` • `GRANULARITY` • `INCREMENTAL` • `INCREMENTAL_STALENESS` • `INCREMENTAL_LEVEL` • `METHOD_OPT` • `NO_INVALIDATE` • `OPTIONS` • `PREFERENCE_OVERRIDES_PARAMETER` • `PUBLISH` • `STALE_PERCENT` • `STAT_CATEGORY` • `TABLE_CACHED_BLOCKS`
`pvalue`	Preference value. If `NULL` is specified, it will set the Oracle default values
`add_sys`	Value `TRUE` will include the Oracle-owned tables

Table 192-124 Statistics Preferences

Preference	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. Consider a synopsis as an internal management structure that samples distinct values. You can specify the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
CASCADE	Determines whether to collect index statistics as part of gathering table statistics.
DEGREE	Determines the degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	Determines the percentage of rows to sample. The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for optimal statistics. This is the default.

Table 192-124 (Cont.) Statistics Preferences

Preference	Description
GLOBAL_TEMP_TABLE_STATS	Controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics. This preference takes two values: • SHARED — All sessions see the same set of statistics • SESSION — Statistics gathered by the GATHER_TABLE_STATS procedure on a global temporary table are session-specific. Thus, the database only uses them for queries issued in the same session as the statistics gathering process. The database deletes session-specific statistics when a session terminates.
GRANULARITY	Determines the granularity of statistics to collect. This preference is only relevant for partitioned tables. The following values are valid: • ALL — Gathers all statistics: subpartition, partition, and global. • AUTO — Determines the granularity based on the partitioning type. This is the default value. • DEFAULT — Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. Use GLOBAL AND PARTITION for this functionality. • GLOBAL — Gathers global statistics. • GLOBAL AND PARTITION — Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. • PARTITION — Gathers partition-level statistics. • SUBPARTITION — Gathers subpartition-level statistics.
INCREMENTAL	Determines whether the global statistics for a partitioned table are maintained without performing a full table scan. When a table is partitioned, an application typically loads data into a new partition. As new partitions are added and data is loaded, global table statistics must be kept up to date. If the following conditions are met, then the database updates the global table statistics by scanning only the changed partitions instead of the entire table: • The INCREMENTAL value for the partitioned table is set to TRUE. • The PUBLISH value for the partitioned table is set to TRUE. • The user specifies AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT and AUTO for GRANULARITY when gathering statistics on the table. If the INCREMENTAL value for the partitioned table was set to FALSE (default value), then the database uses a full table scan to maintain the global statistics. This technique is a much more resource-intensive and time-consuming operation for large tables.

Table 192-124 (Cont.) Statistics Preferences

Preference	Description
INCREMENTAL_LEVEL	Controls which synopses to collect when INCREMENTAL preference is set to TRUE. It takes the following values: • PARTITION — Gathers partition-level synopses. This is the default value. If PARTITION is set on a nonpartitioned table, then the database does not gather synopses. • TABLE — Gathers table-level synopses. Specify this value when you want to exchange this table with a partition. Before the exchange, you can run GATHER_TABLE_STATS on this table with INCREMENTAL set to TRUE and INCREMENTAL_LEVEL to TABLE. The result is that the database gathers table-level synopses on this table. After the exchange, the partition has synopses that come from the table-level synopses of the table before the exchange. You can only use preference value in the SET_TABLE_PREFS procedure: this value is not allowed in the other SET_*_PREFS procedures.

Table 192-124 (Cont.) Statistics Preferences

Preference	Description

INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'.
	The parameter accepts the following values:
	• USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference.
	For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics.
	Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics.
	• USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes.
	For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics.
	• ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist.
	• NULL—A partition or subpartition is considered stale when it has any DML changes.
	For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.
	Note that the following two executions are different:
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	The first execution uses single quotes to set the preference to the value NULL, whereas the second sets the preference to the default, which is ALLOW_MIXED_FORMAT.

Table 192-124 (Cont.) Statistics Preferences

Preference	Description
METHOD_OPT	Controls column statistics collection and histogram creation. When setting preferences at the global, schema, database, or dictionary level, only FOR ALL syntax is allowed: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] The size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} integer — Specifies the number of histogram buckets. The number must be between 1 and 2048. REPEAT — Collects histograms only on the columns that already have histograms. AUTO — Determines the columns on which to collect histograms based on data distribution and the workload of the columns. SKEWONLY — Determines the columns on which to collect histograms based on the data distribution of the columns. The default is FOR ALL COLUMNS SIZE AUTO. You can change the value using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
NO_INVALIDATE	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. If set to TRUE, then the database not invalidate dependent cursors. If set to FALSE, then the procedure invalidates dependent cursors immediately.
OPTIONS	Specifies which objects require statistics to be gathered. Valid values are as follows: • GATHER — Gathers statistics on all objects in the database. This is the default. • GATHER AUTO — Gathers all necessary statistics automatically. This is the default. The database implicitly determines which objects need new statistics and determines how to gather those statistics. When GATHER AUTO is specified, the only additional valid parameters are comp_id, no_invalidate, stattab, statid, and statown; all other parameter settings will be ignored. Also, the database returns a list of objects processed.

Table 192-124 (Cont.) Statistics Preferences

Preference	Description
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • TRUE — Ignores input parameter values, and uses the value of the corresponding preference. • FALSE — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether the database publishes newly gathered statistics after the gathering job completes. You can gather statistics without publishing them immediately. This technique enables you to test new statistics before publishing them.
STALE_PERCENT	Determines the percentage of rows in a table that must change before the statistics on that table are stale and need to be regathered. The valid domain for stale_percent is non-negative numbers. The default value is 10, which means that a table having more than 10% of changes is considered stale.
STAT_CATEGORY	Specifies which statistics to import or export, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics (default) • SYNOPSES — information to support incremental statistics • REALTIME_STATS — specifies only real-time statistics. The value 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.

Security Model

To run this procedure, you must have the SYSDBA role or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privileges.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Invalid or illegal input values

Usage Notes

Both arguments are of type VARCHAR2 and values are enclosed in quotes, even when they represent numbers.

Examples

```
DBMS_STATS.SET_DATABASE_PREFS('CASCADE', 'DBMS_STATS.AUTO_CASCADE');
DBMS_STATS.SET_DATABASE_PREFS('ESTIMATE_PERCENT','9');
DBMS_STATS.SET_DATABASE_PREFS('DEGREE','99');
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set optimizer statistics preferences

192.7.117 SET_GLOBAL_PREFS Procedure

This procedure sets statistics preferences at the global level.

Syntax

```
DBMS_STATS.SET_GLOBAL_PREFS (
    pname     IN   VARCHAR2,
    pvalue    IN   VARCHAR2);
```

Parameters

Table 192-125 SET_GLOBAL_PREFS Procedure Parameters

Parameter	Description
pname	Preference name. The default value for the following preferences can be set: • APPROXIMATE_NDV_ALGORITHM • AUTO_STAT_EXTENSIONS • AUTO_TASK_STATUS • AUTO_TASK_MAX_RUN_TIME • AUTO_TASK_INTERVAL • AUTOSTATS_TARGET • CASCADE • CONCURRENT • DEGREE • ESTIMATE_PERCENT • GLOBAL_TEMP_TABLE_STATS • GRANULARITY • INCREMENTAL • INCREMENTAL_LEVEL • INCREMENTAL_STALENESS • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • STAT_CATEGORY • TABLE_CACHED_BLOCKS • WAIT_TIME_TO_UPDATE_STATS
pvalue	Preference value. If NULL is specified, it will set the Oracle default values

Table 192-126 Global Statistics Preferences

Preference	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. Consider a synopsis as an internal management structure that samples distinct values. You can specify the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
AUTO_TASK_STATUS	Enables or disables the high-frequency automatic optimizer statistics collection. Values are: • ON — Enables high-frequency automatic optimizer statistics collection. • OFF — Disables high-frequency automatic optimizer statistics collection. This is the default.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description
AUTO_TASK_MAX_RUN_TIME	Configures the maximum runtime in seconds of an execution of the the high-frequency automatic statistics collection. The default is 3600 but varies among Oracle Database platforms.
AUTO_TASK_INTERVAL	Specifies the interval in seconds between executions of high-frequency automatic optimizer statistics collection. The minimum value is 60. The default is 900 (equal to 15 minutes).
AUTOSTATS_TARGET	Controls the objects considered for statistics collection. It takes the following values: • 'ALL' - Statistics collected for all objects in system • 'ORACLE' - Statistics collected for all Oracle owned objects • 'AUTO' - Oracle decides on which objects to collect statistics This preference is applicable only for automatic statistics collection.
CASCADE	Determines whether to collect index statistics as part of gathering table statistics.
CONCURRENT	Determines whether statistics are gathered concurrently on multiple objects, or serially, one object at a time. Valid values are: • MANUAL — Concurrency is enabled only for manual statistics gathering. • AUTOMATIC — Concurrency is enabled only for the automatic statistics gathering. • ALL — Concurrency is enabled for both manual and automatic statistics gathering. • OFF — Concurrency is disabled for both manual and automatic statistics. Note the following guidelines: The CONCURRENT preference determines whether statistics are gathered concurrently when the user issues GATHER_*_STATS procedures. DBMS_STATS can collect statistics for a single object in parallel based on the value of the DEGREE parameter. However, parallelism is limited to one object. The CONCURRENT preference extends the scope of parallelism to multiple database objects. This approach is primarily intended for multi-CPU systems, and may not be suitable for small databases on single-CPU computers. To gather statistics concurrently, Resource Manager must be enabled, and the setting for the JOB_QUEUE_PROCESSES initialization parameter must be at least 4.
DEGREE	Determines the degree of parallelism used for gathering statistics.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description
ESTIMATE_PERCENT	Determines the percentage of rows to sample. The valid range is between 0.000001 and 100. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to enable the database to determine the appropriate sample size for optimal statistics. This is the default.
GLOBAL_TEMP_TABLE_STATS	Controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics. This preference takes two values: • SHARED — All sessions see the same set of statistics • SESSION — Statistics gathered by the GATHER_TABLE_STATS procedure on a global temporary table are session-specific. Thus, the database only uses them for queries issued in the same session as the statistics gathering process. The database deletes session-specific statistics when a session terminates.
GRANULARITY	Determines the granularity of statistics to collect. This preference is only relevant for partitioned tables. The following values are valid: • ALL — Gathers all statistics: subpartition, partition, and global. • AUTO — Determines the granularity based on the partitioning type. This is the default value. • DEFAULT — Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. Use GLOBAL AND PARTITION for this functionality. • GLOBAL — Gathers global statistics. • GLOBAL AND PARTITION — Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. • PARTITION — Gathers partition-level statistics. • SUBPARTITION — Gathers subpartition-level statistics.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description
INCREMENTAL	Determines whether the global statistics for a partitioned table are maintained without performing a full table scan. When a table is partitioned, an application typically loads data into a new partition. As new partitions are added and data is loaded, global table statistics must be kept up to date. If the following conditions are met, then the database updates the global table statistics by scanning only the changed partitions instead of the entire table: • The INCREMENTAL value for the partitioned table is set to TRUE. • The PUBLISH value for the partitioned table is set to TRUE. • The user specifies AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT and AUTO for GRANULARITY when gathering statistics on the table. If the INCREMENTAL value for the partitioned table was set to FALSE (default value), then the database uses a full table scan to maintain the global statistics. This technique is a much more resource-intensive and time-consuming operation for large tables.
INCREMENTAL_LEVEL	Controls which synopses to collect when INCREMENTAL preference is set to TRUE. It takes the following values: • PARTITION — Gathers partition-level synopses. This is the default value. If PARTITION is set on a nonpartitioned table, then the database does not gather synopses. • TABLE — Gathers table-level synopses. Specify this value when you want to exchange this table with a partition. Before the exchange, you can run GATHER_TABLE_STATS on this table with INCREMENTAL set to TRUE and INCREMENTAL_LEVEL to TABLE. The result is that the database gathers table-level synopses on this table. After the exchange, the partition has synopses that come from the table-level synopses of the table before the exchange. You can only use preference value in the SET_TABLE_PREFS procedure: this value is not allowed in the other SET_*_PREFS procedures.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description
INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'. The parameter accepts the following values: • USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference. For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics. Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics. • USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes. For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics. • ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist. • NULL—A partition or subpartition is considered stale when it has any DML changes. For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description

	Note that the following two executions are different:
	
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	
	The first execution uses single quotes to set the preference to the value `NULL`, whereas the second sets the preference to the default, which is `ALLOW_MIXED_FORMAT`.

`METHOD_OPT`	Controls column statistics collection and histogram creation. It accepts either of the following options, or both in combination: • `FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause]` `size_clause` is defined as `size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY}` – `integer`: Number of histogram buckets. Must be in the range [1,2048]. – `REPEAT` : Collects histograms only on the columns that already have histograms – `AUTO` : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns – `SKEWONLY` : Oracle determines the columns on which to collect histograms based on the data distribution of the columns - `column_name` : name of a column - `extension` : can be either a column group in the format of `(column_name, colume_name [, ...])` or an expression The default is `FOR ALL COLUMNS SIZE AUTO`.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description
NO_INVALIDATE	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. If set to TRUE, then the database not invalidate dependent cursors. If set to FALSE, then the procedure invalidates dependent cursors immediately.
OPTIONS	Determines the options parameter used in the GATHER_TABLE_STATS procedure. The preference takes the following values: • GATHER — Gathers statistics for all objects in the table. This is the default. • GATHER AUTO — Oracle recommends using GATHER AUTO to gather necessary statistics, such as histograms, after a table has been bulk loaded and acquired online statistics. This is applicable only to tables that are not using INCREMENTAL statistics. For partitioned tables using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO will gather statistics for a table if it is marked stale or has no statistics. In addition, statistics will be gathered for partitions and sub-partitions that are marked stale or have no statistics. For tables not using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO may gather statistics for partitions and sub-partitions, even if they are not marked stale.
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • TRUE — Ignores input parameter values, and uses the value of the corresponding preference. • FALSE — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether the database publishes newly gathered statistics after the gathering job completes. You can gather statistics without publishing them immediately. This technique enables you to test new statistics before publishing them.

Table 192-126 (Cont.) Global Statistics Preferences

Preference	Description
STALE_PERCENT	Determines the percentage of rows in a table that must change before the statistics on that table are stale and need to be regathered. The valid domain for stale_percent is non-negative numbers. The default value is 10, which means that a table having more than 10% of changes is considered stale.
STAT_CATEGORY	Specifies which statistics to import or export, accepting multiple values separated by a comma. Values supported: • OBJECT_STATS — table statistics, column statistics and index statistics (default) • SYNOPSES — information to support incremental statistics • REALTIME_STATS — specifies only real-time statistics. The value 'OBJECT_STATS, SYNOPSES' specifies table statistics, column statistics, index statistics, and synopses.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.
WAIT_TIME_TO_UPDATE_STATS	Specifies the number of minutes before timing out for locks and pins required for updating statistics. It accepts values in the range 0 to 65535. The default value is 15 minutes. The value 0 gets the locks and pins in no-wait mode.

Security Model

To run this procedure, you must have the SYSDBA or both ANALYZE ANY DICTIONARY and ANALYZE ANY system privilege.

Exceptions

ORA-20000: Insufficient privileges

ORA-20001: Invalid or illegal input values

Usage Notes

- This setting is honored only if there is no preference specified for the table to be analyzed.
- Both arguments are of type VARCHAR2 and values are enclosed in quotes, even when they represent numbers.

Example 192-16 Overriding Statistics Preferences at the Global Level

You use the `SET_GLOBAL_PREFS` procedure to set the `ESTIMATE_PERCENT` preference to `5` for every table in the database that does not have a table preference set. Because `sh.costs` does not have a preference set, the global setting applies to this table.

```
SQL> EXEC DBMS_STATS.SET_GLOBAL_PREFS ('ESTIMATE_PERCENT', '5');

PL/SQL procedure successfully completed.
```

You use `SET_TABLE_PREFS` to set the `PREFERENCE_OVERRIDES_PARAMETER` preference to true for the `sh.sales` table only.

```
SQL> EXEC
DBMS_STATS.SET_TABLE_PREFS('sh','sales','PREFERENCE_OVERRIDES_PARAMETER','TRUE
');

PL/SQL procedure successfully completed.
```

A script attempts to set `ESTIMATE_PERCENT` to `10` when gathering statistics for `sh.sales`. However, because `PREFERENCE_OVERRIDES_PARAMETER` is `TRUE` for this table, and because a global preference is defined, Oracle Database gathers statistics using the global setting of `5` rather than the specified setting of `10`:

```
SQL> EXEC DBMS_STATS.GATHER_TABLE_STATS ('sh', 'costs', ESTIMATE_PERCENT=>10);

PL/SQL procedure successfully completed.
```

Example 192-17 Configuring High-Frequency Automatic Optimizer Statistics Collection

Oracle Database 19c introduces high-frequency automatic optimizer statistics collection. This lightweight task periodically gathers statistics for stale objects. The default interval is 15 minutes. In contrast to the automated statistics collection job, the high-frequency task does not perform actions such as purging statistics for non-existent objects or invoking Optimizer Statistics Advisor.

In this example, you enable high-frequency collection, set its maximum run time to a half hour, and then set the frequency interval to 10 minutes:

```
EXEC DBMS_STATS.SET_GLOBAL_PREFS('AUTO_TASK_STATUS','ON');
EXEC DBMS_STATS.SET_GLOBAL_PREFS('AUTO_TASK_MAX_RUN_TIME','1800');
EXEC DBMS_STATS.SET_GLOBAL_PREFS('AUTO_TASK_INTERVAL','600');
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set optimizer statistics preferences

192.7.118 SET_INDEX_STATS Procedures

These procedures set index-related statistics.

The version of this procedure that accepts `ext_stats` sets statistics for use with domain indexes. The statistics type specified is the type to store in the dictionary, in addition to the actual user-defined statistics. If this statistics type is null, then the database stores the statistics type associated with the index or column.

Syntax

```
DBMS_STATS.SET_INDEX_STATS (
   ownname       VARCHAR2,
   indname       VARCHAR2,
   partname      VARCHAR2  DEFAULT NULL,
   stattab       VARCHAR2  DEFAULT NULL,
   statid        VARCHAR2  DEFAULT NULL,
   numrows       NUMBER    DEFAULT NULL,
   numlblks      NUMBER    DEFAULT NULL,
   numdist       NUMBER    DEFAULT NULL,
   avglblk       NUMBER    DEFAULT NULL,
   avgdblk       NUMBER    DEFAULT NULL,
   clstfct       NUMBER    DEFAULT NULL,
   indlevel      NUMBER    DEFAULT NULL,
   flags         NUMBER    DEFAULT NULL,
   statown       VARCHAR2  DEFAULT NULL,
   no_invalidate BOOLEAN   DEFAULT to_no_invalidate_type(
                                    get_param('NO_INVALIDATE')),
   guessq        NUMBER    DEFAULT NULL,
   cachedblk     NUMBER    DEFAULT NULL,
   cachehit      NUMBER    DEFUALT NULL,
   force         BOOLEAN   DEFAULT FALSE);
```

Use the following syntax for user-defined domain index statistics:

```
DBMS_STATS.SET_INDEX_STATS (
   ownname       VARCHAR2,
   indname       VARCHAR2,
   partname      VARCHAR2  DEFAULT NULL,
   stattab       VARCHAR2  DEFAULT NULL,
   statid        VARCHAR2  DEFAULT NULL,
   ext_stats     RAW,
   stattypown    VARCHAR2 DEFAULT NULL,
   stattypname   VARCHAR2 DEFAULT NULL,
   statown       VARCHAR2 DEFAULT NULL,
   no_invalidate BOOLEAN DEFAULT to_no_invalidate_type(
                                    get_param('NO_INVALIDATE')),
   force         BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 192-127 SET_INDEX_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
indname	Name of the index
partname	Name of the index partition in which to store the statistics. If the index is partitioned and if partname is NULL, then the statistics are stored at the global index level.
stattab	User statistics table identifier describing where to store the statistics. If stattab is NULL, then the statistics are stored directly in the dictionary.
statid	Identifier (optional) to associate with these statistics within stattab (Only pertinent if stattab is not NULL)
ext_stats	User-defined statistics
stattypown	Schema of the statistics type
stattypname	Name of the statistics type
numrows	Number of rows in the index (partition)
numlblks	Number of leaf blocks in the index (partition)
numdist	Number of distinct keys in the index (partition)
avglblk	Average integral number of leaf blocks in which each distinct key appears for this index (partition). If not provided, then this value is derived from numlblks and numdist.
avgdblk	Average integral number of data blocks in the table pointed to by a distinct key for this index (partition). If not provided, then this value is derived from clstfct and numdist.
clstfct	See clustering_factor column of the all_indexes view for a description
indlevel	Height of the index (partition)
flags	For internal Oracle use (should be left as NULL)
statown	Schema containing stattab (if different than ownname)
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
guessq	Guess quality. See the pct_direct_access column of the all_indexes view for a description.
cachedblk	Internal use only. Do not set.
cachehit	Internal use only. Do not set.
force	Sets the values even if statistics of the index are locked

Security Model

To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20001`: Invalid input value

`ORA-20005`: Object statistics are locked

Usage Notes

- The Optimizer uses the cached data to estimate number of cached blocks for index or statistics table access. The total cost of the operation will be combined from the I/O cost of reading not cached blocks from disk, the CPU cost of getting cached blocks from the buffer cache, and the CPU cost of processing the data.
- Oracle maintains `cachedblk` and `cachehit` at all times but uses correspondent caching statistics for optimization as part of the table and index statistics only when the user calls `DBMS_STATS.GATHER_[TABLE/INDEX/SCHEMA/DATABASE]_STATS` procedure for auto mode or `DBMS_STATS.GATHER_SYSTEM_STATS` for manual mode. In order to prevent the user from utilizing inaccurate and unreliable data, the optimizer will compute a 'confidence factor' for each `cachehit` and a `cachedblk` for each object. If the 'confidence factor' for the value meets confidence criteria, this value will be used, otherwise the defaults will be used.
- The automatic maintenance algorithm for object caching statistics assumes that there is only one major workload for the system and adjusts statistics to this workload, ignoring other "minor" workloads. If this is not the case, you must use manual mode for maintaining object caching statistics.
- The object caching statistics maintenance algorithm for auto mode prevents you from using statistics in the following situations
  - When not enough data has been analyzed, such as when an object has been recently create
  - When the system does not have one major workload resulting in averages not corresponding to real values.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set artificial statistics

192.7.119 SET_PARAM Procedure

This deprecated procedure sets default values for parameters of `DBMS_STATS` procedures.

> **Note:**
>
> This subprogram has been replaced by improved technology and is maintained only for purposes of backward compatibility. In this case, use the SET_GLOBAL_PREFS Procedure.
>
> See also DBMS_STATS Deprecated Subprograms.

You can use the `GET_PARAM Function` to get the current default value of a parameter.

Syntax

```
DBMS_STATS.SET_PARAM (
   pname       IN    VARCHAR2,
   pval        IN    VARCHAR2);
```

Parameters

Table 192-128 SET_PARAM Procedure Parameters

Parameter	Description
`pname`	The parameter name The default value for following parameters can be set. • `CASCADE` - The default value for `CASCADE` set by `SET_PARAM` is not used by export/import procedures.It is used only by gather procedures. • `DEGREE` • `ESTIMATE_PERCENT` • `METHOD_OPT` • `NO_INVALIDATE` • `GRANULARITY` • `AUTOSTATS_TARGET` - This parameter is applicable only for auto statistics collection. The value of this parameter controls the objects considered for statistics collection (see `pval`)
`pval`	The parameter value. If `NULL` is specified, it will set the default value determined by Oracle. When `pname` is `AUTOSTATS_TARGET`, the following are valid values: • `'ALL'` - Statistics are collected for all objects in the system • `'ORACLE'` - Statistics are collected for all Oracle owned objects • `'AUTO'` - Oracle decides for which objects to collect statistics

Usage Notes

- To run this procedure, you must have the `SYSDBA` or both the `ANALYZE ANY DICTIONARY` and `ANALYZE ANY` system privileges.

- Note that both arguments are of type `VARCHAR2` and the values need to be enclosed in quotes even when they represent numbers.
- Note also the difference between `NULL` and `'NULL'`:
  - When `NULL` is unquoted, this sets the parameter to the value Oracle recommends.
  - In the case of the quoted `'NULL'`, this sets the value of the parameter to `NULL`.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20001`: Invalid or illegal input value

Examples

```
DBMS_STATS.SET_PARAM('CASCADE','DBMS_STATS.AUTO_CASCADE');
DBMS_STATS.SET_PARAM('ESTIMATE_PERCENT','5');
DBMS_STATS.SET_PARAM('DEGREE','NULL');
```

192.7.120 SET_PLSQL_PREFS Procedure

Sets the PL/SQL function preference either for a specific PL/SQL function or for all functions in the same package.

Dynamic sampling is a performance feature that instructs the optimizer to read a sample of data during query compilation to formulate dynamic statistics. Dynamic statistics supplement optimizer statistics such as table and index block counts, table and join cardinalities (estimated number of rows), join column statistics, and GROUP BY statistics. This information helps the optimizer improve plans by making better estimates for predicate cardinality.

The preference value controls whether functions can be executed during dynamic sampling to improve cardinality estimates. High-performance PL/SQL is generally suitable for use in dynamic sampling, but you should exclude long-running PL/SQL functions since they will increase SQL compilation times.

Function-specific preferences set with SET_PLSQL_PREFS are persistent across sessions.

Syntax

```
DBMS_STATS.SET_PLSQL_PREFS(
    user           IN VARCHAR2,
    package_name   IN VARCHAR2,
    function_name  IN VARCHAR2,
    pname          IN VARCHAR2,
    pvalue         IN VARCHAR2
);
```

Parameters

Table 192-129 SET_PLSQL_PREFS Procedure Parameters

Parameter	Description
`user`	Owner of the function. If NULL, the current user is used.

Table 192-129 (Cont.) SET_PLSQL_PREFS Procedure Parameters

Parameter	Description
package_name	The PL/SQL package name. If NULL, this specifies a stand-alone, top-level function.
function_name	Name of the PL/SQL function. To apply the preference to all functions within the package, specify NULL.
pname	The preference that you want to set. Currently only 'dynamic_stats' is accepted.
pvalue	The value of the preference. These are the options for the 'dynamic_stats' preference: • ON - Allow dynamic sampling for this PL/SQL function • OFF - Disallow dynamic sampling for this PL/SQL function • CHOOSE - Leave enabling or disabling dynamic statistics to the discretion of the optimizer. • If pvalue is NULL, the dynamic statistics preference set for this function is deleted and the global preference is used.

Note:

Preference settings have the following order of precedence from higher to lowest:

1. The plsql_function_dynamic_stats hint
2. The session level parameter plsql_function_dynamic_stats
3. The specific PL/SQL function preference that is set using SET_PLSQL_PREFS
4. The global preference that is set using SET_GLOBAL_PLSQL_PREFS

Examples

In this example the owner is the current schema and so the USER parameter is set to NULL. Dynamic statistics for the named function are turned off.

```
EXEC SET_PLSQL_PREFS(NULL, '<PACKAGE_NAME>', '<FUNCTION_NAME>',
'dynamic_stats', 'OFF')
```

In this example, the function_name parameter is set to NULL. Therefore the dynamic statistics preference is set on all functions within the named PL/SQL package.

```
EXEC SET_PLSQL_PREFS('<USER>', '<PACKAGE_NAME>', NULL, 'dynamic_stats', 'ON')
```

In the example below, `pvalue` is set to NULL. The preference set for this function is deleted and the global preference is used. Setting `pvalue` to NULL in this procedure has the same effect as executing `DELETE_PLSQL_PREFS` on the function.

```
EXEC SET_PLSQL_PREFS('<USER>', '<PACKAGE_NAME>',
'<FUNCTION_NAME>','dynamic_stats', NULL)
```

Usage Notes

> **See Also:**
>
> The following privileges enable use of dynamic sampling for PL/SQL functions:
>
> - The owner of the PL/SQL function (or package) can set preferences on those functions.
> - If the session user has SYSDBA privileges, they can set preferences any PL/SQL function (or all functions within a package).
> - If the session user has the ANALYZE ANY DICTIONARY privilege, they can set preferences on any function or package.
> - If the user in the session has the ANALYZE ANY privilege, they can set preferences on any PL/SQL function or package except those owned by SYS.
>
> For additional controls over PL/SQL preferences see Controls for Dynamic Sampling With PL/SQL Functions in the *Database SQL Tuning Guide*. Oracle provides a hint to control PL/SQL dynamic sampling behavior for an individual SQL statement, as well as a session-level parameter.
>
> The *SQL Tuning Guide* also provides details on dynamic statistics at Configuring Options for Dynamic Statistics

192.7.121 SET_GLOBAL_PLSQL_PREFS Procedure

Sets the global preference for PL/SQL functions.

By default, the global preference is set to `CHOOSE`, which lets the optimizer choose the dynamic statistics behavior.

Global preferences set with SET_GLOBAL_PLSQL_PREFS are persistent across sessions.

Syntax

```
DBMS_STATS.SET_GLOBAL_PLSQL_PREFS(
    pname           IN VARCHAR2,
    pvalue          IN VARCHAR2
);
```

Parameters

Table 192-130 SET_GLOBAL_PLSQL_PREFS Procedure Parameters

Parameter	Description
pname	The preference that you want to set. Currently only 'dynamic_stats' is accepted.
pvalue	The value of the preference can either be NULL or one of these three options: • ON - Allow dynamic statistics for PL/SQL functions. • OFF - Disallow dynamic sampling for PL/SQL functions. • CHOOSE - Leave enabling or disabling dynamic statistics to the discretion of the Optimizer. If pvalue is NULL, the global preference is set to the default – CHOOSE.

Examples

Enable dynamic statistics for PL/SQL functions:

```
EXEC DBMS_STATS.SET_PLSQL_GLOBAL_PREFS('dynamic_stats', 'ON')
```

Disable dynamic statistics for PL/SQL functions:

```
EXEC DBMS_STATS.SET_PLSQL_GLOBAL_PREFS('dynamic_stats', 'OFF')
```

Let the optimizer choose whether or not enable the use of dynamic statistics for PL/SQL functions:

```
EXEC DBMS_STATS.SET_PLSQL_GLOBAL_PREFS('dynamic_stats', 'CHOOSE')
EXEC DBMS_STATS.SET_PLSQL_GLOBAL_PREFS('dynamic_stats', NULL)
```

192.7.122 SET_PROCESSING_RATE Procedure

This procedure sets the value of rate of processing for a given operation.

Syntax

```
DBMS_STATS.SET_PROCESSING_RATE (
   opname      IN    VARCHAR2,
   procrate    IN    NUMBER);
```

Parameters

Table 192-131 SET_PROCESSING_RATE Procedure Parameters

Parameter	Description
opname	Name of the operation. Valid values are as follows: • AGGR • ALL • CPU • CPU_ACCESS • CPU_AGGR • CPU_BYTES_PER_SEC • CPU_FILTER • CPU_GBY • CPU_HASH_JOIN • CPU_JOIN • CPU_NL_JOIN • CPU_RANDOM_ACCESS • CPU_SEQUENTIAL_ACCESS • CPU_SM_JOIN • CPU_SORT • HASH • IO • IO_ACCESS • IO_BYTES_PER_SEC • IO_RANDOM_ACCESS • IO_SEQUENTIAL_ACCESS • MEMCMP • MEMCPY
procrate	Processing rate.

Security Model

You must have the OPTIMIZER_PROCESSING_RATE role to run this procedure.

Usage Notes

AUTO DOP uses processing rates to determine the optimal degree of parallelism for a SQL statement.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or illegal input value

192.7.123 SET_SCHEMA_PREFS Procedure

This procedure sets the statistics preferences of all tables owned by the specified user.

Syntax

```
DBMS_STATS.SET_SCHEMA_PREFS (
    ownname    IN   VARCHAR2,
    pname      IN   VARCHAR2,
    pvalue     IN   VARCHAR2);
```

Parameters

Table 192-132 SET_SCHEMA_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
pname	Preference name. You can set the default value for the following preferences: • AUTO_STAT_EXTENSIONS • CASCADE • DEGREE • ESTIMATE_PERCENT • GLOBAL_TEMP_TABLE_STATS • GRANULARITY • INCREMENTAL • INCREMENTAL_LEVEL • INCREMENTAL_STALENESS • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • TABLE_CACHED_BLOCKS
pvalue	Preference value. If NULL is specified, it sets the database default value.

Table 192-133 Statistics Preferences

Preference	Description
AUTO_STAT_EXTENSIONS	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. You can consider a synopsis as an internal management structure that samples distinct values. You can set the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
CASCADE	Determines whether index statistics are collected as part of gathering table statistics.
DEGREE	Determines degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	Determines the percentage of rows to estimate. The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.
GLOBAL_TEMP_TABLE_STATS	This preference takes the following values: • SHARED - All sessions see the same set of statistics • SESSION - Statistics gathered by the GATHER_TABLE_STATS Procedure on a global temporary table are session specific, and hence are only going to be used by the queries issued in the same session as the statistics gathering process. Session-specific statistics are deleted when a session is ended.

Table 192-133 (Cont.) Statistics Preferences

Preference	Description
GRANULARITY	Determines the granularity of statistics to collect (only pertinent if the table is partitioned). Possible values are: 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION'- gathers partition-level statistics 'SUBPARTITION' - gathers subpartition-level statistics.
INCREMENTAL	Determines whether the global statistics of a partitioned table will be maintained without doing a full table scan. With partitioned tables it is very common to load new data into a new partition. As new partitions are added and data loaded, the global table statistics need to be kept up to date. Oracle will update the global table statistics by scanning only the partitions that have been changed instead of the entire table if the following conditions hold: • INCREMENTAL value for the partitioned table is set to TRUE • PUBLISH value for the partitioned table is set to TRUE; • User specifies AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT and AUTO for GRANULARITY when gathering statistics on the table If the INCREMENTAL value for the partitioned table was set to FALSE (default value), a full table scan is used to maintain the global statistics which is a much more resource intensive and time-consuming operation for large tables.

Table 192-133 (Cont.) Statistics Preferences

Preference	Description
INCREMENTAL_LEVEL	Controls which synopses to collect when INCREMENTAL preference is set to TRUE It takes two values: • TABLE - table level synopses are gathered. This is used when you want to exchange this table with a partition. You can run GATHER_TABLE_STATS on this table with INCREMENTAL to TRUE and INCREMENTAL_LEVEL to TABLE before the exchange. The result is that table level synopses are gathered on this table (currently Oracle supports only table level synopses on non-predestined tables). Once the exchange occurs, the partition will have synopses which come from the table level synopses of the table before exchange. This preference value can be only used in the SET_TABLE_PREFS Procedure. It is not allowed in the SET_GLOBAL/DATABASE/SCHEMA_PREFS procedures. • PARTITION - partition level synopses are gathered. This is the default value. If PARTITION is set on a non partitioned table, no synopses are gathered.

Table 192-133 (Cont.) Statistics Preferences

Preference	Description
INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT, USE_LOCKED_STATS, ALLOW_MIXED_FORMAT'. The parameter accepts the following values: • USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference. For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics. Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics. • USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes. Assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics. • ALLOW_MIXED_FORMAT—Partitions with synopses in adaptive sampling format are not considered stale, even when the APPROXIMATE_NDV_ALGORITHM preference is set to HYPERLOGLOG. The database uses existing synopses to derive global NDV. This is the default setting. • NULL—A partition or subpartition is considered stale when it has any DML changes. For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics. Note that the following two executions are different: `EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL'); EXEC DBMS_STATS.SET_TABLE_PREFS`

Table 192-133 (Cont.) Statistics Preferences

Preference	Description
	('sh', 'sales', 'INCREMENTAL_STALENESS', null); The first execution uses single quotes to set the preference to the value NULL, whereas the second sets the preference to the default, which is ALLOW_MIXED_FORMAT.
METHOD_OPT	Controls column statistics collection and histogram creation. It accepts either of the following options, or both in combination: • FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause] size_clause is defined as size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY} – integer: Number of histogram buckets. Must be in the range [1,2048]. – REPEAT : Collects histograms only on the columns that already have histograms – AUTO : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns – SKEWONLY : Oracle determines the columns on which to collect histograms based on the data distribution of the columns - column_name : name of a column - extension : can be either a column group in the format of (column_name, colume_name [, ...]) or an expression The default is FOR ALL COLUMNS SIZE AUTO.
NO_INVALIDATE	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.

Table 192-133 (Cont.) Statistics Preferences

Preference	Description
OPTIONS	Determines the options parameter used in the GATHER_TABLE_STATS procedure. The preference takes the following values: • GATHER — Gathers statistics for all objects in the table. This is the default. • GATHER AUTO — Oracle recommends using GATHER AUTO to gather necessary statistics, such as histograms, after a table has been bulk loaded and acquired online statistics. This is applicable only to tables that are not using INCREMENTAL statistics. For partitioned tables using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO will gather statistics for a table if it is marked stale or has no statistics. In addition, statistics will be gathered for partitions and sub-partitions that are marked stale or have no statistics. For tables not using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO may gather statistics for partitions and sub-partitions, even if they are not marked stale.
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • TRUE — Ignores input parameter values, and uses the value of the corresponding preference. • FALSE — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether newly gathered statistics will be published after the statistics gathering job completes. In releases before Oracle Database 11*g* Release 1 (11.1), when a statistic gathering job completed, the new statistics were automatically published in the dictionary tables. In subsequent releases, you can gather statistics without publishing them immediately. Thus, you can test new statistics before publishing them.
STALE_PERCENT	Determines the percentage of rows in a table that have to change before the statistics on that table are deemed stale and should be regathered. The valid domain for stale_percent is non-negative numbers. The default value is 10%. Note that if you set stale_percent to zero the AUTO STATS gathering job will gather statistics for this table every time a row in the table is modified.
TABLE_CACHED_BLOCKS	Specifies the average number of blocks assumed to be cached in the buffer cache when calculating the index clustering factor. The preference applies only when gathering statistics using DBMS_STATS. Index statistics gathered during CREATE INDEX or REBUILD INDEX operations will use the default value 1.

Security Model

To run this procedure, you must be the schema owner, or have the `SYSDBA` privilege, or have the `ANALYZE ANY` system privilege.

Exceptions

`ORA-20000`: Object does not exist or insufficient privileges

`ORA-20001`: Invalid or illegal input value

Usage Notes

Both arguments are of type `VARCHAR2` and values are enclosed in quotes, even when they represent numbers.

Examples

```
DBMS_STATS.SET_SCHEMA_PREFS('SH','CASCADE','DBMS_STATS.AUTO_CASCADE');
DBMS_STATS.SET_SCHEMA_PREFS('SH','ESTIMATE_PERCENT','9');
DBMS_STATS.SET_SCHEMA_PREFS('SH','DEGREE','99');
```

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set optimizer statistics preferences

192.7.124 SET_SYSTEM_STATS Procedure

This procedure sets systems statistics.

Syntax

```
DBMS_STATS.SET_SYSTEM_STATS (
   pname          VARCHAR2,
   pvalue         NUMBER,
   stattab   IN   VARCHAR2 DEFAULT NULL,
   statid    IN   VARCHAR2 DEFAULT NULL,
   statown   IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 192-134 SET_SYSTEM_STATS Procedure Parameters

Parameter	Description
pname	The parameter name to get, which can have one of the following values: • iotfrspeed—I/O transfer speed in bytes for each millisecond • ioseektim - Seek time + latency time + operating system overhead time, in milliseconds • sreadtim - Average time to read single block (random read), in milliseconds • mreadtim - Average time to read an mbrc block at once (sequential read), in milliseconds • cpuspeed - Average number of CPU cycles for each second, in millions, captured for the workload (statistics collected using 'INTERVAL' or 'START' and 'STOP' options) • cpuspeednw - Average number of CPU cycles for each second, in millions, captured for the no-workload (statistics collected using 'NOWORKLOAD' option. • mbrc - Average multiblock read count for sequential read, in blocks • maxthr - Maximum I/O system throughput, in bytes/second • slavethr - Average slave I/O throughput, in bytes/second
pvalue	Parameter value to get
stattab	Identifier of the user statistics table where the statistics will be obtained. If stattab is null, the statistics will be obtained from the dictionary.
statid	Optional identifier associated with the statistics saved in the stattab
statown	Schema containing stattab (if different from current schema)

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid input value

ORA-20002: Bad user statistics table; may need to be upgraded

ORA-20003: Unable to set system statistics

ORA-20004: Parameter does not exist

Usage Notes

To run this procedure, you need the GATHER_SYSTEM_STATISTICS role.

192.7.125 SET_TABLE_PREFS Procedure

This procedure sets the statistics preferences of the specified table in the specified schema.

Syntax

```
DBMS_STATS.SET_TABLE_PREFS (
    ownname    IN  VARCHAR2,
    tabname    IN  VARCHAR2,
```

```
   pname      IN  VARCHAR2,
   pvalue     IN  VARCHAR2);
```

Parameters

Table 192-135 SET_TABLE_PREFS Procedure Parameters

Parameter	Description
ownname	Owner name
tabname	Table name
pname	Preference name. You can set the default value for following preferences: • APPROXIMATE_NDV_ALGORITHM • AUTO_STAT_EXTENSIONS • CASCADE • DEGREE • ESTIMATE_PERCENT • GRANULARITY • INCREMENTAL • INCREMENTAL_LEVEL • INCREMENTAL_STALENESS • METHOD_OPT • NO_INVALIDATE • OPTIONS • PREFERENCE_OVERRIDES_PARAMETER • PUBLISH • STALE_PERCENT • TABLE_CACHED_BLOCKS
pvalue	Preference value. If NULL is specified, it will set the Oracle default value.

Table 192-136 Statistics Preferences

Preference	Description
APPROXIMATE_NDV_ALGORITHM	Specifies the synopsis generation algorithm. A synopsis is special type of statistic that tracks the number of distinct values (NDV) for each column in a partition. Consider a synopsis as an internal management structure that samples distinct values. You can specify the following preferences: • REPEAT OR HYPERLOGLOG This is the default. If INCREMENTAL is enabled on the table, then the database preserves the format of any existing synopses that use the adaptive sampling algorithm. However, the database creates any new synopses in HyperLogLog format. This approach is attractive when existing performance is acceptable, and you do not want to incur the performance cost of reformatting legacy content. • ADAPTIVE SAMPLING The database uses the adaptive sampling algorithm for all synopses. This is the most conservative option. • HYPERLOGLOG The database uses the HyperLogLog algorithm for all new and stale synopses. In contrast to adaptive sampling, the HyperLogLog algorithm uses a randomization technique. The advantages of HyperLogLog over adaptive sampling are: – The accuracy of the new algorithm is similar to the original algorithm. – The memory required is significantly lower, which typically leads to huge reductions in synopsis size.
AUTO_STAT_EXTENSIONS	Controls the automatic creation of extensions when database statistics are gathered. You can set the following values: • ON — When applicable, a SQL plan directive can trigger the creation of column group statistics based on usage of columns in the predicates in the workload. • OFF— The database does not create column group statistics automatically. The database creates them only when the CREATE_EXTENDED_STATS function is executed, or when extended statistics are specified explicitly in the METHOD_OPT clause of DBMS_STATS. This is the default.
CASCADE	Determines whether to collect index statistics as part of gathering table statistics.

Table 192-136 (Cont.) Statistics Preferences

Preference	Description
CONCURRENT	This preference determines whether statistics will be gathered concurrently on multiple objects, or serially, one object at a time: • 'MANUAL' - Concurrency will be turned on only for manual statistics gathering. • 'AUTOMATIC': Concurrency will be turned on only for the automatic statistics gathering. • 'ALL': Concurrency will be turned on for both manual and automatic statistics gathering. • 'OFF': Concurrency will be turned off for both manual and automatic statistics
DEGREE	Determines degree of parallelism used for gathering statistics.
ESTIMATE_PERCENT	Determines the percentage of rows to estimate. The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.
GLOBAL_TEMP_TABLE_STATS	This controls whether the statistics gathered for a global temporary table should be stored as shared statistics or session statistics. It takes two values: • SHARED - All sessions see the same set of statistics • SESSION - Statistics gathered by the GATHER_TABLE_STATS Procedure on a global temporary table are session specific, and hence are only going to be used by the queries issued in the same session as the statistics gathering process. Session-specific statistics are deleted when a session is ended.
GRANULARITY	Determines granularity of statistics to collect (only pertinent if the table is partitioned). 'ALL' - Gathers all (subpartition, partition, and global) statistics 'AUTO'- Determines the granularity based on the partitioning type. This is the default value. 'DEFAULT' - Gathers global and partition-level statistics. This option is obsolete, and while currently supported, it is included in the documentation for legacy reasons only. You should use the 'GLOBAL AND PARTITION' for this functionality. Note that the default value is now 'AUTO'. 'GLOBAL' - Gathers global statistics 'GLOBAL AND PARTITION' - Gathers the global and partition level statistics. No subpartition level statistics are gathered even if it is a composite partitioned object. 'PARTITION'- Gathers partition-level statistics 'SUBPARTITION' - Gathers subpartition-level statistics.

Table 192-136 (Cont.) Statistics Preferences

Preference	Description
INCREMENTAL	Determines whether or not the global statistics of a partitioned table will be maintained without doing a full table scan. With partitioned tables it is very common to load new data into a new partition. As new partitions are added and data loaded, the global table statistics need to be kept up to date. Oracle will update the global table statistics by scanning only the partitions that have been changed instead of the entire table if the following conditions hold: • INCREMENTAL value for the partitioned table is set to TRUE; • PUBLISH value for the partitioned table is set to TRUE; • User specifies AUTO_SAMPLE_SIZE for ESTIMATE_PERCENT and AUTO for GRANULARITY when gathering statistics on the table. If the INCREMENTAL value for the partitioned table was set to FALSE (default value), a full table scan is used to maintain the global statistics which is a much more resource intensive and time-consuming operation for large tables.
INCREMENTAL_LEVEL	This value controls what synopses to collect when INCREMENTAL preference is set to TRUE It takes two values: • TABLE - table level synopses are gathered. This is used when you want to exchange this table with a partition. You can run GATHER_TABLE_STATS on this table with INCREMENTAL to TRUE and INCREMENTAL_LEVEL to TABLE before the exchange. The result is that table level synopses are gathered on this table (currently Oracle supports only table level synopses on non-predestined tables). Once the exchange occurs, the partition will have synopses which come from the table level synopses of the table before exchange. This preference value can be only used in the SET_TABLE_PREFS Procedure. It is not allowed in the SET_GLOBAL/DATABASE/SCHEMA_PREFS procedures. • PARTITION - partition level synopses are gathered. This is the default value. If PARTITION is set on a non partitioned table, no synopses are gathered.

Table 192-136 (Cont.) Statistics Preferences

Preference	Description
INCREMENTAL_STALENESS	Specifies when a partition or subpartition is considered stale. This parameter takes an enumeration of values, such as 'USE_STALE_PERCENT' and 'USE_LOCKED_STATS'. You can also specify multiple values, such as 'USE_STALE_PERCENT,USE_LOCKED_STATS,ALLOW_MIXED_FORMAT'. The parameter accepts the following values: • USE_STALE_PERCENT—A partition or subpartition is not considered stale when DML changes are below the threshold set by the STALE_PERCENT preference. For example, assume that STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. The partition has 5% DML changes. The database does not regather statistics. Assume a different case in which STALE_PERCENT is 10. You specify USE_STALE_PERCEMENT for INCREMENTAL_STALENESS. However, in this case the partition is locked and has 20% of DML changes. Because the partition is locked, the database does not regather statistics. • USE_LOCKED_STATS—Locked partitions or subpartitions statistics are never considered stale, regardless of DML changes. For example, assume that STALE_PERCENT is 10. You specify 'USE_LOCKED_STATS, USE_STALE_PERCENT'. The partition, which is locked, has 20% DML changes. The partition is not considered stale. The database uses existing statistics to derive global statistics. • ALLOW_MIXED_FORMAT—Adaptive sampling synopses and HyperLogLog synopses are permitted to coexist. • NULL—A partition or subpartition is considered stale when it has any DML changes. For example, assume that STALE_PERCENT is 10. You specify the value 'NULL' for INCREMENTAL_STALENESS. The partition has 5% of DML changes. The database regathers statistics.

Table 192-136 (Cont.) Statistics Preferences

Preference	Description

	Note that the following two executions are different:
	
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', 'NULL');`
	
	`EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'sales', 'INCREMENTAL_STALENESS', null);`
	
	The first execution uses single quotes to set the preference to the value `NULL`, whereas the second sets the preference to the default, which is `ALLOW_MIXED_FORMAT`.

`METHOD_OPT`	The value controls column statistics collection and histogram creation. It accepts either of the following options, or both in combination: • `FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause]` `size_clause` is defined as `size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY}` – `integer`: Number of histogram buckets. Must be in the range [1,2048]. – `REPEAT` : Collects histograms only on the columns that already have histograms – `AUTO` : Oracle determines the columns on which to collect histograms based on data distribution and the workload of the columns – `SKEWONLY` : Oracle determines the columns on which to collect histograms based on the data distribution of the columns - `column_name` : name of a column - `extension` : can be either a column group in the format of `(column_name, colume_name [, ...])` or an expression The default is `FOR ALL COLUMNS SIZE AUTO`.

Table 192-136 (Cont.) Statistics Preferences

Preference	Description
NO_INVALIDATE	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
OPTIONS	Determines the options parameter used in the GATHER_TABLE_STATS procedure. The preference takes the following values: • GATHER — Gathers statistics for all objects in the table. This is the default. • GATHER AUTO — Oracle recommends using GATHER AUTO to gather necessary statistics, such as histograms, after a table has been bulk loaded and acquired online statistics. This is applicable only to tables that are not using INCREMENTAL statistics. For partitioned tables using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO will gather statistics for a table if it is marked stale or has no statistics. In addition, statistics will be gathered for partitions and sub-partitions that are marked stale or have no statistics. For tables not using INCREMENTAL statistics, GATHER_TABLE_STATS with GATHER AUTO may gather statistics for partitions and sub-partitions, even if they are not marked stale.
PREFERENCE_OVERRIDES_PARAMETER	Determines whether to override the input value of a parameter with the preference value of that parameter for a statistics operation. Possible values are: • TRUE — Ignores input parameter values, and uses the value of the corresponding preference. • FALSE — Obeys input parameter values. Specifying this preference does not change the order of precedence of table, global, and default.
PUBLISH	Determines whether the database publishes newly gathered statistics after the gathering job completes. You can gather statistics without publishing them immediately. This technique enables you to test new statistics before publishing them.

Table 192-136 (Cont.) Statistics Preferences

Preference	Description
STALE_PERCENT	Determines the percentage of rows in a table that must change before the statistics on that table are stale and need to be regathered. The valid domain for stale_percent is non-negative numbers. The default value is 10, which means that a table having more than 10% of changes is considered stale.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid or illegal input values

Usage Notes

- To run this procedure, you must connect as owner of the table or have the ANALYZE ANY system privilege.
- All arguments are of type VARCHAR2 and values are enclosed in quotes, even when they represent numbers.

Examples

```
DBMS_STATS.SET_TABLE_PREFS('SH', 'SALES', 'CASCADE', 'DBMS_STATS.AUTO_CASCADE');
DBMS_STATS.SET_TABLE_PREFS('SH', 'SALES', 'ESTIMATE_PERCENT', '9');
DBMS_STATS.SET_TABLE_PREFS('SH', 'SALES', 'DEGREE', '99');
```

Example 192-18 Overriding Statistics Preferences

In this example, legacy scripts set ESTIMATE_PERCENT explicitly rather than using the recommended AUTO_SAMPLE_SIZE. Your goal is to prevent users from using these scripts to set preferences on the sh.costs table.

No preference for ESTIMATE_PERCENT is set for sh.costs or at the global level, so the preference defaults to AUTO_SAMPLE_SIZE:

```
SELECT DBMS_STATS.GET_PREFS ('ESTIMATE_PERCENT', 'sh','costs') AS
"STAT_PREFS" FROM DUAL;

STAT_PREFS
----------
DBMS_STATS.AUTO_SAMPLE_SIZE
```

By default, Oracle Database accepts preferences that are passed to the statistics gathering procedures. To override these parameters, use SET_TABLE_PREFS to set the PREFERENCE_OVERRIDES_PARAMETER preference to TRUE for the costs table only:

```
EXEC DBMS_STATS.SET_TABLE_PREFS ('sh', 'costs',
'PREFERENCE_OVERRIDES_PARAMETER', 'TRUE');
```

A user-created script attempts to set `estimate_percent` to `100` when gathering statistics for `sh.costs`.

```
EXEC DBMS_STATS.GATHER_TABLE_STATS('sh', 'costs', ESTIMATE_PERCENT=>100);
```

However, because `PREFERENCE_OVERRIDES_PARAMETER` is `TRUE` for this table, Oracle Database gathers statistics using `AUTO_SAMPLE_SIZE`, which is the default, rather than the specified value of `100`.

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set optimizer statistics preferences

192.7.126 SET_TABLE_STATS Procedure

This procedure creates artificial table statistics for testing purposes.

Syntax

```
DBMS_STATS.SET_TABLE_STATS (
   ownname         VARCHAR2,
   tabname         VARCHAR2,
   partname        VARCHAR2 DEFAULT NULL,
   stattab         VARCHAR2 DEFAULT NULL,
   statid          VARCHAR2 DEFAULT NULL,
   numrows         NUMBER   DEFAULT NULL,
   numblks         NUMBER   DEFAULT NULL,
   avgrlen         NUMBER   DEFAULT NULL,
   flags           NUMBER   DEFAULT NULL,
   statown         VARCHAR2 DEFAULT NULL,
   no_invalidate   BOOLEAN  DEFAULT to_no_invalidate_type (
                                  get_param('NO_INVALIDATE')),
   cachedblk       NUMBER   DEFAULT NULL,
   cachehit        NUMBER   DEFAULT NULL,
   force           BOOLEAN  DEFAULT FALSE,
   im_imcu_count   NUMBER   DEFAULT NULL,
   im_block_count  NUMBER   DEFAULT NULL,
   scanrate        NUMBER   DEFAULT NULL);
```

Parameters

Table 192-137 SET_TABLE_STATS Procedure Parameters

Parameter	Description
`ownname`	Name of the schema.
`tabname`	Name of the table.

Table 192-137 (Cont.) SET_TABLE_STATS Procedure Parameters

Parameter	Description
partname	Name of the table partition in which to store the statistics. If the table is partitioned and partname is NULL, then the statistics are stored at the global table level.
stattab	Table in which to store the statistics. If stattab is NULL, then the database stores the statistics in the data dictionary.
statid	Identifier (optional) to associate with these statistics within stattab. This identifier is only relevant if stattab is not NULL.
numrows	Number of rows in the table or partition.
numblks	Number of blocks that the table or partition occupies.
avgrlen	Average row length for the table or partition.
flags	For internal use only. Do not set.
statown	Schema containing stattab (if different than ownname).
no_invalidate	Controls the invalidation of dependent cursors when statistics are gathered. The parameter takes the following values: • TRUE: Dependent cursors are not invalidated. • FALSE: Dependent cursors are marked for immediate invalidation. • AUTO: This is the default value. Rolling invalidation is used to invalidate all dependent cursors over a period of time. The performance impact on the database is reduced especially in cases where a large number of cursors are invalidated. You can change the default using the SET_DATABASE_PREFS Procedure, SET_GLOBAL_PREFS Procedure, SET_SCHEMA_PREFS Procedure and SET_TABLE_PREFS Procedure.
cachedblk	For internal use only. Do not set.
cachehit	For internal use only. Do not set.
force	A flag that determines the behavior when statistics are locked. If TRUE, then the procedure sets the values even if the table statistics are locked. By default, the setting is FALSE.
im_imcu_count	The number of In-Memory Compression Units (IMCUs) in the table or partition.
im_block_count	The number of In-Memory blocks in the table or partition.
scanrate	The rate, in MB/s, at which the database scans external tables. This parameter is relevant only for external tables.

Security Model

To invoke this procedure you must be owner of the table, or have the ANALYZE ANY privilege. For objects owned by SYS, you must be either the owner of the table, or have the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Invalid input value

ORA-20002: Bad user statistics table; may need to upgrade it

`ORA-20005`: Object statistics are locked

Usage Notes

For testing purposes, you can manually create artificial statistics for a table, index, or the system using the `DBMS_STATS.SET_*_STATS` procedures. These procedures insert the artificial statistics into the data dictionary directly (when `stattab` is null) or into a user-created table.

> **Note:**
>
> The `DBMS_STATS.SET_*_STATS` procedures are intended for development testing only. Do not use them in a production database. If you set statistics in the data dictionary, then Oracle Database considers the set statistics as the “real” statistics, which means that statistics gathering jobs may not re-gather artificial statistics when they do not meet the criteria for staleness.

The most typical use cases for the `DBMS_STATS.SET_*_STATS` procedures are showing how execution plans change as the numbers of rows or blocks in a table change, or creating realistic statistics for temporary tables.

- The optimizer uses the cached data to estimate number of cached blocks for index or statistics table access. The database calculates the total cost of the operation by combining the I/O cost of reading not cached blocks from disk, the CPU cost of getting cached blocks from the buffer cache, and the CPU cost of processing the data.
- The database maintains `cachedblk` and `cachehit` at all times. However, the database uses the corresponding caching statistics for optimization as part of the table and index statistics only when the user calls the `DBMS_STATS.GATHER_[TABLE/INDEX/SCHEMA/DATABASE]_STATS` procedure for automatic mode or `DBMS_STATS.GATHER_SYSTEM_STATS` for manual mode. To prevent the user from utilizing inaccurate and unreliable data, the optimizer computes a “confidence factor” for each `cachehit` and a `cachedblk` for each object. If the confidence factor for the value meets confidence criteria, then the database uses this value; otherwise, the database uses defaults.
- The automatic maintenance algorithm for object caching statistics assumes that only one major database workload exists. The algorithm adjusts statistics to this workload, ignoring other "minor" workloads. If this assumption is false, then you must use manual mode for maintaining object caching statistics.
- The object caching statistics maintenance algorithm for automatic mode prevents you from using statistics in the following situations:
  - When not enough data has been analyzed, such as when an object has been recently created
  - When the database does not have one major workload, resulting in averages that do not correspond to real values

> **See Also:**
>
> *Oracle Database SQL Tuning Guide* to learn how to set artificial statistics

192.7.127 SHOW_EXTENDED_STATS_NAME Function

This function returns the name of the statistics entry that is created for the user-specified extension. It raises an error if no extension has been created.

Syntax

```
DBMS_STATS.SHOW_EXTENDED_STATS_NAME (
   ownname    VARCHAR2,
   tabname    VARCHAR2,
   extension  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 192-138 SHOW_EXTENDED_STATS_NAME Function Parameters

Parameter	Description
ownname	Owner name of a table
tabname	Name of the table
extension	Can be either a column group or an expression. Suppose the specified table has two column c1, c2. An example column group can be "(c1, c2)" and an example expression can be "(c1 + c2)".

Exceptions

ORA-20000: Object does not exist or insufficient privileges

ORA-20001: Error when processing extension

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.128 TRANSFER_STATS Procedure

This procedure transfers statistics for specified table(s) from a remote database specified by dblink to the local database.

The statistics at the source database are retained. It likewise transfers statistics-related structures such as synopses and DML monitoring information.

Syntax

```
DBMS_STATS.TRANSFER_STATS (
   ownname     IN     VARCHAR2,
   tabname     IN     VARCHAR2,
   dblink      IN     VARCHAR2,   options     IN     NUMBER DEFAULT NULL);
```

Parameters

Table 192-139 TRANSFER_STATS Procedure Parameters

Parameter	Description
ownname	Owner name of a table. If NULL all schemas in the database. If NULL, the procedure will transfer global preferences as well.
tabname	Name of the table. If NULL, all tables in OWNNAME.
dblink	Database link name
options	By default the procedure does not transfer the global preferences. Specifying ADD_GLOBAL_PREFS copies global preferences.

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.129 UNLOCK_PARTITION_STATS Procedure

This procedure enables the user to unlock statistics for a partition.

Syntax

```
DBMS_STATS.UNLOCK_PARTITION_STATS (
    ownname     VARCHAR2,
    tabname     VARCHAR2,
    partname    VARCHAR2);
```

Parameters

Table 192-140 UNLOCK_PARTITION_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema to unlock
tabname	Name of the table
partname	[Sub]Partition name

Usage Notes

To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.

192.7.130 UNLOCK_SCHEMA_STATS Procedure

This procedure unlocks the statistics on all the tables in schema.

Syntax

```
DBMS_STATS.UNLOCK_SCHEMA_STATS (
   ownname     VARCHAR2);
```

Parameters

Table 192-141 UNLOCK_SCHEMA_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema

Usage Notes

- To invoke this procedure you must be owner of the table, or you need the `ANALYZE ANY` privilege. For objects owned by `SYS`, you need to be either the owner of the table, or you need the `ANALYZE ANY DICTIONARY` privilege or the `SYSDBA` privilege.
- When statistics on a table is locked, all the statistics depending on the table, including table statistics, column statistics, histograms and statistics on all dependent indexes, are considered to be locked.
- The `SET_*, DELETE_*, IMPORT_*, GATHER_*` procedures that modify statistics in the dictionary of an individual table, index or column will raise an error if statistics of the object is locked.
- Procedures that operates on multiple objects (such as `GATHER_SCHEMA_STATS`) will skip modifying the statistics of an object if it is locked. Many procedures have force argument to override the lock.
- Neither the UNLOCK_SCHEMA_STATS Procedure nor the UNLOCK_TABLE_STATS Procedure is designed to unlock statistics of corresponding partitions. When you invoke theLOCK_TABLE_STATS Procedure, it sets the statistics lock bit at the table level. In that case, you cannot gather statistics on dependent objects such as partitions and indexes. By the same token, if table statistics are locked, the dependents are locked and you do not need to explicitly invoke the LOCK_PARTITION_STATS Procedure.

192.7.131 UNLOCK_TABLE_STATS Procedure

This procedure unlocks the statistics on the table.

Syntax

```
DBMS_STATS.UNLOCK_TABLE_STATS (
   ownname     VARCHAR2,
   tabname     VARCHAR2);
```

Parameters

Table 192-142 UNLOCK_TABLE_STATS Procedure Parameters

Parameter	Description
ownname	Name of the schema
tabname	Name of the table

Usage Notes

- To invoke this procedure you must be owner of the table, or you need the ANALYZE ANY privilege. For objects owned by SYS, you need to be either the owner of the table, or you need the ANALYZE ANY DICTIONARY privilege or the SYSDBA privilege.
- When statistics on a table is locked, all the statistics depending on the table, including table statistics, column statistics, histograms and statistics on all dependent indexes, are considered to be locked.
- The SET_*, DELETE_*, IMPORT_*, GATHER_* procedures that modify statistics in the dictionary of an individual table, index or column will raise an error if statistics of the object is locked.
- Procedures that operates on multiple objects (such as GATHER_SCHEMA_STATS) will skip modifying the statistics of an object if it is locked. Many procedures have force argument to override the lock.
- Neither the UNLOCK_SCHEMA_STATS Procedure nor the UNLOCK_TABLE_STATS Procedure is designed to unlock statistics of corresponding partitions. When you invoke theLOCK_TABLE_STATS Procedure, it sets the statistics lock bit at the table level. In that case, you cannot gather statistics on dependent objects such as partitions and indexes. By the same token, if table statistics are locked, the dependents are locked and you do not need to explicitly invoke the LOCK_PARTITION_STATS Procedure.

192.7.132 UPGRADE_STAT_TABLE Procedure

This procedure upgrades a user statistics table from an older version.

Syntax

```
DBMS_STATS.UPGRADE_STAT_TABLE (
   ownname    VARCHAR2,
   stattab    VARCHAR2);
```

Parameters

Table 192-143 UPGRADE_STAT_TABLE Procedure Parameters

Parameter	Description
ownname	Name of the schema
stattab	Name of the table

Exceptions

ORA-20000: Unable to upgrade table

Usage Notes

To invoke this procedure you need the privileges to drop and create a table.

193
DBMS_STORAGE_MAP

With the `DBMS_STORAGE_MAP` package, you can communicate with the Oracle background process FMON to invoke mapping operations that populate mapping views. FMON communicates with operating and storage system vendor-supplied mapping libraries.

This chapter contains the following topics:

- Overview
- Operational Notes
- Summary of DBMS_STORAGE_MAP Subprograms

193.1 DBMS_STORAGE_MAP Overview

This terminology and descriptions will help you understand the `DBMS_STORAGE_MAP` API.

- Mapping libraries

  Mapping libraries help you map the components of I/O processing stack elements. Examples of I/O processing components include files, logical volumes, and storage array I/O targets. The mapping libraries are identified in `filemap.ora`.

- Mapping files

  A mapping file is a mapping structure that describes a file. It provides a set of attributes, including file size, number of extents that the file is composed of, and file type.

- Mapping elements and sub-elements

  A mapping element is the abstract mapping structure that describes a storage component within the I/O stack. Examples of elements include mirrors, stripes, partitions, raid5, concatenated elements, and disks—structures that are the mapping building blocks. A mapping sub-element describes the link between an element and the next elements in the I/O mapping stack

- Mapping file extents

  A mapping file extent describes a contiguous chunk of blocks residing on one element. This includes the device offset, the extent size, the file offset, the type (data or parity), and the name of the element where the extent resides. In the case of a raw device or volume, the file is composed of only one file extent component. A mapping file extent is different from Oracle extents.

> **See Also:**
>
> - *Oracle Database Administrator's Guide*for more information
> - *Oracle Database Reference*for `V$MAP` views, including `V$MAP_FILE`, `V$MAP_ELEMENT`, `V$MAP_SUBELEMENT`, `V$MAP_FILE_EXTENT`

193.2 DBMS_STORAGE_MAP Operational Notes

Invoking the `MAP_ELEMENT`, `MAP_FILE`, and `MAP_ALL` functions when mapping information already exists will refresh the mapping, if configuration IDs are supported. If configuration IDs are not supported, invoking these functions again will rebuild the mapping.

> **See Also:**
>
> *Oracle Database Administrator's Guide*for a discussion of the configuration ID, an attribute of the element or file that is changed.

193.3 Summary of DBMS_STORAGE_MAP Subprograms

This table lists the `DBMS_STORAGE_MAP` subprograms and briefly describes them.

Table 193-1 DBMS_STORAGE_MAP Package Subprograms

Subprogram	Description
DROP_ALL Function	Drops all mapping information in the shared memory of the instance
DROP_ELEMENT Function	Drops the mapping information for the element defined by `elemname`
DROP_FILE Function	Drops the file mapping information defined by `filename`
LOCK_MAP Procedure	Locks the mapping information in the shared memory of the instance
MAP_ALL Function	Builds the entire mapping information for all types of Oracle files (except archive logs), including all directed acyclic graph (DAG) elements
MAP_ELEMENT Function	Builds mapping information for the element identified by `elemname`
MAP_FILE Function	Builds mapping information for the file identified by `filename`
MAP_OBJECT Function	Builds the mapping information for the Oracle object identified by the object name, owner, and type
RESTORE Function	Loads the entire mapping information from the data dictionary into the shared memory of the instance
SAVE Function	Saves information needed to regenerate the entire mapping into the data dictionary
UNLOCK_MAP Procedure	Unlocks the mapping information in the shared memory of the instance.

193.3.1 DROP_ALL Function

This function drops all mapping information in the shared memory of the instance.

Syntax

```
DBMS_STORAGE_MAP.DROP_ALL(
      dictionary_update IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 193-2 DROP_ALL Function Parameters

Parameter	Description
dictionary_update	If TRUE, mapping information in the data dictionary is updated to reflect the changes. The default value is TRUE; dictionary_update is an overloaded argument.

193.3.2 DROP_ELEMENT Function

This function drops the mapping information for the element defined by elemname.

Syntax

```
DBMS_STORAGE_MAP.DROP_ELEMENT(
   elemname          IN VARCHAR2,
   cascade           IN BOOLEAN,
   dictionary_update IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 193-3 DROP_ELEMENT Function Parameters

Parameter	Description
elemname	The element for which mapping information is dropped.
cascade	If TRUE, then DROP_ELEMENT is invoked recursively on all elements of the DAG defined by elemname, if possible.
dictionary_update	If TRUE, mapping information in the data dictionary is updated to reflect the changes. The default value is TRUE; dictionary_update is an overloaded argument.

193.3.3 DROP_FILE Function

This function drops the file mapping information defined by filename.

Syntax

```
DBMS_STORAGE_MAP.DROP_FILE(
   filename          IN VARCHAR2,
   cascade           IN BOOLEAN,
   dictionary_update IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 193-4 DROP_FILE Function Parameters

Parameter	Description
filename	The file for which file mapping information is dropped.
cascade	If TRUE, then the mapping DAGs for the elements where the file resides are also dropped, if possible.

Table 193-4 (Cont.) DROP_FILE Function Parameters

Parameter	Description
dictionary_update	If TRUE, mapping information in the data dictionary is updated to reflect the changes. The default value is TRUE; dictionary_update is an overloaded argument.

193.3.4 LOCK_MAP Procedure

This procedure locks the mapping information in the shared memory of the instance.

This is useful when you need a consistent snapshot of the V$MAP tables. Without locking the mapping information, V$MAP_ELEMENT and V$MAP_SUBELEMENT, for example, may be inconsistent.

Syntax

```
DBMS_STORAGE_MAP.LOCK_MAP;
```

193.3.5 MAP_ALL Function

This function builds the entire mapping information for all types of Oracle files (except archive logs), including all directed acyclic graph (DAG) elements. It obtains the latest mapping information because it explicitly synchronizes all mapping libraries.

Syntax

```
DBMS_STORAGE_MAP.MAP_ALL(
   max_num_fileext   IN NUMBER DEFAULT 100,
   dictionary_update IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 193-5 MAP_ALL Function Parameters

Parameter	Description
max_num_fileext	Defines the maximum number of file extents to be mapped. This limits the amount of memory used when mapping file extents. The default value is 100; max_num_fileextent is an overloaded argument.
dictionary_update	If TRUE, mapping information in the data dictionary is updated to reflect the changes. The default value is TRUE; dictionary_update is an overloaded argument.

Usage Notes

You must explicitly call MAP_ALL in a cold startup scenario.

193.3.6 MAP_ELEMENT Function

This function builds mapping information for the element identified by elemname. It may not obtain the latest mapping information if the element being mapped, or any one of the elements

within its I/O stack (if `cascade` is `TRUE`), is owned by a library that must be explicitly synchronized.

Syntax

```
DBMS_STORAGE_MAP.MAP_ELEMENT(
   elemname          IN VARCHAR2,
   cascade           IN BOOLEAN,
   dictionary_update IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 193-6 MAP_ELEMENT Function Parameters

Parameter	Description
`elemname`	The element for which mapping information is built.
`cascade`	If `TRUE`, all elements within the elemname I/O stack DAG are mapped.
`dictionary_update`	If `TRUE`, mapping information in the data dictionary is updated to reflect the changes. The default value is `TRUE`; `dictionary_update` is an overloaded argument.

193.3.7 MAP_FILE Function

This function builds mapping information for the file identified by `filename`. Use this function if the mapping of one particular file has changed. The Oracle database server does not have to rebuild the entire mapping.

Syntax

```
DBMS_STORAGE_MAP.MAP_FILE(
   filename           IN VARCHAR2,
   filetype           IN VARCHAR2,
   cascade            IN BOOLEAN,
   max_num_fileextent IN NUMBER DEFAULT 100,
   dictionary_update  IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 193-7 MAP_FILE Function Parameters

Parameter	Description
`filename`	The file for which mapping information is built.
`filetype`	Defines the type of the file to be mapped. It can be `"DATAFILE"`, `"SPFILE"`, `"TEMPFILE"`, `"CONTROLFILE"`, `"LOGFILE"`, or `"ARCHIVEFILE"`.
`cascade`	Should be `TRUE` only if a storage reconfiguration occurred. For all other instances, such as file resizing (either through an `ALTER SYSTEM` command or DML operations on extended files), `cascade` can be set to `FALSE` because the mapping changes are limited to the file extents only. If `TRUE`, mapping DAGs are also built for the elements where the file resides.

Table 193-7 (Cont.) MAP_FILE Function Parameters

Parameter	Description
max_num_fileextent	Defines the maximum number of file extents to be mapped. This limits the amount of memory used when mapping file extents. The default value is 100; max_num_fileextent is an overloaded argument.
dictionary_update	If TRUE, mapping information in the data dictionary is updated to reflect the changes. The default value is TRUE; dictionary_update is an overloaded argument.

Usage Notes

This function may not obtain the latest mapping information if the file being mapped, or any one of the elements within its I/O stack (if cascade is TRUE), is owned by a library that must be explicitly synchronized.

193.3.8 MAP_OBJECT Function

This function builds the mapping information for the Oracle object identified by the object name, owner, and type.

Syntax

```
DBMS_STORAGE_MAP.MAP_OBJECT(
   objname  IN  VARCHAR2,
   owner    IN  VARCHAR2,
   objtype  IN  VARCHAR2);
```

Parameters

Table 193-8 MAP_OBJECT Function Parameters

Parameter	Description
objname	The name of the object.
owner	The owner of the object.
objtype	The type of the object.

193.3.9 RESTORE Function

This function loads the entire mapping information from the data dictionary into the shared memory of the instance.

You can invoke RESTORE only after a SAVE operation. You must explicitly call RESTORE in a warm startup scenario.

Syntax

```
DBMS_STORAGE_MAP.RESTORE;
```

193.3.10 SAVE Function

This function saves information needed to regenerate the entire mapping into the data dictionary.

Syntax

```
DBMS_STORAGE_MAP.SAVE;
```

193.3.11 UNLOCK_MAP Procedure

This procedure unlocks the mapping information in the shared memory of the instance.

Syntax

```
DBMS_STORAGE_MAP.UNLOCK_MAP;
```

194
DBMS_SYNC_REFRESH

The `DBMS_SYNC_REFRESH` package provides an interface to perform a synchronous refresh of materialized views.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information on using `DBMS_SYNC_REFRESH`

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_SYNC_REFRESH Subprograms

194.1 DBMS_SYNC_REFRESH Overview

Synchronous refresh is a refresh method introduced in Oracle Database Release 12*c*, which enables you to keep a set of tables and the materialized views defined on them to be always in sync.

> **See Also:**
>
> *Oracle Database Data Warehousing Guide* for more information about using synchronous refresh

194.2 DBMS_SYNC_REFRESH Security Model

The execute privilege for this package is granted to `PUBLIC`, so all users can execute the procedures in this package to perform synchronous refresh on objects owned by them. The database administrator can perform synchronous refresh operations on all tables and materialized views in the database.

In general, if a user without the `DBA` privilege wants to use synchronous refresh on another user's table, he must complete privileges to read from and write to that table, that it, the user must have the `SELECT` or `READ`, `INSERT`, `UPDATE`, and `DELETE` privileges on that table or materialized view. A couple of exceptions are:

- `PURGE_REFRESH_STATS` and `ALTER_REFRESH_STATS_RETENTION` Functions

  These two functions implement the purge policy and can be used to change the default retention period. These functions can be only be executed by the database administrator.

- The `CAN_SYNCREF_TABLE` Function

  This is an advisory function which examines the eligibility for sync refresh of all the materialized views associated with a specified table. Hence, this function requires the `SELECT` or `READ` privilege on all materialized views associated with the specified table.

194.3 Summary of DBMS_SYNC_REFRESH Subprograms

This table lists and briefly describes the `DBMS_SYNC_REFRESH` package subprograms.

Table 194-1 DBMS_SYNC_REFRESH Package Subprograms

Subprogram	Description
ABORT_REFRESH Procedure	Aborts a refresh.
ALTER_REFRESH_STATS_RETENTION Procedure	Alters the refresh history retention value, specified in days.
CAN_SYNCREF_TABLE Procedure	Advises on whether a table and its dependent materialized views are eligible for synchronous refresh.
EXECUTE_REFRESH Procedure	Executes synchronous refresh on the synchronous refresh groups.
GET_ALL_GROUP_IDS Function	Returns the group IDs of all the synchronous refresh groups in the database.
GET_GROUP_ID Function	Returns the group ID of a table or materialized view.
GET_GROUP_ID_LIST Function	Returns the group IDs of the tables in a given list of objects (tables or materialized views).
PREPARE_REFRESH Procedure	Prepares the sync refresh groups for refresh.
PREPARE_STAGING_LOG Procedure	Validates and collects statistics on the data in the staging log.
PURGE_REFRESH_STATS Procedure	Purges the refresh history of sync refreshes that took place within a time specified by a timestamp parameter.
REGISTER_MVIEWS	Registers materialized views for synchronous refresh.
REGISTER_PARTITION_OPERATION Procedure	Registers a partition maintenance operation on a partition of a base table.
UNREGISTER_MVIEWS	Unregisters materialized views from synchronous refresh.
UNREGISTER_PARTITION_OPERATION Procedure	Unregisters a partition maintenance operation on a partition of a base table.

194.3.1 ABORT_REFRESH Procedure

This procedure undoes all the changes made by `PREPARE_REFRESH` or `EXECUTE_REFRESH` for the specified sync refresh groups. It helps you to recover to a state where the tables and materialized views are usable and consistent in case they encounter unexpected errors.

This procedure is overloaded.

Syntax

```
DBMS_SYNC_REFRESH.ABORT_REFRESH (
   group_id      IN NUMBER);

DBMS_SYNC_REFRESH.ABORT_REFRESH (
   group_id_list IN DBMS_UTILITY.NUMBER_ARRAY);
```

Parameters

Table 194-2 ABORT_REFRESH Procedure Parameters

Parameter	Description
group_id	The group ID of a sync refresh group.
group_id_list	An array of group IDs of the sync refresh groups to be aborted for sync refresh.

Usage Notes

If called after PREPARE_REFRESH, this procedure drops the outside tables created by it and unlocks the tables and materialized views in the sync refresh group.

If called after EXECUTE_REFRESH fails, this procedure restores the state of tables to before EXECUTE_REFRESH by undoing any partition exchanges which successfully finished.

This procedure releases the locks placed on the tables in the sync refresh group which were placed on them by the PREPARE_REFRESH procedure. See "PREPARE_REFRESH Procedure" for a description of these locks.

ABORT_REFRESH will work only if a PREPARE_REFRESH or EXECUTE_REFRESH statement has failed. It cannot be used after successful runs of those commands, and throws an error in such cases.

194.3.2 ALTER_REFRESH_STATS_RETENTION Procedure

This procedure alters the refresh history retention value, specified in days. It is intended for use in conjunction with PURGE_REFRESH_HISTORY. It also requires the SYSDBA privilege in addition to the privilege to execute it.

Syntax

```
DBMS_SYNC_REFRESH.ALTER_REFRESH_STATS_RETENTION (
   retention    IN NUMBER);
```

Parameters

Table 194-3 ALTER_REFRESH_STATS_RETENTION Procedure Parameters

Parameter	Description
retention	The retention time in days. The refresh history will be retained for at least these many number of days. The valid range is 1 to 365,000. You can use the following values for special purposes: • -1 - Refresh history is never purged by PREPARE_REFRESH. • 0 - Old refresh history is never saved. PREPARE_REFRESH will delete all refresh history. • NULL - Change refresh history retention to default value.

194.3.3 CAN_SYNCREF_TABLE Procedure

This procedure advises on whether a table and its dependent materialized views are eligible for sync refresh. It provides an explanation of its analysis. If not eligible, you can examine the reasons and take appropriate action if possible.

This procedure lists all of the table's dependent materialized views and whether they qualify for sync refresh. Note that a materialized view may qualify for sync refresh even though the base table may not.

The eligibility rules for materialized views for synchronous refresh are discussed in detail in *Oracle Database Data Warehousing Guide.*

You can invoke `CAN_SYNCREF_TABLE` in two ways. The first is to use a table, while the second is to create a `VARRAY`.

Syntax

```
DBMS_SYNC_REFRESH.CAN_SYNCREF_TABLE (
    schema_name     IN VARCHAR2,
    table_name      IN VARCHAR2,
    statement_id    IN VARCHAR2);

DBMS_SYNC_REFRESH.CAN_SYNCREF_TABLE (
    schema_name     IN VARCHAR2,
    table_name      IN VARCHAR2,
    output_array    IN OUT Sys.CanSyncRefTypeArray);
```

Note that only one of `statement_id` or `output_array` need be provided to `CAN_SYNCREF_TABLE`.

Parameters

Table 194-4 CAN_SYNCREF_TABLE Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema of the base table.
`table_name`	The name of the base table.
`statement_id`	A string (`VARCHAR2(30)`) to identify the rows pertaining to an invocation of `CAN_SYNCREF_TABLE` when the output is directed to a table named `SYNCREF_TABLE` in the user's schema.
`output_array`	The output array into which `CAN_SYNCREF_TABLE` records the information on the eligibility of the base table and its dependent materialized views for synchronous refresh.

Using SYNCREF_TABLE

The output of `CAN_SYNCREF_TABLE` can be directed to a table named `SYNCREF_TABLE`. The user is responsible for creating the `SYNCREF_TABLE`; it can be dropped when it is no longer needed. Its structure is as follows:

```
CREATE TABLE SYNCREF_TABLE (
   statement_id    VARCHAR2(30),
   schema_name     VARCHAR2(30),
   table_name      VARCHAR2(30),
   mv_schema_name  VARCHAR2(30),
```

```
   mv_name         VARCHAR2(30),
   eligible        VARCHAR2(1),
   seq_num         NUMBER,
   msg_number      NUMBER,
   message         VARCHAR2(4000));
```

Using a VARRAY

You can save the output of `CAN_SYNCREF_TABLE` in a PL/SQL `VARRAY`. The elements of this array are of type `CanSyncRefMessage`, which is predefined in the `SYS` schema, as shown in the following:

```
TYPE CanSyncRefMessage IS OBJECT (
   schema_name     VARCHAR2(30),
   table_name      VARCHAR2(30),
   mv_schema_name  VARCHAR2(30),
   mv_name         VARCHAR2(30),
   eligible        VARCHAR2(1),
   seq_num         NUMBER,
   msg_number      NUMBER,
   message         VARCHAR2(4000));
```

The array type `CanSyncRefArrayType`, which is a varray of `CanSyncRefMessage` objects, is predefined in the `SYS` schema as follows:

```
TYPE CanSyncRefArrayType AS VARRAY(256) OF CanSyncRefMessage;
```

Each `CanSyncRefMessage` record provides a message concerning the eligibility of the base table or a dependent materialized view for synchronous refresh. The semantics of the fields is the same as that of the corresponding fields in the `SYNCREF_TABLE`. However, the `SYNCREF_TABLE` has a `statement_id` field which is absent in `CanSyncRefMessage` because no `statement_id` is supplied (because it is not required) when `CAN_SYNCREF_TABLE` is called with a `VARRAY` parameter.

194.3.4 EXECUTE_REFRESH Procedure

This procedure executes sync refresh on the sync refresh groups prepared by `DBMS_SYNC_REFRESH.PREPARE_REFRESH`. These groups are identified by their group IDs.

Note this procedure will only perform the refresh on those materialized views that have been registered for synch refresh; any other materialized views will become stale once this procedure completes.

For more information on how to monitor the status of the two synchronous refresh operations, `PREPARE_REFRESH` and `EXECUTE_REFRESH` and how to troubleshoot errors that might occur using the information in the catalog views, refer to "Trouble-Shooting Synchronous Refresh Operations" in *Oracle Database Data Warehousing Guide.*

This procedure is overloaded.

Syntax

```
DBMS_SYNC_REFRESH.EXECUTE_REFRESH (
   group_id    IN NUMBER);

DBMS_SYNC_REFRESH.EXECUTE_REFRESH (
   group_id_list  IN DBMS_UTILITY.NUMBER_ARRAY);
```

Parameters

Table 194-5 EXECUTE_REFRESH Procedure Parameters

Parameter	Description
group_id	The group ID of a sync refresh group.
group_id_list	An array of group IDs of the sync refresh groups to be executed for sync refresh.

Usage Notes

This procedure also releases the locks placed on the tables in the sync refresh group that were placed on them by the PREPARE_REFRESH procedure. See "PREPARE_REFRESH Procedure" for a description of these locks and *Oracle Database Reference* for information regarding the status of the refresh operation after DBMS_SYNC_REFRESH.EXECUTE_REFRESH.

194.3.5 GET_ALL_GROUP_IDS Function

This function returns the group IDs of all the sync refresh groups in the database.

Syntax

```
FUNCTION DBMS_SYNC_REFRESH.GET_ALL_GROUP_IDS
           RETURN DBMS_UTILITY.NUMBER_ARRAY;
```

Parameters

Table 194-6 GET_ALL_GROUP_IDS Function Parameter

Parameter	Description
get_all_group_ids	Returns the group IDs of all the sync refresh groups in the database.

194.3.6 GET_GROUP_ID Function

This function returns the group ID of a materialized view. The group ID identifies the sync refresh group the table belongs to. A sync refresh group is a group of related tables and their dependent materialized views which must be all refreshed together jointly to ensure consistency and correctness.

Syntax

```
DBMS_SYNC_REFRESH.GET_GROUP_ID (
   object_name_list   IN VARCHAR2)
RETURN DBMS_UTILITY.NUMBER_ARRAY;
```

Parameters

Table 194-7 GET_GROUP_ID Function Parameter

Parameter	Description
object_name_list	The name of the materialized view. The name can be schema-qualified.

194.3.7 GET_GROUP_ID_LIST Function

This function returns the group IDs of the tables in a given list of objects (materialized views).

Syntax

```
DBMS_SYNC_REFRESH.GET_GROUP_ID_LIST (
   object_name_list    IN VARCHAR2)
RETURN DBMS_UTILITY.NUMBER_ARRAY;
```

Parameters

Table 194-8 GET_GROUP_ID_LIST Function Parameter

Parameter	Description
object_name_list	A comma-separated list of object names (materialized views). Each name can be schema-qualified.

194.3.8 PREPARE_REFRESH Procedure

This procedure prepares for refresh the sync refresh groups identified by the group ID in the input.

A sync refresh group consists of a set of related tables and all materialized views dependent on those base tables. Note this procedure will only prepare for refresh those dependent materialized views that have been registered for synchronous refresh.

For more information on how to monitor the status of the two synchronous refresh operations, PREPARE_REFRESH and EXECUTE_REFRESH and how to troubleshoot errors that might occur using the information in the catalog views, refer to "Trouble-Shooting Synchronous Refresh Operations" in *Oracle Database Data Warehousing Guide.*

Syntax

```
DBMS_SYNC_REFRESH.PREPARE_REFRESH (
   group_id    IN NUMBER)
RETURN DBMS_UTILITY.NUMBER_ARRAY;
```

Parameters

Table 194-9 PREPARE_REFRESH Procedure Parameters

Parameter	Description
group_id	The group ID of the sync refresh group to be prepared for sync refresh.

Usage Notes

This procedure plans the three phases of the sync refresh operation and executes the steps associated with the prepare phase itself. These steps include identifying the partitions of the fact tables and materialized views that have been changed, and computing their new values as a result of the changes. The new values of the partitions are stored in tables called outside tables that are exchanged into their corresponding partitions at the time of the EXECUTE_REFRESH.

Before running this procedure, the user must run PREPARE_STAGING_LOG on all tables in the group. This is required even for staging logs that do not have changes in them. The user must also register any partition operations on the tables in the group using the REGISTER_PARTITION_OPERATION.

One of the side effects of this procedure is that the tables being prepared are locked in this sense: the staging logs of the tables will be locked to prevent any DMLs from occurring and the registration of partition operations will be disabled. These locks will be in effect until you issue an EXECUTE_REFRESH statement. Alternatively, you can issue an ABORT_REFRESH operation to release these locks. Another side effect of this procedure is that it purges from the catalog records of earlier sync refresh operations; if they are older than the retention period, they are purged.

The degree of parallelism of the prepare refresh job is inherited from the session parameters which you can control with an ALTER SESSION statement.

The group ID of a table can be found using GET_GROUP_ID(table_name). The group IDs of a list of tables can be found with GET_GROUP_ID_LIST(table_name_list). The group IDs of all the lists of tables can be retrieved with GET_ALL_GROUP_IDS.

By default, synchronous refresh does not maintain global indexes belonging to the tables and materialized views in the sync refresh group. If you wish to do so, you can set event 31904, level 64 before executing PREPARE_REFRESH. This will cause the partition exchange DDL statements generated by PREPARE_REFRESH to have the UPDATE INDEXES clause appended to them, and when they are executed by EXECUTE_REFRESH, the global indexes will be maintained.

194.3.9 PREPARE_STAGING_LOG Procedure

This procedure collects statistics on the data in the staging log of the base table and validates the data in the log.

It can be run in several different modes ranging from the enforced mode in which strict checking of the data is done to trusted mode in which no checking is done. You should run this procedure after loading the staging log and before running PREPARE_REFRESH.

In the enforced mode, which is the default, this procedure will fill in the missing values of the columns of the rows being deleted or updated. An error is thrown if any violations of the staging-log rules are found. You can query the view USER_SR_STLOG_EXCEPTIONS to get details on the exceptions.

The notion of the staging log key is described in *Oracle Database Data Warehousing Guide.*

In the enforced mode, this procedure processes each delete/update row in the staging log as follows:

- It verifies the existence of the row in the base table using the key.
- For the rows being deleted (DMLTYPE$$ is 'D'), it verifies a row with this key exists in the base table; if non-null non-key values are supplied in the staging log, it verifies the values match the corresponding columns in the base table; else an exception is logged in the exceptions table. If the values of any of the non-key columns are missing, it fills in those values from the row in the base table.
- For the rows being updated (DMLTYPE$$ is 'UO' or 'UN'), it verifies a row with this key exists in the base table. In the old values row (DMLTYPE$$ is 'UO'), it makes the same check and does the same processing as with rows being deleted. In the new values row (DMLTYPE$$ is 'UN'), it checks that at least the value of one the columns differs from its old value; else an exception is logged.

- In the new values row (DMLTYPE$$ is 'UN'), a null value in a column is interpreted as having the same value as the old value of the column except if the old value is non-null and the new value is null in which case, the new value of the column is interpreted as being null. This requires that the user must provide the old value of columns which are being updated to NULL.

In the default enforced mode, this procedure verifies that each key is specified for at most once for a delete or update operation. This means that the user, when doing the change consolidation, must consolidate delete-insert of the same row into an update operation with rows 'UO' and 'UN;' multiple updates must be consolidated into a single update; and null changes such as an insert-update-delete of the same row must not appear in the staging log.

The checking done in the enforced mode can be time-consuming. If you are confident in the integrity of the data, you can choose a lower level of checking. You can choose to:

- trust all the insert rows (DMLTYPE$$ is 'I') by choosing the psl_mode of DBMS_SYNC_REFRESH.INSERT_TRUSTED
- trust all the delete rows (DMLTYPE$$ is 'D') by choosing the psl_mode of DBMS_SYNC_REFRESH.DELETE_TRUSTED
- trust all the update rows (DMLTYPE$$ is 'UO' or 'UN') by choosing the psl_mode of DBMS_SYNC_REFRESH.UPDATE_TRUSTED
- trust all three types of DMLs by choosing the psl_mode of DBMS_SYNC_REFRESH.TRUSTED.

In addition, you can specify the psl_mode as a bitmask of the flags described above. For example, DBMS_SYNC_REFRESH.INSERT_TRUSTED + DBMS_SYNC_REFRESH_DELETE_TRUSTED will treat inserts and deletes to be trusted but not updates.

Syntax

```
DBMS_SYNC_REFRESH.PREPARE_STAGING_LOG (
   schema_name      IN VARCHAR2,
   base_table_name  IN VARCHAR2,
   psl_mode         IN NUMBER DEFAULT
   DBMS_SYNC_REFRESH.ENFORCED);
```

Parameters

Table 194-10 PREPARE_STAGING_LOG Procedure Parameters

Parameter	Description
schema_name	The name of the schema of the base table.
base_table_name	The name of the base table.
psl_mode	The mode in which staging log preparation should be done. The possible values are: • DBMS_SYNC_REFRESH.ENFORCED (the default) • DBMS_SYNC_REFRESH.INSERT_TRUSTED • DBMS_SYNC_REFRESH.DELETE_TRUSTED • DBMS_SYNC_REFRESH.UPDATE_TRUSTED • DBMS_SYNC_REFRESH.TRUSTED

194.3.10 PURGE_REFRESH_STATS Procedure

This procedure purges the refresh history of sync refreshes that took place before the value specified by the `BEFORE_TIMESTAMP` parameter.

This procedure requires the `SYSDBA` privilege in addition to the privilege to execute it.

Syntax

```
DBMS_SYNC_REFRESH.PURGE_REFRESH_STATS (
   before_timestamp IN TIMESTAMP WITH TIME ZONE);
```

Parameters

Table 194-11 PURGE_REFRESH_STATS Procedure Parameter

Parameter	Description
`before_timestamp`	Records of sync refreshes saved before this timestamp are purged. If `NULL`, it uses the purging policy used by automatic purge. The automatic purge deletes all history older than (current time - refresh - history retention). The refresh history retention value can be changed using `ALTER_REFRESH_STATS_RETENTION`. The default is 31 days.

194.3.11 REGISTER_MVIEWS

This procedure registers a list of materialized views for synchronous refresh.

It checks each materialized view in the list for eligibility and places it in the sync refresh group it belongs to. A sync refresh group is a set of related tables and materialized views defined on top of them. Two tables are considered related if there is a referential constraint between them.

The eligibility rules of materialized views for synchronous refresh are described in detail in *Oracle Database Data Warehousing Guide.* The principal requirements are that the materialized view must be partitioned and its partition key must be derivable from the partition key of its fact table. The materialized view definition must specify the `USING TRUSTED CONSTRAINTS` clause because sync refresh trusts the foreign key and primary key relationships to perform various refresh optimizations. The materialized view's refresh policy must be specified as `ON DEMAND`.

You have an option to register only some of the materialized views associated with a table, and leave some unregistered. Oracle Corporation does not recommend this, and in such a case, the user has to maintain the unregistered ones using the PCT or complete refresh methods.

A staging table must have been created for each base table of each materialized view in the materialized view list (`mv_list`), or else an error is thrown.

If any of the materialized views are not eligible for sync refresh, an error is thrown and the registration of all materialized views in the materialized view list fails.

Syntax

```
DBMS_SYNC_REFRESH.REGISTER_MVIEWS (
   mv_list   IN VARCHAR2);
```

Parameter

Table 194-12 REGISTER_MVIEWS Procedure Parameter

Parameter	Description
mv_list	A comma-delimited list of materialized views to register. These names are optionally schema-qualified.

194.3.12 REGISTER_PARTITION_OPERATION Procedure

This procedure registers a partition-maintenance operation (PMOP) on a partition of a base table.

Syntax

```
DBMS_SYNC_REFRESH.REGISTER_PARTITION_OPERATION (
   partition_op                IN VARCHAR2,
   schema_name                 IN VARCHAR2,
   base_table_name             IN VARCHAR2,
   partition_name              IN VARCHAR2,
   outside_partn_table_schema  IN VARCHAR2,
   outside_partn_table_name    IN VARCHAR2);
```

Parameters

Table 194-13 REGISTER_PARTITION_OPERATION Procedure Parameters

Parameter	Description
partition_op	The name of the partition operation (DROP, EXCHANGE, or TRUNCATE).
schema_name	The name of the schema of the base table.
base_table_name	The name of the base table.
partition_name	The name of the partition to be changed; either exchanged with the outside partition table or dropped or truncated.
outside_partn_table_schema	The name of the schema of the outside partition table (required for EXCHANGE only).
outside_partn_table_name	The name of the outside partition table (required for EXCHANGE only).

Usage Notes

The three kinds of change operations that may be specified on partitions are DROP, TRUNCATE, and EXCHANGE.

If DROP is specified, then the partition will be dropped from the base table at the time of EXECUTE_REFRESH. If TRUNCATE is specified, then the data from the partition will be deleted but the partition itself will not be dropped. These operations provide a more efficient way of specifying the deletes of all the rows in a partition than specifying them individually in the staging log.

If EXCHANGE is specified, then the contents of the outside table is exchanged with contents of the specified partition of EXECUTE_REFRESH. This provides an alternative method to the user of providing the changes to the base tables instead of populating the staging log.

194.3.13 UNREGISTER_MVIEWS

This procedure unregisters a list of materialized views from synchronous refresh. Once a materialized view is unregistered, it can be maintained by the user with any of the traditional refresh methods, such as complete or PCT, refresh.

Syntax

```
DBMS_SYNC_REFRESH.UNREGISTER_MVIEWS (
   mv_list   IN VARCHAR20;
```

Parameter

Table 194-14 UNREGISTER_MVIEWS Parameter

Parameter	Description
mv_list	A comma-delimited list of materialized views to unregister. These names are optionally schema-qualified.

194.3.14 UNREGISTER_PARTITION_OPERATION Procedure

This procedure unregisters a partition-maintenance operation (PMOP) that had been previously registered with REGISTER_PARTITION_OPERATION on a base table. The three kinds of change operations that can be specified on partitions are DROP, TRUNCATE, and EXCHANGE.

Syntax

```
DBMS_SYNC_REFRESH.UNREGISTER_PARTITION_OPERATION (
   partition_op     IN VARCHAR2,
   schema_name      IN VARCHAR2,
   base_table_name  IN VARCHAR2,
   partition_name   IN VARCHAR2);
```

Parameters

Table 194-15 UNREGISTER_PARTITION_OPERATION Procedure Parameters

Parameter	Description
partition_op	The name of the partition operation (DROP, EXCHANGE, or TRUNCATE).
schema_name	The name of the schema of the base table.
base_table_name	The name of the base table.
partition_name	The name of the partition to be changed; either exchanged with the outside partition table or dropped or truncated.

195
DBMS_TABLE_DATA

The `DBMS_TABLE_DATA` package provides procedures that can be used to retrieve the byte value of a single column, or a series of columns, in any table. These procedures, along with a publicly documented meta-data structure for columns, can be used to retrieve the row content for row data on which the hash or user signature is computed.

This chapter contains the following topics:

- DBMS_TABLE_DATA Overview
- DBMS_TABLE_DATA Security Model
- Summary of DBMS_TABLE_DATA Subprograms

> **See Also:**
>
> - *Oracle Database Administrator's Guide*
> - *Oracle Database Concepts*
> - *Oracle Database SQL Language Reference*
> - *Oracle Database Reference*

195.1 DBMS_TABLE_DATA Overview

The `DBMS_TABLE_DATA` package can be used to retrieve the byte value of a single column, or a series of columns, in any table.

These procedures, along with a publicly documented meta-data structure for columns, can be used to retrieve the `row content` for row data on which the hash or user signature is computed.

The procedures provide platform-neutral binary values for columns.

The procedures give you fine-grained control over the scope of a digital signature or cryptographic hash in terms of the columns you extract and the order in which you assemble the results.

195.2 DBMS_TABLE_DATA Security Model

The `DBMS_TABLE_DATA` package is owned by `SYS` and is installed as part of database installation. The routines in the package are run with invokers' rights (run with the privileges of the current user). The current user needs `SELECT` privileges on the table to extract column data.

195.3 Summary of DBMS_TABLE_DATA Subprograms

The DBMS_TABLE_DATA package uses GET_BYTES_FOR_COLUMN, GET_BYTES_FOR_COLUMNS, and GET_BYTES_FOR_ROW subprograms to perform various functions.

Table 195-1 DBMS_TABLE_DATA Package Subprograms

Subprogram	Description
GET_BYTES_FOR_COLUMN Procedure	Returns the column_data in bytes for the particular column with name column_name for row_id row in the particular table identified by schema_name.table_name.
GET_BYTES_FOR_COLUMNS Procedure	Returns a concatenated array of column byte values in column_data in the order of columns specified in the column_names array.
GET_BYTES_FOR_ROW Procedure	Returns a concatenated array of column byte values in column_data in the order of column positions for the particular row identified by row_id.

195.3.1 GET_BYTES_FOR_COLUMN Procedure

This procedure returns the column_data in bytes for the particular column with name column_name for row_id row in the particular table identified by schema_name.table_name.

Syntax

```
DBMS_TABLE_DATA.GET_BYTES_FOR_COLUMN(
   schema_name          IN VARCHAR2,
   table_name              IN VARCHAR2,
   row_id                  IN ROWID,
   column_name             IN VARCHAR2,
   column_data             IN OUT BLOB);
```

Parameters

Table 195-2 GET_BYTES_FOR_COLUMN Procedure Parameters

Parameter	Description
schema_name	The name of the schema.
table_name	The name of the table.
row_id	The row id.
column_name	The column name.
column_data	The data in the column.

Usage Notes

All arguments are required.

195.3.2 GET_BYTES_FOR_COLUMNS Procedure

This procedure returns a concatenated array of column byte values in `column_data` in the order of columns specified in the `column_names` array.

Syntax

```
DBMS_TABLE_DATA.GET_BYTES_FOR_COLUMNS(
   schema_name          IN VARCHAR2,
   table_name              IN VARCHAR2,
   row_id                  IN ROWID,
   column_names            IN VARCHAR2,
   column_data             IN OUT BLOB);
```

Parameters

Table 195-3 GET_BYTES_FOR_COLUMNS Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the table.
`row_id`	The row id.
`column_names`	The names of the columns.
`column_data`	The data in the columns.

Usage Notes

All arguments are required.

195.3.3 GET_BYTES_FOR_ROW Procedure

This procedure returns a concatenated array of column byte values in `column_data` in the order of column positions for the particular row identified by `row_id`.

Syntax

```
DBMS_TABLE_DATA.GET_BYTES_FOR_ROW(
   schema_name          IN VARCHAR2,
   table_name              IN VARCHAR2,
   row_id                  IN ROWID,
   column_data             IN OUT BLOB);
```

Parameters

Table 195-4 GET_BYTES_FOR_ROW Procedure Parameters

Parameter	Description
`schema_name`	The name of the schema.
`table_name`	The name of the table.
`row_id`	The row id.
`column_data`	The data in the row.

Usage Notes

All arguments are required.

196
DBMS_TDB

The `DBMS_TDB` package reports whether a database can be transported between platforms using the RMAN `CONVERT DATABASE` command.

The package verifies that databases on the current host platform are of the same endian format as the destination platform, and that the state of the current database does not prevent transport of the database.

> **See Also:**
>
> *Oracle Database Backup and Recovery User's Guide* regarding database transport using `CONVERT DATABASE`

This chapter contains the following topics:

- Overview
- DBMS_TDB Security Model
- Constants
- Views
- Operational Notes
- Summary of DBMS_TDB Subprograms

196.1 DBMS_TDB Overview

In many cases, Oracle supports transporting databases between platforms which have the same endian format. However, even when the endian formats are the same, a database must undergo a conversion process to move from one platform to another. There are also preconditions required for the process of transporting a database, such as having the database to be transported open read-only.

The `DBMS_TDB` package serves two purposes:

- Confirming that Oracle supports transporting a database from a given source platform to a given target platform
- Determining whether a database to be transported has been properly prepared for transport, and if not, identifying the condition that prevents database transport

The actual conversion is performed using the Recovery Manager `CONVERT DATABASE` command. For a complete discussion of the requirements for transporting a database, the process of converting a database for transport across platforms, and examples of the use of the `DBMS_TDB` subprograms in the conversion process, see *Oracle Database Backup and Recovery User's Guide*.

196.2 DBMS_TDB Security Model

Use of this package requires the DBA privilege.

196.3 DBMS_TDB Constants

The DBMS_TDB package defines several enumerated constants that should be used for specifying parameter values. Enumerated constants must be prefixed with the package name, for example, DBMS_TDB.SKIP_NONE.

The DBMS_TDB package uses the constants shown in Table 196-1.

Table 196-1 DBMS_TDB Constants

Name	Type	Value	Description
SKIP_NONE	NUMBER	0	Check all files when checking whether a database is ready for transport.
SKIP_OFFLINE	NUMBER	2	Skip files in offline tablespaces when checking whether a database is ready for transport.
SKIP_READONLY	NUMBER	3	Skip files in read-only tablespaces when checking whether a database is ready for transport.

196.4 DBMS_TDB Views

The DBMS_TDB package uses the V$DB_TRANSPORTABLE_PLATFORM view.

This view is described in *Oracle Database Reference*.

- V$DB_TRANSPORTABLE_PLATFORM, which specifies which combinations of source and target platforms support database transport

196.5 DBMS_TDB Operational Notes

The following notes apply to DBMS_TDB.

- The subprograms in this package are useful both in determining whether the desired cross-platform database conversion is possible, and in checking whether the database is ready for conversion. See *Oracle Database Backup and Recovery User's Guide* for details on the different uses of these subprograms are used in the conversion process.
- The subprograms in this package return simple TRUE or FALSE results to indicate whether database transport is possible. Use the subprograms with SERVEROUTPUT ON for informative messages about why transport is not possible.

196.6 Summary of DBMS_TDB Subprograms

This table lists the DBMS_TDB subprograms and briefly describes them.

Table 196-2 DBMS_TDB Package Subprograms

Subprogram	Description
CHECK_DB Function	Checks whether a database can be transported to a target platform
CHECK_EXTERNAL Function	Checks whether a database has external tables, directory or BFILEs

196.6.1 CHECK_DB Function

This function checks whether a database can be transported to a target platform. It tests whether transport is supported at all for a given source and destination platform, and whether the database is currently in the correct state for transport.

You can specify whether to skip checking parts of the database that are read-only or offline, if you do not plan to transport them.

The function is overloaded. The different functionality of each form of syntax is presented along with the definition.

Syntax

```
DBMS_TDB.CHECK_DB (
     target_platform_name   IN VARCHAR2,
     skip_option            IN  NUMBER)
   RETURN BOOLEAN;

DBMS_TDB.CHECK_DB (
     target_platform_name   IN VARCHAR2)
   RETURN BOOLEAN;

DBMS_TDB.CHECK_DB
   RETURN BOOLEAN;
```

Parameters

Table 196-3 CHECK_DB Function Parameters

Parameter	Description
`target_platform_name`	The name of the destination platform, as it appears in `V$DB_TRANSPORTABLE_PLATFORM`.
`skip_option`	Specifies which, if any, parts of the database to skip when checking whether the database can be transported. Supported values are listed in Table 196-1.

Return Values

If the database cannot be transported to the target platform or is not ready to be transported, returns `FALSE`. If the database is ready for transport, returns `TRUE`.

Usage Notes

- If `SERVEROUTPUT` is `ON`, then the output will contain the reasons why the database cannot be transported and how to fix the problems. For details on possible reasons and fixes, see Table 196-4.

Table 196-4 Reasons for CHECK_DB Function to Return FALSE

Cause	Action
Unrecognized target platform name.	Check `V$DB_TRANSPORTABLE_PLATFORM` for recognized platform names.
Target platform has a different endian format.	Conversion is not supported.
Database is not open read-only.	Open database read-only and retry.
There are active or in-doubt transactions in the database.	Open the database read-write. After the active transactions are rolled back, open the database read-only and retry the operation. This situation can occur if users flash back the database and open it read only. The active transactions will be rolled back when the database is opened read-write.
Deferred transaction rollback needs to be done.	Open the database read-write and bring online the necessary tablespaces. Once the deferred transaction rollback is complete, open the database read-only and retry the operation.
Database compatibility version is below 10.0.0.	Change the `COMPATIBLE` initialization parameter to 10.0.0 or higher, open the database read-only, and retry the operation.
Some tablespaces have not been open read-write with compatibility version is 10.0.0 or higher.	Change the `COMPATIBLE` initialization parameter to 10.0.0 or higher, then open the affected tablespaces read-write. Shut down the database, open it read-only, and retry the operation.

Examples

This example illustrates the use of `CHECK_DB` with a database that is open read-write:

```
SQL> SET SERVEROUTPUT ON
SQL> DECLARE
       db_ready BOOLEAN;
     BEGIN
       db_ready := DBMS_TDB.CHECK_DB('Microsoft Windows IA (32-bit)');
     END;
     /

Database is not open READ ONLY. Please open database READ ONLY and retry.

PL/SQL procedure successfully completed.
```

196.6.2 CHECK_EXTERNAL Function

This function determines whether a database has external tables, directories, or BFILEs.

Syntax

```
DBMS_TDB.CHECK_EXTERNAL
   RETURN BOOLEAN;
```

Return Values

If the database has external tables, directories, or BFILEs, return `TRUE`. Otherwise, return `FALSE`.

Usage Notes

- If SERVEROUTPUT is ON, then the function will output the names of the external tables, directories, and BFILEs in the database.
- The database must be open read-write.

Examples

This example illustrates the use of CHECK_EXTERNAL with a database that has several external tables, directories, and BFILEs:

```
SQL> SET SERVEROUTPUT ON
SQL> DECLARE
        external BOOLEAN;
     BEGIN
       external := DBMS_TDB.CHECK_EXTERNAL;
     END;
     /
The following external tables exist in the database:
SH.SALES_TRANSACTIONS_EXT
The following directories exist in the database:
SYS.MEDIA_DIR, SYS.DATA_FILE_DIR, SYS.LOG_FILE_DIR, SYS.DATA_PUMP_DIR
The following BFILEs exist in the database:
PM.PRINT_MEDIA

PL/SQL procedure successfully completed.
```

197
DBMS_TF

The DBMS_TF package contains utilities for Polymorphic Table Functions (PTF) implementation. You can use DBMS_TF subprograms to consume and produce data, and get information about its execution environment.

You must be familiar with the Polymorphic Table Function (PTF) concepts, syntax and semantics.

> **See Also:**
>
> - *Oracle Database PL/SQL Language Reference* for an overview of Polymorphic Table Function (PTF) concepts
> - *Oracle Database PL/SQL Language Reference* for more information about `CREATE FUNCTION PIPELINED` clause syntax and semantics

This chapter contains the following topics:

- DBMS_TF Overview
- DBMS_TF Security Model
- DBMS_TF Constants
- DBMS_TF Operational Notes
- DBMS_TF Execution Flow
- DBMS_TF Restrictions
- DBMS_TF Examples
- DBMS_TF Data Structures
- Summary of DBMS_TF Subprograms

197.1 DBMS_TF Overview

The DBMS_TF package contains types, constants, and subprograms that can be used by Polymorphic Table Functions (PTFs).

Polymorphic Table Functions (PTFs) need various services from the database to implement their functionality. PTFs need a mechanism to get rows from the database and send back new rows, for instance. The DBMS_TF package provides these server and client interfaces utilities.

197.2 DBMS_TF Security Model

PUBLIC is granted the EXECUTE privilege on package DBMS_TF. Its subprograms execute with invoker's rights privileges.

197.3 DBMS_TF Constants

This topic describes useful constants defined in the DBMS_TF package.

The `DBMS_TF` package defines several enumerated constants that should be used for specifying parameter values or types. Enumerated constants must be prefixed with the package name, for example, `DBMS_TF.TYPE_DATE`.

Table 197-1 DBMS_TF Supported Types

Name	Description
`TYPE_BINARY_DOUBLE`	Type code for `BINARY_DOUBLE`
`TYPE_BINARY_FLOAT`	Type code for `BINARY_FLOAT`
`TYPE_BLOB`	Type code for `BLOB`
`TYPE_BOOLEAN`	Type code for `BOOLEAN`
`TYPE_CHAR`	Type code for `CHAR`
`TYPE_CLOB`	Type code for `CLOB`
`TYPE_DATE`	Type code for `DATE`
`TYPE_INTERVAL_DS`	Type code for `INTERVAL_DS`
`TYPE_INTERVAL_YM`	Type code for `INTERVAL_YM`
`TYPE_NUMBER`	Type code for `NUMBER`
`TYPE_ROWID`	Type code for `ROWID`
`TYPE_RAW`	Type code for `RAW`
`TYPE_TIMESTAMP`	Type code for `TIMESTAMP`
`TYPE_TIMESTAMP_TZ`	Type code for `TIMESTAMP_TZ`
`TYPE_VARCHAR2`	Type code for `VARCHAR2`

Additional constants are defined for use with specific subprograms.

See Also:

- Table 197-3 for more information about CSTORE related constants
- Table 197-4 for more information about predefined PTF method names
- Table 197-6 for more information about XSTORE related constants
- Supported Types Collections for more information about predefined collections of supported types

197.4 DBMS_TF Operational Notes

These operational notes describe the client and the server-side interfaces, and detail the compilation and execution statement management of Polymorphic Table Functions (PTF).

197.4.1 PTF Client Interface

The Polymorphic Table Function (PTF) implementation client interface is a set of subprograms with fixed names that every PTF must provide.

The PTF client interface can have up to four subprograms as follow :

- `DESCRIBE` function (Required)
- `OPEN` procedure (Optional)
- `FETCH_ROWS` procedure (Optional)
- `CLOSE` procedure (Optional)

The function `DESCRIBE` is invoked during SQL cursor compilation.

The procedures `OPEN`, `FETCH_ROWS`, and `CLOSE` are invoked during query execution.

The arguments to the implementation functions must match the PTF function with the following modifications:

1. Arguments of the type `TABLE` and `COLUMNS` are skipped for the execution procedures `OPEN`, `FETCH_ROWS`, and `CLOSE`.
2. The `TABLE` and `COLUMNS` arguments have descriptor types for the `DESCRIBE` function.
3. Scalar arguments that are not available during compilation are passed as NULL values (when using bind variables for instance). During execution, the actual values are passed in.

DESCRIBE Function

The `DESCRIBE` function is invoked to determine the type of rows (row shape) produced by the Polymorphic Table Function (PTF). It returns a `DBMS_TF.DESCRIBE_T` table.

The function `DESCRIBE` is invoked during SQL cursor compilation when a SQL query references a PTF. The SQL compiler locates the `DESCRIBE` function defined in the PTF implementation package. All the argument values from the query calling the PTF are passed to the `DESCRIBE` function. Like any PLSQL function, the `DESCRIBE` function can be overloaded and can have arguments default values.

The arguments of the PTF function and `DESCRIBE` function must match, but with the type of any `TABLE` argument replaced with the `DBMS_TF.TABLE_T` descriptor type, and the type of any `COLUMNS` argument replaced with `DBMS_TF.COLUMN_T` descriptor.

The `DESCRIBE` function indicates which columns must be kept by the database and passed unchanged as the PTF output (Pass-Through columns). In addition, the `DESCRIBE` function indicates any input columns that the PTF will use for its computation (Read columns).

Finally, the `DESCRIBE` function returns the list of any new columns that the PTF will create (or `NULL` if no new columns are being produced) using the `DBMS_TF.DESCRIBE_T` descriptor.

OPEN Procedure

The `OPEN` procedure purpose is to initialize and allocate any execution specific state. The `OPEN` procedure is most useful when you implement a Table Semantics PTF. The function typically calls the `GET_XID` function to get a unique ID for managing the execution state.

`OPEN` procedure is generally invoked before calling the `FETCH_ROWS` procedure.

FETCH_ROWS Procedure

The `FETCH_ROWS` procedure produces an output rowset that it sends to the database. The number of invocations of this function and the size of each rowset are data dependent and determined during query execution.

CLOSE Procedure

The `CLOSE` procedure is called at the end of the PTF execution. The procedure releases resources associated with the PTF execution state.

Example 197-1 Noop Polymorphic Table Function Example

This example creates a PTF called noop. This PTF returns the input rows as the output rows without any modification or filtering. Noop is one of the smallest PTF you can write.

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Noop Polymorphic Table Function

To implement the noop PTF, you first create the implementation package noop_package.

```
CREATE PACKAGE noop_package AS
   FUNCTION describe(t IN OUT DBMS_TF.TABLE_T)
     RETURN DBMS_TF.DESCRIBE_T;

   PROCEDURE fetch_rows;
END noop_package;
```

The `DESCRIBE` function does not produce any new columns and hence, returns `NULL`. Executing `FETCH_ROWS` also results in `NULL`.

```
CREATE PACKAGE BODY noop_package AS
   FUNCTION describe(t IN OUT DBMS_TF.TABLE_T)
     RETURN DBMS_TF.DESCRIBE_T AS
     BEGIN
       RETURN NULL;
     END;

   PROCEDURE fetch_rows AS
   BEGIN
     RETURN;
   END;
END noop_package;
```

The noop PTF is defined to execute the noop_package when it is invoked.

```
CREATE FUNCTION noop (t TABLE)
  RETURN TABLE PIPELINED ROW POLYMORPHIC USING noop_package;
```

The PTF can be invoked in queries. For example:

```
SELECT *
FROM   NOOP(emp)
WHERE deptno = 10;

      7782 CLARK      MANAGER         7839 09-JUN-81       2450        10
      7839 KING       PRESIDENT            17-NOV-81       5000        10
      7934 MILLER     CLERK           7782 23-JAN-82       1300        10

WITH e
     AS (SELECT *
         FROM   emp
                NATURAL JOIN dept
         WHERE dname = 'SALES')
SELECT t.*
FROM   NOOP(e) t;

  30       7499 ALLEN      SALESMAN        7698 20-FEB-81       1600       300
SALES          CHICAGO
  30       7521 WARD       SALESMAN        7698 22-FEB-81       1250       500
SALES          CHICAGO
  30       7654 MARTIN     SALESMAN        7698 28-SEP-81       1250      1400
SALES          CHICAGO
  30       7698 BLAKE      MANAGER         7839 01-MAY-81       2850
SALES          CHICAGO
  30       7844 TURNER     SALESMAN        7698 08-SEP-81       1500         0
SALES          CHICAGO
  30       7900 JAMES      CLERK           7698 03-DEC-81        950
SALES          CHICAGO
```

197.4.1.1 DESCRIBE Only Polymorphic Table Function

A Polymorphic Table Function (PTF) can have a DESCRIBE function only.

A PTF which does not have any runtime methods (Open/Fetch_Rows/Close) is used only at cursor compilation time with no runtime row source allocated. The explain plan output of a Describe-Only PTF will not show any rows for the PTF.

197.4.2 PTF Server Side Interface

The `DBMS_TF` package provides the server side interface needed for Polymorphic Table Functions (PTFs) implementation to read and write information in the database.

This topic contains a partial list of types and subprograms used for the PTF server side implementation.

Table 197-2 Summary of Commonly Used Types and Subprograms in PTF Server Side Interface

NAME	DESCRIPTION
COLUMN_METADATA_T	Column metadata record
COLUMN_T	Column descriptor record
TABLE_T	Table descriptor record

Table 197-2 (Cont.) Summary of Commonly Used Types and Subprograms in PTF Server Side Interface

NAME	DESCRIPTION
COLUMNS_T	Collection containing column names
COLUMNS_NEW_T	Collection for new columns
TAB_<typ>_T	Collection for each supported types, where <typ> is described in "Supported Types Collections "
ROW_SET_T	Data for a rowset record
GET_COL Procedure	Fetches data for a specified (input) column
PUT_COL Procedure	Returns data for a specified (new) column
GET_ROW_SET Procedure	Fetches the input rowset of column values
PUT_ROW_SET Procedure	Returns data for ALL (new) columns
SUPPORTED_TYPE Function	Verifies if a type is supported by DBMS_TF subprograms
GET_XID Function	Returns a unique execution ID to index PTF state in a session

See Also:

- DBMS_TF Data Structures for the complete list of types
- Summary of DBMS_TF Subprograms for the complete list of subprograms

197.4.3 Read Columns

Read columns are a set of table columns that the Polymorphic Table Function (PTF) processes when executing the FETCH_ROWS procedure.

The PTF indicates the read columns inside DESCRIBE by annotating them in the input table descriptor, TABLE_T. Only the indicated read columns will be fetched and thus available for processing during FETCH_ROWS.

The PTF invocation in a query will typically use the COLUMNS operator to indicate which columns the query wants the PTF to read, and this information is passed to the DESCRIBE function which then in turn sets the COLUMN_T.FOR_READ boolean flag.

Only scalar SQL data types are allowed for the read columns.

The Echo Polymorphic Table Function Example takes a table and a list of columns and produces new columns with the same values.

197.4.4 Pass-Through Columns

Pass-through columns are passed from the input table of the Polymorphic Table Function (PTF) to the output, without any modifications.

The DESCRIBE function indicates the pass-through columns by setting the COLUMN_T.PASS_THROUGH boolean flag on the input table descriptor, DBMS_TF.TABLE_T.

All columns in the `Row Semantics` PTF are marked as pass-through by default. For `Table Semantics` PTF, the default value for pass-through is set to false. For the `Table Semantics` PTF, the partitioning columns are always pass-through and this cannot be changed by the DESCRIBE function.

Note, the notions of Pass-Through and Read are orthogonal, and indicating a column as one has no implication for the other.

197.4.5 State Management

The database manages the compilation and execution states of the polymorphic table functions (PTF).

The database fulfills the PTF conductor role. As such, it is responsible for the PTF compilation state and execution state.

1. Compilation State : This is the immutable state that is generated by `DESCRIBE` which is needed before execution.
2. Execution State: This is the state used by the execution time procedures (`OPEN`, `FETCH_ROWS`, and `CLOSE`) of a `Table semantics` PTF.

The most common use of compilation state is to keep track of the columns to be read and the new columns that are to be produced. The PTF Server interface provides functions that can be used to achieve this: `GET_ENV`, and `GET_ROW_SET`. The PTF author who defines, documents, and implements the PTF can rely on the database to manage the PTF states. The PTF author should not attempt to use the session state (such as PL/SQL package global variables) to store any compilation state. Problems can arise because in a given session all cursors using the PTF will share that state, and other sessions executing the PTF cursor will not see the original compilation state.

Since the execution state is session and cursor private, a `Table Semantics` PTF can use package globals for storing execution state, but with the provision that the PTF uses the database provided unique execution ID to identify that state. The `GET_XID` function guarantees to provide an execution unique ID for the PTF's execution procedures, where this ID remains constant for all the execution functions of a PTF.

197.4.5.1 CSTORE Compilation State Management

The CSTORE is the PTF compilation state management interface.

The CSTORE enables Polymorphic Table Functions (PTF) to store the compilation state in the SQL cursor.

The CSTORE interface is used to store key-value pairs during cursor compilation through the `DBMS_TF.DESCRIBE_T` record.

The compilation state information is retrieved during execution procedures such as `OPEN`, `FETCH_ROWS` and `CLOSE`.

CSTORE Subprograms

The CSTORE interface consists of the following subprograms.

Name	Description

CSTORE_GET procedure	Fetches item of specified type. If not found, the OUT value remains unchanged.
CSTORE_EXISTS function	If an item with the given key exists in the CSTORE, this function returns TRUE.

CSTORE Supported Types

The DBMS_TF.DESCRIBE_T supports specifying key-value pairs for these scalar types: VARCHAR2, NUMBER, DATE, BOOLEAN.

Table 197-3 DBMS_TF CSTORE Scalar Supported Types

Name	Description
CSTORE_TYPE_VARCHAR2	CSTORE VARCHAR2 type code
CSTORE_TYPE_NUMBER	CSTORE NUMBER type code
CSTORE_TYPE_DATE	CSTORE DATE type code
CSTORE_TYPE_BOOLEAN	CSTORE BOOLEAN type code

Collections For Compilation Storage

These predefined collection types are used for compilation state management.

```
TYPE CSTORE_CHR_T IS TABLE OF VARCHAR2(32767)  INDEX BY VARCHAR2(32767);
TYPE CSTORE_NUM_T IS TABLE OF NUMBER           INDEX BY VARCHAR2(32767);
TYPE CSTORE_BOL_T IS TABLE OF BOOLEAN          INDEX BY VARCHAR2(32767);
TYPE CSTORE_DAT_T IS TABLE OF DATE             INDEX BY VARCHAR2(32767);
```

DBMS_TF Method Names

The method names are also stored in the DBMS_TF.DESCRIBE_T record. These predefined values for the method names can be customized by the PTF author.

See Method Name Overrides for more information about changing the default method names

Table 197-4 DBMS_TF Method Names Constants

Name	Type	Value	Description
CLOSE	DBMS_QUOTED_ID	'CLOSE'	Predefined index value for the method named CLOSE
FETCH_ROWS	DBMS_QUOTED_ID	'FETCH_ROWS'	Predefined index value for the method named FETCH_ROWS
OPEN	DBMS_QUOTED_ID	'OPEN'	Predefined index value for the method named OPEN

197.4.5.2 XSTORE Execution State Management

XSTORE is the PTF execution state management interface.

The `XSTORE` key-value interface simplifies the implementation of `Table Semantics` PTFs by providing automatic state management capabilities when the keys are strings and values are of commonly used scalar types.

The database automatically manages the deletion of all execution states allocated using this interface.

XSTORE Subprograms

The execution state management interface consists of the following subprograms.

Table 197-5 DBMS_TF XSTORE Subprograms

Name	Description
`XSTORE_CLEAR` procedure	Removes all key-value pairs from the XSTORE execution state
`XSTORE_EXISTS` function	Returns TRUE if an item with a given key exists in the XSTORE
`XSTORE_GET` procedure	Gets the associated value for a given key stored in the XSTORE
`XSTORE_REMOVE` procedure	Removes an item associated with the given key and key_type
`XSTORE_SET` procedure	Sets the value for the given key for PTF Execution State Management

XSTORE Predefined Types

The XSTORE supports specifying key-value pairs for these scalar types: `VARCHAR2`, `NUMBER`, `DATE`, and `BOOLEAN`.

Table 197-6 DBMS_TF XSTORE Scalar Supported Types

Name	Description
`XSTORE_TYPE_VARCHAR2`	XSTORE `VARCHAR2` type code
`XSTORE_TYPE_NUMBER`	XSTORE `NUMBER` type code
`XSTORE_TYPE_DATE`	XSTORE `DATE` type code
`XSTORE_TYPE_BOOLEAN`	XSTORE `BOOLEAN` type code

197.4.6 Method Name Overrides

When multiple polymorphic table function (PTF) implementations are in the same package, you can override the default runtime method names (`OPEN`, `FETCH_ROWS`, and `CLOSE`) with your PTF specific names.

To override a method name, the application can specify the new method names using `DBMS_TF METHOD_NAMES` collection (see DESCRIBE_T Record Type).

> **See Also:**
>
> Table 197-4

Example 197-2 DBMS_TF Method Name Overrides

This example shows how to change the default method name of the noop_p PTF fetch_rows method to noop_fetch.

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at DBMS_TF Method Name Overrides

Create the PTF implementation package noop_p.

```
CREATE PACKAGE noop_p AS
   FUNCTION describe(tab IN OUT DBMS_TF.table_t)
               RETURN DBMS_TF.describe_t;
   PROCEDURE noop_fetch;
END noop_p;
```

To provide a method name override, you can specify the new method names using DBMS_TF.Method_Names collection. The FETCH_ROWS method name is changed to 'Noop_Fetch'. The procedure noop_fetch to implement this method is defined in the package.

```
 CREATE OR replace PACKAGE BODY noop_p
AS
   FUNCTION describe(tab IN OUT DBMS_TF.table_t)
               RETURN DBMS_TF.describe_t AS
      methods DBMS_TF.methods_t := DBMS_TF.methods_t(DBMS_TF.fetch_rows =>
'Noop_Fetch');
     BEGIN
          RETURN DBMS_TF.describe_t(method_names => methods);
   END;
   PROCEDURE noop_fetch AS
     BEGIN
         RETURN;
   END;
END noop_p;
```

The noop PTF is defined to execute the noop_p when it is invoked.

```
CREATE FUNCTION noop (t TABLE) RETURN TABLE PIPELINED ROW POLYMORPHIC USING
noop_p;
```

The PTF is invoked in the FROM clause of a query block.

```
SELECT *
FROM noop(scott.emp)
WHERE deptno =10;
```

197.4.7 Using the COLUMNS Pseudo-Operator

The `COLUMNS` pseudo-operator is an addition to the SQL expression language.

Use the `COLUMNS` pseudo-operator to specify the arguments when invoking a Polymorphic Table Function (PTF) in the `FROM` clause. The `COLUMNS` pseudo-operator arguments specify the list of column names, or the list of column names with associated types.

> **See Also:**
>
> *Oracle Database PL/SQL Language Reference* for more information about the COLUMNS pseudo-operator syntax and semantics

197.4.8 Query Transformations

About predicate, projection and partitioning.

The pass-through columns of a `Row Semantics` PTF, and the `PARTITION BY` key columns of a `Table Semantics` PTF can be used for projection and predicate pushdown.

Example 197-3 Query Transformations

This example illustrates the predicate and projection pushdown for a `Row Semantics` PTF.

This query calls the echo PTF created in Echo Polymorphic Table Function Example.

```
SELECT empno, ename, sal, comm, echo_sal
FROM echo(emp, COLUMNS(sal,comm))
WHERE deptno = 30
  AND echo_sal > 1000;


     EMPNO ENAME             SAL       COMM   ECHO_SAL
---------- ---------- ---------- ---------- ----------
      7499 ALLEN            1600        300       1600
      7521 WARD             1250        500       1250
      7654 MARTIN           1250       1400       1250
      7698 BLAKE            2850                  2850
      7844 TURNER           1500          0       1500
```

Conceptually, this query will get rewritten as:

```
WITH t AS (SELECT empno, ename, sal, comm
FROM emp
WHERE deptno=30)
SELECT empno, ename, sal, comm, echo_sal
FROM echo(t, COLUMNS(sal, comm))
WHERE echo_sal > 1000;
```

197.4.9 Parallel Execution

A key benefit of Polymorphic Table Functions (PTFs) is that their execution can be parallelized.

Row and table semantic PTFs execute in parallel differently.

Row Semantics PTF

Under `Row Semantics` PTF, the parallel query executes with the same degree of parallelism (DOP) as it would if the PTF were not present. The DOP is driven by the child row source.

Provided that the DOP on table emp has been set to 5, the following is an example that shows this parallelization:

```
EXPLAIN PLAN FOR
SELECT * FROM echo(emp, COLUMNS(ename, job))
WHERE deptno != 20;
-------------------------------------------------
| Id | Operation                     | Name     |
-------------------------------------------------
| 0 | SELECT STATEMENT                |          |
| 1 | PX COORDINATOR                  |          |
| 2 | PX SEND QC (RANDOM)             | :TQ10000 |
| 3 | POLYMORPHIC TABLE FUNCTION      | ECHO     |
| 6 | PX BLOCK ITERATOR               |          |
|* 7 | TABLE ACCESS FULL              | EMP      |
-------------------------------------------------
Predicate Information (identified by operation id):
-------------------------------------------------
5 - filter("EMP"."DEPTNO"<>20)
```

Table Semantics PTF

`Table Semantics` PTF requires its input table rows to be redistributed using the `PARTITION BY` key. The parallel execution is determined by the `PARTITION BY` clause specified in the query.

197.5 DBMS_TF Execution Flow

Query executions invoking Polymorphic Table Functions (PTF) follow this execution model and data transfers flow.

The PTF execution procedures (`OPEN`, `FETCH_ROWS` and `CLOSE`) are called by the database during query execution.

The PTF execution follows this flow :

1. `OPEN` (if present)
2. `FETCH_ROWS` (can be invoked multiple times)
3. `CLOSE` (if present)

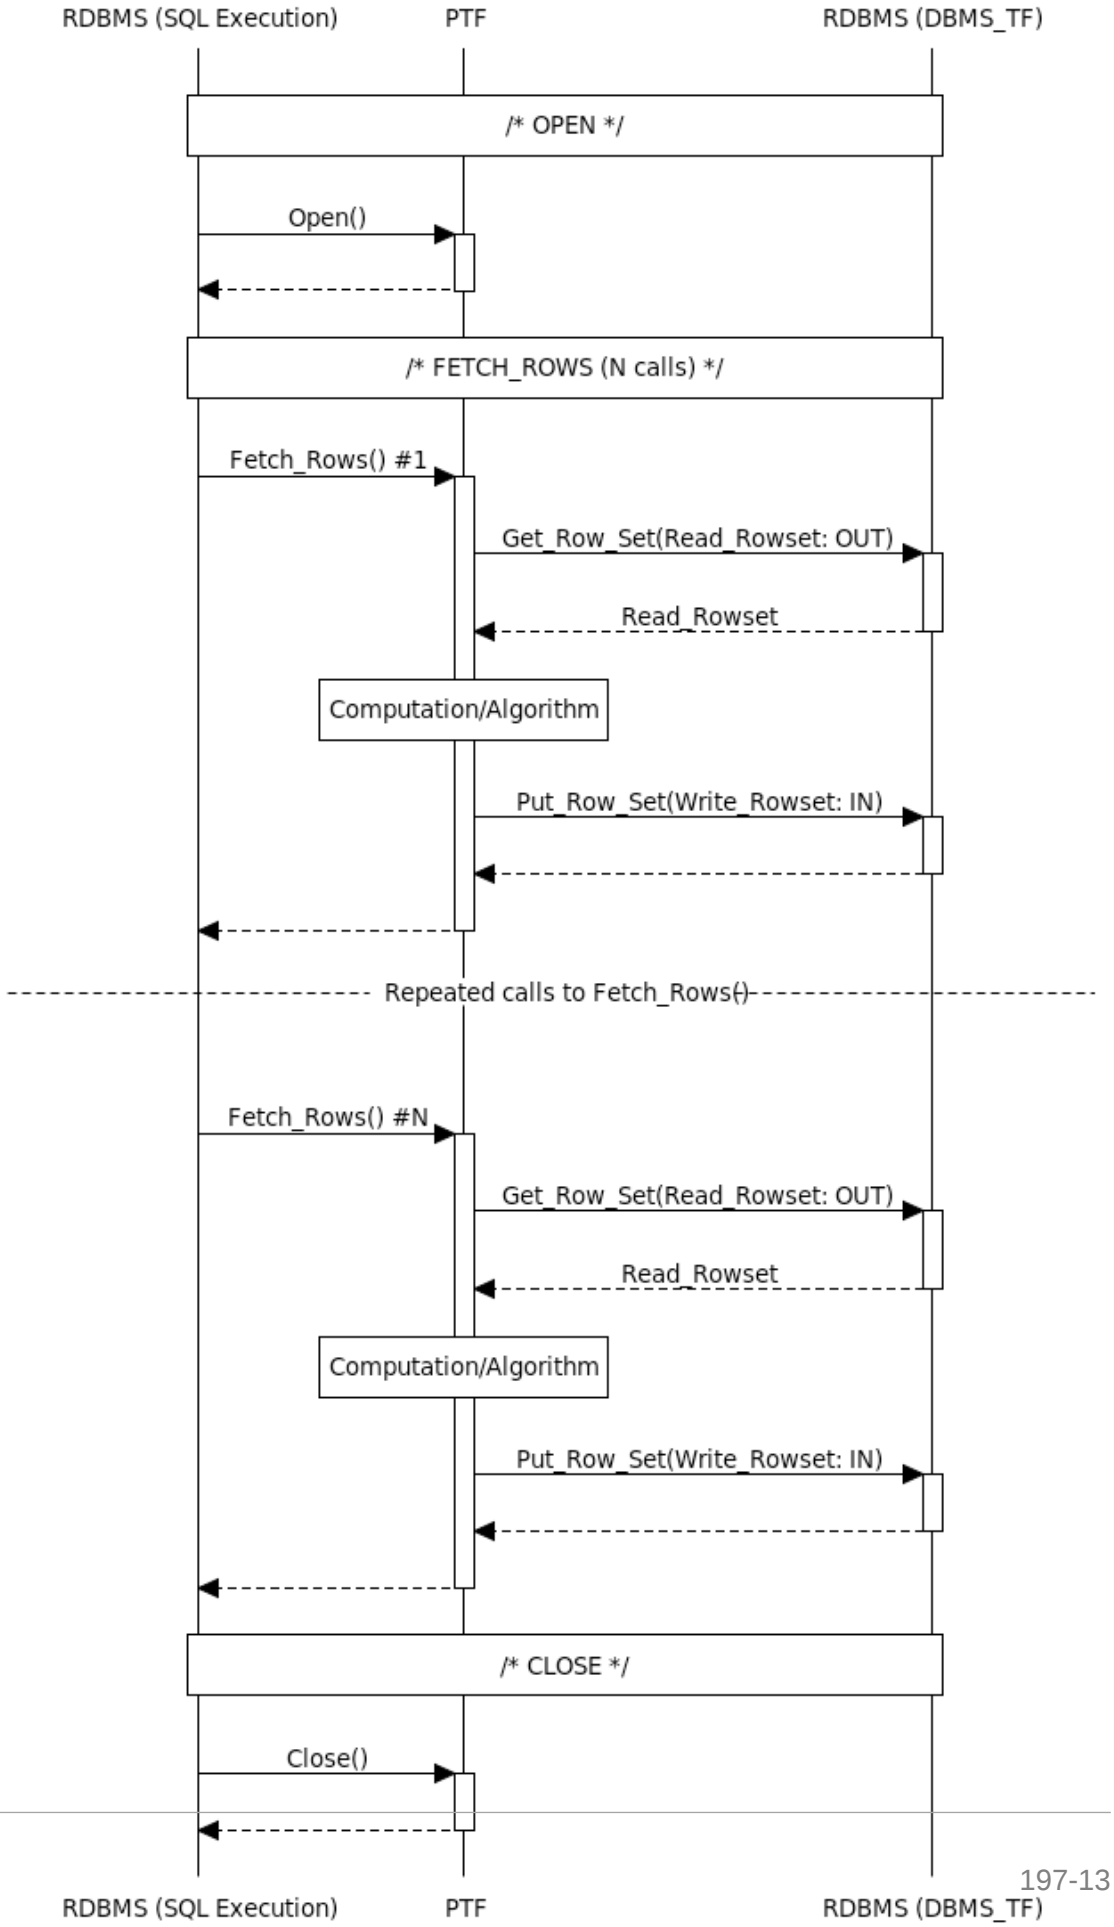
RDBMS (SQL Execution)
PTF
RDBMS (DBMS_TF)
/* OPEN */
Open()
/* FETCH_ROWS (N calls) */
Fetch_Rows() #1
Get_Row_Set(Read_Rowset: OUT)
Read_Rowset
Computation/Algorithm
Put_Row_Set(Write_Rowset: IN)
Repeated calls to Fetch_Rows()
Fetch_Rows() #N
Get_Row_Set(Read_Rowset: OUT)
Read_Rowset
Computation/Algorithm
Put_Row_Set(Write_Rowset: IN)
/* CLOSE */
Close()
RDBMS (SQL Execution)
PTF
RDBMS (DBMS_TF)

The FETCH_ROWS procedure reads the data for a rowset (collection of rows), and produces an output rowset.

Each call to FETCH_ROWS is associated with a rowset which is a data collection of input rows that is expected to be processed by the PTF.

The GET_ROW_SET or GET_COL is used to read the input rowset.

The PUT_ROW_SET or PUT_COL is used to produce an output rowset, that is written back to the database.

PUT_ROW_SET is used to set all the new columns in a single call.

The ROWSET_T record holds data for multiple columns. When the PTF algorithm is more suited toward producing a single output column at a time, you can use PUT_COL to produce a single column. A given column can only be produced once within a call to FETCH_ROWS.

For a Row Semantics PTF, the FETCH_ROWS procedure will return the new rows using the PTF Server interface before returning back to the database.

197.6 DBMS_TF Restrictions

These restrictions apply to Polymorphic Table Functions (PTFs) and using the DBMS_TF package.

Type Restrictions

A Polymorphic Table Function (PTF) can operate on a table with columns of any SQL types. However, read and new columns are restricted to scalar types. The read and new columns are used in the PUT_ROW_SET, PUT_COL, GET_ROW_SET and GET_COL procedures. All SQL types can be used with pass-through columns. The DESCRIBE function can determine the supported types using the DBMS_TF.SUPPORTED_TYPE function.

PTF Invocation and Execution Restrictions

Polymorphic table functions cannot be nested in the FROM clause of a query. Nesting PTF is only allowed using WITH clause.

Nesting table function with polymorphic table function is only allowed using CURSOR expressions. A PTF cannot be specified as an argument of a table function.

You cannot select a rowid from a Polymorphic Table Function (PTF).

The PARTITION BY and the ORDER BY clause can only be specified on an argument of a Table Semantics PTF.

The PTF execution methods OPEN, FETCH_ROWS, and CLOSE must be invoked in the polymorphic table function execution context only.

You cannot invoke the DESCRIBE method directly.

This example shows ten PTF nested invocation.

```
WITH t0
     AS (SELECT /*+ parallel */ *
         FROM   noop(dept)),
     t1
```

```
      AS (SELECT *
          FROM   noop(t0)),
      t2
      AS (SELECT *
          FROM   noop(t1)),
      t3
      AS (SELECT *
          FROM   noop(t2)),
      t4
      AS (SELECT *
          FROM   noop(t3)),
      t5
      AS (SELECT *
          FROM   noop(t4)),
      t6
      AS (SELECT *
          FROM   noop(t5)),
      t7
      AS (SELECT *
          FROM   noop(t6)),
      t8
      AS (SELECT *
          FROM   noop(t7)),
      t9
      AS (SELECT *
          FROM   noop(t8))
SELECT *
FROM   noop(t9)
WHERE  deptno = 10;

        10 ACCOUNTING     NEW YORK
```

197.7 DBMS_TF Examples

These examples use DBMS_TF subprograms.

Summary of DBMS_TF Examples

These examples are incomplete and for demonstration purpose only.

- Example 197-1, "Noop Polymorphic Table Function"
- Echo Polymorphic Table Function Example
- Example 197-2, "DBMS_TF Method Name Overrides"
- Example 197-3, "Query Transformations"
- Example 197-5 , "DBMS_TF.COLUMN_TYPE_NAME Example"
- Example 197-6 , "DBMS_TF.COL_TO_CHAR Example"
- Example 197-7, "DBMS_TF.CSTORE_EXISTS Example"
- Example 197-8, "DBMS_TF.GET_COL Example"
- Example 197-9, "DBMS_TF.GET_ENV Example"
- Example 197-10, "DBMS_TF.GET_ROW_SET Example"
- Example 197-12, "DBMS_TF.GET_XID Example"

- Rand_col Polymorphic Table Function Example, (DBMS_TF.PUT_COL Example)
- Stack Polymorphic Table Function Example
- Split Polymorphic Table Function Example, (DBMS_TF.GET_ROW_SET and PUT_ROW_SET Example)
- Example 197-14, "DBMS_TF.PUT_ROW_SET Example"
- Example 197-16, "Replicate : DBMS_TF.ROW_REPLICATION Example"
- Example 197-17, "DBMS_TF.ROW_TO_CHAR Example"
- Example 197-18, "DBMS_TF.TRACE Example"
- Row_num Polymorphic Table Function Example, (DBMS_TF.XSTORE_GET and XSTORE_SET Example)

In other books :

- *Oracle PL/SQL Language Reference* , "Skip_col Polymorphic Table Function Example"

  Oracle PL/SQL Language Reference, "To_doc Polymorphic Table Function Example"

197.7.1 Echo Polymorphic Table Function Example

The echo PTF takes in a table and a list of columns and produces new columns with same values.

This PTF returns all the columns in the input table tab, and adds to it the columns listed in cols but with the column names prefixed with "ECHO_".

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Echo Polymorphic Table Function

The echo PTF can appear in the FROM clause of the query. The COLUMNS operator is used to specify columns, for example:

```
SELECT *
FROM echo(scott.dept, COLUMNS(dname, loc));

   DEPTNO DNAME          LOC           ECHO_DNAME     ECHO_LOC
--------- -------------- ------------- -------------- -------------
        10 ACCOUNTING     NEW YORK      ACCOUNTING     NEW YORK
        20 RESEARCH       DALLAS        RESEARCH       DALLAS
        30 SALES          CHICAGO       SALES          CHICAGO
        40 OPERATIONS     BOSTON        OPERATIONS     BOSTON
```

A PTF consists of the following :

- PTF implementation package specification : The specification must have the DESCRIBE method. The OPEN, FETCH_ROWS and CLOSE methods are optional.
- PTF implementation package body: The DESCRIBE method may have a new-columns parameter (the additional columns created by this PTF), which is followed by the PTF functions parameters.
- PTF Function: The PTF function has a reference to the implementation package.

The echo_package package specification defines the DESCRIBE and FETCH_ROWS methods.

```
CREATE PACKAGE echo_package
AS
  prefix DBMS_ID := 'ECHO_';
  FUNCTION describe(
    tab  IN OUT DBMS_TF.TABLE_T,
    cols IN DBMS_TF.COLUMNS_T)
  RETURN DBMS_TF.DESCRIBE_T;
  PROCEDURE fetch_rows;
END echo_package;
```

The echo_package package body contains the PTF implementation.

```
CREATE PACKAGE BODY echo_package
AS
  FUNCTION describe(tab  IN OUT DBMS_TF.TABLE_T,
                    cols IN DBMS_TF.COLUMNS_T)
  RETURN DBMS_TF.DESCRIBE_T
  AS
    new_cols DBMS_TF.COLUMNS_NEW_T;
    col_id   PLS_INTEGER := 1;
  BEGIN
      FOR I IN 1 .. tab.COLUMN.COUNT LOOP
          FOR J IN 1 .. cols.COUNT LOOP
              IF ( tab.COLUMN(i).description.name = cols(j) ) THEN
                IF ( NOT
DBMS_TF.SUPPORTED_TYPE(tab.COLUMN(i).description.TYPE) )
                THEN
                  RAISE_APPLICATION_ERROR(-20102, 'Unsupported column type ['
||
                  TAB.COLUMN(i).description.TYPE||']');
                END IF;

                TAB.COLUMN(i).for_read := TRUE;
                NEW_COLS(col_id) := TAB.COLUMN(i).description;
                NEW_COLS(col_id).name := prefix ||
TAB.COLUMN(i).description.name;
                col_id := col_id + 1;

                EXIT;
              END IF;
          END LOOP;
      END LOOP;

      /* Verify all columns were found */
      IF ( col_id - 1 != cols.COUNT ) THEN
        RAISE_APPLICATION_ERROR(-20101, 'Column mismatch ['||col_id - 1||'],
['||cols.COUNT||']');
      END IF;

      RETURN DBMS_TF.DESCRIBE_T(new_columns => new_cols);
  END;
  PROCEDURE FETCH_ROWS
```

```
  AS
    ROWSET DBMS_TF.ROW_SET_T;
  BEGIN
      DBMS_TF.GET_ROW_SET(rowset);
      DBMS_TF.PUT_ROW_SET(rowset);
  END;
END echo_package;
```

The PTF echo references the implementation package echo_package.

```
CREATE FUNCTION echo(tab TABLE,
                     cols COLUMNS)
  RETURN TABLE
  PIPELINED ROW POLYMORPHIC USING echo_package;
```

Example 197-4 Using the Echo PTF in Queries

This example selects all employees in department 20. The resulting rows have three new columns `ECHO_ENAME`, `ECHO_HIREDATE`, and `ECHO_SAL`.

```
SELECT *
FROM   echo(scott.emp, COLUMNS(ename, sal, hiredate))
WHERE  deptno = 20;
```

```
EMPNO ENAME JOB      MGR  HIREDATE  SAL  COMM DEPTNO ECHO_ENAME ECHO_HIRE ECHO_SAL
----- ----- -------- ---- --------- ---- ---- ------ ---------- --------- --------
7369  SMITH CLERK    7902 17-DEC-80  800          20 SMITH      17-DEC-80      800
7566  JONES MANAGER  7839 02-APR-81 2975          20 JONES      02-APR-81     2975
7788  SCOTT ANALYST  7566 19-APR-87 3000          20 SCOTT      19-APR-87     3000
7876  ADAMS CLERK    7788 23-MAY-87 1100          20 ADAMS      23-MAY-87     1100
7902  FORD  ANALYST  7566 03-DEC-81 3000          20 FORD       03-DEC-81     3000
```

Using subquery `W`, display `ENAME`, `ECHO_LOC` and `DNAME` columns for all employees in department 30 with a salary greater than 1000.

```
WITH w
     AS (SELECT e.*,
                dname,
                loc
         FROM   scott.emp e,
                scott.dept d
         WHERE  e.deptno = d.deptno)
SELECT ename,
       echo_loc,
       dname
FROM   echo(w, COLUMNS(sal, dname, loc, hiredate))
WHERE  deptno = 30
       AND echo_sal > 1000;
```

```
ENAME      ECHO_LOC       DNAME
---------- -------------- --------
ALLEN      CHICAGO        SALES
WARD       CHICAGO        SALES
MARTIN     CHICAGO        SALES
```

```
BLAKE      CHICAGO         SALES
TURNER     CHICAGO         SALES
```

Using subquery `W`, display `ENAME` and `DNAME` columns for all employees with a salary greater than 1000.

```
WITH w
     AS (SELECT e.*,
                dname,
                loc
         FROM   scott.emp e,
                scott.dept d
         WHERE  e.deptno = d.deptno)
SELECT echo_ename,
       dname
FROM   echo(w, COLUMNS(loc, deptno, dname, ename)) e
WHERE  ename IN (SELECT echo_ename
                 FROM   echo(scott.emp, COLUMNS(sal, deptno, ename, hiredate))
                 WHERE  deptno = e.echo_deptno
                        AND sal > 1000);


ECHO_ENAME DNAME
---------- ----------
ALLEN      SALES
MILLER     ACCOUNTING
CLARK      ACCOUNTING
WARD       SALES
ADAMS      RESEARCH
TURNER     SALES
SCOTT      RESEARCH
BLAKE      SALES
JONES      RESEARCH
KING       ACCOUNTING
FORD       RESEARCH
MARTIN     SALES
```

197.8 DBMS_TF Data Structures

The DBMS_TF package defines these RECORD types, TABLE types and subtype.

RECORD Types

- COLUMN_DATA_T Record Type
- COLUMN_METADATA_T Record Type
- COLUMN_T Record Type
- DESCRIBE_T Record Type
- ENV_T Record Type
- PARALLEL_ENV_T Record Type
- TABLE_T Record Type

TABLE Types

- Supported Types Collections (TAB_<typ>_T)

- COLUMNS_NEW_T Table Type
- COLUMNS_T Table Type
- COLUMNS_WITH_TYPE_T Table Type
- TABLE_COLUMNS_T Table Type
- ROW_SET_T Table Type

Types

- XID_T Subtype

CSTORE and XSTORE Data Structures

The compilation and execution state management interfaces use data structures internally.

See Collections For Compilation Storage for more information.

197.8.1 Supported Types Collections

Each supported type has a corresponding predefined collection defined.

Syntax

```
TYPE TAB_BOOLEAN_T       IS TABLE OF BOOLEAN       INDEX BY PLS_INTEGER;

TYPE TAB_BINARY_FLOAT_T  IS TABLE OF BINARY_FLOAT  INDEX BY PLS_INTEGER;

TYPE TAB_BINARY_DOUBLE_T IS TABLE OF BINARY_DOUBLE INDEX BY PLS_INTEGER;

TYPE TAB_BLOB_T          IS TABLE OF BLOB          INDEX BY PLS_INTEGER;

TYPE TAB_CHAR_T          IS TABLE OF CHAR(32767)   INDEX BY PLS_INTEGER;

TYPE TAB_CLOB_T          IS TABLE OF CLOB          INDEX BY PLS_INTEGER;

TYPE TAB_DATE_T          IS TABLE OF DATE          INDEX BY PLS_INTEGER;

TYPE TAB_INTERVAL_YM_T   IS TABLE OF YMINTERVAL_UNCONSTRAINED   INDEX BY PLS_INTEGER;

TYPE TAB_INTERVAL_DS_T   IS TABLE OF DSINTERVAL_UNCONSTRAINED   INDEX BY PLS_INTEGER;

TYPE TAB_NATURALN_T      IS TABLE OF NATURALN      INDEX BY PLS_INTEGER;

TYPE TAB_NUMBER_T        IS TABLE OF NUMBER        INDEX BY PLS_INTEGER;

TYPE TAB_RAW_T           IS TABLE OF RAW(32767)    INDEX BY PLS_INTEGER;

TYPE TAB_ROWID_T         IS TABLE OF ROWID         INDEX BY PLS_INTEGER;

TYPE TAB_VARCHAR2_T      IS TABLE OF VARCHAR2(32767) INDEX BY PLS_INTEGER;

TYPE TAB_TIMESTAMP_T     IS TABLE OF TIMESTAMP_UNCONSTRAINED INDEX BY PLS_INTEGER;

TYPE TAB_TIMESTAMP_TZ_T  IS TABLE OF TIMESTAMP_TZ_UNCONSTRAINED INDEX BY PLS_INTEGER;

TYPE TAB_TIMESTAMP_LTZ_T IS TABLE OF TIMESTAMP_LTZ_UNCONSTRAINED INDEX BY PLS_INTEGER;
```

> **See Also:**
>
> Table 197-1 for more information about the DBMS_TF supported types

197.8.2 COLUMN_DATA_T Record Type

Data for a single column (variant record).

Exactly one variant field is active in the record. The description includes information about the column type that is active.

See Table 197-1 for the list of supported types.

Syntax

```
TYPE COLUMN_DATA_T IS RECORD
 ( description            COLUMN_METADATA_T,
   tab_varchar2           TAB_VARCHAR2_T,
   tab_number             TAB_NUMBER_T,
   tab_date               TAB_DATE_T,
   tab_binary_float       TAB_BINARY_FLOAT_T,
   tab_binary_double      TAB_BINARY_DOUBLE_T,
   tab_raw                TAB_RAW_T,
   tab_char               TAB_CHAR_T,
   tab_clob               TAB_CLOB_T,
   tab_blob               TAB_BLOB_T,
   tab_timestamp          TAB_TIMESTAMP_T,
   tab_timestamp_tz       TAB_TIMESTAMP_TZ_T,
   tab_interval_ym        TAB_INTERVAL_YM_T,
   tab_interval_ds        TAB_INTERVAL_DS_T,
   tab_timestamp_ltz      TAB_TIMESTAMP_LTZ_T,
   tab_rowid              TAB_ROWID_T,
   tab_boolean            TAB_BOOLEAN_T);
```

Fields

Table 197-7 COLUMN_DATA_T Fields

Field	Description
description	The tag defines the metadata for the column indicating which variant field is active.
tab_varchar2	Variant field
tab_number	Variant field
tab_date	Variant field
tab_binary_float	Variant field
tab_binary_double	Variant field
tab_raw	Variant field
tab_char	Variant field
tab_clob	Variant field
tab_blob	Variant field
tab_timestamp	Variant field

Table 197-7 (Cont.) COLUMN_DATA_T Fields

Field	Description
tab_timestamp_tz	Variant field
tab_interval_ym	Variant field
tab_interval_ds	Variant field
tab_timestamp_ltz	Variant field
tab_rowid	Variant field
tab_boolean_t	Variant field

197.8.3 COLUMN_METADATA_T Record Type

This type contains metadata about an existing table column or a new column produced by PTF.

Syntax

```
TYPE COLUMN_METADATA_T IS RECORD
  ( type               PLS_INTEGER,
    max_len            PLS_integer DEFAULT -1,
    name               VARCHAR2(32767),
    name_len           PLS_INTEGER,
    precision          PLS_INTEGER,
    scale              PLS_INTEGER,
    charsetid          PLS_INTEGER,
    charsetform        PLS_INTEGER,
    collation          PLS_INTEGER );
```

Fields

Table 197-8 COLUMN_METADATA_T Fields

Field	Description
type	Internal Oracle typecode for the column's type
max_len	Maximum length of a column. If it is less than the maximum allowed length then that value will be used, if it is NULL or zero, zero will be used. If it is less than zero, then maximum allowed length will be used. If types (like date,float), does not care about length, then this value will be ignored.
name	Name of the column
name_len	Length of the name
precision	The precision, or the maximum number of significant decimal digits (for numeric data types)
scale	Scale, or the number of digits from the decimal point to the least significant digit (for numeric data types)
charsetid	Character set id (internal Oracle code, applies to string types)
charsetform	Character set form (internal Oracle code, applies to string types)
collation	Collation id (internal Oracle code, applies to string types)

197.8.4 COLUMN_T Record Type

The column descriptor record for the type COLUMN_METADATA_T that contains PTF specific attributes.

Syntax

```
TYPE column_t IS RECORD (
    description            COLUMN_METADATA_T,
    pass_through           BOOLEAN,
    for_read               BOOLEAN);
```

Fields

Table 197-9 COLUMN_T Fields

Field	Description
description	Column metadata
pass_through	Is this a pass through column
for_read	Is this column read by the PTF

197.8.5 DESCRIBE_T Record Type

The return type from the DESCRIBE method of PTF.

Syntax

```
TYPE DESCRIBE_T          IS RECORD
  ( NEW_COLUMNS        COLUMNS_NEW_T DEFAULT COLUMNS_NEW_T(),
    CSTORE_CHR         CSTORE_CHR_T  DEFAULT CSTORE_CHR_T(),
    CSTORE_NUM         CSTORE_NUM_T  DEFAULT CSTORE_NUM_T(),
    CSTORE_BOL         CSTORE_BOL_T  DEFAULT CSTORE_BOL_T(),
    CSTORE_DAT         CSTORE_DAT_T  DEFAULT CSTORE_DAT_T(),
    METHOD_NAMES       METHODS_T     DEFAULT METHODS_T());
```

Fields

Table 197-10 DESCRIBE_T Fields

Field	Description
NEW_COLUMNS	New columns description that will be produced by the PTF
CSTORE_CHR	CStore array key type : VARCHAR2 (optional)
CSTORE_NUM	CStore array key type : NUMBER (optional)
CSTORE_BOL	CStore array key type : BOOLEAN (optional)
CSTORE_DAT	CStore array key type : DATE (optional)
METHOD_NAMES	Method names, if user wants to override OPEN, FETCH_ROWS, CLOSE methods

197.8.6 ENV_T Record Type

This record contains metadata about the polymorphic table function execution state.

Syntax

```
TYPE ENV_T IS RECORD (
  get_columns      TABLE_METADATA_T,
  put_columns      TABLE_METADATA_T,
  ref_put_col      REFERENCED_COLS_T,
  parallel_env     PARALLEL_ENV_T,
  query_optim      BOOLEAN,
  row_count        PLS_INTEGER,
  row_replication  BOOLEAN,
  row_insertion    BOOLEAN);
```

Fields

Table 197-11 ENV_T Fields

Field	Description
get_columns	Metadata about the columns read by PTF GET_COL procedure
put_columns	Metadata about columns sent back to database by PUT_COL procedure
ref_put_col	TRUE if the put column was referenced in the query
parallel_env	Parallel execution information (when a query runs in parallel)
query_optim	Is this execution for query optimization? TRUE, if the query was running on behalf of optimizer
row_count	Number of rows in current row set
row_replication	Is Row Replication Enabled?
row_insertion	Is Row Insertion Enabled?

197.8.7 PARALLEL_ENV_T Record Type

The record contains metadata specific to polymorphic table functions parallel execution.

Syntax

```
TYPE PARALLEL_ENV_T         IS RECORD
  ( instance_id       PLS_INTEGER,
    session_id        PLS_INTEGER,
    server_grp        PLS_INTEGER,
    server_set_no     PLS_INTEGER,
    no_slocal_servers PLS_INTEGER,
    global_server_no  PLS_INTEGER,
    no_local_servers  PLS_INTEGER,
    local_server_no   PLS_INTEGER );
```

Fields

Table 197-12 PARALLEL_ENV_T Fields

Field	Description
instance_id	QC instance ID
session_id	QC session ID
server_grp	Server group

Table 197-12 (Cont.) PARALLEL_ENV_T Fields

Field	Description
server_set_no	Server set number
no_slocal_servers	Number of sibling servers (including self)
global_server_no	Global server number (base 0)
no_local_servers	Number of sibling servers running on instance
local_server_no	Local server number (base 0)

197.8.8 TABLE_T Record Type

The DESCRIBE function input table descriptor argument is of TABLE_T record type.

Syntax

```
TYPE TABLE_T IS RECORD(
  column                TABLE_COLUMNS_T,
  schema_name           DBMS_id,
  package_name          DBMS_id,
  ptf_name              DBMS_id);
```

Fields

Table 197-13 TABLE_T Fields

Field	Description
column	Column information
schema_name	The PTF schema name
package_name	The PTF implementation package name
ptf_name	The PTF name invoked

197.8.9 COLUMNS_NEW_T Table Type

Collection for new columns

Syntax

```
TYPE COLUMNS_NEW_T IS TABLE OF COLUMN_METADATA_T INDEX BY PLS_INTEGER;
```

197.8.10 COLUMNS_T Table Type

Collection containing column names

Syntax

```
TYPE COLUMNS_T IS TABLE OF DBMS_QUOTED_ID;
```

197.8.11 COLUMNS_WITH_TYPE_T Table Type

Collection containing columns metadata

Syntax

```
TYPE COLUMNS_WITH_TYPE_T IS TABLE OF COLUMN_METADATA_T;
```

197.8.12 TABLE_COLUMNS_T Table Type

A collection of columns(COLUMN_T)

Syntax

```
TYPE TABLE_COLUMNS_T IS TABLE OF COLUMN_T;
```

197.8.13 ROW_SET_T Table Type

Data for a rowset

Syntax

```
TYPE ROW_SET_T IS TABLE OF COLUMN_DATA_T INDEX BY PLS_INTEGER;
```

197.8.14 XID_T Subtype

The XID_T subtype is defined to store the execution unique ID returned by function GET_XID.

Syntax

```
SUBTYPE XID_T IS VARCHAR2(1024);
```

197.9 Summary of DBMS_TF Subprograms

This summary briefly describes the `DBMS_TF` package subprograms.

Table 197-14 DBMS_TF Subprograms

Subprogram	Description
`COLUMN_TYPE_NAME` Function	Returns the type name of the specified column type
`COL_TO_CHAR` Function	Returns the string representation of the specified column
`CSTORE_EXISTS` Function	Returns TRUE if an item with a given key exists in the PTF Compilation State management Store
`CSTORE_GET` Procedure	Gets item(s) of specified type from the PTF Compilation State management Store
`GET_COL` Procedure	Gets read column values
`GET_ENV` Function	Returns information about the PTF runtime environment
`GET_ROW_SET` Procedure	Gets read set of column values in the collection
`GET_XID` Function	Returns a unique execution id that can be used by the PTF to index any cursor execution specific runtime state
`PUT_COL` Procedure	Puts column values in the database

Table 197-14 (Cont.) DBMS_TF Subprograms

Subprogram	Description
`PUT_ROW_SET` Procedure	Puts the collection read set of column values in the database
`ROW_REPLICATION` Procedure	Sets the row replication factor
`ROW_TO_CHAR` Function	Returns the string representation of a row in a rowset
`SUPPORTED_TYPE` Function	Returns TRUE if a specified type is supported by PTF infrastructure
`TRACE` Procedure	Prints data structures to help development and problem diagnosis
`XSTORE_CLEAR` Procedure	Removes all key-value pairs from XStore
`XSTORE_EXISTS` Procedure	Returns TRUE if the key has an associated value
`XSTORE_GET` Procedure	Gets a key-value store for PTF Execution State Management
`XSTORE_REMOVE` Procedure	Removes any value associated with the given key
`XSTORE_SET` Procedure	Sets the value for the given key store for PTF Execution State Management

197.9.1 COLUMN_TYPE_NAME Function

Returns the type name for the specified column type.

Syntax

```
FUNCTION COLUMN_TYPE_NAME(
   col COLUMN_METADATA_T)
   RETURN VARCHAR2;
```

Parameters

Table 197-15 DBMS_TF.COLUMN_TYPE_NAME Function Parameters

Parameter	Description
`col`	The column metadata. See COLUMN_METADATA_T Record Type

Return Values

Returns the column type converted as text.

Example 197-5 DBMS_TF.COLUMN_TYPE_NAME Example

This example shows an application type check that invokes `COLUMN_TYPE_NAME` to compare the column type and raise an application error if the column type is not `VARCHAR2`.

```
FUNCTION describe(
   tab   IN OUT DBMS_TF.table_t,
   cols  IN     DBMS_TF.columns_t)
   RETURN DBMS_TF.describe_t
AS
```

```
  new_cols DBMS_TF.columns_new_t;
  col_id    PLS_INTEGER := 1;
BEGIN
    FOR i IN 1 .. tab.count LOOP
       FOR j IN 1 .. cols.count LOOP
          IF (tab(i).description.name = cols(j)) THEN
               IF (DBMS_TF.column_type_name(tab(i).description.type) != 'VARCHAR2') THEN
                  raise_application_error(-20102,
                   'Unsupported column type ['||tab(i).description.type||']');
               END IF;
             tab(i).for_read       := true;
             new_cols(col_id)      := tab(i).description;
             new_cols(col_id).name := 'ECHO_'|| tab(i).description.name;
             col_id                := col_id + 1;
              EXIT;
         END IF;
       END LOOP;
    END LOOP;

    --  Verify all columns were found
    IF (col_id - 1 != cols.count) THEN
       raise_application_error(-20101,
         'Column mismatch ['||col_id-1||'], ['||cols.count||']');
    END IF;

    RETURN DBMS_TF.describe_t(new_columns => new_cols);
END;
```

197.9.2 COL_TO_CHAR Function

Returns the string representation of the specified column.

Syntax

```
FUNCTION COL_TO_CHAR(
   col   COLUMN_DATA_T,
   rid   PLS_INTEGER,
   quote VARCHAR2 DEFAULT '"')
  RETURN VARCHAR2;
```

Parameters

Table 197-16 DBMS_TF.COL_TO_CHAR Function Parameters

Parameter	Description
col	The column whose value is to be converted
rid	Row number
quote	Quotation mark to use for non-numeric values

Return Values

The string representation of a column data value.

Example 197-6 DBMS_TF.COL_TO_CHAR Example

```
PROCEDURE Fetch_Rows AS
    rowset DBMS_TF.rROW_SET_T;
    str    VARCHAR2(32000);
```

```
BEGIN
     DBMS_TF.GET_ROW_SET(rowset);
     str := DBMS_TF.COL_TO_CHAR(rowset(1), 1)
END;
```

197.9.3 CSTORE_EXISTS Function

Returns TRUE if an item with a given key exists in the Store PTF Compilation State.

Syntax

```
FUNCTION CSTORE_EXISTS
  (key      IN VARCHAR2,
   key_type IN PLS_INTEGER default NULL)
   return BOOLEAN;
```

Parameters

Table 197-17 CSTORE_EXISTS Function Parameters

Parameter	Description
`key`	A unique character key
`key_type`	The type of key (optional) Default : NULL

Return Values

Returns TRUE if the key has an associated value. When the key_type is NULL (default), it returns TRUE if the key has an associated value of any of the supported type.

When a key_type parameter value is passed, it returns TRUE if the key and specified type of key has an associated value. Otherwise, it returns FALSE.

Example 197-7 DBMS_TF.CSTORE_EXISTS Example

This code excerpt checks if an item with the key exists before reading it from the compilation store.

```
IF (DBMS_TF.CSTORE_EXISTS('min'||j)) THEN
    DBMS_TF.CSTORE_GET('min'||j, min_col);
END IF;
```

197.9.4 CSTORE_GET Procedure

You can use the CSTORE_GET procedure to get the associated value for a given key stored for PTF Compilation State.

CSTORE is the PTF compilation state management interface. The CSTORE interface is used to set and store key-value pairs during cursor compilation through the `DBMS_TF.DESCRIBE` function.

You can get the PTF compilation state during runtime procedures such as `OPEN`, `FETCH_ROWS` and `CLOSE`.

This procedure is overloaded. The `DESCRIBE_T` supports specifying key-value pairs for these scalar types: `VARCHAR2`, `NUMBER`, `DATE`, `BOOLEAN`.

See Table 197-3 for more information.

Syntax

Get the value associated with the key in the value out variable. The value type returned is one of the supported scalar types.

```
PROCEDURE CSTORE_GET(
   key   IN     VARCHAR2,
   value IN OUT VARCHAR2);


PROCEDURE CSTORE_GET(
   key   IN     VARCHAR2,
   value IN OUT NUMBER);


PROCEDURE CSTORE_GET(
   key   IN     VARCHAR2,
   value IN OUT DATE);


PROCEDURE CSTORE_GET(
   key   IN     VARCHAR2,
   value IN OUT BOOLEAN);
```

When no specific key is passed as an input parameter, the entire collection of key values for that type that exist in the CSTORE is returned.

```
PROCEDURE CSTORE_GET(key_value OUT CSTORE_CHR_T);

PROCEDURE CSTORE_GET(key_value OUT CSTORE_NUM_T);

PROCEDURE CSTORE_GET(key_value OUT CSTORE_BOL_T);

PROCEDURE CSTORE_GET(key_value OUT CSTORE_DAT_T);
```

Parameters

Table 197-18 DBMS_TF.CSTORE_GET Procedure Parameters

Parameter	Description
key	A unique character key
value	Value corresponding to the key for supported types
key_value	Key value

197.9.5 GET_COL Procedure

Get Read Column Values

Syntax

```
PROCEDURE GET_COL(
   columnId NUMBER,
   collection IN OUT NOCOPY <datatype>);
```

Where <*datatype*> can be any one of the supported types.

See Table 197-1 for the list of supported types.

Parameters

Table 197-19 GET_COL Procedure Parameters

Parameter	Description
`columnid`	The id for the column
`collection`	The data for the column

Usage Notes

This procedure is used to get the read column values in the collection of scalar type.

The column numbers are in the get column order as created in `DESCRIBE` method of PTF.

For the same ColumnId, `GET_COL` and `PUT_COL` may correspond to different column.

Example 197-8 DBMS_TF.GET_COL Example

This example is an excerpt of a fetch_rows procedure defined in the PTF implementation package.

```
PROCEDURE fetch_rows
IS
  col1 DBMS_TF.TAB_CLOB_T;
  col2 DBMS_TF.TAB_CLOB_T;
  out1 DBMS_TF.TAB_CLOB_T;
  out2 DBMS_TF.TAB_CLOB_T;
BEGIN
    DBMS_TF.GET_COL(1, col1);
    DBMS_TF.GET_COL(2, col2);

    FOR I IN 1 .. col1.COUNT LOOP
        out1(i) := 'ECHO-' || col1(i);
    END LOOP;

    FOR I IN 1 .. col2.COUNT LOOP
        out2(i) := 'ECHO-' || col2(i);
    END LOOP;

    DBMS_TF.PUT_COL(1, out1);
    DBMS_TF.PUT_COL(2, out2);
END;
```

Note, invoking the DBMS_TF APIs directly is not allowed. An error is raised if an attempt is made to execute these procedures out of context.

```
exec fetch_rows

ERROR at line 1:
ORA-62562: The API Get_Col can be called only during execution time of a polymorphic
table function.
```

197.9.6 GET_ENV Function

Returns information about the PTF runtime environment

Syntax

```
FUNCTION GET_ENV
         RETURN ENV_T;
```

Return Values

Returns information about the PTF runtime environment.

Example 197-9 DBMS_TF.GET_ENV Example

This line shows how you could initialize a local variable env of type ENV_T with the PTF execution information in a `FETCH_ROWS` implementation procedure.

```
env          DBMS_TF.ENV_T  := DBMS_TF.GET_ENV();
```

197.9.7 GET_ROW_SET Procedure

Get Read Column Values

The `FETCH_ROW` procedure can call the `GET_ROW_SET` procedure to read the input rowset set of column values in the collection of supported scalar type. This procedure is overloaded.

Syntax

```
PROCEDURE GET_ROW_SET(
  rowset    OUT NOCOPY ROW_SET_T);

PROCEDURE GET_ROW_SET(
  rowset    OUT NOCOPY ROW_SET_T,
  row_count OUT        PLS_INTEGER);

PROCEDURE GET_ROW_SET(
  rowset    OUT NOCOPY ROW_SET_T,
  row_count OUT        PLS_INTEGER,
  col_count OUT        PLS_INTEGER);
```

Parameters

Table 197-20 GET_ROW_SET Procedure Parameters

Parameter	Description
`rowset`	The collection of data and metadata
`row_count`	The number of rows in the columns
`col_count`	The number of columns

Example 197-10 DBMS_TF.GET_ROW_SET Example

This example is an excerpt from a PTF implementation package for demonstration purpose.

```
PROCEDURE fetch_rows(new_name IN VARCHAR2 DEFAULT 'PTF_CONCATENATE')
AS
  rowset      DBMS_TF.ROW_SET_T;
  accumulator DBMS_TF.TAB_VARCHAR2_T;
  row_count   PLS_INTEGER;

  FUNCTION get_value(col  PLS_INTEGER,
                     ROW  PLS_INTEGER)
  RETURN VARCHAR2
  AS
    col_type  PLS_INTEGER := rowset(col).description.TYPE;
  BEGIN
      CASE col_type
        WHEN DBMS_TF.TYPE_VARCHAR2 THEN
          RETURN NVL(rowset(col).TAB_VARCHAR2 (ROW), 'empty');
        ELSE
          RAISE_APPLICATION_ERROR(-20201, 'Non-Varchar Type='||col_type);
      END CASE;
  END;
BEGIN
    DBMS_TF.GET_ROW_SET(rowset, row_count);

    IF ( rowset.count = 0 ) THEN
      RETURN;
    END IF;

    FOR row_num IN 1 .. row_count LOOP
        accumulator(row_num) := 'empty';
    END LOOP;

    FOR col_num IN 1 .. rowset.count LOOP
        FOR row_num IN 1 .. row_count LOOP
            accumulator(row_num) := accumulator(row_num) ||
get_value(col_num, row_num);
        END LOOP;
    END LOOP;
    -- Pushout the accumulator
    DBMS_TF.PUT_COL(1, accumulator);
END;
```

197.9.7.1 Stack Polymorphic Table Function Example

The stack PTF example unpivots the non-null values of the specified numeric columns by converting each column value into a new row.

Example 197-11 Stack Polymorphic Table Function Example

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Stack Polymorphic Table Function

Create the PTF implementation package stack_p.

The parameters are :

- tab - Input table
- col - The names of numeric (input) table columns to stack

```
CREATE PACKAGE stack_p AS

  FUNCTION  describe(tab  IN OUT dbms_tf.table_t,
                     col         dbms_tf.columns_t)
            RETURN dbms_tf.describe_t;

  PROCEDURE fetch_rows;

END stack_p;
```

Create the PTF implementation package body stack_p.

This PTF produces two new columns, `COLUMN_NAME` and `COLUMN_VALUE`, where the former contains the name of the unpivoted column and the latter contains the numeric value of that column. Additionally, the unpivoted columns are removed from the PTF's output.

```
CREATE PACKAGE BODY stack_p AS

  FUNCTION  describe(tab  IN OUT dbms_tf.table_t,
                     col         dbms_tf.columns_t)
            RETURN dbms_tf.describe_t  AS
  BEGIN
    FOR i IN 1 .. tab.column.count LOOP
      FOR j IN 1 .. col.count LOOP
        IF (tab.column(i).description.name = col(j) AND
            tab.column(i).description.TYPE = dbms_tf.type_number) THEN
          tab.column(i).pass_through := false;
          tab.column(i).for_read     := true;
        END IF;
      END LOOP;
    END LOOP;

    RETURN dbms_tf.describe_t(
             new_columns => dbms_tf.columns_new_t(
               1 => dbms_tf.column_metadata_t(name => 'COLUMN_NAME',
                                              TYPE => dbms_tf.type_varchar2),
               2 => dbms_tf.column_metadata_t(name => 'COLUMN_VALUE',
                                              TYPE => dbms_tf.type_number)),
             row_replication => true);
```

```
  END;

  PROCEDURE fetch_rows  AS
    env    dbms_tf.env_t := dbms_tf.get_env();
    rowset dbms_tf.row_set_t;
    colcnt PLS_INTEGER;
    rowcnt PLS_INTEGER;
    repfac dbms_tf.tab_naturaln_t;
    namcol dbms_tf.tab_varchar2_t;
    valcol dbms_tf.tab_number_t;
  BEGIN
    dbms_tf.get_row_set(rowset, rowcnt, colcnt);

    FOR i IN 1 .. rowcnt LOOP
      repfac(i) := 0;
    END LOOP;

    FOR r IN 1 .. rowcnt LOOP
      FOR c IN 1 .. colcnt LOOP
        IF rowset(c).tab_number(r) IS NOT NULL THEN
          repfac(r)                    := repfac(r) + 1;
          namcol(nvl(namcol.last+1,1)) :=
            INITCAP(regexp_replace(env.get_columns(c).name, '^"|"$'));
          valcol(NVL(valcol.last+1,1)) := rowset(c).tab_number(r);
        END IF;
      END LOOP;
    END LOOP;

    dbms_tf.row_replication(replication_factor => repfac);
    dbms_tf.put_col(1, namcol);
    dbms_tf.put_col(2, valcol);

  END;

END stack_p;
```

Create the standalone PTF named stack. Specify exactly one formal argument of type `TABLE`, specify the return type of the PTF as `TABLE`, specify a Row Semantics PTF type, and indicate the PTF implementation package to use is stack_p.

```
CREATE FUNCTION stack(tab TABLE,
                      col columns)
                   RETURN TABLE
PIPELINED ROW POLYMORPHIC USING stack_p;
```

For all employees in departments 10 and 30, report values of columns `MGR`, `SAL`, and `COMM` ordered by department number and employee name.

```
SELECT   deptno, ename, column_name, column_value
FROM     stack(scott.emp, COLUMNS(mgr, sal, comm))
WHERE    deptno IN (10, 30)
ORDER BY deptno, ename;
```

```
    DEPTNO ENAME      COLUMN_NAME     COLUMN_VALUE
---------- ---------- --------------- ------------
        10 CLARK      Mgr                     7839
        10 CLARK      Sal                     2450
        10 KING       Sal                     5000
        10 MILLER     Sal                     1300
        10 MILLER     Mgr                     7782
        30 ALLEN      Comm                     300
        30 ALLEN      Mgr                     7698
        30 ALLEN      Sal                     1600
        30 BLAKE      Mgr                     7839
        30 BLAKE      Sal                     2850
        30 JAMES      Sal                      950
        30 JAMES      Mgr                     7698
        30 MARTIN     Comm                    1400
        30 MARTIN     Mgr                     7698
        30 MARTIN     Sal                     1250
        30 TURNER     Comm                       0
        30 TURNER     Sal                     1500
        30 TURNER     Mgr                     7698
        30 WARD       Comm                     500
        30 WARD       Mgr                     7698
        30 WARD       Sal                     1250
```

197.9.8 GET_XID Function

Returns a unique execution id that can be used by the PTF to index any cursor-execution specific runtime state.

Syntax

```
FUNCTION GET_XID
      RETURN XID_T;
```

Return Values

A unique execution id that can be used by the PTF to index any cursor-execution specific runtime state.

Example 197-12 DBMS_TF.GET_XID Example

This is an excerpt of code showing an invocation of GET_XID to initialize a local variable indexed using the execution id to a zero value.

```
PROCEDURE open IS
BEGIN
  xst(DBMS_TF.GET_XID()) := 0;
END;
```

197.9.9 PUT_COL Procedure

Put Column Values

Syntax

```
PROCEDURE PUT_COL(
   columnid NUMBER,
   collection IN <datatype>);
```

Where <*datatype*> can be any one of the supported types.

See Table 197-1 for the list of supported types.

Parameters

Table 197-21 PUT_COL Procedure Parameters

Parameter	Description
`columnid`	The id for the column
`collection`	The data for the column

Usage Notes

This procedure is used to put the read column values in the collection of scalar type.

The collection of scalar type should be of supported type only.

The column numbers are in the get column order as created in `DESCRIBE` method of PTF.

For the same columnid, `GET_COL` and `PUT_COL` may correspond to different column.

197.9.9.1 Rand_col Polymorphic Table Function Example

The rand_col PTF appends specified number of random-valued columns to the output.

Example 197-13 Rand_col Polymorphic Table Function Example

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Rand_col Polymorphic Table Function

This rand_col PTF example appends col_count number of random-valued columns to the output. Optionally, the caller can restrict the random values to a numeric range by specifying [low, high]. The new columns are named "RAND_<*n*>"

Create the PTF implementation package rand_col_p.

The parameters are :

- tab : Input table
- col_count (optional) : Number of random-valued columns to generate [Default = 1]
- low (optional) : Lower bound for the random numbers [Default = Null]
- high (optional) : Upper bound for the random numbers [Default = Null]

```
CREATE PACKAGE rand_col_p AS

  FUNCTION describe(tab       IN OUT DBMS_TF.table_t,
                    col_count        NATURALN DEFAULT 1,
                    low              NUMBER   DEFAULT NULL,
                    high             NUMBER   DEFAULT NULL)
```

```
          RETURN DBMS_TF.describe_t;

  PROCEDURE fetch_rows(col_count NATURALN DEFAULT 1,
                       low       NUMBER   DEFAULT NULL,
                       high      NUMBER   DEFAULT NULL);

END rand_col_p;
```

Create the PTF implementation package body rand_col_p.

The parameter col_count is a 'shape-determining' parameter and thus must be a constant (no binds, correlations, or expressions). By defining the type of col_count to be `NATURALN`, which has an implicit `NOT NULL` constraint, we guarantee that a cursor with non-constant value for this parameter will get a compilation error.

```
CREATE PACKAGE BODY rand_col_p AS
  col_name_prefix CONSTANT dbms_id := 'RAND_';

  FUNCTION describe(tab       IN OUT DBMS_TF.table_t,
                    col_count        NATURALN DEFAULT 1,
                    low              NUMBER   DEFAULT NULL,
                    high             NUMBER   DEFAULT NULL)
           RETURN DBMS_TF.describe_t
  AS
    cols DBMS_TF.columns_new_t;
  BEGIN
    FOR i IN 1 .. col_count LOOP
     cols(i):= DBMS_TF.column_metadata_t(name=>col_name_prefix||i,
TYPE=>DBMS_TF.type_number);
    END LOOP;

    RETURN DBMS_TF.describe_t(new_columns => cols);
  END;

  PROCEDURE fetch_rows(col_count NATURALN DEFAULT 1,
                       low       NUMBER   DEFAULT NULL,
                       high      NUMBER   DEFAULT NULL)
  AS
    row_count CONSTANT PLS_INTEGER := DBMS_TF.get_env().row_count;
    col       DBMS_TF.tab_number_t;
  BEGIN
    FOR c IN 1 .. col_count LOOP
      FOR i IN 1 .. row_count LOOP
        col(i) :=  CASE WHEN (low IS NULL OR high IS NULL)
                        THEN dbms_random.VALUE
                         ELSE dbms_random.VALUE(low,
high)
                    END;
      END LOOP;
      DBMS_TF.put_col(c, col);
    END LOOP;
  END;

END rand_col_p;
```

Create the standalone rand_col PTF. Specify exactly one formal argument of type TABLE, specify the return type of the PTF as TABLE, specify a Row Semantics PTF type, and indicate the PTF implementation package to use is rand_col_p.

```
CREATE FUNCTION rand_col(tab TABLE,
                         col_count NATURALN DEFAULT 1,
                         low       NUMBER   DEFAULT NULL,
                         high      NUMBER   DEFAULT NULL)
                  RETURN TABLE
                  PIPELINED ROW POLYMORPHIC USING rand_col_p;
```

Invoke the rand_col PTF to display all columns of table SCOTT.DEPT with one produced RAND_1 column.

```
  SELECT *
    FROM rand_col(scott.dept);
```

```
    DEPTNO DNAME          LOC              RAND_1
---------- -------------- ------------- ----------
        10 ACCOUNTING     NEW YORK      .738666262
        20 RESEARCH       DALLAS        .093256312
        30 SALES          CHICAGO       .992944835
        40 OPERATIONS     BOSTON        .397948124
```

Invoke the rand_col PTF to display all columns of table SCOTT.DEPT with two produced RAND_1 and RAND_2 columns.

```
  SELECT *
    FROM rand_col(scott.dept, col_count => 2);
```

```
    DEPTNO DNAME          LOC              RAND_1     RAND_2
---------- -------------- ------------- ---------- ----------
        10 ACCOUNTING     NEW YORK      .976521361 .209802028
        20 RESEARCH       DALLAS        .899577891  .10050334
        30 SALES          CHICAGO       .277238362 .110736583
        40 OPERATIONS     BOSTON        .989839995 .164822363
```

For all employees for which their job is not being a SALESMAN, display the employee name, job, and produce three RAND columns generating random values between —10 and 10.

```
  SELECT ename, job, rand_1, rand_2, rand_3
    FROM   rand_col(scott.emp, col_count => 3, low => -10, high => +10)
    WHERE  job != 'SALESMAN';
```

```
ENAME      JOB           RAND_1     RAND_2     RAND_3
---------- --------- ---------- ---------- ----------
SMITH      CLERK     8.91760464 6.67366638 -9.2789076
JONES      MANAGER   6.78612961 -1.8617958  6.5282227
BLAKE      MANAGER   7.59545803 5.22269017 -2.7966401
CLARK      MANAGER   -6.4747304 -7.3650276 3.28388872
SCOTT      ANALYST   6.80492435 -3.2271045 -.97099797
KING       PRESIDENT -9.3161177 6.27762154 -1.8184785
ADAMS      CLERK     -1.6618848 3.13119089 8.06363075
JAMES      CLERK     2.86918245 -3.5187936 -.72913809
FORD       ANALYST   6.67038328 -7.4989893 1.99072598
MILLER     CLERK     -2.1574578 -8.5082989 -.56046716
```

197.9.10 PUT_ROW_SET Procedure

Writes a collection of new column values in the database.

You can use this procedure to write all new columns in a collection of rows in the database.

This procedure is overloaded. Rows are not replicated by default. You can use the `ROW_REPLICATION` procedure to set the replication factor.

Syntax

This syntax is used when rows are not replicated.

```
PROCEDURE PUT_ROW_SET(
   rowset IN ROW_SET_T);
```

This syntax is used when the replication factor is a constant.

```
  PROCEDURE PUT_ROW_SET(
   rowset             IN   ROW_SET_T,
   replication_factor IN   NATURALN);
```

This syntax is used when the replication factor is specified as an array with multiple values.

```
PROCEDURE PUT_ROW_SET(
   rowset             IN   ROW_SET_T,
   replication_factor IN   TAB_NATURALN_T);
```

Parameters

Table 197-22 PUT_ROW_SET Procedure Parameters

Parameter	Description
`rowset`	The collection of data and metadata
`replication_factor`	The replication factor per row

Example 197-14 DBMS_TF.PUT_ROW_SET Example

This code excerpt fetches a collection of rows and writes all new columns back to the database without any processing.

```
PROCEDURE fetch_rows
AS
  rowset DBMS_TF.ROW_SET_T;
BEGIN
    DBMS_TF.GET_ROW_SET(rowset);
    DBMS_TF.PUT_ROW_SET(rowset);
END;
```

197.9.10.1 Split Polymorphic Table Function Example

The split PTF example splits each row of the input table into specified pieces.

Example 197-15 Split Polymorphic Table Function Example

This PTF example splits each row of the input table into cnt pieces dividing the values of the split columns.

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Split Polymorphic Table Function

Create the PTF implementation package split_p.

The parameters are :

- tab - Input table
- col - The names of numeric (input) table columns to split
- cnt - The number of times each input row is to be split

```
CREATE PACKAGE split_p AS

  FUNCTION  describe(tab  IN OUT DBMS_TF.table_t,
                     col         DBMS_TF.columns_t,
                     cnt         NATURALN)
            RETURN DBMS_TF.describe_t;

  PROCEDURE fetch_rows(cnt NATURALN);

END split_p;
```

Create the PTF implementation package body split_p. Each row of the input table is split into cnt pieces dividing the values of the split columns.

```
CREATE PACKAGE BODY split_p AS

  FUNCTION  describe(tab  IN OUT DBMS_TF.Table_t,
                     col         DBMS_TF.Columns_t,
                     cnt         NATURALN)
            RETURN DBMS_TF.describe_t
  AS
    new_cols DBMS_TF.columns_new_t;
    col_id   PLS_INTEGER := 1;
  BEGIN
    FOR i IN 1 .. tab.column.count LOOP
      FOR j IN 1 .. col.count LOOP
        IF (tab.column(i).description.name = col(j) AND
            tab.column(i).description.TYPE = DBMS_TF.type_number) THEN
          tab.column(i).pass_through := FALSE;
          tab.column(i).for_read     := TRUE;
          new_cols(col_id) := tab.column(i).description;
          col_id := col_id + 1;
        END IF;
      END LOOP;
    END LOOP;
```

```
    RETURN DBMS_TF.describe_t(new_columns=>new_cols, row_replication=>true);
  END;
  PROCEDURE fetch_rows(cnt NATURALN)
  AS
    inp_rs DBMS_TF.row_set_t;
    out_rs DBMS_TF.row_set_t;
    rows   PLS_INTEGER;
  BEGIN
    DBMS_TF.get_row_set(inp_rs, rows);

    FOR c IN 1 .. inp_rs.count() LOOP
      FOR r IN 1 .. rows LOOP
        FOR i IN 1 .. cnt LOOP
          out_rs(c).tab_number((r-1)*cnt+i) := inp_rs(c).tab_number(r)/cnt;
        END LOOP;
      END LOOP;
    END LOOP;

    DBMS_TF.put_row_set(out_rs, replication_factor => cnt);
  END;

END split_p;
```

Create the standalone PTF named split. Specify exactly one formal argument of type `TABLE`, specify the return type of the PTF as `TABLE`, specify a Row Semantics PTF type, and indicate the PTF implementation package to use is split_p.

```
CREATE FUNCTION split(tab TABLE, col columns, cnt NATURALN)
                  RETURN TABLE
PIPELINED ROW POLYMORPHIC USING split_p;
```

For all employees in department 30, display the `ENAME`, `SAL`, and `COMM` columns. Invoke the split PTF with the `COLUMNS` pseudo-operator to divide the value of `SAL` and `COMM` by 2 for each replicated row returned by the query. Each row is replicated twice.

```
SELECT ename, sal, comm
FROM   split(scott.emp, COLUMNS(sal, comm), cnt => 2)
WHERE  deptno=30;
```

```
ENAME             SAL                COMM
---------- ---------- --------------------
ALLEN             800                 150
ALLEN             800                 150
WARD              625                 250
WARD              625                 250
MARTIN            625                 700
MARTIN            625                 700
BLAKE            1425
BLAKE            1425
TURNER            750                   0
TURNER            750                   0
JAMES             475
JAMES             475
```

197.9.11 ROW_REPLICATION Procedure

Sets the row replication factor either as a fixed value or as a value per row.

This procedure is overloaded. A `Row Semantics` polymorphic table function will either produce a single output row for a given input row (one-to-one), or it can produce more output rows for a given input rows (one-to-many), or it can produce no output rows (one-to-none).

Syntax

Sets the row replication factor as a fixed value.

```
PROCEDURE ROW_REPLICATION(
  replication_factor IN NATURALN);
```

Sets the row replication factor as a value per row.

```
PROCEDURE ROW_REPLICATION(
  replication_factor IN TAB_NATURALN_T);
```

Parameters

Table 197-23 ROW_REPLICATION Procedure Parameters

Parameter	Description
`replication_factor`	The replication factor per row

Example 197-16 Replicate Polymorphic Table Function Example

This example creates a PTF that replicates each input row by the replication_factor that is given as a parameter.

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Replicate Polymorphic Table Function

Create the PTF implementation package replicate_p.

```
CREATE PACKAGE replicate_p
AS

    FUNCTION Describe(tab IN OUT DBMS_TF.TABLE_T,
                      replication_factor NATURAL)
        RETURN DBMS_TF.describe_t;

    PROCEDURE Fetch_Rows(replication_factor NATURALN);

END replicate_p;
```

Create the PTF implementation package body replicate_p. The PTF replicates each input row by the replication_factor that is given as a parameter.

```
CREATE PACKAGE body replicate_p
AS

    FUNCTION Describe(tab IN OUT DBMS_TF.Table_t
                    , replication_factor NATURAL)
        RETURN DBMS_TF.describe_t AS
    BEGIN
        RETURN DBMS_TF.describe_t(row_replication => True);
    END;

    PROCEDURE Fetch_Rows(replication_factor NATURALN)
    AS
    BEGIN
        DBMS_TF.ROW_REPLICATION(replication_factor);
    END;
END replicate_p;
```

Create a standalone PTF named replicate. Specify exactly one formal argument of type TABLE, specify the return type of the PTF as TABLE, specify a Row Semantics PTF type, and indicate the PTF implementation package to use is replicate_p.

```
CREATE FUNCTION replicate(tab TABLE,
                          replication_factor NATURAL)
    RETURN TABLE PIPELINED ROW POLYMORPHIC USING replicate_p;
```

This example sets the replication_factor to 2 which results in doubling the number of rows.

```
 SELECT *
 FROM replicate(dept, replication_factor => 2);


   DEPTNO DNAME          LOC
--------- -------------- -------------
        10 ACCOUNTING     NEW YORK
        10 ACCOUNTING     NEW YORK
        20 RESEARCH       DALLAS
        20 RESEARCH       DALLAS
        30 SALES          CHICAGO
        30 SALES          CHICAGO
        40 OPERATIONS     BOSTON
        40 OPERATIONS     BOSTON
```

This example sets the replication_factor to zero.

```
SELECT *
FROM replicate(dept, replication_factor => 0);


no rows selected
```

Count the number of employees in each department. Invoke the replicate PTF to report from the SCOTT.EMP table to set the replication_factor to 1000000.

```
SELECT deptno, COUNT(*)
FROM replicate(scott.emp, 1e6)
GROUP BY deptno;


    DEPTNO   COUNT(*)
---------- ----------
        30    6000000
        10    3000000
        20    5000000
```

This sets the replication_factor to 1000000000.

```
SELECT COUNT(*)
FROM replicate(dual, 1e9);


  COUNT(*)
----------
1000000000
```

197.9.12 ROW_TO_CHAR Function

The ROW_TO_CHAR function converts a row data value to a string representation.

Syntax

```
FUNCTION ROW_TO_CHAR(
   rowset ROW_SET_T,
   rid    PLS_INTEGER,
   format PLS_INTEGER DEFAULT FORMAT_JSON)
   RETURN VARCHAR2;
```

Parameters

Table 197-24 DBMS_TF.ROW_TO_CHAR Function Parameters

Parameter	Description
rowset	The rowset whose value is to be converted
rid	Row number
format	The string format (default is FORMAT_JSON)

Usage Notes

Only the JSON format is supported.

Return Values

The string representation in JSON format.

Example 197-17 DBMS_TF.ROW_TO_CHAR Example

```
 PROCEDURE Fetch_Rows as
   rowset DBMS_TF.ROW_SET_T;
   str    VARCHAR2(32000);
```

```
BEGIN
  DBMS_TF.GET_ROW_SET(rowset);
  str := DBMS_TF.ROW_TO_CHAR(rowset, 1)
END;
```

197.9.13 SUPPORTED_TYPE Function

This function tests if a specified type is supported with polymorphic table functions.

Syntax

```
FUNCTION SUPPORTED_TYPE(
   type_id PLS_INTEGER)
   RETURN BOOLEAN;
```

Parameters

Table 197-25 DBMS_TF.SUPPORTED_TYPE Function Parameters

Parameter	Description
type_id	The type

Return Values

Returns TRUE if the type_id is a scalar supported by `PUT_COL` and `GET_COL`.

> **See Also:**
>
> Echo Polymorphic Table Function Example for an example of DBMS_TF.SUPPORTED_TYPE use.

197.9.14 TRACE Procedure

Prints data structures to help development and problem diagnosis.

This procedure is overloaded.

Syntax

```
PROCEDURE TRACE(
   msg          VARCHAR2,
   with_id      BOOLEAN   DEFAULT FALSE,
   separator    VARCHAR2  DEFAULT NULL,
   prefix       VARCHAR2  DEFAULT NULL);

PROCEDURE TRACE(
   rowset       IN ROW_SET_T);

PROCEDURE TRACE(
   env          IN ENV_T);

PROCEDURE TRACE(
   columns_new  IN COLUMNS_NEW_T);

PROCEDURE trace(
   cols         IN COLUMNS_T);
```

```
PROCEDURE trace(
   columns_with_type  IN COLUMNS_WITH_TYPE_T);

PROCEDURE trace(
   tab          IN TABLE_T);

PROCEDURE trace(
   col          IN COLUMN_METADATA_T);
```

Parameters

Table 197-26 TRACE Procedure Parameters

Parameter	Description
msg	Custom user tracing message
with_id	Include the unique execution ID in the trace?
separator	Specify a string to use to separate values
prefix	Specify a string to prefix the actual values
rowset	Data for a rowset
env	Metadata about the polymorphic table function execution state
columns_new	Collection for new columns
cols	Collection containing column names
columns_with_type	Collection containing columns metadata
tab	Table descriptor
col	Metadata about an existing table column or a new column produced

Example 197-18 DBMS_TF.TRACE Example

This example adds tracing to a fetch_rows procedure.

```
PROCEDURE fetch_rows
AS
   rowset DBMS_TF.ROW_SET_T;
BEGIN
   DBMS_TF.TRACE('IDENTITY_PACKAGE.Fetch_Rows()', with_id => TRUE);
   DBMS_TF.TRACE(rowset);
   DBMS_TF.GET_ROW_SET(rowset);
   DBMS_TF.TRACE(rowset);
   DBMS_TF.PUT_ROW_SET(rowset);
   DBMS_TF.TRACE(DBMS_TF.GET_ENV);
END;
```

197.9.15 XSTORE_CLEAR Procedure

Removes all key-value pairs from the XSTORE execution state.

Syntax

```
PROCEDURE XSTORE_CLEAR;
```

197.9.16 XSTORE_EXISTS Function

Returns TRUE if an item with a given key exists in the XSTORE.

Syntax

```
FUNCTION XSTORE_EXISTS(
   key      IN VARCHAR2,
   key_type IN PLS_INTEGER DEFAULT NULL)
  RETURN BOOLEAN;
```

Parameters

Table 197-27 DBMS_TF.XSTORE_EXISTS Function Parameters

Parameter	Description
`key`	A unique character key
`key_type`	The type of key (optional). Default : NULL

Return Values

Returns TRUE if the key has an associated value. When the key_type is NULL (default), it returns TRUE if the key has an associated value of any of the supported type.

When a key_type parameter value is passed, it returns TRUE if the key and specified type of key has an associated value. Otherwise, it returns FALSE.

> **See Also:**
>
> Table 197-6 for more information about supported key types.

197.9.17 XSTORE_GET Procedure

You can use the XSTORE_GET procedure to get the associated value for a given key stored for PTF Execution State Management.

XStore is the PTF execution state management interface. The XStore interface is used to set and store key-value pairs during PTF execution.

This procedure is overloaded. The XStore supports specifying key-value pairs for these scalar types: `VARCHAR2`, `NUMBER`, `DATE`, `BOOLEAN`.

See Table 197-6 for more information about supported key types.

Syntax

```
PROCEDURE XSTORE_GET(
   key   IN VARCHAR2,
   value IN OUT VARCHAR2);

PROCEDURE XSTORE_GET(
   key   IN VARCHAR2,
   value IN OUT NUMBER);
```

```
PROCEDURE XSTORE_GET(
   key   IN VARCHAR2,
   value IN OUT DATE);

PROCEDURE XSTORE_GET(
   key   IN VARCHAR2,
   value IN OUT BOOLEAN);
```

Parameters

Table 197-28 DBMS_TF.XSTORE_GET Procedure Parameters

Parameter	Description
key	A unique character key
value	Value corresponding to the key for supported types

Usage Notes

If the key is not found, the value is unchanged.

197.9.17.1 Row_num Polymorphic Table Function Example

The row_num PTF example appends a sequence column to a table.

Example 197-19 Row_num Polymorphic Table Function Example

> **Live SQL:**
>
> You can view and run this example on Oracle Live SQL at Row_num Polymorphic Table Function

Create the PTF implementation package row_num_p.

The parameters are :

- tab - The input table
- ini - The initial value (Default = 1)
- inc - The amount to increment (Default = 1)

```
CREATE PACKAGE row_num_p IS
  FUNCTION describe(tab IN OUT dbms_tf.table_t,
                    ini NUMBER DEFAULT 1,
                    inc NUMBER DEFAULT 1)
           RETURN dbms_tf.describe_t;

  PROCEDURE fetch_rows(ini NUMBER DEFAULT 1, inc NUMBER DEFAULT 1);
END;
```

This PTF accepts any input table and appends the sequence column `ROW_ID` to the table. The sequence values start with the specified value (ini) and each time it is incremented by the specified value (inc).

```
CREATE PACKAGE BODY row_num_p IS
  FUNCTION describe(tab IN OUT dbms_tf.table_t,
                    ini NUMBER DEFAULT 1,
                    inc NUMBER DEFAULT 1)
           RETURN dbms_tf.describe_t AS
  BEGIN
    RETURN dbms_tf.describe_t(new_columns =>
             dbms_tf.columns_new_t(1 =>
               dbms_tf.column_metadata_t(name => 'ROW_ID',
                                         TYPE => dbms_tf.type_number)));
  END;

  PROCEDURE fetch_rows(ini NUMBER DEFAULT 1, inc NUMBER DEFAULT 1) IS
    row_cnt CONSTANT PLS_INTEGER := dbms_tf.get_env().row_count;
    rid     NUMBER               := ini;
    col     dbms_tf.tab_number_t;
  BEGIN
    dbms_tf.xstore_get('rid', rid);
    FOR i IN 1 .. row_cnt LOOP col(i) := rid + inc*(i-1); END LOOP;
    dbms_tf.put_col(1, col);
    dbms_tf.xstore_set('rid', rid + inc*row_cnt);
  END;

 END;
```

Create a standalone polymorphic table function named row_num. Specify exactly one formal argument of type `TABLE`, specify the return type of the PTF as `TABLE`, specify a Table Semantics PTF type, and indicate the PTF implementation package to use is row_num_p.

```
CREATE FUNCTION row_num(tab TABLE,
                        ini NUMBER DEFAULT 1,
                        inc NUMBER DEFAULT 1)
                        RETURN TABLE
PIPELINED TABLE POLYMORPHIC USING row_num_p;
```

The row_num PTF invocation reporting from the `SCOTT.DEPT` table produces a new column `ROW_ID` with value starting at 1 and incremented by 1 in the row set.

```
SELECT * FROM row_num(scott.dept);


DEPTNO DNAME          LOC              ROW_ID
------ -------------- ------------- ----------
    10 ACCOUNTING     NEW YORK              1
    20 RESEARCH       DALLAS                2
    30 SALES          CHICAGO               3
    40 OPERATIONS     BOSTON                4
```

The row_num PTF invocation reporting from the SCOTT.DEPT table produces a new column ROW_ID with value starting at 100 and incremented by 1 in the row set.

```
SELECT * FROM row_num(scott.dept, 100);


DEPTNO DNAME          LOC              ROW_ID
------ -------------- ------------- ----------
    10 ACCOUNTING     NEW YORK            100
    20 RESEARCH       DALLAS              101
    30 SALES          CHICAGO             102
    40 OPERATIONS     BOSTON              103
```

The row_num PTF invocation reporting from the SCOTT.DEPT table produces a new column ROW_ID with value starting at 0 and decremented by 1 in the row set.

```
SELECT * FROM row_num(scott.dept, ini => 0, inc => -1);


DEPTNO DNAME          LOC              ROW_ID
------ -------------- ------------- ----------
    10 ACCOUNTING     NEW YORK             0
    20 RESEARCH       DALLAS              -1
    30 SALES          CHICAGO             -2
    40 OPERATIONS     BOSTON              -3
```

The row_num PTF invocation reporting from the SCOTT.EMP table produces a new column ROW_ID with value starting at 0 and incremented by 0.25 in the row set which is partitioned by department number and ordered by employee name.

```
SELECT deptno, ename, job, sal, row_id
  FROM   row_num(scott.emp PARTITION BY deptno ORDER BY ename, ini => 0, inc
=> 0.25)
WHERE  deptno IN (10, 30);


DEPTNO ENAME      JOB              SAL     ROW_ID
------ ---------- --------- ---------- ----------
    10 CLARK      MANAGER         2450          0
    10 KING       PRESIDENT       5000        .25
    10 MILLER     CLERK           1300         .5
    30 ALLEN      SALESMAN        1600          0
    30 BLAKE      MANAGER         2850        .25
    30 JAMES      CLERK            950         .5
    30 MARTIN     SALESMAN        1250        .75
    30 TURNER     SALESMAN        1500          1
    30 WARD       SALESMAN        1250       1.25
```

197.9.18 XSTORE_REMOVE Procedure

Removes an item associated with the given key and key_type.

Syntax

```
PROCEDURE XSTORE_REMOVE(
   key      IN VARCHAR2,
   key_type IN PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 197-29 DBMS_TF.XSTORE_REMOVE Function Parameters

Parameter	Description
key	A unique character key
key_type	The type of key to remove (optional)

Usage Notes

When a key_type parameter value is passed, it removes the associated item for the key and specified type of key.

197.9.19 XSTORE_SET Procedure

Sets the value for the given key for PTF Execution State Management.

You can use this procedure to store and item key-value pair in the XStore. This procedure is overloaded. The XStore supports specifying key-value pairs for these scalar types: VARCHAR2, NUMBER, DATE, BOOLEAN.

Syntax

```
PROCEDURE XSTORE_SET (
   key   IN VARCHAR2,
   value IN VARCHAR2);

PROCEDURE XSTORE_SET (
   key   IN VARCHAR2,
   value IN NUMBER);

PROCEDURE XSTORE_SET (
   key   IN VARCHAR2,
   value IN DATE);

PROCEDURE XSTORE_SET (
   key   IN VARCHAR2,
   value IN BOOLEAN);
```

Parameters

Table 197-30 DBMS_TF.XSTORE_SET Procedure Parameters

Parameter	Description
key	A unique character key
value	Value corresponding to the key for supported types

Usage Notes

If an item for a given key already exists, the value is replaced.

198
DBMS_TNS

The DBMS_TNS package provides the RESOLVE_TNSNAME function to resolve a TNS name and return the corresponding Oracle Net8 connection string.

This chapter contains the following topics:

- DBMS_TNS Overview
- DBMS_TNS Security Model
- Summary of DBMS_TNS Subprograms

198.1 DBMS_TNS Overview

The DBMS_TNS package contains one function, RESOLVE_TNSNAME, which returns the resolved connect string from configured sources that have been defined in the names.directory_path parameter in the sqlnet.ora file.

The RESOLVE_TNSNAME function in this package helps you to identify interconnected databases and trace the system change number (SCN) flow across distributed databases.

> **See Also:**
>
> - *Oracle Database Administrator's Guide* for more information about finding database link information
> - *Oracle Database Reference* for information about the ALL_DB_LINKS data dictionary view

198.2 DBMS_TNS Security Model

You must be granted the EXECUTE privilege on the DBMS_TNS package.

198.3 Summary of DBMS_TNS Subprograms

This table lists the DBMS_TNS subprograms and briefly describes them.

Table 198-1 DBMS_TNS Parameters

Subprogram	Description
RESOLVE_TNSNAME Function	Returns the resolved connect string from any configured source (tnsnames, LDAP, Easy Connect) as indicated in the sqlnet.ora names.directory_path parameter)

198.3.1 RESOLVE_TNSNAME Function

This function returns the resolved connect string from any configured source (for example, TNS names, Lightweight Directory Access Protocol (LDAP), or Oracle Easy Connect) as indicated in the `sqlnet.ora names.directory_path` parameter.

Syntax

```
DBMS_TNS.RESOLVE_TNSNAME (
   tns_name VARCHAR2)
  RETURN VARCHAR2;
```

RESOLVE_TNSNAME Function Parameters

Parameter	Description
`tns_name`	Database address for establishing the connection

Usage Notes

- Database administrators and system management tools can explicitly call the `DBMS_TNS.RESOLVE_TNSNAME function` to resolve a `tns_name`:
- Because the resolved value can change over time, Oracle does not recommend storing or caching this value, because these kinds of values can become invalid over time.
- You can use the `DBMS_TNS.RESOLVE_TNSNAME` function to resolve TNS names that were referred to by other data dictionary views such as `ALL_DB_LINKS` and `USER_DB_LINKS`. In addition, you can use it to find any TNS name entry in the `tnsnames.ora` file.

Example

```
SELECT DB_LINK, DBMS_TNS.RESOLVE_TNSNAME(HOST) FROM DBA_DB_LINKS;

DB_LINK
--------------------------------------------------------------------------------
--
DBMS_TNS.RESOLVE_TNSNAME(HOST)
--------------------------------------------------------------------------------
--
US.EXAMPLE.COM
(DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=may))
(CONNECT_DATA=(SERVICE_NAME=us.example.com)(CID=(PROGRAM=oracle)(HOST=juno)
(USER=psmith))
))
```

199
DBMS_TRACE

The DBMS_TRACE package contains the interface to trace PL/SQL functions, procedures, and exceptions.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Restrictions
- Operational Notes
- Summary of DBMS_TRACE Subprograms

199.1 DBMS_TRACE Overview

DBMS_TRACE provides subprograms to start and stop PL/SQL tracing in a session. Oracle collects the trace data as the program executes and writes it to database tables.

A typical session involves:

1. (Optional) Limit tracing to specific subprograms and choose a tracing level.

   Tracing all subprograms and exceptions in a large program can produce huge amounts of data that are difficult to manage.
2. Starting PL/SQL tracing in session (DBMS_TRACE.SET_PLSQL_TRACE).
3. Running an application to be traced.
4. Stopping PL/SQL tracing in session (DBMS_TRACE.CLEAR_PLSQL_TRACE).

After you have collected data with Trace, you can query the database tables that contain the performance data and analyze it in the same way that you analyze the performance data from Profiler.

199.2 DBMS_TRACE Security Model

This package must be created under SYS.

199.3 DBMS_TRACE Constants

DBMS_TRACE defines constants to use when specifying parameter values.

These constants are shown in the following table.

Table 199-1 DBMS_TRACE Event Constants

Name	Type	Value	Description
TRACE_ALL_CALLS	INTEGER	1	Traces calls or returns
TRACE_ENABLED_CALLS	INTEGER	2	
TRACE_ALL_EXCEPTIONS	INTEGER	4	Traces exceptions
TRACE_ENABLED_EXCEPTIONS	INTEGER	8	Traces exceptions and handlers
TRACE_LIMIT	INTEGER	16	Save only the last few records. This allows tracing up to a problem area, without filling the database up with masses of irrelevant information. If event 10940 is set, the limit is 1023*(the value of event 10940). This can be overridden by the use of "TRACE_LIMIT" flag.
TRACE_ALL_SQL	INTEGER	32	Traces SQL statements
TRACE_ENABLED_SQL	INTEGER	64	Traces SQL statements at PL/SQL level. This does not invoke SQL Trace
TRACE_ALL_LINES	INTEGER	128	Traces each line
TRACE_ENABLED_LINES	INTEGER	256	
TRACE_PAUSE	INTEGER	4096	Pauses tracing
TRACE_RESUME	INTEGER	8192	Resume tracing
TRACE_STOP	INTEGER	16384	Stops tracing
NO_TRACE_ADMINISTRATIVE	INTEGER	32768	Prevents tracing of 'administrative events such as • PL/SQL Trace Tool started • Trace flags changed • PL/SQL Virtual Machine started • PL/SQL Virtual Machine stopped
NO_TRACE_HANDLED_EXCEPTIONS	INTEGER	65536	Prevents tracing of handled exceptions

Table 199-2 DBMS_TRACE Version Constants

Name	Type	Value	Description
TRACE_MINOR_VERSION	INTEGER	0	
TRACE_MAJOR_VERSION	INTEGER	1	

Oracle recommends using the symbolic form for all these constants.

199.4 DBMS_TRACE Restrictions

You cannot use PL/SQL tracing in a shared server environment.

199.5 DBMS_TRACE Operational Notes

Certain operational notes apply to `DBMS_TRACE`.

These are described in the following sections:

- Controlling Data Volume
- Creating Database Tables to Collect DBMS_TRACE Output
- Collecting Trace Data
- Collected Data
- Trace Control

Controlling Data Volume

Profiling large applications may produce a large volume of data. You can control the volume of data collected by *enabling* specific program units for trace data collection.

You can enable a program unit by compiling it debug. This can be done in one of two ways:

```
alter session set plsql_debug=true;
create or replace ... /* create the library units - debug information will be generated
*/
```

or:

```
/* recompile specific library unit with debug option */
alter [PROCEDURE | FUNCTION | PACKAGE BODY]  <libunit-name> compile debug;
```

> **Note:**
>
> You cannot use the second method for anonymous blocks.

You can limit the amount of storage used in the database by retaining only the most recent 8,192 records (approximately) by including `TRACE_LIMIT` in the `TRACE_LEVEL` parameter of the `SET_PLSQL_TRACE` procedure.

Creating Database Tables to Collect DBMS_TRACE Output

You must create database tables into which the `DBMS_TRACE` package writes output. Otherwise, the data is not collected. To create these tables, run the script `TRACETAB.SQL`. The tables this script creates are owned by SYS.

Collecting Trace Data

The PL/SQL features you can trace are described in the script `DBMSPBT.SQL`. Some of the key tracing features are:

- Tracing Calls
- Tracing Exceptions
- Tracing SQL
- Tracing Lines

Additional features of DBMS_TRACE also allow pausing and resuming trace, and limiting the output.

Tracing Calls

Two levels of call tracing are available:

- Level 1: Trace all calls. This corresponds to the constant `TRACE_ALL_CALLS`.
- Level 2: Trace calls to enabled program units only. This corresponds to the constant `TRACE_ENABLED_CALLS`.

Enabling cannot be detected for remote procedure calls (RPCs); hence, RPCs are only traced with level 1.

Tracing Exceptions

Two levels of exception tracing are available:

- Level 1: Trace all exceptions. This corresponds to `TRACE_ALL_EXCEPTIONS`.
- Level 2: Trace exceptions raised in enabled program units only. This corresponds to `TRACE_ENABLED_EXCEPTIONS`.

Tracing SQL

Two levels of SQL tracing are available:

- Level 1: Trace all SQL. This corresponds to the constant `TRACE_ALL_SQL`.
- Level 2: Trace SQL in enabled program units only. This corresponds to the constant `TRACE_ENABLED_SQL`.

Tracing Lines

Two levels of line tracing are available:

- Level 1: Trace all lines. This corresponds to the constant `TRACE_ALL_LINES`.
- Level 2: Trace lines in enabled program units only. This corresponds to the constant `TRACE_ENABLED_LINES`.

When tracing lines, Oracle adds a record to the database each time the line number changes. This includes line number changes due to procedure calls and returns.

> **Note:**
>
> For all types of tracing, level 1 overrides level 2. For example, if both level 1 and level 2 are enabled, then level 1 takes precedence.

Collected Data

If tracing is requested only for enabled program units, and if the current program unit is not enabled, then no trace data is written.

When tracing calls, both the call and return are traced. The check for whether tracing is "enabled" passes if either the called routine or the calling routine is "enabled".

Call tracing will always output the program unit type, program unit name, and line number for both the caller and the callee. It will output the caller's stack depth. If the caller is enabled, the caller's name will also be output. If the callee is enabled, the callee's name will also be output.

Exception tracing writes out the line number. Raising the exception shows information on whether the exception is user-defined or pre-defined. It also shows the exception number in the case of pre-defined exceptions. Both the place where the exceptions are raised and their handler is traced. The check for tracing being "enabled" is done independently for the place where the exception is raised and the place where the exception is handled. Enabling `NO_TRACE_HANDLED_EXCEPTIONS` limits data collection to unhandled exceptions

All calls to `DBMS_TRACE.SET_PLSQL_TRACE` and `DBMS_TRACE.CLEAR_PLSQL_TRACE` place a special trace record in the database. Therefore, it is always possible to determine when trace settings were changed.

Trace Control

As well as determining which items are collected, you can pause and resume the trace process. No information is gathered between the time that tracing is paused and the time that it is resumed. The constants `TRACE_PAUSE` and `TRACE_RESUME` are used to accomplish this. Trace records are generated to indicate that the trace was paused/resumed.

It is also possible to retain only the last 8,192 trace events of a run by using the constant `TRACE_LIMIT`. This allows tracing to be turned on without filling up the database. When tracing stops, the last 8,192 records are saved. The limit is approximate, since it is not checked on every trace record. At least the requested number of trace records will be generated; up to 1,000 additional records may be generated. At least the requested number of trace records will be generated; up to 1,000 additional records may be generated. The 8,192 record limit can be changed. Setting event 10940 to level *n* changes the record limit to 1024 * *n*.

Enabling `NO_TRACE_ADMINISTRATIVE` prevents the generation of such administrative event records as `PL/SQL Trace Tool started`, `Trace flags changed`, `PL/SQL Virtual Machine started`, and `PL/SQL Virtual Machine stopped`.

199.6 Summary of DBMS_TRACE Subprograms

This table lists the `DBMS_TRACE` subprograms and briefly describes them.

Table 199-3 DBMS_TRACE Package Subprograms

Subprogram	Description
CLEAR_PLSQL_TRACE Procedure	Stops trace data dumping in session
GET_PLSQL_TRACE_LEVEL Function	Gets the trace level
PLSQL_TRACE_VERSION Procedure	Gets the version number of the trace package
SET_PLSQL_TRACE Procedure	Starts tracing in the current session

199.6.1 CLEAR_PLSQL_TRACE Procedure

This procedure disables trace data collection.

Syntax

```
DBMS_TRACE.CLEAR_PLSQL_TRACE;
```

199.6.2 GET_PLSQL_TRACE_LEVEL Function

This procedure returns the current trace level as the sum of one or more DBMS_TRACE constants.

See Table 199-1 for a list of the constants.

Syntax

```
DBMS_TRACE.GET_PLSQL_TRACE_LEVEL
  RETURN BINARY_INTEGER;
```

199.6.3 PLSQL_TRACE_VERSION Procedure

This procedure gets the version number of the trace package. It returns the major and minor version number of the DBMS_TRACE package.

Syntax

```
DBMS_TRACE.PLSQL_TRACE_VERSION (
   major OUT BINARY_INTEGER,
   minor OUT BINARY_INTEGER);
```

Parameters

Table 199-4 PLSQL_TRACE_VERSION Procedure Parameters

Parameter	Description
major	Major version number of DBMS_TRACE.
minor	Minor version number of DBMS_TRACE.

199.6.4 SET_PLSQL_TRACE Procedure

This procedure enables PL/SQL trace data collection.

Syntax

```
DBMS_TRACE.SET_PLSQL_TRACE (
   trace_level INTEGER);
```

Parameters

Table 199-5 *SET_PLSQL_TRACE Procedure Parameters*

Parameter	Description
trace_level	You must supply one or more of the constants as listed in Table 199-1. By summing the constants, you can enable tracing of multiple PL/SQL language features simultaneously. The control constants "TRACE_PAUSE", "TRACE_RESUME" and "TRACE_STOP" should not be used in combination with other constants. Also see DBMS_TRACE Operational Notes: Collecting Trace Data for more information.

200

DBMS_TRANSACTION

The DBMS_TRANSACTION package provides access to SQL transaction statements from stored procedures.

> **See Also:**
>
> *Oracle Database SQL Language Reference*

This chapter contains the following topics:

- DBMS_TRANSACTION Security Model
- Summary of DBMS_TRANSACTION Subprograms

200.1 DBMS_TRANSACTION Security Model

This package runs with the privileges of calling user, rather than the package owner `SYS`.

200.2 Summary of DBMS_TRANSACTION Subprograms

This table lists the `DBMS_TRANSACTION` subprograms and briefly describes them.

Table 200-1 DBMS_TRANSACTION Package Subprograms

Subprogram	Description
ADVISE_COMMIT Procedure	Equivalent to the SQL statement: `ALTER SESSION ADVISE COMMIT`
ADVISE_NOTHING Procedure	Equivalent to the SQL statement: `ALTER SESSION ADVISE NOTHING`
ADVISE_ROLLBACK Procedure	Equivalent to the SQL statement: `ALTER SESSION ADVISE ROLLBACK`
COMMIT Procedure	Equivalent to the SQL statement: `COMMIT`
COMMIT_COMMENT Procedure	Equivalent to the SQL statement: `COMMIT COMMENT <text>`
COMMIT_FORCE Procedure	Equivalent to the SQL statement: `COMMIT FORCE <text>, <number>"`
LOCAL_TRANSACTION_ID Function	Returns the local (to instance) unique identifier for the current transaction
GET_TRANSACTION_ID Function	Returns the transaction identifier for the Sessionless transaction or the XA transaction in the session

Table 200-1 (Cont.) DBMS_TRANSACTION Package Subprograms

Subprogram	Description
GET_TRANSACTION_TIMEOUT Function	Returns the timeout, in seconds, of the XA transaction branch or the Sessionless transaction within a session
GET_TRANSACTION_TYPE Function	Returns the type of the transaction active in the session
PURGE_LOST_DB_ENTRY Procedure	Enables removal of incomplete transactions from the local site when the remote database is destroyed or re-created before recovery completes
PURGE_MIXED Procedure	Deletes information about a given mixed outcome transaction
READ_ONLY Procedure	Equivalent to the SQL statement: `SET TRANSACTION READ ONLY`
READ_WRITE Procedure	equivalent to the SQL statement: `SET TRANSACTION READ WRITE`
ROLLBACK Procedure	Equivalent to the SQL statement: `ROLLBACK`
ROLLBACK_FORCE Procedure	Equivalent to the SQL statement: `ROLLBACK FORCE <text>`
ROLLBACK_SAVEPOINT Procedure	Equivalent to the SQL statement: `ROLLBACK TO SAVEPOINT <savepoint_name>`
SAVEPOINT Procedure	Equivalent to the SQL statement: `SAVEPOINT <savepoint_name>`
START_TRANSACTION Function	Starts a local (vanilla) transaction or a Sessionless transaction
STEP_ID Function	Returns local (to local transaction) unique positive integer that orders the DML operations of a transaction
SUSPEND_TRANSACTION Procedure	Suspends the Sessionless transaction that is active within a session
USE_ROLLBACK_SEGMENT Procedure	Equivalent to the SQL statement: `SET TRANSACTION USE ROLLBACK SEGMENT <rb_seg_name>`

200.2.1 ADVISE_COMMIT Procedure

This procedure is equivalent to the SQL statement: `ALTER SESSION ADVISE COMMIT`

Syntax

```
DBMS_TRANSACTION.ADVISE_COMMIT;
```

200.2.2 ADVISE_NOTHING Procedure

This procedure is equivalent to the SQL statement: `ALTER SESSION ADVISE NOTHING`

Syntax

```
DBMS_TRANSACTION.ADVISE_NOTHING;
```

200.2.3 ADVISE_ROLLBACK Procedure

This procedure is equivalent to the SQL statement: `ALTER SESSION ADVISE ROLLBACK`

Syntax

```
DBMS_TRANSACTION.ADVISE_ROLLBACK;
```

200.2.4 COMMIT Procedure

This procedure is equivalent to the SQL statement: `COMMIT`

This procedure is included for completeness, the functionality being already implemented as part of PL/SQL.

Syntax

```
DBMS_TRANSACTION.COMMIT;
```

200.2.5 COMMIT_COMMENT Procedure

This procedure is equivalent to the SQL statement: `COMMIT COMMENT <text>`

Syntax

```
DBMS_TRANSACTION.COMMIT_COMMENT (
   cmnt VARCHAR2);
```

Parameters

Table 200-2 COMMIT_COMMENT Procedure Parameters

Parameter	Description
cmnt	Comment to associate with this commit.

200.2.6 COMMIT_FORCE Procedure

This procedure is equivalent to the SQL statement: `COMMIT FORCE <text>, <number>"`

Syntax

```
DBMS_TRANSACTION.COMMIT_FORCE (
   xid VARCHAR2,
   scn VARCHAR2 DEFAULT NULL);
```

Parameters

Table 200-3 COMMIT_FORCE Procedure Parameters

Parameter	Description
xid	Local or global transaction ID.
scn	System change number.

200.2.7 GET_TRANSACTION_ID Function

This function returns the transaction identifier in the session.

This function returns the hexadecimal representations of the transaction identifier (in varchar2 format) of the active XA transaction or Sessionless transaction in the session. If there are none of these types of transactions active within the session, it returns NULL.

If it's an XA transaction, the function returns a hex representation of format ID and GTRID separated by a dot (<format id>.<GTRID>).

If it is a Sessionless transaction, the function returns a hex representation of GTRID. Call GET_TRANSACTION_TYPE to get whether the transaction is of type XA or Sessionless.

Syntax

```
GET_TRANSACTION_ID
   RETURN VARCHAR2;
```

200.2.8 GET_TRANSACTION_TIMEOUT Function

This function returns the timeout, in seconds, of the XA transaction branch or the Sessionless transaction within a session.

If a local transaction is active within the session, the function returns UB4MAXVAL.

If no transaction is active within the session, the function returns NULL.

Syntax

```
GET_TRANSACTION_TIMEOUT
   RETURN NUMBER;
```

200.2.9 GET_TRANSACTION_TYPE Function

This function returns the type of the transaction active in the session.

The type of the transaction returned can be one of the following: TRANSACTION_TYPE_LOCAL, TRANSACTION_TYPE_SESSIONLESS, TRANSACTION_TYPE_XA, or NULL (if no transaction is active within the session).

Syntax

```
GET_TRANSACTION_TYPE
   RETURN PLS_INTEGER;
```

200.2.10 LOCAL_TRANSACTION_ID Function

This function returns the local (to instance) unique identifier for the current transaction. It returns null if there is no current transaction.

Syntax

```
DBMS_TRANSACTION.LOCAL_TRANSACTION_ID (
   create_transaction BOOLEAN := FALSE)
  RETURN VARCHAR2;
```

Parameters

Table 200-4 LOCAL_TRANSACTION_ID Function Parameters

Parameter	Description
`create_transaction`	If true, then start a transaction if one is not currently active.

200.2.11 PURGE_LOST_DB_ENTRY Procedure

Procedure `PURGE_LOST_DB_ENTRY` purges entries that control database recovery from a local site.

When a failure occurs during commit processing, automatic recovery consistently resolves the results at all sites involved in the transaction. However, if the remote database is destroyed or re-created before recovery completes, then the entries used to control recovery in `DBA_2PC_PENDING` and associated tables are never removed, and recovery will periodically retry. Procedure `PURGE_LOST_DB_ENTRY` enables removal of such transactions from the local site.

Syntax

```
DBMS_TRANSACTION.PURGE_LOST_DB_ENTRY (
   xid VARCHAR2);
```

Parameters

Table 200-5 PURGE_LOST_DB_ENTRY Procedure Parameters

Parameter	Description
`xid`	Must be set to the value of the `LOCAL_TRAN_ID` column in the `DBA_2PC_PENDING` table.

Usage Notes

> **WARNING:**
>
> PURGE_LOST_DB_ENTRY should *only* be used when the other database is lost or has been re-created. Any other use may leave the other database in an unrecoverable or inconsistent state.

Before automatic recovery runs, the transaction may show up in `DBA_2PC_PENDING` as state "collecting", "committed", or "prepared". If the DBA has forced an in-doubt transaction to have a particular result by using "commit force" or "rollback force", then states "forced commit" or "forced rollback" may also appear. Automatic recovery normally deletes entries in any of these states. The only exception is when recovery finds a forced transaction which is in a state inconsistent with other sites in the transaction; in this case, the entry is left in the table and the `MIXED` column has the value 'yes'.

However, under certain conditions, it may not be possible for automatic recovery to run. For example, a remote database may have been permanently lost. Even if it is re-created, it gets a new database ID, so that recovery cannot identify it (a possible symptom is `ORA-02062`). In this case, the DBA may use the procedure `PURGE_LOST_DB_ENTRY` to clean up the entries in any state other than "prepared". The DBA does not need to be in any particular hurry to resolve these entries, because they are not holding any database resources.

The following table indicates what the various states indicate about the transaction and what the DBA actions should be:

Table 200-6 PURGE_LOST_DB_ENTRY Procedure States

State of Column	State of Global Transaction	State of Local Transaction	Normal DBA Action	Alternative DBA Action
Collecting	Rolled back	Rolled back	None	PURGE_LOST_DB_ENTRY (See Note 1)
Committed	Committed	Committed	None	PURGE_LOST_DB_ENTRY (See Note 1)
Prepared	Unknown	Prepared	None	FORCE COMMIT or ROLLBACK
Forced commit	Unknown	Committed	None	PURGE_LOST_DB_ENTRY (See Note 1)
Forced rollback	Unknown	Rolled back	None	PURGE_LOST_DB_ENTRY (See Note 1)
Forced commit (mixed)	Mixed	Committed	(See Note 2)	
Forced rollback (mixed)	Mixed	Rolled back	(See Note 2)	

> **Note:**
>
> Use only if significant reconfiguration has occurred so that automatic recovery cannot resolve the transaction. Examples are total loss of the remote database, reconfiguration in software resulting in loss of two-phase commit capability, or loss of information from an external transaction coordinator such as a TP monitor.

> **Note:**
>
> Examine and take any manual action to remove inconsistencies; then use the procedure `PURGE_MIXED`.

200.2.12 PURGE_MIXED Procedure

This procedure deletes information about a given mixed outcome transaction

When in-doubt transactions are forced to commit or rollback (instead of letting automatic recovery resolve their outcomes), there is a possibility that a transaction can have a mixed outcome; some sites commit, and others rollback. Such inconsistency cannot be resolved automatically by Oracle. However, Oracle flags entries in `DBA_2PC_PENDING` by setting the `MIXED` column to a value of 'yes'.

Oracle never automatically deletes information about a mixed outcome transaction. When the application or DBA is certain that all inconsistencies that might have arisen as a result of the mixed transaction have been resolved, this procedure can be used to delete the information about a given mixed outcome transaction.

Syntax

```
DBMS_TRANSACTION.PURGE_MIXED (
   xid VARCHAR2);
```

Parameters

Table 200-7 PURGE_MIXED Procedure Parameters

Parameter	Description
`xid`	Must be set to the value of the `LOCAL_TRAN_ID` column in the `DBA_2PC_PENDING` table.

200.2.13 READ_ONLY Procedure

This procedure is equivalent to the SQL statement `SET TRANSACTION READ ONLY`.

Syntax

```
DBMS_TRANSACTION.READ_ONLY;
```

200.2.14 READ_WRITE Procedure

This procedure is equivalent to the SQL statement:

```
SET TRANSACTION READ WRITE
```

Syntax

```
DBMS_TRANSACTION.READ_WRITE;
```

200.2.15 ROLLBACK Procedure

This procedure is equivalent to the SQL statement `ROLLBACK`.

This procedure is included for completeness, the functionality being already implemented as part of PL/SQL.

Syntax

```
DBMS_TRANSACTION.ROLLBACK;
```

200.2.16 ROLLBACK_FORCE Procedure

This procedure is equivalent to the SQL statement `ROLLBACK FORCE <text>`.

Syntax

```
DBMS_TRANSACTION.ROLLBACK_FORCE (
   xid VARCHAR2);
```

Parameters

Table 200-8 ROLLBACK_FORCE Procedure Parameters

Parameter	Description
`xid`	Local or global transaction ID.

200.2.17 ROLLBACK_SAVEPOINT Procedure

This procedure is equivalent to the SQL statement `ROLLBACK TO SAVEPOINT <savepoint_name>`.

This procedure is included for completeness, the functionality being already implemented as part of PL/SQL.

Syntax

```
DBMS_TRANSACTION.ROLLBACK_SAVEPOINT (
   savept VARCHAR2);
```

Parameters

Table 200-9 ROLLBACK_SAVEPOINT Procedure Parameters

Parameter	Description
savept	Savepoint identifier.

200.2.18 SAVEPOINT Procedure

This procedure is equivalent to the SQL statement `SAVEPOINT <savepoint_name>`.

This procedure is included for completeness, the feature being already implemented as part of PL/SQL.

Syntax

```
DBMS_TRANSACTION.SAVEPOINT (
   savept VARCHAR2);
```

Parameters

Table 200-10 SAVEPOINT Procedure Parameters

Parameter	Description
savept	Savepoint identifier.

200.2.19 START_TRANSACTION Function

This function starts a local (vanilla) transaction or a Sessionless transaction depending on the `transaction_type` being `TRANSACTION_TYPE_LOCAL` or `TRANSACTION_TYPE_SESSIONLESS`.

If starting a local transaction, this function returns the `XID <usn>.<slot>.<wrap>` just like `DBMS_TRANSACTION.local_transaction_id()` does. If starting a Sessionless transaction, this function returns the hexadecimal representation of GTRID of the started transaction.

Syntax

```
DBMS_TRANSACTION.START_TRANSACTION(
   xid              in raw default null,
   transaction_type in pls_integer default TRANSACTION_TYPE_LOCAL,
   timeout          in pls_integer default 60,
   flag             in pls_integer default 0
) RETURN VARCHAR2;
```

The following table describes the parameters of the `DBMS_TRANSACTION.START_TRANSACTION` function.

> **Note:**
>
> The `xid`, `timeout`, and `flag` parameters apply only if the `TRANSACTION_TYPE` is `TRANSACTION_TYPE_SESSIONLESS`.

Table 200-11 START_TRANSACTION Function Parameters

Parameter	Description
`xid`	GTRID of the Sessionless transaction to be newly started or resumed. It can be raw byte array of size up to 64 bytes. If `NULL`, the server generates a GTRID.
`timeout`	The duration, in seconds, that this transaction can be resumed after it is suspended. **Note:** As of Oracle Database 23ai, Version 23.6, the `timeout` parameter is only used if the `flag` is `TRANSACTION_NEW`.
`flag`	The flag can be either `TRANSACTION_NEW` for starting a new transaction or `TRANSACTION_RESUME` for resuming a suspended transaction.

200.2.20 STEP_ID Function

This function returns local (to local transaction) unique positive integer that orders the DML operations of a transaction.

Syntax

```
DBMS_TRANSACTION.STEP_ID
   RETURN NUMBER;
```

200.2.21 SUSPEND_TRANSACTION Procedure

This procedure suspends the Sessionless transaction that is active within a session.

If the active transaction is not a Sessionless transaction, an `ORA-26202` error is raised. If there is no transaction active within the session, the suspend transaction operation is a no-op.

Syntax

```
DBMS_TRANSACTION.SUSPEND_TRANSACTION
```

200.2.22 USE_ROLLBACK_SEGMENT Procedure

This procedure is equivalent to the SQL statement `SET TRANSACTION USE ROLLBACK SEGMENT <rb_seg_name>`.

Syntax

```
DBMS_TRANSACTION.USE_ROLLBACK_SEGMENT (
   rb_name VARCHAR2);
```

Parameters

Table 200-12 USE_ROLLBACK_SEGMENT Procedure Parameters

Parameter	Description
rb_name	Name of rollback segment to use.

201
DBMS_TRANSFORM

The DBMS_TRANSFORM package provides an interface to the message format transformation features of Oracle Advanced Queuing.

This chapter contains the following topic:

- Summary of DBMS_TRANSFORM Subprograms

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide*for more on message format transformations.

201.1 Summary of DBMS_TRANSFORM Subprograms

This table lists the DBMS_TRANSFORM subprograms and briefly describes them.

Table 201-1 DBMS_TRANSFORM Package Subprograms

Subprograms	Description
CREATE_TRANSFORMATION Procedure	Creates a transformation that maps an object of the source type to an object of the destination type
DROP_TRANSFORMATION Procedure	Drops the given transformation
MODIFY_TRANSFORMATION Procedure	Modifies an existing transformation

201.1.1 CREATE_TRANSFORMATION Procedure

This procedure creates a transformation that maps an object of the source type to an object of the target type. The transformation expression can be a SQL expression or a PL/SQL function. It must return an object of the target type.

Syntax

```
DBMS_TRANSFORM.CREATE_TRANSFORMATION (
    schema               VARCHAR2(30),
    name                 VARCHAR2(30),
    from_schema          VARCHAR2(30),
    from_type            VARCHAR2(30),
    to_schema            VARCHAR2(30),
    to_type              VARCHAR2(30),
    transformation       VARCHAR2(4000));
```

Parameters

Table 201-2 CREATE_TRANSFORMATION Procedure Parameters

Parameter	Description
schema	Specifies the schema of the transformation.
name	Specifies the name of the transformation.
from_schema	Specifies the schema of the source type.
from_type	Specifies the source type.
to_schema	Specifies the target type schema.
to_type	Specifies the target type.
transformation	Specifies the transformation expression, returning an object of the target type. The expression must be a function returning an object of the target type or a constructor expression for the target type. You can choose not to specify a transformation expression and instead specify transformations for attributes of the target type using MODIFY_TRANSFORMATION.

Usage Notes

- The transformation expression must be a SQL expression or a PL/SQL function returning the type of the specified attribute of the target type.
- To create, modify or drop transformations, a user must be granted execute privileges on DBMS_TRANSFORM. The user must also have execute privileges on the user defined types that are the source and destination types of the transformation. In addition, the user must also have execute privileges on any PLSQL function being used in the transformation function.
- The transformation cannot write database state (perform DML) or commit or rollback the current transaction.
- The transformation must be a SQL function with source type as input type, returning an object of the target type. It could also be a SQL expression of target type, referring to a source type. All references to the source type must be of the form *source.user_data*.
- Both source and target types must be non-scalar database types. A null transformation expression maps to a null target object.

For using the transformation at enqueue and dequeue time, the login user invoking the operation must have execute privileges on the PLSQL functions used by the transformation. For propagation, the owning schema of the queue must have these privileges.

201.1.2 DROP_TRANSFORMATION Procedure

This procedure drops the given transformation.

Syntax

```
DBMS_TRANSFORM.DROP_TRANSFORMATION (
```

schema VARCHAR2(30),

```
    name        VARCHAR2(30));
```

Parameters

Table 201-3 DROP_TRANSFORMATION Procedure Parameters

Parameter	Description
schema	Specifies the schema of the transformation.
name	Specifies the name of the transformation.

201.1.3 MODIFY_TRANSFORMATION Procedure

This procedure modifies the transformation expression for the given transformation.

Syntax

```
DBMS_TRANSFORM.MODIFY_TRANSFORMATION (
   schema             VARCHAR2(30),
   name               VARCHAR2(30),
   attribute_number   INTEGER,
   transformation     VARCHAR2(4000));
```

Parameters

Table 201-4 MODIFY_TRANSFORMATION Procedure Parameters

Parameter	Description
schema	Specifies the schema of the transformation.
name	Specifies the name of the transformation.
attribute_number	The attribute of the target type for which the new transformation expression is being specified. When specifying the new transformation as a single expression of the target type, specify a value of 0.
transformation	The transformation expression must be a SQL expression or a PL/SQL function returning the type of the specified attribute of the target type. If the attribute_number is 0, then the expression must be a PL/SQL function returning an object of the target type or a constructor expression for the target type.

Usage Notes

- If the new transformation is a single expression of the target type, it may be specified with an attribute_number of 0. The new transformation may also be specified for each attribute of the target type.
- You can use this procedure to define the transformation as a separate expression for each attribute of the target type. For large transformations, this representation may be more readable and allow the application of fine grain control over the transformation. If the transformation expression was left unspecified for some of the attributes of the target type, they are evaluated to null when the transformation is applied.

202
DBMS_TSDP_MANAGE

The DBMS_TSDP_MANAGE package provides an interface to import and manage sensitive columns and sensitive column types in the database, and is used in conjunction with the DBMS_TSDP_PROTECT package with regard to transparent sensitive data protection (TSDP) policies.

DBMS_TSDP_MANAGE is available with the Enterprise Edition only.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_TSDP_MANAGE Subprograms

Related Topics

- DBMS_TSDP_PROTECT
  The DBMS_TSDP_PROTECT package provides an interface to configure transparent sensitive data protection (TSDP) policies in conjunction with the DBMS_TSDP_MANAGE package.

> **See Also:**
>
> *Oracle Database Security Guide*

202.1 DBMS_TSDP_MANAGE Overview

The DBMS_TSDP_MANAGE package lets you manage sensitive columns and sensitive types in the Oracle database.

The identified sensitive columns are classified based on the sensitive types. By Using the DBMS_TSDP_PROTECT package to create a policy that protects data for a given class based on a column type rather than the data itself, you can then manage security for these types in a uniform fashion and modify the settings to accommodate changing compliance regulations.

You also can export the policies to other databases, when you perform a full export using Data Pump. You cannot export the policy itself, but an export of the database will include the TSDP policies.

202.2 DBMS_TSDP_MANAGE Security Model

All procedures are executed with invoker's rights. The DBMS_TSDP_MANAGE package is owned by SYS.

The EXECUTE privilege on this package should be granted as appropriate. Typically, an application database administrator should be granted the EXECUTE privilege for this package, while the DBMS_TSDP_PROTECT package would be governed by a security administrator.

202.3 Summary of DBMS_TSDP_MANAGE Subprograms

This table lists the `DBMS_TSDP_MANAGE` subprograms and briefly describes them.

Table 202-1 DBMS_TSDP_MANAGE Package Subprograms

Subprogram	Description
ADD_SENSITIVE_COLUMN Procedure	Adds a column to the sensitive column list
ADD_SENSITIVE_TYPE Procedure	Creates and adds a sensitive column type to the list of sensitive column types in the database
ALTER_SENSITIVE_COLUMN Procedure	Alters the sensitive type and/or the comment of a column in the sensitive column list.
DROP_SENSITIVE_COLUMN Procedure	Removes columns from the sensitive column list
DROP_SENSITIVE_TYPE Procedure	Drops a sensitive column type from the list sensitive column types in the database
DROP_SENSITIVE_TYPE_SOURCE Procedure	Drops sensitive column types corresponding to a source from the list sensitive column types in the database
IMPORT_DISCOVERY_RESULT Procedure	Imports sensitive columns from an external source. This can be an Application Data Model (ADM) from an Oracle Enterprise Manager Cloud Control instance
IMPORT_SENSITIVE_TYPES Procedure	Imports a list of sensitive column types from a source
REMOVE_DISCOVERY_RESULT Procedure	Removes sensitive columns corresponding to an Application Data Model (ADM) from an Oracle Enterprise Manager Cloud Control instance.

202.3.1 ADD_SENSITIVE_COLUMN Procedure

This procedure adds a column to the sensitive column list.

Syntax

```
DBMS_TSDP_MANAGE.ADD_SENSITIVE_COLUMN (
   schema_name        IN VARCHAR2,
   table_name         IN VARCHAR2,
   column_name        IN VARCHAR2,
   sensitive_type     IN VARCHAR2,
   user_comment       IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 202-2 *ADD_SENSITIVE_COLUMN Procedure Parameters*

Parameter	Description
`schema_name`	Schema to which the column belongs
`table_name`	Table containing the column
`column_name`	Sensitive column name
`sensitive_type`	Identifier of the sensitive column type

Table 202-2 (Cont.) *ADD_SENSITIVE_COLUMN Procedure Parameters*

Parameter	Description
user_comment	User comment regarding the sensitive column

Examples

Add a column SAL in SCOTT.EMP:

```
DBMS_TSDP_MANAGE.ADD_SENSITIVE_COLUMN (
  schema_name        => 'SCOTT',
  table_name         => 'EMP',
  column_name        => 'SAL',
  sensitive_type     => 'SALARY_TYPE',
  user_comment       => 'Salary column');
```

202.3.2 ALTER_SENSITIVE_COLUMN Procedure

This procedure alters the Sensitive Type and/or the Comment of a Column in the sensitive column list.

Syntax

```
DBMS_TSDP_MANAGE.ALTER_SENSITIVE_COLUMN (
   schema_name        IN VARCHAR2,
   table_name         IN VARCAHR2,
   column_name        IN VARCAHR2,
   sensitive_type     IN VARCAHR2,
   user_comment       IN VARCAHR2 DEFAULT NULL);
```

Parameters

Table 202-3 *ALTER_SENSITIVE_COLUMN Procedure Parameters*

Parameter	Description
schema_name	Schema to which the column belongs
table_name	Table containing the column
column_name	Sensitive column name
sensitive_type	Identifier of the sensitive column type
user_comment	User comment regarding the sensitive column

Examples

Alter the column SAL in SCOTT.EMP that is listed in the sensitive column list:

```
DBMS_TSDP_MANAGE.ALTER_SENSITIVE_COLUMN (
  schema_name        => 'SCOTT',
  table_name         => 'EMP',
  column_name        => 'SAL',
  sensitive_type     => 'FINANCE_Type',
  user_comment       => 'Finance Type. Earlier categorized as Salary Type');
```

202.3.3 ADD_SENSITIVE_TYPE Procedure

This procedure creates and adds a sensitive column type to the list sensitive column types in the database.

Syntax

```
DBMS_TSDP_MANAGE.ADD_SENSITIVE_TYPE (
   sensitive_type     IN  VARCHAR2,
   user_comment       IN  VARCAHR2 DEFAULT NULL);
```

Parameters

Table 202-4 *ADD_SENSITIVE_TYPE Procedure Parameters*

Parameter	Description
sensitive_type	Name of the sensitive column type
user_comment	User comment regarding the sensitive column

Examples

Add a sensitive column type called SALARY_TYPE that is intended to be associated with columns containing salary data:

```
DBMS_TSDP_MANAGE.ADD_SENSITIVE_TYPE (
  sensitive_type     => 'SALARY_TYPE',
  user_comment       => 'Salary data');
```

202.3.4 DROP_SENSITIVE_COLUMN Procedure

This procedure removes columns from the sensitive column list.

Syntax

```
DBMS_TSDP_MANAGE.DROP_SENSITIVE_COLUMN (
   schema_name        IN VARCHAR2 DEFAULT '%',
   table_name         IN VARCAHR2 DEFAULT '%',
   column_name        IN VARCAHR2 DEFAULT '%');
```

Parameters

Table 202-5 *DROP_SENSITIVE_COLUMN Procedure Parameters*

Parameter	Description
schema_name	Schema to which the column belongs
table_name	Table containing the column
column_name	Sensitive column name

Examples

Remove column SAL in SCOTT.EMP from the sensitive column list:

```
DBMS_TSDP_MANAGE.DROP_SENSITIVE_COLUMN (
  schema_name         => 'SCOTT',
  table_name          => 'EMP',
  column_name         => 'SAL');
```

202.3.5 DROP_SENSITIVE_TYPE Procedure

This procedure drops a sensitive column type from the list sensitive column types in the database.

Syntax

```
DBMS_TSDP_MANAGE.DROP_SENSITIVE_TYPE (
   sensitive_type     IN VARCHAR2);
```

Parameters

Table 202-6 *DROP_SENSITIVE_TYPE Procedure Parameters*

Parameter	Description
sensitive_type	Name of the sensitive column type to be dropped

Examples

To drop SALARY_TYPE:

```
DBMS_TSDP_MANAGE.DROP_SENSITIVE_TYPE (
  sensitive_type     => 'SALARY_TYPE');
```

202.3.6 DROP_SENSITIVE_TYPE_SOURCE Procedure

This procedure drops sensitive column types corresponding to a source from the list sensitive column types in the database.

Syntax

```
DBMS_TSDP_MANAGE.DROP_SENSITIVE_TYPE_SOURCE (
   source      IN VARCHAR2);
```

Parameters

Table 202-7 *DROP_SENSITIVE_TYPE_SOURCE Procedure Parameters*

Parameter	Description
source	Name of the source

Examples

To drop all sensitive column types corresponding to an Application Data Model (ADM) from an Oracle Enterprise Manager Cloud Control instance, ADM_DEMO:

```
DBMS_TSDP_MANAGE.DROP_SENSITIVE_TYPE_SOURCE (
   source      =>  'ADM_DEMO');
```

202.3.7 IMPORT_DISCOVERY_RESULT Procedure

This procedure can be used to import sensitive columns, along with the associated sensitive types, from an external source. The external source can be an Application Data Model (ADM) instance from Oracle Enterprise Manager Cloud Control.

Syntax

```
DBMS_TSDP_MANAGE.IMPORT_DISCOVERY_RESULT (
  discovery_result         IN  CLOB,
  discovery_source         IN  VARCHAR2,
  force                    IN  FORCE DEFAULT FALSE);

DBMS_TSDP_MANAGE.IMPORT_DISCOVERY_RESULT (
  discovery_result         IN  XMLTYPE,
  discovery_source         IN  VARCHAR2,
  force                    IN  FORCE DEFAULT FALSE);
```

Parameters

Table 202-8 *IMPORT_DISCOVERY_RESULT Procedure Parameters*

Parameter	Description
discovery_result	List of sensitive columns, along with the optional list of (the definitions of) the sensitive column types in XML format (possibly as a CLOB).
discovery_source	Source of the import. The discovery_sourcename identifies the list of imported sensitive columns. In case of ADM, this should be the ADM name.
force	Specifies if the discovery result should be imported or not when the discovery result contains columns sensitive columns that are already identified as sensitive by another source. • FALSE (default) - the discovery result will not be imported in case of conflicting columns. None of the columns and the sensitive types are imported. • TRUE - the discovery result is imported and the attributes of the conflicting columns is set based on the incoming discovery result

Examples

Import the list of sensitive columns of ADM instance, ADM_Demo:

```
DBMS_TSDP_MANAGE.IMPORT_DISCOVERY_RESULT (
  discovery_results        =>  xml_adm_result,
  discovery_source         => 'ADM_Demo');
```

202.3.8 IMPORT_SENSITIVE_TYPES Procedure

This procedure imports a list of sensitive column types from a source.

Syntax

```
DBMS_TSDP_MANAGE.IMPORT_SENSITIVE_TYPES (
  sensitive_types        IN CLOB,
  source                 IN VARCHAR2);
```

```
DBMS_TSDP_MANAGE.IMPORT_SENSITIVE_TYPES (
  sensitive_types        IN XMLTYPE,
  source                 IN VARCHAR2);
```

Parameters

Table 202-9 *IMPORT_SENSITIVE_TYPES Procedure Parameters*

Parameter	Description
sensitive_types	List of sensitive column types in XML Format (possibly as a CLOB)
source	Source of the import. The source identifies the list of imported sensitive column types. In case of Application Data Model (ADM) from an Oracle Enterprise Manager Cloud Control instance, this should be the ADM name.

Examples

Import the list of sensitive column types of ADM instance, ADM_Demo:

```
DBMS_TSDP_MANAGE.IMPORT_SENSITIVE_TYPES (
  sensitive_types        =>  xml_adm_result,
  source                 => 'ADM_Demo');
```

202.3.9 REMOVE_DISCOVERY_RESULT Procedure

This procedure removes sensitive columns corresponding to an Application Data Model (ADM) from an Oracle Enterprise Manager Cloud Control instance.

Syntax

```
DBMS_TSDP_MANAGE.REMOVE_DISCOVERY_RESULT (
  discovery_source     IN  VARCHAR2);
```

Parameters

Table 202-10 *REMOVE_DISCOVERY_RESULT Procedure Parameters*

Parameter	Description
discovery_source	Source of the import. In case of ADM, this should be the ADM name, the results of which is to be removed.

Examples

Remove the sensitive columns corresponding to ADM instance, ADM_Demo:

```
DBMS_TSDP_MANAGE.REMOVE_DISCOVERY_RESULT (
  discovery_source    => 'ADM_Demo');
```

203
DBMS_TSDP_PROTECT

The `DBMS_TSDP_PROTECT` package provides an interface to configure transparent sensitive data protection (TSDP) policies in conjunction with the DBMS_TSDP_MANAGE package.

`DBMS_TSDP_PROTECT` is available with the Enterprise Edition only.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Data Structures
- Summary of DBMS_TSDP_PROTECT Subprograms

Related Topics

- DBMS_TSDP_MANAGE
  The `DBMS_TSDP_MANAGE` package provides an interface to import and manage sensitive columns and sensitive column types in the database, and is used in conjunction with the DBMS_TSDP_PROTECT package with regard to transparent sensitive data protection (TSDP) policies.

> **See Also:**
>
> *Oracle Database Security Guide*

203.1 DBMS_TSDP_PROTECT Overview

Use the `DBMS_TSDP_PROTECT` package to create transparent sensitive data protection policies, configure protection by associating the policies with sensitive types, and to enable and disable the configured protection.

Sensitive types can be added using the DBMS_TSDP_MANAGE package.

203.2 DBMS_TSDP_PROTECT Security Model

All procedures are executed with invoker's rights. Typically, a security administrator should have the `EXECUTE` privilege for this package.

203.3 DBMS_TSDP_PROTECT Constants

DBMS_TSDP_PROTECT defines the TSDP_PARAM_MAX constant for use when specifying parameter values.

This constant is described in the following table.

Table 203-1 DBMS_TDSP_PROTECT Constants - Compression Types

Constant	Type	Value	Description
TSDP_PARAM_MAX	INTEGER	4000	Maximum length of the parameter value that can be specified in FEATURE_OPTIONS

203.4 DBMS_TSDP_PROTECT Data Structures

The DBMS_TSDP_PROTECT package defines two TABLE types.

Table Types

- FEATURE_OPTIONS Table Type

  POLICY_CONDITIONS Table Type

203.4.1 FEATURE_OPTIONS Table Type

The following type is an associative array of VARCHAR2(TSDP_PARAM_MAX) that is indexed by VARCHAR2(M_IDEN).

Syntax

```
TYPE FEATURE_OPTIONS IS TABLE OF VARCHAR2(TSDP_PARAM_MAX)
INDEX BY VARCHAR2(M_IDEN);
```

203.4.2 POLICY_CONDITIONS Table Type

The following type is an associative array of VARCHAR2(TSDP_PARAM_MAX) that is indexed by PLS_INTEGER.

Syntax

```
TYPE POLICY_CONDITIONS IS TABLE OF VARCHAR2(TSDP_PARAM_MAX)
INDEX BY PLS_INTEGER;
```

203.5 Summary of DBMS_TSDP_PROTECT Subprograms

This table lists the DBMS_TSDP_PROTECT subprograms and briefly describes them.

Table 203-2 DBMS_TSDP_PROTECT Package Subprograms

Subprogram	Description
ADD_POLICY Procedure	Creates a TSDP policy
ALTER_POLICY Procedure	Alters a TDSP policy
ASSOCIATE_POLICY Procedure	Associates or disassociates a TSDP policy with a sensitive column type
DISABLE_PROTECTION_COLUMN Procedure	Disables protection for columns
DISABLE_PROTECTION_SOURCE Procedure	Disables protection based on the source of truth for the sensitive columns
DISABLE_PROTECTION_TYPE Procedure	Disables protection for a sensitive column type
DROP_POLICY Procedure	Removes a TDSP policy
ENABLE_PROTECTION_COLUMN Procedure	Enables protection for columns
ENABLE_PROTECTION_SOURCE Procedure	Enables protection based on the source of truth for the sensitive columns
ENABLE_PROTECTION_TYPE Procedure	Enables protection for a sensitive column type

203.5.1 ADD_POLICY Procedure

This procedure creates a TDSP policy.

Syntax

```
DBMS_TSDP_PROTECT.ADD_POLICY (
   policy_name              IN VARCHAR2,
   security_feature         IN PLS_INTEGER,
   policy_enable_options    IN FEATURE_OPTIONS,
   policy_apply_condition   IN POLICY_CONDITION DEFAULT TSDP$default_condition);
```

Parameters

Table 203-3 *ADD_POLICY Procedure Parameters*

Parameter	Description
`policy_name`	Name of the policy being created. The maximum length for this identifier is `M_IDEN`. This follows the Oracle naming convention.
`security_feature`	Oracle security feature with which the policy is associated. Allowed values: • `DBMS_TSDP_PROTECT.REDACT` • `DBMS_TSDP_PROTECT.VPD` • `DBMS_TSDP_PROTECT.UNIFIED_AUDIT` • `DBMS_TSDP_PROTECT.FINE_GRAINED_AUDIT` • `DBMS_TSDP_PROTECT.COLUMN_ENCRYPTION`
`policy_enable_options`	Initialized with the parameter-value pairs corresponding to the `security_feature` setting

Table 203-3 (Cont.) *ADD_POLICY Procedure Parameters*

Parameter	Description
policy_apply_condition	Initialized with the property-value pairs that must be satisfied in order to apply the corresponding policy_enable_options. This is an associative array with Property as the key (PLS_INTEGER). Example: example_policy_condition(*Property*)= *property_value*. Permissible values for Property: • DBMS_TSDP_PROPERTY.DATATYPE • DBMS_TSDP_PROPERTY.LENGTH • DBMS_TSDP_PROPERTY.PARENT_SCHEMA • DBMS_TSDP_PROPERTY.PARENT_TABLE

Usage Notes

To create the TDSP policy, you must include the procedure in an anonymous block that defines the type of security feature that will use the policy and conditions to test when the policy is enabled. For more information, see *Oracle Database Security Guide*.

Examples

Create a policy PARTIAL_MASK_POLICY:

```
DECLARE
  redact_feature_options DBMS_TSDP_PROTECT.FEATURE_OPTIONS;
  policy_conditions DBMS_TSDP_PROTECT.POLICY_CONDITIONS;
BEGIN
  redact_feature_options ('expression') :=
   'SYS_CONTEXT(''USERENV'',''SESSION_USER'')      ='APPUSER'';
  redact_feature_options ('function_type')        := 'DBMS_REDACT.PARTIAL';
  redact_feature_options ('function_parameters')  := 'STR, VVVVVVVVV,VVVVVVVVV, *,  1,
6';
  policy_conditions(DBMS_TSDP_PROTECT.DATATYPE)   := 'VARCHAR2';
 DBMS_TSDP_PROTECT.ADD_POLICY
  ('PARTIAL_MASK_POLICY',  DBMS_TSDP_PROTECT.REDACT, redact_feature_options,
policy_conditions);
END;
```

203.5.2 ALTER_POLICY Procedure

This procedure alters an existing TDSP policy

Syntax

```
DBMS_TSDP_PROTECT.ALTER_POLICY (
   policy_name              IN VARCHAR2,
   policy_enable_options    IN FEATURE_OPTIONS,
   policy_apply_condition   IN POLICY_CONDITION default TSDP$default_condition);
```

Parameters

Table 203-4 *ALTER_POLICY Procedure Parameters*

Parameter	Description
policy_name	Name of the policy to alter

Table 203-4 (Cont.) *ALTER_POLICY Procedure Parameters*

Parameter	Description
`policy_enable_options`	Initialized with the parameter-value pairs corresponding to the security feature
`policy_apply_condition`	Initialized with the property-value pairs that must be satisfied in order to apply the corresponding `policy_enable_options`. This is an associative array with Property as the key (`PLS_INTEGER`). Example: `example_policy_condition(`*`Property`*`)=`*`property_value`*. Permissible values for Property: • `DBMS_TSDP_PROPERTY.DATATYPE` • `DBMS_TSDP_PROPERTY.LENGTH` • `DBMS_TSDP_PROPERTY.PARENT_SCHEMA` • `DBMS_TSDP_PROPERTY.PARENT_TABLE`

Usage Notes

- If the `policy_apply_condition` matches an existing condition for the policy, then the corresponding enable options are updated with `policy_enable_options`.
- If the `policy_apply_condition` does not match any existing condition for the policy, the combination of `policy_enable_options` and `policy_apply_condition` is added to the policy.

Examples

Add a new combination of `policy_apply_condition` and `policy_enable_options` to an existing policy `PARTIAL_MASK_POLICY`:

```
DECLARE
  redact_feature_options DBMS_TSDP_PROTECT.FEATURE_OPTIONS;
  policy_conditions DBMS_TSDP_PROTECT.POLICY_CONDITIONS;
BEGIN
   redact_feature_options ('expression') :=
       'SYS_CONTEXT(''USERENV'',''SESSION_USER'')=''APPUSER''';
   redact_feature_options ('function_type') := 'DBMS_REDACT.PARTIAL';
   redact_feature_options ('function_parameters') := 'STR, VVVVVVVVV,VVVVVVVVV, *,
        1, 6';
   policy_conditions (DBMS_TSDP_PROTECT.DATATYPE) := 'VARCHAR2';
   DBMS_TSDP_PROTECT.ALTER_POLICY ('PARTIAL_MASK_POLICY',  redact_feature_options,
policy_conditions);
END;
```

203.5.3 ASSOCIATE_POLICY Procedure

This procedure associates or disassociates a TSDP policy with a sensitive column type.

Syntax

```
DBMS_TSDP_PROTECT.ASSOCIATE_POLICY (
   policy_name        IN  VARCHAR2,
   sensitive_type     IN  VARCHAR2,
   associate          IN  BOOLEAN DEFAULT TRUE);
```

Parameters

Table 203-5 *ASSOCIATE_POLICY Procedure Parameters*

Parameter	Description
policy_name	Name of the TDSP policy
sensitive_type	Name of the sensitive column type:
associate	Associate or Disassociate. TRUE implies Associate

Usage Notes

Both the policy and the sensitive column type should exist in the database.

Examples

Associate PARTIAL_MASK_POLICY with SSN_TYPE:

```
DBMS_TSDP_PROTECT.ASSOCIATE_POLICY ('PARTIAL_MASK_POLICY', 'SSN_TYPE');
```

203.5.4 DISABLE_PROTECTION_COLUMN Procedure

This procedure disables protection for columns.

Syntax

```
DBMS_TSDP_PROTECT.DISABLE_PROTECTION_COLUMN (
   schema_name        IN  VARCHAR2 DEFAULT '%',
   table_name         IN  VARCHAR2 DEFAULT '%',
   column_name        IN  VARCHAR2 DEFAULT '%',
   policy_name        IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 203-6 *DISABLE_PROTECTION_COLUMN Procedure Parameters*

Parameter	Description
schema_name	Name of the schema containing the column
table_name	Table containing the column
column_name	Column name
policy_name	Optional policy name. If given, only this policy is disabled.

Examples

Disable TSDP policies associated with the corresponding sensitive column types for columns that reside in schema with name like %PAYROLL%, table name like EMP%, and column name like SAL%:

```
EXEC DBMS_TSDP_PROTECT.DISABLE_PROTECTION_COLUMN ('%PAYROLL%', 'EMP%', 'SAL%');
```

203.5.5 DISABLE_PROTECTION_SOURCE Procedure

This procedure disables protection based on the source of truth for the sensitive columns.

Syntax

```
DBMS_TSDP_PROTECT.DISABLE_PROTECTION_SOURCE (
   discovery_sourcename      IN  VARCHAR2);
```

Parameters

Table 203-7 *DISABLE_PROTECTION_SOURCE Procedure Parameters*

Parameter	Description
discovery_sourcename	Name of the discovery source. This could be the Application Data Model (ADM) name or the database user.

Examples

Disable protection for all columns corresponding to `ADM_Demo`:

```
DBMS_TSDP_PROTECT.DISABLE_PROTECTION_SOURCE ('ADM_Demo');
```

203.5.6 DISABLE_PROTECTION_TYPE Procedure

This procedure disables protection for a sensitive column type.

Syntax

```
DBMS_TSDP_PROTECT.DISABLE_PROTECTION_TYPE (
   sensitive_type         IN VARCHAR2);
```

Parameters

Table 203-8 *DISABLE_PROTECTION_TYPE Procedure Parameters*

Parameter	Description
sensitive_type	Name of the sensitive column type

Examples

Disable protection for all columns identified by `SSN_TYPE`:

```
DBMS_TSDP_PROTECT.DISABLE_PROTECTION_TYPE ('SSN_TYPE');
```

203.5.7 DROP_POLICY Procedure

This procedure removes a TDSP policy or one of its condition-enable_options combinations.

Syntax

```
DBMS_TSDP_PROTECT.DROP_POLICY (
   policy_name              IN VARCHAR2,
   policy_apply_condition   IN POLICY_CONDITIONS);
```

```
DBMS_TSDP_PROTECT.DROP_POLICY (
   policy_name              IN VARCHAR2);
```

Parameters

Table 203-9 *DROP_POLICY Procedure Parameters*

Parameter	Description
policy_name	Name of the policy to drop
policy_apply_condition	To be initialized with the relevant condition

Usage Notes

- The combination of policy_conditions and policy_enable_options can be dropped from a TSDP policy by giving the policy_apply_condition parameter. The default condition-default options combination can also be dropped (if it exists for the policy) by passing an empty associative array of type DBMS_TSDP_PROTECT.POLICY_CONDITION.
- If the condition-enable_options combination that is being dropped is the last condition-enable_options combination for the policy, the policy itself is dropped.
- A policy can be completely dropped by using the overloaded of the procedure that takes only policy_name.
- A policy or one of its conditions can be dropped only if the policy is not associated with any sensitive column type. This also means that a policy that is being dropped is not enabled on any column (object).

Examples

Dropping the condition-enable_options combination based on a specific condition:

```
DECLARE
   policy_conditions DBMS_TSDP_PROTECT.POLICY_CONDITIONS;
BEGIN
   policy_conditions (DBMS_TSDP_PROTECT.DATATYPE) := 'VARCHAR2';
   DBMS_TSDP_PROTECT.DROP_POLICY ('PARTIAL_MASK_POLICY', policy_conditions);
END;
```

The default condition-enable_options combination can be dropped by passing an empty associative array of type DBMS_TSDP_PROTECT.POLICY_CONDITIONS for the policy_apply_condition parameter:

```
DECLARE
   policy_conditions DBMS_TSDP_PROTECT.POLICY_CONDITIONS;
BEGIN
  DBMS_TSDP_PROTECT.DROP_POLICY ('redact_partial_cc', policy_conditions);
END;
```

Dropping a TSDP policy:

```
BEGIN
   DBMS_TSDP_PROTECT.DROP_POLICY(
   policy_name => 'PARTIAL_MASK_POLICY');
END;
```

203.5.8 ENABLE_PROTECTION_COLUMN Procedure

This procedure enables protection for columns.

Syntax

```
DBMS_TSDP_PROTECT.ENABLE_PROTECTION_COLUMN (
   schema_name         IN  VARCHAR2 DEFAULT '%',
   table_name          IN  VARCHAR2 DEFAULT '%',
   column_name         IN  VARCHAR2 DEFAULT '%',
   policy_name         IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 203-10 *ENABLE_PROTECTION_COLUMN Procedure Parameters*

Parameter	Description
schema_name	Name of the schema containing the column
table_name	Table containing the column
column_name	Column name
policy_name	Optional policy name. If given, only this policy is enabled.

Usage Notes

- Only a TSDP Policy that is associated with the sensitive column type of the sensitive column can be enabled using this Procedure.
- LIKE condition is used for schema_name, table_name and column_name. AND semantics is followed.

Examples

Enable TSDP policies associated with the corresponding sensitive column types for columns that reside in schema with name like %PAYROLL%, table name like EMP%, and column name like SAL%:

```
DBMS_TSDP_PROTECT.ENABLE_PROTECTION_COLUMN ('%PAYROLL%', 'EMP%', 'SAL%');
```

203.5.9 ENABLE_PROTECTION_SOURCE Procedure

This procedure enables protection based on the source of truth for the sensitive columns.

Syntax

```
DBMS_TSDP_PROTECT.ENABLE_PROTECTION_SOURCE (
   discovery_sourcename      IN  VARCHAR2);
```

Parameters

Table 203-11 *ENABLE_PROTECTION_SOURCE Procedure Parameters*

Parameter	Description
discovery_sourcename	Name of the discovery source. This could be the Application Data Model (ADM) name or the database user.

Examples

Enable protection for all columns corresponding to ADM_Demo:

```
DBMS_TSDP_PROTECT.ENABLE_PROTECTION_SOURCE ('ADM_Demo');
```

203.5.10 ENABLE_PROTECTION_TYPE Procedure

This procedure enables protection for a sensitive column type.

Syntax

```
DBMS_TSDP_PROTECT.ENABLE_PROTECTION_TYPE (
   sensitive_type         IN VARCHAR2);
```

Parameters

Table 203-12 *ENABLE_PROTECTION_TYPE Procedure Parameters*

Parameter	Description
sensitive_type	Name of the sensitive column type

Examples

Enable protection for all columns identified by SSN_TYPE:

```
DBMS_TSDP_PROTECT.ENABLE_PROTECTION_TYPE ('SSN_TYPE');
```

204
DBMS_TTS

The DBMS_TTS package checks if the transportable set is self-contained. All violations are inserted into a temporary table that can be selected from the view `TRANSPORT_SET_VIOLATIONS`.

This chapter contains the following topics:

- Security Model
- Exceptions
- Operational Notes
- Summary of DBMS_TTS Subprograms

> **See Also:**
>
> - *Oracle Database Administrator's Guide*
> - *Oracle Database Upgrade Guide*

204.1 DBMS_TTS Security Model

Only users having the `execute_catalog_role` can execute this procedure. This role is initially only assigned to user SYS.

204.2 DBMS_TTS Exceptions

The `DBMS_TTS` package creates exceptions for missing or invalid transportable tablespaces.

```
ts_not_found  EXCEPTION;
PRAGMA exception_init(ts_not_found, -29304);
ts_not_found_num NUMBER := -29304;

invalid_ts_list  EXCEPTION;
PRAGMA exception_init(invalid_ts_list, -29346);
invalid_ts_list_num NUMBER := -29346;

sys_or_tmp_ts     EXCEPTION;
PRAGMA exception_init(sys_or_tmp_ts, -29351);
sys_or_tmp_ts_num NUMBER := -29351;
```

204.3 DBMS_TTS Operational Notes

With respect to transportable tablespaces, disabled and enabled referential integrity constraints are handled differently.

- A disabled referential integrity constraint does not violate the transportability rules and is dropped during the import phase.

- An enabled referential integrity constraint violates the transportability rules if it references a table in a tablespace outside the transportable set.

204.4 Summary of DBMS_TTS Subprograms

The two procedures listed in the table are designed to be called by database administrators.

Table 204-1 DBMS_TTS Package Subprograms

Subprogram	Description
DOWNGRADE Procedure	Downgrades transportable tablespace-related data
TRANSPORT_SET_CHECK Procedure	Checks if a set of tablespaces (to be transported) is self-contained

204.4.1 DOWNGRADE Procedure

This procedure downgrades transportable tablespace related data in releases earlier than Oracle Database 10g.

Syntax

```
DBMS_TTS.DOWNGRADE;
```

> **Note:**
>
> This procedure has no effect in Oracle Database 10g and later releases. It is included only for backward compatibility to support references to it in existing scripts.

204.4.2 TRANSPORT_SET_CHECK Procedure

This procedure checks if a set of tablespaces (to be transported) is self-contained. After calling this procedure, the user may select from a view to see a list of violations, if there are any.

Syntax

```
DBMS_TTS.TRANSPORT_SET_CHECK (
   ts_list          IN CLOB,
   incl_constraints IN BOOLEAN DEFAULT FALSE,
   full_check       IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 204-2 TRANSPORT_SET_CHECK Procedure Parameters

Parameter	Description
`ts_list`	List of one or more tablespaces, separated by comma.
`incl_constraints`	`TRUE` if you want to count in referential integrity constraints when examining if the set of tablespaces is self-contained. (The `incl_constraints` parameter is a default so that `TRANSPORT_SET_CHECK` will work if it is called with only the `ts_list` argument.)

Table 204-2 (Cont.) TRANSPORT_SET_CHECK Procedure Parameters

Parameter	Description
full_check	Indicates whether a full or partial dependency check is required. If TRUE, treats all IN and OUT pointers (dependencies) and captures them as violations if they are not self-contained in the transportable set. The parameter should be set to TRUE for TSPITR or if a strict version of transportable is desired. By default the parameter is set to FALSE. It will only consider OUT pointers as violations.

Examples

If the view does not return any rows, then the set of tablespaces is self-contained. For example,

```
SQLPLUS> EXECUTE DBMS_TTS.TRANSPORT_SET_CHECK('foo,bar', TRUE);
SQLPLUS> SELECT * FROM TRANSPORT_SET_VIOLATIONS;
```

205
DBMS_TYPES

The DBMS_TYPES package consists of constants, which represent the built-in and user-defined types.

This chapter contains the following topics:

- Constants
- Exceptions

205.1 DBMS_TYPES Constants

The DBMS_TYPES package defines several constants to use when specifying parameter values.

These constants are listed in following table .

Table 205-1 ***DBMS_TYPES Constants***

Constant	Description
NO_DATA	Is only relevant if PieceWise is called, for a collection or any dataset. Denotes the end of collection/any dataset when all the elements have been accessed
SUCCESS	The operation succeeded
TYPECODE_BDOUBLE	A NUMBER type
TYPECODE_BFILE	A BFILE type
TYPECODE_BFLOAT	A NUMBER type
TYPECODE_BLOB	A BLOB type
TYPECODE_CFILE	A CFILE type
TYPECODE_CHAR	A CHAR type
TYPECODE_CLOB	A CLOB type
TYPECODE_DATE	A DATE type
TYPECODE_INTERVAL_DS	An INTERVAL_DS type
TYPECODE_INTERVAL_YM	A INTERVAL_YM type
TYPECODE_MLSLABEL	An MLSLABEL type
TYPECODE_NAMEDCOLLECTION	A named collection (VARRAY/nested table) type
TYPECODE_NCHAR	A NCHAR type
TYPECODE_NCLOB	A NCLOB type
TYPECODE_NUMBER	A NUMBER type
TYPECODE_NVARCHAR2	A NVARCHAR2 type
TYPECODE_OBJECT	An OBJECT type
TYPECODE_OPAQUE	An OPAQUE type

Table 205-1 (Cont.) *DBMS_TYPES Constants*

Constant	Description
TYPECODE_RAW	A RAW type
TYPECODE_REF	A REF type
TYPECODE_TABLE	A nested table collection type
TYPECODE_TIMESTAMP	A TIMESTAMP type
TYPECODE_TIMESTAMP_LTZ	A TIMESTAMP_LTZ type
TYPECODE_TIMESTAMP_TZ	A TIMESTAMP_TZ type
TYPECODE_UROWID	A UROWID type
TYPECODE_VARCHAR2	A VARCHAR2 type
TYPECODE_VARCHAR	A VARCHAR type
TYPECODE_VARRAY	A VARRAY collection type

205.2 DBMS_TYPES Exceptions

DBMS_TYPES throws these exceptions.

- INVALID_PARAMETERS
- INCORRECT_USAGE
- TYPE_MISMATCH

206
DBMS_UMF

The `DBMS_UMF` package provides an interface for deploying the Remote Management Framework (RMF) for an Oracle Database. The RMF is used for collecting performance statistics for an Oracle Database.

> **See Also:**
>
> *Oracle Database Performance Tuning Guide* for more information about configuring the RMF for an Oracle Database.

This chapter contains the following topic:

- Summary of DBMS_UMF Subprograms

206.1 Summary of DBMS_UMF Subprograms

This topic lists the `DBMS_UMF` subprograms in alphabetical order and briefly describes them.

Table 206-1 DBMS_UMF Package Subprograms

Subprogram	Description
CONFIGURE_NODE Procedure	Configures a node in the RMF topology
CREATE_LINK Procedure	Creates a database link between two nodes in the RMF topology
CREATE_TOPOLOGY Procedure	Creates the RMF topology
DROP_LINK Procedure	Removes a database link between two nodes in the RMF topology
DROP_TOPOLOGY Procedure	Deletes the RMF topology
ENABLE_SERVICE Procedure	Enables a service on a node in the RMF topology
GET_NODE_ID_LOCAL Function	Returns the node ID of a node in the RMF topology
GET_NODE_NAME_LOCAL Function	Returns the node name of a node in the RMF topology
GET_TARGET_ID Function	Returns the destination ID in the RMF topology
GET_TOPOLOGY_NAME_LOCAL Function	Returns the RMF topology name of the local node
QUERY_LINK_INFO Procedure	Returns the information about a database link in the RMF topology
QUERY_NODE_INFO Procedures	Returns the information about a node in the RMF topology
REGISTER_NODE Function and Procedure	Registers a node in the RMF topology
SWITCH_DESTINATION Procedure	Designates a source node as a destination node in the RMF topology

Table 206-1 (Cont.) DBMS_UMF Package Subprograms

Subprogram	Description
UNCONFIGURE_NODE Procedure	Resets the configuration of a node in the RMF topology
UNREGISTER_NODE Procedure	Removes the registration of a node in the RMF topology

206.1.1 CONFIGURE_NODE Procedure

This procedure configures a node that needs to be registered with the RMF topology. This procedure must be executed on the node that needs to be configured.

Syntax

```
DBMS_UMF.CONFIGURE_NODE(
   node_name          IN VARCHAR2 DEFAULT NULL,
   dblink_to_target   IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 206-2 CONFIGURE_NODE Procedure Parameters

Parameter	Description
node_name	Name of the node. Each node in the RMF topology must be assigned a unique name. If a name is not provided for a node in this procedure, then the value of the initialization parameter DB_UNIQUE_NAME is assigned as the name for the node by default.
dblink_to_target	Database link from this node to the destination node.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.2 CREATE_LINK Procedure

This procedure creates database links between two nodes in the RMF topology. This procedure must be executed only on the destination node.

Syntax

```
DBMS_UMF.CREATE_LINK(
   topology_name      IN VARCHAR2,
   node_a_name        IN VARCHAR2,
   node_b_name        IN VARCHAR2,
   dblink_a_to_b      IN VARCHAR2,
   dblink_b_to_a      IN VARCHAR2);
```

Parameters

Table 206-3 CREATE_LINK Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology.
node_a_name	Name of the first node.
node_b_name	Name of the second node.
dblink_a_to_b	Database link from the first node to the second node.
dblink_b_to_b	Database link from the second node to the first node.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.3 CREATE_TOPOLOGY Procedure

This procedure creates the RMF topology and designates the node on which it is executed as the destination node for that topology.

Syntax

```
DBMS_UMF.CREATE_TOPOLOGY(
   topology_name IN VARCHAR2);
```

Parameters

Table 206-4 CREATE_TOPOLOGY Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.4 DROP_LINK Procedure

This procedure removes the database links between two nodes in the RMF topology. This procedure must be executed only on the destination node.

Syntax

```
DBMS_UMF.DROP_LINK(
   topology_name  IN VARCHAR2,
   node_a_name    IN VARCHAR2,
   node_b_name    IN VARCHAR2);
```

Parameters

Table 206-5 DROP_LINK Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology.
node_a_name	Name of the first node.
node_b_name	Name of the second node.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.5 DROP_TOPOLOGY Procedure

This procedure deletes the RMF topology. This procedure must be executed only on the destination node.

Syntax

```
DBMS_UMF.DROP_TOPOLOGY(
   topology_name IN VARCHAR2);
```

Parameters

Table 206-6 DROP_TOPOLOGY Procedure Parameters

Parameter	Description
topology_name	Name of the topology to delete.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.6 ENABLE_SERVICE Procedure

This procedure enables a service, such as the AWR service, on a node in the RMF topology. This procedure must be executed only on the destination node.

Syntax

```
DBMS_UMF.ENABLE_SERVICE(
   topology_name      IN VARCHAR2,
   node_name          IN VARCHAR2,
   service_type       IN NUMBER);
```

Parameters

Table 206-7 ENABLE_SERVICE Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology.
node_name	Name of the node on which a specific service needs to be enabled.
service_type	Numeric constant identifying the service. The only allowed value for this parameter is UMF_SERVICE_TYPE_AWR, which is Automatic Workload Repository (AWR) service.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.7 GET_NODE_ID_LOCAL Function

This function returns the node ID of the node in the RMF topology on which this function is executed.

Syntax

```
DBMS_UMF.GET_NODE_ID_LOCAL(
   topology_name  IN  VARCHAR2 DEFAULT NULL)
 RETURN NUMBER;
```

Parameters

Table 206-8 GET_NODE_ID_LOCAL Function Parameters

Parameter	Description
topology_name	Name of the RMF topology with which the node is registered.

Return Value

Returns the node ID of the node in the RMF topology on which this function is executed.

Security Model

Only the database users SYS and SYS$UMF can execute this function.

206.1.8 GET_NODE_NAME_LOCAL Function

This function returns the name of the node in the RMF topology on which this function is executed.

Syntax

```
DBMS_UMF.GET_NODE_NAME_LOCAL RETURN VARCHAR2;
```

Return Value

Returns the name of the node in the RMF topology on which this function is executed.

Security Model

Only the database users `SYS` and `SYS$UMF` can execute this procedure.

206.1.9 GET_TARGET_ID Function

This function returns the ID of the destination node in the RMF topology. This function can be executed on any node in the RMF topology.

Syntax

```
DBMS_UMF.GET_TARGET_ID(
   topology_name  IN  VARCHAR2)
 RETURN NUMBER;
```

Parameters

Table 206-9 GET_TARGET_ID Function Parameters

Parameter	Description
`topology_name`	Name of the RMF topology.

Return Value

Returns the ID of the destination node in the RMF topology.

Security Model

Only the database users `SYS` and `SYS$UMF` can execute this procedure.

206.1.10 GET_TOPOLOGY_NAME_LOCAL Function

This function returns the name of the active RMF topology of the node on which this function is executed.

Syntax

```
DBMS_UMF.GET_TOPOLOGY_NAME_LOCAL RETURN VARCHAR2;
```

Return Value

Returns the name of the active RMF topology of the node on which this function is executed.

Security Model

Only the database users `SYS` and `SYS$UMF` can execute this procedure.

206.1.11 QUERY_LINK_INFO Procedure

This procedure returns the name of the database link between two nodes in the RMF topology. This procedure can be executed on any node in the RMF topology.

Syntax

```
DBMS_UMF.QUERY_LINK_INFO(
   topology_name      IN   VARCHAR2,
   from_node_id       IN   NUMBER,
   to_node_id         IN   NUMBER,
   link_name          OUT  VARCHAR2);
```

Parameters

Table 206-10 QUERY_LINK_INFO Procedure Parameters

Parameter	Description
`topology_name`	Name of the RMF topology.
`from_node_id`	Node ID of the first node.
`to_node_id`	Node ID of the second node.
`link_name`	Name of the database link from the first node to the second node returned by the procedure.

Security Model

Only the database users `SYS` and `SYS$UMF` can execute this procedure.

206.1.12 QUERY_NODE_INFO Procedures

This procedure returns information about a node in the RMF topology. This procedure can be executed on any node in the RMF topology.

Syntax

```
DBMS_UMF.QUERY_NODE_INFO(
   topology_name      IN    VARCHAR2  DEFAULT NULL,
   node_name          IN    VARCHAR2,
   node_id            OUT   NUMBER);

DBMS_UMF.QUERY_NODE_INFO(
   node_id            IN    NUMBER,
```

```
   topology_name       OUT   VARCHAR2,
   node_name           OUT   VARCHAR2);
```

Parameters

Table 206-11 QUERY_NODE_INFO Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology with which the node is registered.
node_name	Name of the node.
node_id	Identifier of the node.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.13 REGISTER_NODE Function and Procedure

This function and procedure registers a node with the RMF topology. This procedure and function must be executed only on the destination node in the RMF topology.

Syntax

```
DBMS_UMF.REGISTER_NODE(
   topology_name           IN  VARCHAR2,
   node_name               IN  VARCHAR2,
   dblink_to_node          IN  VARCHAR2 DEFAULT NULL,
   dblink_from_node        IN  VARCHAR2 DEFAULT NULL,
   as_source               IN  VARCHAR2 DEFAULT 'TRUE',
   as_candidate_target     IN  VARCHAR2 DEFAULT 'FALSE');

DBMS_UMF.REGISTER_NODE(
   topology_name           IN  VARCHAR2,
   node_name               IN  VARCHAR2,
   dblink_to_node          IN  VARCHAR2 DEFAULT NULL,
   dblink_from_node        IN  VARCHAR2 DEFAULT NULL,
   as_source               IN  VARCHAR2 DEFAULT 'TRUE',
   as_candidate_target     IN  VARCHAR2 DEFAULT 'FALSE');
   node_id                 OUT VARCHAR2);

DBMS_UMF.REGISTER_NODE(
   topology_name           IN  VARCHAR2,
   node_name               IN  VARCHAR2,
   dblink_to_node          IN  VARCHAR2 DEFAULT NULL,
   dblink_from_node        IN  VARCHAR2 DEFAULT NULL,
   as_source               IN  VARCHAR2 DEFAULT 'TRUE',
   as_candidate_target     IN  VARCHAR2 DEFAULT 'FALSE')
 RETURN NUMBER;
```

Parameters

Table 206-12 REGISTER_NODE Function and Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology.
node_name	Name of the node to register.
dblink_to_node	Name for the database link from the destination to the node.
dblink_from_node	Name for the database link from the node to the destination.
as_source	Set to TRUE, if the node is a source, else set to FALSE.
as_candidate_target	Set to TRUE, if the node is a candidate destination, else set to FALSE.
node_id	Node ID returned by the procedure.

Return Value

Returns the node ID of the registered node.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.14 SWITCH_DESTINATION Procedure

This procedure makes the candidate destination as the new destination in the RMF topology. This procedure must be executed only on the candidate destination node.

Syntax

```
DBMS_UMF.SWITCH_DESTINATION(
   topology_name      IN VARCHAR2,
   force_switch       IN BOOLEAN DEFAULT TRUE);
```

Parameters

Table 206-13 SWITCH_DESTINATION Procedure Parameters

Parameter	Description
topology_name	Name of the RMF topology.

Table 206-13 (Cont.) SWITCH_DESTINATION Procedure Parameters

Parameter	Description
force_switch	If set to FALSE, the execution of this procedure fails in the following situations: • Candidate destination is a read-only database, and hence it cannot become the new destination. • Candidate destination does not have database links to one or more sources in the topology. • Candidate destination is unable to get the latest AWR data from the old destination. If set to TRUE, the execution of this procedure fails in the following situation: • Candidate destination is a read-only database, and hence it cannot become the new destination.

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.15 UNCONFIGURE_NODE Procedure

This procedure removes the configuration details of the node on which this procedure is executed.

Syntax

```
DBMS_UMF.UNCONFIGURE_NODE;
```

Security Model

Only the database users SYS and SYS$UMF can execute this procedure.

206.1.16 UNREGISTER_NODE Procedure

This procedure removes a node for the RMF topology. This procedure must be executed only on the destination node.

Syntax

```
DBMS_UMF.UNREGISTER_NODE(
   topology_name    IN VARCHAR2,
   node_name        IN VARCHAR2);
```

Parameters

Table 206-14 UNREGISTER_NODE Procedure Parameters

Parameter	Description
`topology_name`	Name of the RMF topology.
`node_name`	Name of the node which needs to be removed from the topology.

Security Model

Only the database users `SYS` and `SYS$UMF` can execute this procedure.

207
DBMS_UNDO_ADV

The Undo Advisor assists in correctly sizing the undo tablespace and sets the low threshold value of the Undo Retention Period for any Oracle Flashback requirements. The Undo Advisor can also be used to estimate the undo tablespace required for migration from manual to automatic undo management before actually creating the new undo tablespace.

The `DBMS_UNDO_ADVISOR` package provides subprograms to manage the execution of the Undo Advisor feature.

The Undo Advisor relies for its analysis on data collected in the Automatic Workload Repository (AWR). Thus, it is important that the AWR have adequate workload statistics available so that the Undo Advisor can make accurate recommendations. For newly created databases, adequate statistics may not be available immediately. In such cases, continue to use the default auto-extending undo tablespace until at least one workload cycle completes.

This chapter contains the following topic:

- Summary of DBMS_UNDO_ADV Subprograms

207.1 Summary of DBMS_UNDO_ADV Subprograms

This topic lists the `DBMS_UNDO_ADV` subprograms in alphabetical order and briefly describes them.

Table 207-1 DBMS_UNDO_ADV Package Subprograms

Subprogram	Description
BEST_POSSIBLE_RETENTION Function	Returns the best possible value for the `undo_retention` parameter that the current undo tablespace can satisfy in order to maximize the usage of the current undo tablespace based on the historical information of the given period.
LONGEST QUERY Function	This function returns the duration of the longest query, in seconds, for a given period.
RBU_MIGRATION Function	Estimates the undo tablespace needed for migration from manual to automatic undo management.
REQUIRED_RETENTION Function	Returns the required value for the `undo_retention` parameter to satisfy the longest query based on undo statistics available for a given period.
REQUIRED_UNDO_SIZE Function	Returns the required undo tablespace size in MB to satisfy a certain undo retention value based on undo statistics available for a given period.
UNDO_ADVISOR Function	Uses the advisor framework to check if there is any problem with the current instance. It also provides recommendations.
UNDO_AUTOTUNE Function	Determines whether auto-tuning of undo retention is enabled for the current undo tablespace.

Table 207-1 (Cont.) DBMS_UNDO_ADV Package Subprograms

Subprogram	Description
UNDO_HEALTH Function	Checks whether there is any problem with the current setting of undo retention and undo tablespace size based on the historical information of a given period and provides recommendations to fix the problem.
UNDO_INFO Function	Retrieves information about the undo tablespace of the current instance.

207.1.1 BEST_POSSIBLE_RETENTION Function

This function returns the best possible value for the `undo_retention` parameter that the current undo tablespace can satisfy in order to maximize the usage for the current undo tablespace based on the historical information of a given period.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.BEST_POSSIBLE_RETENTION()
RETURN NUMBER;
```

Viewing the output using start time and end time:

```
DBMS_UNDO_ADV.BEST_POSSIBLE_RETENTION(
    start_time IN DATE,
    end_time   IN DATE)
RETURN NUMBER;
```

Viewing the output using begin and end AWR snapshot ID:

```
DBMS_UNDO_ADV.BEST_POSSIBLE_RETENTION(
    begin_snap IN NUMBER,
    end_snap   IN NUMBER)
RETURN NUMBER;
```

Parameters

Table 207-2 BEST_POSSIBLE_RETENTION Function Parameters

Parameter	Description
`start_time`	Start time of the given period.
`end_time`	End time of the given period.
`begin_snap`	Begin snapshot identifier. It is based on historical information in AWR from the `begin_snap` identifier.
`end_snap`	End snapshot identifier. It is based on historical information in AWR until the `end_snap` identifier.

207.1.2 RBU_MIGRATION Function

This function estimates the undo tablespace needed for migration from manual to automatic undo management.

If you are currently using manual undo management (rollback segments) to manage undo space, then Oracle recommends migrating to automatic undo management. You must first create an undo tablespace, before opening a newly upgraded database. The required size of undo tablespace depends on the system workload and Flashback requirements.

> **Note:**
>
> The `RBU_MIGRATION` function should be called only when `undo_management = manual`.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.RBU_MIGRATION()
RETURN NUMBER;
```

Viewing the output using the Start Time and End Time:

```
DBMS_UNDO_ADV.RBU_MIGRATION(
    START_TIME IN DATE,
    END_TIME   IN DATE)
RETURN NUMBER;
```

Parameters

Table 207-3 RBU_MIGRATION Function Parameters

Parameter	Description
`START_TIME`	Start time of the given period.
`END_TIME`	End time of the given period.

207.1.3 LONGEST_QUERY Function

This function returns the duration of the longest query, in seconds, for a given period.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.LONGEST_QUERY()
RETURN NUMBER;
```

Viewing the output using start time and end time:

```
DBMS_UNDO_ADV.LONGEST_QUERY(
    start_time IN DATE,
    end_time   IN DATE)
RETURN NUMBER;
```

Viewing the output using begin and end AWR snapshot ID:

```
DBMS_UNDO_ADV.LONGEST_QUERY(
    begin_snap IN NUMBER,
    end_snap   IN NUMBER)
RETURN NUMBER;
```

Parameters

Table 207-4 LONGEST_QUERY Function Parameters

Parameter	Description
`start_time`	Start time for the specified period.
`end_time`	End time for the specified period.
`begin_snap`	Begin snapshot identifier. It is based on historical information in AWR from the `begin_snap` identifier.
`end_snap`	End snapshot identifier. It is based on historical information in AWR until the `end_snap` identifier.

207.1.4 REQUIRED_RETENTION Function

This function returns the required value for the `undo_retention` parameter to satisfy the longest query based on undo statistics available for a given period.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.REQUIRED_RETENTION()
RETURN NUMBER;
```

Viewing the output using start time and end time:

```
DBMS_UNDO_ADV.REQUIRED_RETENTION(
    start_time IN DATE,
    end_time   IN DATE)
RETURN NUMBER;
```

Viewing the output using begin and end AWR snapshot ID:

```
DBMS_UNDO_ADV.REQUIRED_RETENTION(
    begin_snap IN NUMBER,
```

```
    end_snap    IN NUMBER)
RETURN NUMBER;
```

Parameters

Table 207-5 REQUIRED_RETENTION Function Parameters

Parameter	Description
start_time	Start time of the given period.
end_time	End time of the given period.
begin_snap	Begin snapshot identifier. It is based on historical information in AWR from the begin_snap identifier.
end_snap	End snapshot identifier. It is based on historical information in AWR until the end_snap identifier.

207.1.5 REQUIRED_UNDO_SIZE Function

This function returns the required undo tablespace size (in MB) to satisfy certain undo retention value based on undo statistics available for a given period.

> **Note:**
>
> Zero will be returned if the information about the given period is not available.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.REQUIRED_UNDO_SIZE(
    retention IN NUMBER)
RETURN NUMBER;
```

Viewing the output using start time and end time:

```
DBMS_UNDO_ADV.REQUIRED_UNDO_SIZE(
    retention  IN NUMBER,
    start_time IN DATE,
    end_time   IN DATE)
RETURN NUMBER;
```

Viewing the output using begin and end AWR snapshot ID:

```
DBMS_UNDO_ADV.REQUIRED_UNDO_SIZE(
    retention  IN NUMBER,
    begin_snap IN NUMBER,
    end_snap   IN NUMBER)
RETURN NUMBER;
```

Parameters

Table 207-6 REQUIRED_UNDO_SIZE Function Parameters

Parameter	Description
retention	Retention value you want to set for the undo_retention init.ora parameter.
start_time	Start time of the given period.
end_time	End time of the given period.
begin_snap	Begin snapshot identifier. It is based on historical information in AWR from the begin_snap identifier.
end_snap	End snapshot identifier. It is based on historical information in AWR until the end_snap identifier.

207.1.6 UNDO_ADVISOR Function

This function uses the advisor framework to check if there is any problem with the current instance and provide recommendations.

> **Note:**
>
> This function should be used when undo_management is set to auto.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.UNDO_ADVISOR(
    instance_id IN NUMBER)
RETURN VARCHAR2;
```

Viewing the output using start time and end time:

```
DBMS_UNDO_ADV.UNDO_ADVISOR(
    start_time  IN DATE,
    end_time    IN DATE,
    instance_id IN NUMBER)
RETURN VARCHAR2;
```

Viewing the output using begin and end AWR snapshot ID:

```
DBMS_UNDO_ADV.UNDO_ADVISOR(
    begin_snap  IN NUMBER,
    end_snap    IN NUMBER,
    instance_id IN NUMBER)
RETURN VARCHAR2;
```

Parameters

Table 207-7 UNDO_ADVISOR Function Parameters

Parameter	Description
start_time	Start time of the given period.
end_time	End time of the given period.
begin_snap	Begin snapshot identifier. It is based on historical information in AWR from the begin_snap identifier.
end_snap	End snapshot identifier. It is based on historical information in AWR until the end_snap identifier.
instance_id	Instance ID of the current instance.

207.1.7 UNDO_AUTOTUNE Function

This function finds out whether the auto-tuning of undo retention is enabled for the current undo tablespace.

Syntax

```
DBMS_UNDO_ADV.UNDO_AUTOTUNE(
     chk OUT BOOLEAN)
RETURN BOOLEAN;
```

Parameters

Table 207-8 UNDO_AUTOTUNE Function Parameters

Parameter	Description
chk	TRUE if auto-tuning of undo retention is enabled, FALSE otherwise.

207.1.8 UNDO_HEALTH Function

Checks whether there is any problem with the current setting of undo retention and undo tablespace size based on the historical information of a given period and provides recommendations to fix the problem.

If the return value is 0, no problem is found. Otherwise, parameter prob and reco are the problem and recommendation on fixing the problem.

Syntax

Viewing the output using the historical information in memory:

```
DBMS_UNDO_ADV.UNDO_HEALTH(
    prob OUT VARCHAR2,
    reco OUT VARCHAR2,
    rtn1 OUT VARCHAR2,
```

```
    retn OUT NUMBER,
    utbs OUT NUMBER);
```

Viewing the output using start time and end time:

```
DBMS_UNDO_ADV.UNDO_HEALTH(
    prob OUT VARCHAR2,
    reco OUT VARCHAR2,
    rtn1 OUT VARCHAR2,
    retn OUT NUMBER,
    utbs OUT NUMBER)
RETURN NUMBER;
```

Viewing the output using begin and end AWR snapshot ID:

```
DBMS_UNDO_ADV.UNDO_HEALTH(
    begin_snap IN  NUMBER,
    end_snap   IN  NUMBER,
    prob       OUT VARCHAR2,
    reco       OUT VARCHAR2,
    rtn1       OUT VARCHAR2,
    retn       OUT NUMBER,
    utbs       OUT NUMBER)
RETURN NUMBER;
```

Parameters

Table 207-9 UNDO_HEALTH Function Parameters

Parameter	Description
start_time	Start time of the given period.
end_time	End time of the given period.
begin_snap	Begin snapshot identifier. It is based on historical information in AWR from the begin_snap identifier.
end_snap	End snapshot identifier. It is based on historical information in AWR until the end_snap identifier.
prob	Problem that is being diagnosed. For example, long running query may fail or undo tablespace cannot satisfy undo_retention.
reco	Recommendation for fixing the problem.
rtn1	Rationale for the recommendation.
retn	The numerical value of retention if the recommendation is to change retention.
utbs	The numerical value of undo tablespace size (in MB) if the recommendation is to change undo tablespace size.

207.1.9 UNDO_INFO Function

This function retrieves information about the undo tablespace of the current instance.

This function returns the undo tablespace name, maximum possible size for the undo tablespace, current undo retention value, and boolean values to verify if it is auto-extensible or if the undo tablespace has guaranteed undo retention.

Syntax

```
DBMS_UNDO_ADV.UNDO_INFO(
    tbs_name       OUT VARCHAR2,
    tbs_size       OUT NUMBER,
    tbs_autoextend OUT BOOLEAN,
    tbs_retention  OUT NUMBER,
    tbs_guarantee  OUT BOOLEAN)
RETURN BOOLEAN;
```

Parameters

Table 207-10 UNDO_INFO Function Parameters

Parameter	Description
tbs_name	Name of the current undo tablespace the instance is using.
tbs_size	The size of the undo tablespace in MB, if the undo tablespace is fixed-sized. If the tablespace is auto_extensible, it is the maximum possible size of the undo tablespace in MB.
tbs_autoextend	TRUE if the undo tablespace is extensible, FALSE otherwise.
tbs_retention	The value of the undo_retentioninit.ora parameter.
tbs_guarantee	TRUE if the undo tablespace has guaranteed retention, FALSE otherwise.

208
DBMS_USER_CERTS

The DBMS_USER_CERTS package allows you add and delete certificates.

This chapter contains the following topics:

- DBMS_USER_CERTS Overview
- DBMS_USER_CERTS Security Model
- Summary of DBMS_USER_CERTS Subprograms

> **See Also:**
>
> - *Oracle Database Administrator's Guide*
> - *Oracle Database Concepts*
> - *Oracle Database SQL Language Reference*
> - *Oracle Database Reference*
> - For information on hidden columns in blockchain tables, see *Hidden Columns in Blockchain Tables*

208.1 DBMS_USER_CERTS Overview

The DBMS_USER_CERTS package allows you add and delete certificates.

The DBMS_USER_CERTS package allows you to:

- add certificates
- delete certificates

One use case is to attach a digital signature to a row in a blockchain table and be able to verify the digital signature later. A certificate used to sign one or more rows in blockchain tables and which has since expired should remain in the database until all those rows have been deleted from the blockchain tables.

208.2 DBMS_USER_CERTS Security Model

The DBMS_USER_CERTS package is owned by SYS and is installed as part of database installation.

- Both ADD_CERTIFICATE() and DROP_CERTIFICATE() sub-programs can be executed by any database user. The EXECUTE privilege on DBMS_USER_CERTS package is granted to PUBLIC role.
- Using ADD_CERTIFICATE(), a database user can add multiple certificates, each one would be identified uniquely using a GUID.

- Only SYSDBA or the user who owns the certificate, can drop it using DROP_CERTIFICATE().

208.3 Summary of DBMS_USER_CERTS Subprograms

The DBMS_USER_CERTS package uses ADD_CERTIFICATE, ADD_COPY, and DROP_CERTIFICATE subprograms to add, copy, and delete X.509 certificates which are used for signature verification for blockchain tables by the current user.

Table 208-1 DBMS_USER_CERTS Package Subprograms

Subprogram	Description
ADD_CERTIFICATE Procedure	Adds X.509 certificates which are used for signature verification of blockchain tables.
ADD_COPY Procedure	Adds a certificate to the database and assign the certificate a specific global unique identifier (GUID).
DROP_CERTIFICATE Procedure	Drops a certificate that is used for signature verification of blockchain tables.

208.3.1 ADD_CERTIFICATE Procedure

This procedure can be used by the current user to add an X.509 certificate that is used for signature verification of blockchain tables.

Syntax

```
DBMS_USER_CERTS.ADD_CERTIFICATE(
   x509_cert                 IN  BLOB,
   cert_id                   OUT RAW);
```

Parameters

Table 208-2 ADD_CERTIFICATE Procedure Parameters

Parameter	Description
x509_cert	The X.509 certificate used for signature verification of blockchain tables.
cert_id	The Global Unique Identifier (GUID) for the certificate.

208.3.2 ADD_COPY Procedure

This procedure enables you add a certificate to the database and assign the certificate a specific global unique identifier (GUID).

This procedure is used when a user needs to make a copy of the certificate that was previously created in one database and add the copy to another database while preserving its GUID. When recording blockchain table signatures in one database and verifying the signatures in another, the certificate GUID must be preserved. One scenario is when you use Oracle Data Pump to copy a blockchain table between databases. Another scenario is when you use Oracle GoldenGate to replicate rows in a blockchain table between databases.

Syntax

```
DBMS_USER_CERTS.ADD_COPY(
   x509_cert IN   BLOB,
```

```
   cert_id   IN   RAW,
   username  IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 208-3 ADD_COPY Procedure Parameters

Parameter	Description
x509_cert	The X.509 certificate used for signature verification of blockchain tables.
cert_id	The Global Unique Identifier (GUID) for the certificate.
username	The name of the user who will own the certificate. A NULL value defaults to the current user.

Usage Note

If username is not NULL, either username must refer to the current user, or the current user must be SYSDBA.

208.3.3 DROP_CERTIFICATE Procedure

This procedure can be used by the current user to drop a certificate that is used for signature verification of blockchain tables.

Syntax

```
DBMS_USER_CERTS.DROP_CERTIFICATE(
   cert_id                    IN  RAW);
```

Parameters

Table 208-4 DROP_CERTIFICATE Procedure Parameters

Parameter	Description
cert_id	The Global Unique Identifier (GUID) of the certificate.

209
DBMS_USERDIAG

The DBMS_USERDIAG package allows you to perform a limited set of diagnosis operations on the PDB.

This chapter contains the following topics:

- DBMS_USERDIAG Overview
- Summary of DBMS_USERDIAG Subprograms

Enabling traces on cloud requires additional configuration as documented here: Perform SQL Tracing on Autonomous Database.

209.1 DBMS_USERDIAG Overview

The DBMS_USERDIAG package is for diagnosis and allows you to set up a trace within a PDB.

The DBMS_USERDIAG package utilizes underneath DBMS_SYSTEM functionality but with a narrow set of definitions which restrict arbitrary event settings.

For a given PDB, the DBMS_USERDIAG package allows you:

- to enable the SQL trace at a given level.
- to disable the SQL trace.
- to check the SQL trace.

Most of the regular diagnostic mechanisms have been restricted outside of a given PDB using lockdown profiles, so that arbitrary events cannot be enabled from user sessions in a shared tenancy in CBD deployments in cloud instances. In particular, "alter session set events" statement is blocked in cloud deployments, because it can be misused to set events/actions which may change code-path execution or simulate errors.

209.2 Summary of DBMS_USERDIAG Subprograms

The DBMS_USERDIAG package uses ENABLE_SQL_TRACE_EVENT and CHECK_SQL_TRACE_EVENT subprograms to enable, disable, and monitor sql_trace events.

Table 209-1 DBMS_USERDIAG Package Subprograms

Subprogram	Description
CHECK_SQL_TRACE_EVENT Procedure	Checks the current sql_trace event and retrieves the level.
ENABLE_SQL_TRACE_EVENT Procedure	Enables sql_trace event at a given level in user session, and generates SQL traces into respective process trace files.
GET_CALL_ERROR_MSG Function	Obtains the error message if the last call to DBMS_USERDIAG API returned an error.
GET_CALL_STATUS Function	Obtains the status of the last call to the DBMS_USERDIAG API.

Table 209-1 (Cont.) DBMS_USERDIAG Package Subprograms

Subprogram	Description
SET_EXCEPTION_MODE Procedure	Raises an exception on any error or silently ignores the same (default).
SET_TRACEFILE_IDENTIFIER Procedure	Sets a custom trace file identifier for the active trace file in the current ADR home.
TRACE Procedure	Traces message to the trace file or alert log.

209.2.1 CHECK_SQL_TRACE_EVENT Procedure

This procedure can be used by the current user to check the current `sql_trace` event and retrieves the level.

Syntax

```
DBMS_USERDIAG.CHECK_SQL_TRACE_EVENT(
   sql_id                IN VARCHAR2       DEFAULT NULL
   sys                   IN BINARY_INTEGER DEFAULT NULL
   level                 OUT BINARY_INTEGER);
```

Parameters

Table 209-2 CHECK_SQL_TRACE_EVENT Procedure Parameters

Parameter	Description
`level`	Required. Output of the current level at which `sql_trace` event is enabled.
`sql_id`	`sql_trace` in scope of a given `sql_id` alone. This is currently ignored.
`sys`	Optional. Set event instance-wide or PDB-wide. Default 0. This is currently ignored.

209.2.2 ENABLE_SQL_TRACE_EVENT Procedure

This procedure enables `sql_trace` event (`sql_trace/10046`) at a given `level` in user session, and generates SQL traces into respective process trace files.

Optionally a `sql_id` can also be specified in which case tracing is enabled for that `sql_id` alone. Default is to enable events without sql scope. Default is to set event in current session alone. Event can be disabled by setting `disable` to non-zero value (default zero).

Syntax

```
DBMS_USERDIAG.ENABLE_SQL_TRACE_EVENT(
   level                 IN  BINARY_INTEGER DEFAULT 1,
   sid                   IN  BINARY_INTEGER DEFAULT 0,
   ser                   IN  BINARY_INTEGER DEFAULT 0,
   binds                 IN  BINARY_INTEGER DEFAULT 0,
   waits                 IN  BINARY_INTEGER DEFAULT 0,
   plan_stat             IN  VARCHAR2       DEFAULT NULL,
   sql_id                IN  VARCHAR2       DEFAULT NULL,
```

```
  disable               IN  BINARY_INTEGER DEFAULT 0,
  sys                   IN  BINARY_INTEGER DEFAULT 0);
```

Parameters

Table 209-3 ENABLE_SQL_TRACE_EVENT Procedure Parameters

Parameter	Description
level	Optional. Specifies the level associated with sql_trace event. Other settings augment the level value. For example, binds=>true adds to whatever level value has been provided by the user. This is kept for ease of usage.
sid	Optional. Target session identifier. sid, ser pair must be validated to be within current PDB id.
ser	Optional. Target serial number.
binds	Optional. A non-zero value traces bind values. Maps to level 4.
waits	Optional. A non-zero value traces waits. Maps to level 8.
plan_stat	Optional. Allowed values: FIRST_EXECUTION (default), NEVER, ALL_EXECUTIONS, and ADAPTIVE
sql_id	Enable sql_trace in scope of a given sql_id alone
disable	Optional. Non-zero value disables the already set sql_trace event.
sys	Optional. Set event instance-wide or PDB-wide. Default 0.

209.2.3 GET_CALL_ERROR_MSG Function

This function is used to obtain the error message if the last call to DBMS_USERDIAG API returned an error.

Syntax

```
DBMS_USERDIAG.GET_CALL_ERROR_MSG RETURN VARCHAR2;
```

Return Value

If the previous call to DBMS_USERDIAG was unsuccessful, the VARCHAR2 contains the error message associated with it.

If the previous call to DBMS_USERDIAG was successful, the value returned is NULL.

209.2.4 GET_CALL_STATUS Function

This function is used to obtain the status of the last call to the DBMS_USERDIAG API.

Syntax

```
DBMS_USERDIAG.GET_CALL_STATUS RETURN NUMBER;
```

Return Value

Returns a number for the status of the last call to the DBMS_USERDIAG API.

If the previous call was successful, the value of this call is `NOERROR(0)`.

If the previous call was unsuccessful, the value of this call is the error code number.

209.2.5 SET_EXCEPTION_MODE Procedure

This procedure raises an exception on any error or silently ignores the same (default).

Syntax

```
DBMS_USERDIAG.SET_EXCEPTION_MODE(
   exc_mode IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 209-4 SET_EXCEPTION_MODE Procedure Parameters

Parameter	Description
`exc_mode`	When `TRUE`, an exception is raised on any error. When `FALSE` (default), errors are ignored.

209.2.6 SET_TRACEFILE_IDENTIFIER Procedure

This procedure is used to set a custom trace file identifier for the active trace file in the current ADR home.

Syntax

```
DBMS_USERDIAG.SET_TRACEFILE_IDENTIFIER(
   trc_identifier IN VARCHAR2);
```

Parameters

Table 209-5 SET_TRACEFILE_IDENTIFIER Procedure Parameters

Parameter	Description
`trc_identifier`	Specifies the trace identifier for the active trace file in the current ADR home.

209.2.7 TRACE Procedure

This procedure is used to trace message to the trace file or alert log.

Syntax

```
DBMS_USERDIAG.TRACE(
   message IN VARCHAR2,
   alert   IN BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 209-6 TRACE Procedure Parameters

Parameter	Description
`message`	Specifies the text or message to be traced.
`alert`	When `TRUE`, the alert logs are included with the trace file for the trace operation. When `FALSE` (default), the alert logs are ignored and only the trace file is used for the trace operation.

210
DBMS_UTILITY

The `DBMS_UTILITY` package provides various utility subprograms.

This chapter contains the following topics:

- DBMS_UTILITY Deprecated Subprograms
- Security Model
- Constants
- Exceptions
- Data Structures
- Summary of DBMS_UTILITY Subprograms

210.1 DBMS_UTILITY Deprecated Subprograms

These `DBMS_UTILITY` subprograms are deprecated in Oracle Database 12*c* release 12.2.

> **Note:**
>
> Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

- GET_PARAMETER_VALUE Function

  Query `v$_parameter` directly to find the value of an `init.ora` parameter.
- ANALYZE_PART_OBJECT Procedure

  Use `DBMS_STATS` to gather statistics.
- GET_DEPENDENCY Procedure

  There is no replacement for this subprogram. However, you can directly query the dictionary views.

There are no replacements for these subprograms.

210.2 DBMS_UTILITY Security Model

`DBMS_UTILITY` runs with the privileges of the calling user for the `NAME_RESOLVE` procedure and the `COMPILE_SCHEMA` procedure. This is necessary so that the SQL works correctly.

The package does not run as `SYS`. The privileges are checked using `DBMS_DDL`.

Related Topics

- NAME_RESOLVE Procedure
  This procedure resolves the given name, including synonym translation and authorization checking as necessary.
- COMPILE_SCHEMA Procedure
  This procedure compiles all procedures, functions, packages, views and triggers in the specified schema.

210.3 DBMS_UTILITY Constants

The `DBMS_UTILITY` package defines one constant to use when specifying parameter values.

This constant is shown in the following table.

Table 210-1 DBMS_UTILITY Constants

Name	Type	Value	Description
`INV_ERROR_ON_RESTRICTIONS`	`PLS_INTEGER`	`1`	This constant is the only legal value for the `p_option_flags` parameter of the `INVALIDATE` subprogram

210.4 DBMS_UTILITY Exceptions

This table lists the exceptions raised by `DBMS_UTILITY`.

Table 210-2 Exceptions Raised by DBMS_UTILITY

Exception	Error Code	Description
`INV_NOT_EXIST_OR_NO_PRIV`	-24237	Raised by the `INVALIDATE` subprogram when the `object_id` argument is `NULL` or invalid, or when the caller does not have `CREATE` privileges on the object being invalidated
`INV_MALFORMED_SETTINGS`	-24238	Raised by the `INVALIDATE` subprogram if a compiler setting is specified more than once in the `p_plsql_object_settings` parameter
`INV_RESTRICTED_OBJECT`	-24239	Raised by the `INVALIDATE` subprogram when different combinations of conditions pertaining to the `p_object_id` parameter are contravened

210.5 DBMS_UTILITY Data Structures

The `DBMS_UTILITY` package defines a single `RECORD` type and `TABLE` types.

Record Types

- INSTANCE_RECORD Record Type

Table Types

- DBLINK_ARRAY TABLE Type

- INDEX_TABLE_TYPE Table Type
- INSTANCE_TABLE Table Type
- LNAME_ARRAY Table Type
- NAME_ARRAY Table Type
- NUMBER_ARRAY Table Type
- UNCL_ARRAY Table Type

210.5.1 DBMS_UTILITY INSTANCE_RECORD Record Type

This type describes a list of active instance number-name pairs.

Syntax

```
TYPE INSTANCE_RECORD IS RECORD (
     inst_number   NUMBER,
     inst_name     VARCHAR2(60));
```

Fields

Table 210-3 INSTANCE_RECORD Record Type Fields

Field	Description
inst_number	Active instance number
inst_name	Instance name

210.5.2 DBMS_UTILITY DBLINK_ARRAY TABLE Type

This type stores a list of database links.

Syntax

```
TYPE DBLINK_ARRAY IS TABLE OF VARCHAR2(128) INDEX BY BINARY_INTEGER;
```

210.5.3 DBMS_UTILITY INDEX_TABLE_TYPE Table Type

This type describes the order in which generated objects are returned to a user.

Syntax

```
TYPE INDEX_TABLE_TYPE IS TABLE OF BINARY_INTEGER INDEX BY BINARY_INTEGER;
```

210.5.4 DBMS_UTILITY INSTANCE_TABLE Table Type

This type describes a table of `INSTANCE_RECORD` Record Type.

Syntax

```
TYPE INSTANCE_TABLE IS TABLE OF INSTANCE_RECORD INDEX BY BINARY_INTEGER;
```

Usage Notes

The starting index of `INSTANCE_TABLE` Is 1; `INSTANCE_TABLE` Is Dense.

Related Topics

- DBMS_UTILITY INSTANCE_RECORD Record Type
  This type describes a list of active instance number-name pairs.

210.5.5 DBMS_UTILITY LNAME_ARRAY Table Type

This type stores lists of `LONG NAME` including fully qualified attribute names.

Syntax

```
TYPE LNAME_ARRAY IS TABLE OF VARCHAR2(4000) INDEX BY BINARY_INTEGER;
```

210.5.6 DBMS_UTILITY NAME_ARRAY Table Type

This type stores lists of `NAME`.

Syntax

```
TYPE NAME_ARRAY IS TABLE OF VARCHAR2(30) INDEX BY BINARY_INTEGER;
```

210.5.7 DBMS_UTILITY NUMBER_ARRAY Table Type

This type describes the order in which generated objects are returned to users.

Syntax

```
TYPE NUMBER_ARRAY IS TABLE OF NUMBER INDEX BY BINARY_INTEGER;
```

210.5.8 DBMS_UTILITY UNCL_ARRAY Table Type

This type stores lists of "`user`"."`name`"."`column`"`@link`

Syntax

```
TYPE UNCL_ARRAY IS TABLE OF VARCHAR2(227) INDEX BY BINARY_INTEGER;
```

210.6 Summary of DBMS_UTILITY Subprograms

This table lists the `DBMS_UTILITY` subprograms and briefly describes them.

Table 210-4 *DBMS_UTILITY Package Subprograms*

Subprogram	Description
ACTIVE_INSTANCES Procedure	Returns the active instance
ANALYZE_DATABASE Procedure	Analyzes all the tables, clusters and indexes in a database
ANALYZE_PART_OBJECT Procedure	Analyzes the given tables and indexes This procedure is deprecated from the `DBMS_UTILITY` package with Oracle Database 12c release 12.2 and later. Use `DBMS_STATS` to gather statistics.
ANALYZE_SCHEMA Procedure	Analyzes all the tables, clusters and indexes in a schema
CANONICALIZE Procedure	Canonicalizes a given string

Table 210-4 (Cont.) *DBMS_UTILITY Package Subprograms*

Subprogram	Description
COMMA_TO_TABLE Procedures	Converts a comma-delimited list of names into a PL/SQL table of names
COMPILE_SCHEMA Procedure	Compiles all procedures, functions, packages, views and triggers in the specified schema
CREATE_ALTER_TYPE_ERROR_TABLE Procedure	Creates an error table to be used in the `EXCEPTION` clause of the `ALTER TYPE` statement
CURRENT_INSTANCE Function	Returns the current connected instance number
DATA_BLOCK_ADDRESS_BLOCK Function	Gets the block number part of a data block address
DATA_BLOCK_ADDRESS_FILE Function	Gets the file number part of a data block address
DB_VERSION Procedure	Returns version information for the database
EXEC_DDL_STATEMENT Procedure	Executes the DDL statement in `parse_string`
EXPAND_SQL_TEXT Procedure	Recursively replaces any view references in the input SQL query with the corresponding view subquery
FORMAT_CALL_STACK Function	Formats the current call stack
FORMAT_ERROR_BACKTRACE Function	Formats the backtrace from the point of the current error to the exception handler where the error has been caught
FORMAT_ERROR_STACK Function	Formats the current error stack
GET_CPU_TIME Function	Returns the current CPU time in 100th's of a second
GET_DEPENDENCY Procedure	Shows the dependencies on the object passed in. This procedure is deprecated from the `DBMS_UTILITY` package with Oracle Database 12c release 12.2 and later. There is no replacement for this subprogram.
GET_ENDIANNESS Function	Gets the endianness of the database platform
GET_HASH_VALUE Function	Computes a hash value for the given string
GET_PARAMETER_VALUE Function	Gets the value of specified init.ora parameter. This function is deprecated from the `DBMS_UTILITY` package with Oracle Database 12c release 12.2 and later. You can query `v$_parameter` directly.
GET_SQL_HASH Function	Computes a hash value for the given string using MD5 algorithm
GET_TIME Function	Returns the current time in 100th's of a second
GET_TZ_TRANSITIONS Procedure	Returns time zeone transitions by `regionid` from the `timezone.dat` file
INVALIDATE Procedure	Invalidates a database object and (optionally) modifies its PL/SQL compiler parameter settings
IS_BIT_SET Function	Checks the bit setting for the given bit in the given `RAW` value
IS_CLUSTER_DATABASE Function	Determines if the database is running in cluster database mode
MAKE_DATA_BLOCK_ADDRESS Function	Creates a data block address given a file number and a block number
NAME_RESOLVE Procedure	Resolves the given name
NAME_TOKENIZE Procedure	Calls the parser to parse the given name

Table 210-4 (Cont.) *DBMS_UTILITY Package Subprograms*

Subprogram	Description
OLD_CURRENT_SCHEMA Function	Returns the session value from `SYS_CONTEXT ('USERENV', 'CURRENT_SCHEMA')`
OLD_CURRENT_USER Function	Returns the session value from `SYS_CONTEXT ('USERENV', 'CURRENT_USER')`
PORT_STRING Function	Returns a string that uniquely identifies the version of Oracle and the operating system
SQLID_TO_SQLHASH Function	Converts a SQL ID into a hash value
TABLE_TO_COMMA Procedures	Converts a PL/SQL table of names into a comma-delimited list of names
VALIDATE Procedure	Makes invalid database objects valid
WAIT_ON_PENDING_DML Function	Waits until all transactions (other than the caller's own) that have locks on the listed tables and began prior to the specified SCN have either committed or been rolled back

210.6.1 ACTIVE_INSTANCES Procedure

This procedure returns the active instance.

Syntax

```
DBMS_UTILITY.ACTIVE_INSTANCES (
   instance_table   OUT INSTANCE_TABLE,
   instance_count   OUT NUMBER);
```

Parameters

Table 210-5 ACTIVE_INSTANCES Procedure Parameters

Procedure	Description
`instance_table`	Contains a list of the active instance numbers and names. When no instance is up, the list is empty.
`instance_count`	Number of active instances

210.6.2 ANALYZE_DATABASE Procedure

This procedure analyzes all the tables, clusters and indexes in a database.

Syntax

```
DBMS_UTILITY.ANALYZE_DATABASE (
   method             IN  VARCHAR2,
   estimate_rows      IN  NUMBER DEFAULT NULL,
   estimate_percent   IN  NUMBER DEFAULT NULL,
   method_opt         IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 210-6 ANALYZE_DATABASE Procedure Parameters

Parameter	Description
method	One of ESTIMATE, COMPUTE or DELETE. If ESTIMATE then either estimate_rows or estimate_percent must be nonzero.
estimate_rows	Number of rows to estimate
estimate_percent	Percentage of rows to estimate. If estimate_rows is specified ignore this parameter.
method_opt	Method options of the following format: [FOR TABLE] [FOR ALL [INDEXED] COLUMNS] [SIZE n] [FOR ALL INDEXES]

Exceptions

ORA-20000: Insufficient privileges for some object in this database

210.6.3 ANALYZE_PART_OBJECT Procedure

This procedure is equivalent to SQL: "ANALYZE TABLE|INDEX [<schema>.]<object_name> PARTITION <pname> [<command_type>] [<command_opt>] [<sample_clause>]

> **Note:**
>
> This subprogram has been deprecated and replaced by improved technology. It is maintained only for purposes of backward compatibility. As an alternative, you can use DBMS_STATS to gather statistics.

Syntax

```
DBMS_UTILITY.ANALYZE_PART_OBJECT (
   schema        IN VARCHAR2 DEFAULT NULL,
   object_name   IN VARCHAR2 DEFAULT NULL,
   object_type   IN CHAR     DEFAULT 'T',
   command_type  IN CHAR     DEFAULT 'E',
   command_opt   IN VARCHAR2 DEFAULT NULL,
   sample_clause IN VARCHAR2 DEFAULT 'sample 5 percent ');
```

Parameters

Table 210-7 ANALYZE_PART_OBJECT Procedure Parameters

Parameter	Description
schema	Schema of the object_name
object_name	Name of object to be analyzed, must be partitioned
object_type	Type of object, must be T (table) or I (index)

Table 210-7 (Cont.) ANALYZE_PART_OBJECT Procedure Parameters

Parameter	Description
command_type	Must be V (validate structure)
command_opt	Other options for the command type. For C, E it can be FOR table, FOR all LOCAL indexes, FOR all columns or combination of some of the 'for' options of analyze statistics (table). For V, it can be CASCADE when object_type is T.
sample_clause	Sample clause to use when command_type is 'E'

Usage Notes

For each partition of the object, run in parallel using job queues.

210.6.4 ANALYZE_SCHEMA Procedure

This procedure analyzes all the tables, clusters and indexes in a schema.

Syntax

```
DBMS_UTILITY.ANALYZE_SCHEMA (
   schema             IN  VARCHAR2,
   method             IN  VARCHAR2,
   estimate_rows      IN  NUMBER DEFAULT NULL,
   estimate_percent   IN  NUMBER DEFAULT NULL,
   method_opt         IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 210-8 ANALYZE_SCHEMA Procedure Parameters

Parameter	Description
schema	Name of the schema
method	One of ESTIMATE, COMPUTE or DELETE. If ESTIMATE then either estimate_rows or estimate_percent must be nonzero.
estimate_rows	Number of rows to estimate
estimate_percent	Percentage of rows to estimate. If estimate_rows is specified ignore this parameter.
method_opt	Method options of the following format: [FOR TABLE] [FOR ALL [INDEXED] COLUMNS] [SIZE n] [FOR ALL INDEXES]

Exceptions

ORA-20000: Insufficient privileges for some object in this schema

210.6.5 CANONICALIZE Procedure

This procedure canonicalizes the given string. The procedure handles a single reserved or key word (such as 'table'), and strips off white spaces for a single identifier so that ' table ' becomes TABLE.

Syntax

```
DBMS_UTILITY.CANONICALIZE(
   name           IN    VARCHAR2,
   canon_name     OUT   VARCHAR2,
   canon_len      IN    BINARY_INTEGER);
```

Parameters

Table 210-9 CANONICALIZE Procedure Parameters

Parameter	Description
`name`	String to be canonicalized
`canon_name`	Canonicalized string
`canon_len`	Length of the string (in bytes) to canonicalize

Return Values

Returns the first `canon_len` bytes in `canon_name`.

Usage Notes

- If `name` is NULL, `canon_name` becomes NULL.
- If `name` is not a dotted name, and if `name` begins and ends with a double quote, remove both quotes. Alternatively, convert to upper case with NLS_UPPER. Note that this case does not include a name with special characters, such as a space, but is not doubly quoted.
- If name is a dotted name (such as a."b".c), for each component in the dotted name in the case in which the component begins and ends with a double quote, no transformation will be performed on this component. Alternatively, convert to upper case with NLS_UPPER and apply begin and end double quotes to the capitalized form of this component. In such a case, each canonicalized component will be concatenated together in the input position, separated by ".".
- Any other character after a[.b]* will be ignored.
- The procedure does not handle cases like 'A B.'

Examples

- `a` becomes `A`
- `"a"` becomes `a`
- `"a".b` becomes `"a"."B"`
- `"a".b,c.f` becomes `"a"."B"` with`",c.f"` ignored.

210.6.6 COMMA_TO_TABLE Procedures

These procedures convert a comma-delimited list of names into a PL/SQL table of names. The second version supports fully-qualified attribute names.

Syntax

```
DBMS_UTILITY.COMMA_TO_TABLE (
   list   IN  VARCHAR2,
   tablen OUT BINARY_INTEGER,
   tab    OUT uncl_array);

DBMS_UTILITY.COMMA_TO_TABLE (
   list   IN  VARCHAR2,
   tablen OUT BINARY_INTEGER,
   tab    OUT lname_array);
```

Parameters

Table 210-10 COMMA_TO_TABLE Procedure Parameters

Parameter	Description
`list`	Comma separated list of list of 'names', where a name should have the following format for the first overloading: `a [. b [. c ]][ @ d ]` and the following format for the second overloading: `a [. b]*` where `a`, `b`, `c`, `d` are simple identifiers (quoted or unquoted).
`tablen`	Number of tables in the PL/SQL table
`tab`	PL/SQL table which contains list of names

Return Values

A PL/SQL table is returned, with values `1..n` and `n+1 is null`.

Usage Notes

- The `list` must be a non-empty comma-delimited list: Anything other than a comma-delimited list is rejected. Commas inside double quotes do not count.
- Entries in the comma-delimited list cannot include multibyte characters.
- The values in `tab` are copied from the original list, with no transformations.
- The procedure fails if the string between separators is longer than 30 bytes.

210.6.7 COMPILE_SCHEMA Procedure

This procedure compiles all procedures, functions, packages, views and triggers in the specified schema.

Syntax

```
DBMS_UTILITY.COMPILE_SCHEMA (
   schema          IN VARCHAR2,
   compile_all     IN BOOLEAN DEFAULT TRUE,
   reuse_settings  IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 210-11 COMPILE_SCHEMA Procedure Parameters

Parameter	Description
schema	Name of the schema
compile_all	If TRUE, will compile everything within the schema regardless of whether it is VALID If FALSE, will compile only INVALID objects
reuse_settings	Indicates whether the session settings in the objects should be reused, or whether the current session settings should be adopted instead

Exceptions

Table 210-12 COMPILE_SCHEMA Procedure Exceptions

Exception	Description
ORA-20000	Insufficient privileges for some object in this schema
ORA-20001	Cannot recompile SYS objects
ORA-20002	Maximum iterations exceeded. Some objects may not have been recompiled.

Usage Notes

- Note that this subprogram is a wrapper for the RECOMP_SERIAL Procedure included with the UTL_RECOMP package.
- After calling this procedure, you should select from view ALL_OBJECTS for items with status of INVALID to see if all objects were successfully compiled.
- To see the errors associated with INVALID objects, you may use the Enterprise Manager command:

```
SHOW ERRORS <type> <schema>.<name>
```

210.6.8 CREATE_ALTER_TYPE_ERROR_TABLE Procedure

This procedure creates an error table to be used in the EXCEPTION clause of the ALTER TYPE statement.

Syntax

```
DBMS_UTILITY.CREATE_ALTER_TYPE_ERROR_TABLE(
   schema_name      IN     VARCHAR2,
   table_name       IN     VARCHAR2);
```

Parameters

Table 210-13 CREATE_ALTER_TYPE_ERROR_TABLE Procedure Parameters

Parameter	Description
schema_name	Name of the schema
table_name	Name of the table created

Exceptions

An error is returned if the table already exists.

210.6.9 CURRENT_INSTANCE Function

This function returns the current connected instance number. It returns NULL when connected instance is down.

Syntax

```
DBMS_UTILITY.CURRENT_INSTANCE
   RETURN NUMBER;
```

210.6.10 DATA_BLOCK_ADDRESS_BLOCK Function

This function gets the block number part of a data block address.

Syntax

```
DBMS_UTILITY.DATA_BLOCK_ADDRESS_BLOCK (
   dba NUMBER)
  RETURN NUMBER;
```

Parameters

Table 210-14 DATA_BLOCK_ADDRESS_BLOCK Function Parameters

Parameter	Description
dba	Data block address

Pragmas

```
pragma restrict_references(data_block_address_block, WNDS, RNDS, WNPS, RNPS);
```

Return Values

Block offset of the block.

Usage Notes

This function should not be used with datablocks which belong to bigfile tablespaces.

210.6.11 DATA_BLOCK_ADDRESS_FILE Function

This function gets the file number part of a data block address.

Syntax

```
DBMS_UTILITY.DATA_BLOCK_ADDRESS_FILE (
   dba NUMBER)
  RETURN NUMBER;
```

Parameters

Table 210-15 DATA_BLOCK_ADDRESS_FILE Function Parameters

Parameter	Description
dba	Data block address

Pragmas

```
pragma restrict_references (data_block_address_file, WNDS, RNDS, WNPS, RNPS);
```

Return Values

File that contains the block.

Usage Notes

This function should not be used with datablocks which belong to bigfile tablespaces.

210.6.12 DB_VERSION Procedure

This procedure returns version information for the database.

Syntax

```
DBMS_UTILITY.DB_VERSION (
   version       OUT VARCHAR2,
   compatibility OUT VARCHAR2);
```

Parameters

Table 210-16 DB_VERSION Procedure Parameters

Parameter	Description
version	A string which represents the internal software version of the database (for example, 7.1.0.0.0). The length of this string is variable and is determined by the database version.
compatibility	The compatibility setting of the database determined by the "compatible" init.ora parameter. If the parameter is not specified in the init.ora file, then NULL is returned.

210.6.13 EXEC_DDL_STATEMENT Procedure

This procedure executes the DDL statement in `parse_string`.

Syntax

```
DBMS_UTILITY.EXEC_DDL_STATEMENT (
   parse_string IN VARCHAR2);
```

Parameters

Table 210-17 EXEC_DDL_STATEMENT Procedure Parameters

Parameter	Description
`parse_string`	DDL statement to be executed

210.6.14 EXPAND_SQL_TEXT Procedure

This procedure recursively replaces any view references in the input SQL query with the corresponding view subquery.

Syntax

```
DBMS_UTILITY.EXPAND_SQL_TEXT (
   input_sql_text     IN           CLOB,
   output_sql_text    OUT NOCOPY   CLOB);
```

Parameters

Table 210-18 EXPAND_SQL_TEXT Procedure Parameters

Parameter	Description
`input_sql_text`	Input SQL query text
`output_sql_text`	View-expanded query text

Exceptions

Table 210-19 EXPAND_SQL_TEXT Procedure Exceptions

Exception	Description
`ORA-00942`	Current user does not have select privileges on all the views and tables recursively referenced in the `input_sql_text`
`ORA-24251`	`input_sql_text` is not a `SELECT` statement
`ORA-00900`	Input is not valid
`ORA-29477`	Input LOB size exceeds maximum size of 4GB -1

Usage Notes

The expanded and merged SQL statement text is copied to output_sql_text on successful completion. The resulting query text only contains references to underlying tables and is semantically equivalent with some caveats:

- If there are invoker rights functions called from any of the views, they may be called as a different user in the resulting query text if the view owner is different from the user who will eventually compile/run the expanded SQL text.
- The VPD policy expands differently if there is a function supplied to generate the dynamic `WHERE` clause. This function would return differently, for example, if the userid caused the expansion to be different.
- If there are references to remote objects, results are undetermined.

210.6.15 FORMAT_CALL_STACK Function

This function formats the current call stack. This can be used on any stored procedure or trigger to access the call stack. This can be useful for debugging.

Syntax

```
DBMS_UTILITY.FORMAT_CALL_STACK
  RETURN VARCHAR2;
```

Pragmas

```
pragma restrict_references(format_call_stack,WNDS);
```

Return Values

This returns the call stack, up to 2000 bytes.

210.6.16 FORMAT_ERROR_BACKTRACE Function

This function displays the call stack at the point where an exception was raised, even if the subprogram is called from an exception handler in an outer scope.

The output is similar to the output of the `SQLERRM` function, but not subject to the same size limitation.

Syntax

```
DBMS_UTILITY.FORMAT_ERROR_BACKTRACE
  RETURN VARCHAR2;
```

Return Values

The backtrace string. A `NULL` string is returned if no error is currently being handled.

Examples

```
CREATE OR REPLACE PROCEDURE Log_Errors ( i_buff in varchar2 ) IS
  g_start_pos integer := 1;
  g_end_pos  integer;

  FUNCTION Output_One_Line RETURN BOOLEAN IS
  BEGIN
```

```
    g_end_pos := Instr ( i_buff, Chr(10), g_start_pos );

    CASE g_end_pos > 0
      WHEN true THEN
        DBMS_OUTPUT.PUT_LINE ( Substr ( i_buff, g_start_pos,
g_end_pos-g_start_pos ) );
        g_start_pos := g_end_pos+1;
        RETURN TRUE;

      WHEN FALSE THEN
        DBMS_OUTPUT.PUT_LINE ( Substr ( i_buff, g_start_pos,
(Length(i_buff)-g_start_pos)+1 ) );
        RETURN FALSE;
    END CASE;
  END Output_One_Line;

BEGIN
  WHILE Output_One_Line() LOOP NULL;
  END LOOP;
END Log_Errors;
/

Set Doc Off
Set Feedback off
Set Echo Off

CREATE OR REPLACE PROCEDURE P0 IS
  e_01476 EXCEPTION; pragma exception_init ( e_01476, -1476 );
BEGIN
  RAISE e_01476;
END P0;
/
Show Errors

CREATE OR REPLACE PROCEDURE P1 IS
BEGIN
  P0();
END P1;
/
SHOW ERRORS

CREATE OR REPLACE PROCEDURE P2 IS
BEGIN
  P1();
END P2;
/
SHOW ERRORS

CREATE OR REPLACE PROCEDURE P3 IS
BEGIN
  P2();
END P3;
/
SHOW ERRORS

CREATE OR REPLACE PROCEDURE P4 IS
  BEGIN P3(); END P4;
/
CREATE OR REPLACE PROCEDURE P5 IS
  BEGIN P4(); END P5;
/
SHOW ERRORS
```

```
CREATE OR REPLACE PROCEDURE Top_Naive IS
BEGIN
  P5();
END Top_Naive;
/
SHOW ERRORS

CREATE OR REPLACE PROCEDURE Top_With_Logging IS
  -- NOTE: SqlErrm in principle gives the same info as Format_Error_Stack.
  -- But SqlErrm is subject to some length limits,
  -- while Format_Error_Stack is not.
BEGIN
  P5();
EXCEPTION
  WHEN OTHERS THEN
    Log_Errors ( 'Error_Stack...' || Chr(10) ||
      DBMS_UTILITY.FORMAT_ERROR_STACK() );
    Log_Errors ( 'Error_Backtrace...' || Chr(10) ||
      DBMS_UTILITY.FORMAT_ERROR_BACKTRACE() );
    DBMS_OUTPUT.PUT_LINE ( '----------' );
END Top_With_Logging;
/
SHOW ERRORS

--------------------------------------------------------------------------------

Set ServerOutput On
call Top_Naive()
  /*
  ERROR at line 1:
  ORA-01476: divisor is equal to zero
  ORA-06512: at "U.P0", line 4
  ORA-06512: at "U.P1", line 3
  ORA-06512: at "U.P2", line 3
  ORA-06512: at "U.P3", line 3
  ORA-06512: at "U.P4", line 2
  ORA-06512: at "U.P5", line 2
  ORA-06512: at "U.TOP_NAIVE", line 3
  */
  ;

Set ServerOutput On
call Top_With_Logging()
  /*
  Error_Stack...
  ORA-01476: divisor is equal to zero
  Error_Backtrace...
  ORA-06512: at "U.P0", line 4
  ORA-06512: at "U.P1", line 3
  ORA-06512: at "U.P2", line 3
  ORA-06512: at "U.P3", line 3
  ORA-06512: at "U.P4", line 2
  ORA-06512: at "U.P5", line 2
  ORA-06512: at "U.TOP_WITH_LOGGING", line 6
  ----------
  */
  ;

/*
  ORA-06512:
  Cause:
```

```
    Backtrace message as the stack is
    unwound by unhandled exceptions.
  Action:
    Fix the problem causing the exception
    or write an exception handler for this condition.
    Or you may need to contact your application administrator
    or database administrator.
*/
```

210.6.17 FORMAT_ERROR_STACK Function

This function formats the current error stack. This can be used in exception handlers to look at the full error stack.

Syntax

```
DBMS_UTILITY.FORMAT_ERROR_STACK
  RETURN VARCHAR2;
```

Return Values

This returns the error stack, up to 2000 bytes.

210.6.18 GET_CPU_TIME Function

This function returns a measure of current CPU processing time in hundredths of a second. The difference between the times returned from two calls measures the CPU processing time (not the total elapsed time) between those two points.

Syntax

```
  DBMS_UTILITY.GET_CPU_TIME
   RETURN NUMBER;
```

Return Values

Time is the number of 100th's of a second from some arbitrary epoch.

Usage Notes

The amount of work performed is calculated by measuring the difference between a start point and end point for a particular operation.

210.6.19 GET_DEPENDENCY Procedure

This **deprecated procedure** shows the dependencies on the object passed in.

> **Note:**
>
> This subprogram has been deprecated and replaced in Oracle Database 12c release 12.2 and later. Oracle recommends that you do not use deprecated subprograms. It is maintained only for purposes of backward compatibility.

Syntax

```
DBMS_UTILITY.GET_DEPENDENCY
 type      IN     VARCHAR2,
 schema    IN     VARCHAR2,
 name      IN     VARCHAR2);
```

Parameters

Table 210-20 GET_DEPENDENCY Procedure Parameters

Parameter	Description
type	Type of the object, for example if the object is a table give the type as 'TABLE'
schema	Schema name of the object
name	Name of the object

Usage Notes

This procedure uses the DBMS_OUTPUT package to display results, and so you must declare SET SERVEROUTPUT ON if you wish to view dependencies. Alternatively, any application that checks the DBMS_OUTPUT output buffers can invoke this subprogram and then retrieve the output by means of DBMS_OUTPUT subprograms such as GET_LINES.

210.6.20 GET_ENDIANNESS Function

This function gets the endianness of the database platform.

Syntax

```
DBMS_UTILITY.GET_ENDIANNESS
 RETURN NUMBER;
```

Return Values

A NUMBER value indicating the endianness of the database platform: 1 for big-endian or 2 for little-endian.

210.6.21 GET_HASH_VALUE Function

This function computes a hash value for the given string.

Syntax

```
DBMS_UTILITY.GET_HASH_VALUE (
   name      VARCHAR2,
   base      NUMBER,
   hash_size NUMBER)
  RETURN NUMBER;
```

Parameters

Table 210-21 GET_HASH_VALUE Function Parameters

Parameter	Description
name	String to be hashed.
base	Base value for the returned hash value at which to start
hash_size	Desired size of the hash table

Pragmas

```
pragma restrict_references(get_hash_value, WNDS, RNDS, WNPS, RNPS);
```

Return Values

A hash value based on the input string. For example, to get a hash value on a string where the hash value should be between 1000 and 3047, use 1000 as the base value and 2048 as the hash_size value. Using a power of 2 for the hash_size parameter works best.

210.6.22 GET_PARAMETER_VALUE Function

This deprecated function gets the value of specified init.ora parameter.

> **Note:**
>
> This subprogram has been deprecated and replaced by improved technology. It is maintained only for purposes of backward compatibility. As an alternative, you can query v$_parameter directly.

Syntax

```
DBMS_UTILITY.GET_PARAMETER_VALUE (
   parnam     IN        VARCHAR2,
   intval     IN OUT    BINARY_INTEGER,
   strval     IN OUT    VARCHAR2,
   listno     IN        BINARY_INTEGER DEFAULT 1)
  RETURN BINARY_INTEGER;
```

Parameters

Table 210-22 GET_PARAMETER_VALUE Function Parameters

Parameter	Description
parnam	Parameter name
intval	Value of an integer parameter or the value length of a string parameter
strval	Value of a string parameter
listno	List item number. If retrieving parameter values for a parameter that can be specified multiple times to accumulate values, use this parameter to get each individual parameter.

Return Values

Parameter type:

- 0 if parameter is an `INTEGER`/`BOOLEAN` parameter
- 1 if parameter is a string/file parameter

Usage Notes

- To execute the this function, you must have the `SELECT` privilege on the V$PARAMETER dynamic view.

Examples

```
DECLARE
  parnam VARCHAR2(256);
  intval BINARY_INTEGER;
  strval VARCHAR2(256);
  partyp BINARY_INTEGER;
BEGIN
  partyp := dbms_utility.get_parameter_value('max_dump_file_size',
                                             intval, strval);
  dbms_output.put('parameter value is: ');
  IF partyp = 1 THEN
    dbms_output.put_line(strval);
  ELSE
    dbms_output.put_line(intval);
  END IF;
  IF partyp = 1 THEN
    dbms_output.put('parameter value length is: ');
    dbms_output.put_line(intval);
  END IF;
  dbms_output.put('parameter type is: ');
  IF partyp = 1 THEN
    dbms_output.put_line('string');
  ELSE
    dbms_output.put_line('integer');
  END IF;
END;
```

210.6.23 GET_SQL_HASH Function

This function computes a hash value for the given string using MD5 algorithm.

Syntax

```
Dbms_utility.get_sql_hash (
   name          IN   VARCHAR2,
   hash          OUT  RAW,
   pre10ihash    OUT  NUMBER)
  RETURN NUMBER;
```

Pragmas

Pragma Restrict_references(`Get_sql_hash`, `Wnds`, `Rnds`, `Wnps`, `Rnps`);

Parameters

Table 210-23 GET_SQL_HASH Procedure Parameters

Parameter	Description
`name`	String to be hashed
`hash`	Stores all 16 bytes of returned hash value
`pre10ihash`	Stores the pre 10i database version hash value

Return Values

A hash value (last 4 bytes) based on the input string. the MD5 hash algorithm computes a 16 byte hash value, but we only return the last 4 bytes so that we can return an actual number. Use the `hash` parameter to get all 16 bytes and pre10*i* hash parameter to store the pre 10*i* hash value of 4 bytes.

210.6.24 GET_TIME Function

This function determines the current time in hundredths of a second. This subprogram is primarily used for determining elapsed time. The subprogram is called twice – at the beginning and end of some process – and then the first (earlier) number is subtracted from the second (later) number to determine the time elapsed.

Syntax

```
DBMS_UTILITY.GET_TIME
  RETURN NUMBER;
```

Return Values

Time is the number of hundredths of a second from the point in time at which the subprogram is invoked.

Usage Notes

Numbers are returned in the range -2147483648 to 2147483647 depending on platform and machine, and your application must take the sign of the number into account in determining the interval. For instance, in the case of two negative numbers, application logic must allow that the first (earlier) number will be larger than the second (later) number which is closer to zero. By the same token, your application should also allow that the first (earlier) number be negative and the second (later) number be positive.

210.6.25 GET_TZ_TRANSITIONS Procedure

This procedure returns time zone transitions by `regionid` from the `timezone.dat` file.

Syntax

```
DBMS_UTILITY.GET_TZ_TRANSITIONS
   regionid      IN     NUMBER,
   transitions   OUT    MAXRAW);
```

Parameters

Table 210-24 GET_TZ_TRANSITIONS Procedure Parameters

Parameter	Description
regionid	Number corresponding to the region
transitions	Raw bytes from the timezone.dat file

Exceptions

Table 210-25 GET_TZ_TRANSITIONS Procedure Exceptions

Exception	Description
ORA-6502: PL/SQL: NUMERIC OR VALUE ERROR	For an invalid regionid

210.6.26 INVALIDATE Procedure

This procedure invalidates a database object and (optionally) modifies its PL/SQL compiler parameter settings. It also invalidates any objects that (directly or indirectly) depend on the object being invalidated.

Syntax

```
DBMS_UTILITY.INVALIDATE (
   p_object_id              NUMBER,
   p_plsql_object_settings  VARCHAR2 DEFAULT NULL,
   p_option_flags           PLS_INTEGER DEFAULT 0);
```

Parameters

Table 210-26 INVALIDATE Procedure Parameters

Parameter	Description
p_object_id	ID number of object to be invalidated. This is the same as the value of the OBJECT_ID column from ALL_OBJECTS. If the object_id argument is NULL or invalid then the exception inv_not_exist_or_no_priv is raised. The caller of this procedure must have create privileges on the object being invalidated else the inv_not_exist_or_no_priv exception is raised.

Table 210-26 (Cont.) INVALIDATE Procedure Parameters

Parameter	Description
p_plsql_object_settings	Optional parameter that ignored if the object specified by p_object_id is not a PL/SQL object. If no value is specified for this parameter then the PL/SQL compiler settings are left unchanged, that is, equivalent to REUSE SETTINGS. If a value is provided, it must specify the values of the PL/SQL compiler settings separated by one or more spaces. Each setting can be specified only once else inv_malformed_settings exception will be raised. The setting values are changed only for the object specified by p_object_id and do not affect dependent objects that may be invalidated. The setting names and values are case insensitive. If a setting is omitted and REUSE SETTINGS is specified, then if a value was specified for the compiler setting in an earlier compilation of this library unit, Oracle Database uses that earlier value. If a setting is omitted and REUSE SETTINGS was not specified or no value has been specified for the parameter in an earlier compilation, then the database will obtain the value for that setting from the session environment.
p_option_flags	Optional parameter defaults to zero (no flags). Option flags supported by invalidate. • inv_error_on_restrictions (see Constants): The subprogram imposes various restrictions on the objects that can be invalidated. For example, the object specified by p_object_id cannot be a table. By default, invalidate quietly returns on these conditions (and does not raise an exception). If the caller sets this flag, the exception inv_restricted_object is raised.

Exceptions

Table 210-27 INVALIDATE Exceptions

Exception	Description
INV_NOT_EXIST_OR_NO_PRIV	Raised when the object_id argument is NULL or invalid, or when the caller does not have CREATE privileges on the object being invalidated
INV_MALFORMED_SETTINGS	Raised if a compiler setting is specified more than once in the p_plsql_object_settings parameter
INV_RESTRICTED_OBJECT	Raised when different combinations of conditions pertaining to the p_object_id parameter are contravened

Usage Notes

The object type (object_type column from ALL_OBJECTS) of the object specified by p_object_id must be a PROCEDURE, FUNCTION, PACKAGE, PACKAGE BODY, TRIGGER, TYPE, TYPE BODY, LIBRARY, VIEW, OPERATOR, SYNONYM, or JAVA CLASS. If the object is not one of these types and the flag inv_error_on_restrictions is specified in p_option_flags then the exception inv_restricted_object is raised, else no action is taken.

If the object specified by p_object_id is the package specification of STANDARD, DBMS_STANDARD, or specification or body of DBMS_UTILITY and the flag

inv_error_on_restrictions is specified in p_option_flags then the exception inv_restricted_object is raised, else no action is taken.

If the object specified by p_object_id is an object type specification and there exist tables which depend on the type and the flag inv_error_on_restrictions is specified in p_option_flags then the exception inv_restricted_object is raised, else no action is taken.

Examples

Example 1

```
DBMS_UTILITY.INVALIDATE (1232, 'PLSQL_OPTIMIZE_LEVEL = 2 REUSE SETTINGS');
```

Assume that the object_id 1232 refers to the procedure remove_emp in the HR schema. Then the above call will mark the remove_emp procedure invalid and change it's PLSQL_OPTIMIZE_LEVEL compiler setting to 2. The values of other compiler settings will remain unchanged since REUSE SETTINGS is specified.

Objects that depend on hr.remove_emp will also get marked invalid. Their compiler parameters will not be changed.

Example 2

```
DBMS_UTILITY.INVALIDATE (40775, 'plsql_code_type = native');
```

Assume that the object_id 40775 refers to the type body leaf_category_typ in the OE schema. Then the above call will mark the type body invalid and change its PLSQL_CODE_TYPE compiler setting to NATIVE. The values of other compiler settings will be picked up from the current session environment since REUSE SETTINGS has not been specified.

Since no objects can depend on bodies, there are no cascaded invalidations.

Example 3

```
DBMS_UTILITY.INVALIDATE (40796);
```

Assume that the object_id 40796 refers to the view oc_orders in the OE schema. Then the above call will mark the oc_orders view invalid.

Objects that depend on oe.oc_orders will also get marked invalid.

210.6.27 IS_BIT_SET Function

This function checks the bit setting for the given bit in the given RAW value.

Syntax

```
DBMS_UTILITY.IS_BIT_SET (
   r     IN    RAW,   n     IN    NUMBER)
 RETURN NUMBER;
```

Parameters

Table 210-28 IS_BET_SET Function Parameters

Parameter	Description
r	RAW source

Table 210-28 (Cont.) IS_BET_SET Function Parameters

Parameter	Description
n	Bit in r to check

Return Values

This function returns 1 if bit n in raw r is set, zero otherwise. Bits are numbered high to low with the lowest bit being bit number 1.

210.6.28 IS_CLUSTER_DATABASE Function

This function finds out if this database is running in cluster database mode.

Syntax

```
DBMS_UTILITY.IS_CLUSTER_DATABASE
  RETURN BOOLEAN;
```

Return Values

This function returns TRUE if this instance was started in cluster database mode; FALSE otherwise.

210.6.29 MAKE_DATA_BLOCK_ADDRESS Function

This function creates a data block address given a file number and a block number.

A data block address is the internal structure used to identify a block in the database. This function is useful when accessing certain fixed tables that contain data block addresses.

Syntax

```
DBMS_UTILITY.MAKE_DATA_BLOCK_ADDRESS (
   file  NUMBER,
   block NUMBER)
  RETURN NUMBER;
```

Parameters

Table 210-29 MAKE_DATA_BLOCK_ADDRESS Function Parameters

Parameter	Description
file	File that contains the block
block	Offset of the block within the file in terms of block increments

Pragmas

```
pragma restrict_references (make_data_block_address, WNDS, RNDS, WNPS, RNPS);
```

Return Values

Data block address.

210.6.30 NAME_RESOLVE Procedure

This procedure resolves the given name, including synonym translation and authorization checking as necessary.

Syntax

```
DBMS_UTILITY.NAME_RESOLVE (
   name          IN  VARCHAR2,
   context       IN  NUMBER,
   schema        OUT VARCHAR2,
   part1         OUT VARCHAR2,
   part2         OUT VARCHAR2,
   dblink        OUT VARCHAR2,
   part1_type    OUT NUMBER,
   object_number OUT NUMBER);
```

Parameters

Table 210-30 NAME_RESOLVE Procedure Parameters

Parameter	Description
`name`	Name of the object. This can be of the form [[a.]b.]c[@d], where a, b, c are SQL identifier and d is a dblink. No syntax checking is performed on the dblink. If a dblink is specified, or if the name resolves to something with a dblink, then object is not resolved, but the `schema`, `part1`, `part2` and `dblink` `OUT` parameters are filled in. a, b and c may be delimited identifiers, and may contain Globalization Support (NLS) characters (single and multibyte).
`context`	Must be an integer between 0 and 9. • 0 - table • 1 - PL/SQL (for 2 part names) • 2 - sequences • 3 - trigger • 4 - Java Source • 5 - Java resource • 6 - Java class • 7 - type • 8 - Java shared data • 9 - index
`schema`	Schema of the object: c. If no schema is specified in `name`, then the `schema` is determined by resolving the name.
`part1`	First part of the name. The type of this name is specified `part1_type` (synonym or package).
`part2`	If this is non-`NULL`, then this is a subprogram name. If part1 is non-`NULL`, then the subprogram is within the package indicated by part1. If part1 is `NULL`, then the subprogram is a top-level subprogram.
`dblink`	If this is non-`NULL`, then a database link was either specified as part of `name` or `name` was a synonym which resolved to something with a database link. In this case, if further name translation is desired, then you must call the `DBMS_UTILITY.NAME_RESOLVE` procedure on this remote node.

Table 210-30 (Cont.) NAME_RESOLVE Procedure Parameters

Parameter	Description
part1_type	Type of part1 is: • 5 - synonym • 7 - procedure (top level) • 8 - function (top level) • 9 - package
object_number	Object identifier

Exceptions

All errors are handled by raising exceptions. A wide variety of exceptions are possible, based on the various syntax error that are possible when specifying object names.

210.6.31 NAME_TOKENIZE Procedure

This procedure calls the parser to parse the given name as `a [. b [. c ]][@ dblink ]`.

It strips double quotes, or converts to uppercase if there are no quotes. It ignores comments of all sorts, and does no semantic analysis. Missing values are left as `NULL`.

Syntax

```
DBMS_UTILITY.NAME_TOKENIZE (
   name     IN  VARCHAR2,
   a        OUT VARCHAR2,
   b        OUT VARCHAR2,
   c        OUT VARCHAR2,
   dblink   OUT VARCHAR2,
   nextpos OUT BINARY_INTEGER);
```

Parameters

Table 210-31 NAME_RESOLVE Procedure Parameters

Parameter	Description
name	Input name, consisting of SQL identifiers (for example, scott.foo@dblink)
a	Output for the first token of the name
b	Output for the second token of the name (if applicable)
c	Output for the third token of the name (if applicable)
dblink	Output for the dblink of the name
nextpos	Next position after parsing the input name

210.6.32 OLD_CURRENT_SCHEMA Function

This function returns the session value from `sys_context` ('`userenv`', '`current_schema`').

Syntax

```
DBMS_UTILITY.OLD_CURRENT_SCHEMA
 RETURN VARCHAR2;
```

210.6.33 OLD_CURRENT_USER Function

This function returns the session value from `sys_context` ('`userenv`', '`current_user`').

Syntax

```
DBMS_UTILITY.OLD_CURRENT_USER
 RETURN VARCHAR2;
```

210.6.34 PORT_STRING Function

This function returns a string that identifies the operating system and the `TWO TASK PROTOCOL` version of the database. For example, "`VAX/VMX-7.1.0.0`"

The maximum length is port-specific.

Syntax

```
DBMS_UTILITY.PORT_STRING
   RETURN VARCHAR2;
```

Pragmas

```
pragma restrict_references(port_string, WNDS, RNDS, WNPS, RNPS);
```

210.6.35 SQLID_TO_SQLHASH Function

This function converts a SQL ID into a hash value.

Syntax

```
DBMS_UTILITY.SQLID_TO_SQLHASH (
   sql_id    IN   VARCHAR2)
 RETURN NUMBER;
```

Parameters

Table 210-32 SQLID_TO_SQLHASH Function Parameters

Parameter	Description
`sql_id`	SQL ID of a SQL statement. Must be `VARCHAR2(13)`.

210.6.36 TABLE_TO_COMMA Procedures

This procedure converts a PL/SQL table of names into a comma-delimited list of names.

This takes a PL/SQL table, `1..n`, terminated with `n+1 null`. The second version supports fully-qualified attribute names.

Syntax

```
DBMS_UTILITY.TABLE_TO_COMMA (
   tab    IN  UNCL_ARRAY,
   tablen OUT BINARY_INTEGER,
   list   OUT VARCHAR2);

DBMS_UTILITY.TABLE_TO_COMMA (
   tab    IN  lname_array,
   tablen OUT BINARY_INTEGER,
   list   OUT VARCHAR2);
```

Parameters

Table 210-33 TABLE_TO_COMMA Procedure Parameters

Parameter	Description
`tab`	PL/SQL table which contains list of table names
`tablen`	Number of tables in the PL/SQL table
`list`	Comma separated list of tables

Return Values

A comma-delimited list and the number of elements found in the table.

210.6.37 VALIDATE Procedure

This procedure makes invalid database objects valid.

Syntax

```
DBMS_UTILITY.VALIDATE(
    object_id       NUMBER);

DBMS_UTILITY.VALIDATE(
   owner          VARCHAR2,
   objname        VARCHAR2,
   namespace      NUMBER,   edition_name   := SYS_CONTEXT ('USERENV',
'CURRENT_EDITION'));
```

Parameters

Table 210-34 VALIDATE Procedure Parameters

Parameter	Description
`owner`	Name of the user who owns the object. Same as the `OWNER` field in `ALL_OBJECTS`.

Table 210-34 (Cont.) VALIDATE Procedure Parameters

Parameter	Description
objname	Name of the object to be validated. Same as the OBJECT_NAME field in ALL_OBJECTS.
namespace	Namespace of the object. Same as the namespace field in obj$. Equivalent numeric values are as follows: • 1 — TABLE/PROCEDURE/TYPE • 2 — BODY • 3 — TRIGGER • 4 — INDEX • 5 — CLUSTER • 8 — LOB • 9 — DIRECTORY • 10 — QUEUE • 11 — REPLICATION OBJECT GROUP • 12 — REPLICATION PROPAGATOR • 13 — JAVA SOURCE • 14 — JAVA RESOURCE • 58 — (Data Mining) MODEL
edition_name	[Note: Currently not operable. Reserved for future use]

Usage Notes

- No errors are raised if the object does not exist or is already valid or is an object that cannot be validated.
- If the object being validated is not actual in the specified edition, the subprogram automatically switches into the edition in which the object is actual prior to validation. That is, a call to VALIDATE will not actualize the object in the specified edition.
- The INVALIDATE Procedure invalidates a database object and optionally changes its PL/SQL compiler parameter settings. The object to be invalidated is specified by its object_id. The subprogram automatically switches to the edition in which the object is actual prior to invalidation. That is, a call to INVALIDATE will not actualize the object in the current edition.

210.6.38 WAIT_ON_PENDING_DML Function

This function waits until all transactions (other than the caller's own) that have locks on the listed tables and began prior to the specified scn have either committed or been rolled back.

Syntax

```
DBMS_UTILITY.WAIT_ON_PENDING_DML (
    tables     IN       VARCHAR2,
    timeout    IN       BINARY_INTEGER,
    scn        IN OUT   NUMBER)
  RETURN BOOLEAN;
```

Parameters

Table 210-35 WAIT_ON_PENDING_DML Function Parameters

Parameter	Description
tables	Comma-separated list of one or more table names. The list must be valid for COMMA_TO_TABLE Procedures, and each item valid to the NAME_RESOLVE Procedure. Neither column specifiers nor DBLINK (database link) specifiers are allowed in the names, and each name must resolve to an existing table in the local database.
timeout	Maximum number of seconds to wait, totalled across all tables/ transactions. A NULL or negative value will cause a very long wait.
scn	SCN prior to which transactions must have begun to be considered relevant to this request. If the value is NULL or not recognized as a meaningful scn on input, the most current SCN across all instances will be used and will be set into the passed argument as an output. If a meaningful value is passed in, its value will be preserved in the output.

Return Values

TRUE if all relevant transactions have committed or been rolled back, FALSE if the timeout occurred prior to all relevant transactions committing or being rolled back

211
DBMS_VECTOR

The `DBMS_VECTOR` package provides APIs to support common operations with Oracle AI Vector Search, such as extracting chunks or embeddings from user data, generating text for a given prompt, creating a vector index, or reporting on index accuracy.

These functions accept their respective input parameters in JSON format.

Related Topics

- *Oracle Database AI Vector Search User's Guide*

211.1 Summary of DBMS_VECTOR Subprograms

This table lists the `DBMS_VECTOR` subprograms and briefly describes them.

Table 211-1 ***DBMS_VECTOR Package Subprograms***

Subprogram	Description
ONNX Model Related Procedures:	
These procedures enable you to load an ONNX model into Oracle Database and drop the ONNX model.	
LOAD_ONNX_MODEL	Loads an ONNX model into the database
LOAD_ONNX_MODEL_CLOUD	Loads an ONNX model from object storage into the database
DROP_ONNX_MODEL Procedure	Drops the ONNX model
Chainable Utility (UTL) Functions:	
These functions are a set of modular and flexible functions within vector utility PL/SQL packages. You can chain these together to automate end-to-end data transformation and similarity search operations.	
UTL_TO_CHUNKS	Splits data into smaller pieces or chunks
UTL_TO_EMBEDDING and UTL_TO_EMBEDDINGS	Converts data to one or more vector embeddings
UTL_TO_GENERATE_TEXT	Generates text for a prompt (input string) or an image
Credential Helper Procedures:	
These procedures enable you to securely manage authentication credentials in the database. You require these credentials to enable access to third-party service providers for making REST calls.	
CREATE_CREDENTIAL	Creates a credential name
DROP_CREDENTIAL	Drops an existing credential name
Data Access Functions:	
These functions enable you to retrieve data, create index, and perform simple similarity search operations.	
CREATE_INDEX	Creates a vector index
REBUILD_INDEX	Rebuilds a vector index
GET_INDEX_STATUS	Describes the status of a vector index creation
ENABLE_CHECKPOINT	Enables the Checkpoint feature for a vector index user and index name

Table 211-1 (Cont.) *DBMS_VECTOR Package Subprograms*

Subprogram	Description
DISABLE_CHECKPOINT	Disables the Checkpoint feature for a vector index user and index name
INDEX_VECTOR_MEMORY_ADVISOR	Determines the vector memory size that is needed for a vector index
QUERY	Performs a similarity search query
RERANK	Reorders search results for more relevant output
Accuracy Reporting Function:	
These functions enable you to determine the accuracy of existing search indexes and to capture accuracy values achieved by approximate searches performed by past workloads.	
INDEX_ACCURACY_QUERY	Verifies the accuracy of a vector index
INDEX_ACCURACY_REPORT	Captures accuracy values achieved by approximate searches

> **Note:**
>
> `DBMS_VECTOR` is a lightweight package that does not support text processing or summarization operations. Therefore, the `UTL_TO_TEXT` and `UTL_TO_SUMMARY` chainable utility functions and all the chunker helper procedures are available only in the advanced `DBMS_VECTOR_CHAIN` package.

211.1.1 CREATE_CREDENTIAL

Use the `DBMS_VECTOR.CREATE_CREDENTIAL` credential helper procedure to create a credential name for storing user authentication details in Oracle Database.

Purpose

To securely manage authentication credentials in the database. You require these credentials to enable access during REST API calls to your chosen third-party service provider, such as Cohere, Google AI, Hugging Face, Oracle Cloud Infrastructure (OCI) Generative AI, OpenAI, or Vertex AI.

A credential name holds authentication parameters, such as user name, password, access token, private key, or fingerprint.

Note that if you are using Oracle Database as the service provider, then you do not need to create a credential.

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

```
DBMS_VECTOR.CREATE_CREDENTIAL (
    CREDENTIAL_NAME     IN VARCHAR2,
    PARAMS              IN JSON DEFAULT NULL
);
```

CREDENTIAL_NAME

Specify a name of the credential that you want to create for holding authentication parameters.

PARAMS

Specify authentication parameters in JSON format, based on your chosen service provider.

Generative AI requires the following authentication parameters:

```
{
"user_ocid"       : "<user ocid>",
"tenancy_ocid"    : "<tenancy ocid>",
"compartment_ocid": "<compartment ocid>",
"private_key"     : "<private key>",
"fingerprint"     : "<fingerprint>"
}
```

Cohere, Google AI, Hugging Face, OpenAI, and Vertex AI require the following authentication parameter:

```
{ "access_token": "<access token>" }
```

Table 211-2 Parameter Details

Parameter	Description
`user_ocid`	Oracle Cloud Identifier (OCID) of the user, as listed on the User Details page in the OCI console.
`tenancy_ocid`	OCID of your tenancy, as listed on the Tenancy Details page in the OCI console.
`compartment_ocid`	OCID of your compartment, as listed on the Compartments information page in the OCI console.

Table 211-2 (Cont.) Parameter Details

Parameter	Description

`private_key`	OCI private key.
	Note: The generated private key may appear as:
	
	`-----BEGIN RSA PRIVATE KEY-----`
	`<private key string>`
	`-----END RSA PRIVATE KEY-----`
	
	You pass the `<private key string>` value (excluding the `BEGIN` and `END` lines), either as a single line or as multiple lines.

`fingerprint`	Fingerprint of the OCI profile key, as listed on the User Details page under API Keys in the OCI console.
`access_token`	Access token obtained from your third-party service provider.

Required Privilege

You need the `CREATE CREDENTIAL` privilege to call this API.

Examples

- For Generative AI:

```
declare
  jo json_object_t;
begin
  jo := json_object_t();

jo.put('user_ocid','ocid1.user.oc1..aabbalbbaa1112233aabbaabb1111222aa1111b
b');

jo.put('tenancy_ocid','ocid1.tenancy.oc1..aaaaalbbbb1112233aaaabbaa1111222a
aa111a');

jo.put('compartment_ocid','ocid1.compartment.oc1..ababalabab1112233abababab
1111222aba11ab');
  jo.put('private_key','AAAaaaBBB11112222333...AAA111AAABBB222aaa1a/+');
  jo.put('fingerprint','01:1a:a1:aa:12:a1:12:1a:ab:12:01:ab:a1:12:ab:1a');
  dbms_vector.create_credential(
    credential_name   => 'OCI_CRED',
    params            => json(jo.to_string));
end;
/
```

- For Cohere:

```
declare
  jo json_object_t;
begin
  jo := json_object_t();
  jo.put('access_token', 'A1Aa0abA1AB1a1Abc123ab1A123ab123AbcA12a');
  dbms_vector.create_credential(
```

```
        credential_name   => 'COHERE_CRED',
        params            => json(jo.to_string));
    end;
    /
```

End-to-end examples:

To run end-to-end example scenarios using this procedure, see Use LLM-Powered APIs to Generate Summary and Text.

211.1.2 CREATE_INDEX

Use the `DBMS_VECTOR.CREATE_INDEX` procedure to create an index.

Syntax

```
DBMS_VECTOR.CREATE_INDEX (
    IDX_NAME                    IN VARCHAR2,
    TABLE_NAME                  IN VARCHAR2,
    IDX_VECTOR_COL              IN VARCHAR2,
    IDX_INCLUDE_COLS            IN VARCHAR2 DEFAULT NULL,
    IDX_PARTITIONING_SCHEME     IN VARCHAR2 default 'LOCAL',
    IDX_ORGANIZATION            IN VARCHAR2,
    IDX_DISTANCE_METRIC         IN VARCHAR2 DEFAULT COSINE,
    IDX_ACCURACY                IN NUMBER DEFAULT 90,
    IDX_PARAMETERS              IN CLOB,
    IDX_PARALLEL_CREATION       IN NUMBER DEFAULT 1
);
```

Parameters

Specify the input parameters in JSON format.

Table 211-3 CREATE_INDEX (IN) Parameters of DBMS_VECTOR

Parameter	Description
`table_name`	Table on which to create the index.
`idx_name`	Name of the index.
`idx_vector_col`	Vector column on which to create the index.
`idx_partitioning_scheme`	Partitioning scheme. `LOCAL` or `GLOBAL`.
`idx_organization`	Index organization. Either `NEIGHBOR PARTITIONS` or `INMEMORY NEIGHBOR GRAPH`.
`idx_distance_metric`	Distance computation metric. Defaults to `COSINE`. Can also be `MANHATTAN`, `HAMMING`, `DOT`, `EUCLIDEAN`, `L2_SQUARED`, `EUCLIDEAN_SQUARED`.
`idx_accuracy`	Target accuracy.
`idx_parameters`	Takes in CLOB with type of index and associated parameters. Type, neighbor partitions for IVF, neighbors/efConstruction for HNSW) .
`idx_parallel_creation`	Number of parallel threads used for index construction.

211.1.3 DISABLE_CHECKPOINT

Use the `DISABLE_CHECKPOINT` procedure to disable the Checkpoint feature for a given Hierarchical Navigable Small World (HNSW) index user and HNSW index name. This operation purges all older checkpoints for the HNSW index. It also disables the creation of future checkpoints as part of the HNSW graph refresh.

Syntax

```
DBMS_VECTOR.DISABLE_CHECKPOINT('INDEX_USER',['INDEX_NAME']);
```

INDEX_USER

Specify the user name of the HNSW vector index owner.

INDEX_NAME

Specify the name of the HNSW vector index for which you want to disable the Checkpoint feature.

The `INDEX_NAME` clause is optional. If you do not specify the index name, then this procedure disables the Checkpoint feature for all HNSW vector indexes under the given user.

Examples

- Using both the index name and index user:

  ```
  DBMS_VECTOR.DISABLE_CHECKPOINT('VECTOR_USER','VIDX1');
  ```

- Using only the index user:

  ```
  DBMS_VECTOR.DISABLE_CHECKPOINT('VECTOR_USER');
  ```

Related Topics

- *Oracle Database AI Vector Search User's Guide*
- ENABLE_CHECKPOINT
  Use the `ENABLE_CHECKPOINT` procedure to enable the Checkpoint feature for a given Hierarchical Navigable Small World (HNSW) index user and HNSW index name.

211.1.4 DROP_CREDENTIAL

Use the `DBMS_VECTOR.DROP_CREDENTIAL` credential helper procedure to drop an existing credential name from the data dictionary.

Syntax

```
DBMS_VECTOR.DROP_CREDENTIAL (
    CREDENTIAL_NAME      IN VARCHAR2
);
```

CREDENTIAL_NAME

Specify the credential name that you want to drop.

Examples

- For Generative AI:

```
exec dbms_vector.drop_credential('OCI_CRED');
```

- For Cohere:

```
exec dbms_vector.drop_credential('COHERE_CRED');
```

211.1.5 ENABLE_CHECKPOINT

Use the `ENABLE_CHECKPOINT` procedure to enable the Checkpoint feature for a given Hierarchical Navigable Small World (HNSW) index user and HNSW index name.

> **Note:**
>
> - This procedure only allows the index to create checkpoints. The checkpoint is created as part of the next HNSW graph refresh.
> - By default, HNSW checkpointing is enabled. If required, you can disable it using the `DBMS_VECTOR.DISABLE_CHECKPOINT` procedure.

Syntax

```
DBMS_VECTOR.ENABLE_CHECKPOINT('INDEX_USER',['INDEX_NAME']);
```

INDEX_USER

Specify the user name of the HNSW vector index owner.

INDEX_NAME

Specify the name of the HNSW vector index for which you want to enable the Checkpoint feature.

The `INDEX_NAME` clause is optional. If you do not specify the index name, then this procedure enables the Checkpoint feature for all HNSW vector indexes under the given user.

Examples

- Using both the index name and index user:

```
DBMS_VECTOR.ENABLE_CHECKPOINT('VECTOR_USER','VIDX1');
```

- Using only the index user:

```
DBMS_VECTOR.ENABLE_CHECKPOINT('VECTOR_USER');
```

Related Topics

- *Oracle Database AI Vector Search User's Guide*

- DISABLE_CHECKPOINT
  Use the `DISABLE_CHECKPOINT` procedure to disable the Checkpoint feature for a given Hierarchical Navigable Small World (HNSW) index user and HNSW index name. This operation purges all older checkpoints for the HNSW index. It also disables the creation of future checkpoints as part of the HNSW graph refresh.

211.1.6 DROP_ONNX_MODEL Procedure

This procedure deletes the specified ONNX model.

Syntax

```
DBMS_VECTOR.DROP_ONNX_MODEL (model_name IN VARCHAR2,
                                force      IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 211-4 DROP_ONNX_MODEL Procedure Parameters

Parameter	Description
`model_name`	Name of the machine learning ONNX model in the form [*schema_name*.]*model_name*. If you do not specify a schema, then your own schema is used.
`force`	Forces the machine learning ONNX model to be dropped even if it is invalid. An ONNX model may be invalid if a serious system error interrupted the model build process.

Usage Note

To drop an ONNX model, you must be the owner or you must have the `DB_DEVELOPER_ROLE`.

Example

You can use the following command to delete a valid ONNX model named `doc_model` that exists in your schema.

```
BEGIN
  DBMS_VECTOR.DROP_ONNX_MODEL(model_name => 'doc_model');
END;
/
```

211.1.7 GET_INDEX_STATUS

Use the `GET_INDEX_STATUS` procedure to query the status of a vector index creation.

Syntax

```
DBMS_VECTOR.GET_INDEX_STATUS ('USER_NAME','INDEX_NAME');
```

USER_NAME

Specify the user name of the vector index owner.

INDEX_NAME

Specify the name of the vector index. You can query the index creation status for both Hierarchical Navigable Small World (HNSW) indexes and Inverted File Flat (IVF) indexes.

Usage Notes

- You can use the `GET_INDEX_STATUS` procedure only during a vector index creation.
- The `Percentage` value is shown in the output only for Hierarchical Navigable Small World (HNSW) indexes (and not for Inverted File Flat (IVF) indexes).
- Along with the `DB_DEVELOPER_ROLE` privilege, you must have read access to the `VECSYS.VECTOR$INDEX$BUILD$` table.
- You can use the following query to view all auxiliary tables:

```
select IDX_AUXILIARY_TABLES from vecsys.vector$index;
```

  - For HNSW indexes:

    `rowid_vid_map` stores the mapping between a row ID and vector ID. `shared_journal_change_log` stores the DML changes that are yet to be incorporated into an HNSW graph.
  - For IVF indexes:

    `centroids` stores the location for each centroid. `centroid_partitions` stores the best centroid for each vector.
- The possible values of `Stage` for HNSW vector indexes are:

Value	Description
`HNSW Index Initialization`	Initialization phase for the HNSW vector index creation
`HNSW Index Auxiliary Tables Creation`	Creation of the internal auxiliary tables for the HNSW Neighbor Graph vector index
`HNSW Index Graph Allocation`	Allocation of memory from the vector memory pool for the HNSW graph
`HNSW Index Loading Vectors`	Loading of the base table vectors into the vector pool memory
`HNSW Index Graph Construction`	Creation of the multi-layered HNSW graph with the previously loaded vectors
`HNSW Index Creation Completed`	HNSW vector index creation finished

- The possible values of `Stage` for IVF vector indexes are:

Value	Description
`IVF Index Initialization`	Initialization phase for the IVF vector index creation
`IVF Index Centroids Creation`	The K-means clustering phase that computes the cluster centroids on a sample of base table vectors
`IVF Index Centroid Partitions Creation`	Centroids assignment phase for the base table vectors
`IVF Index Creation Completed`	IVF vector index creation completed

Example

```
exec DBMS_VECTOR.GET_INDEX_STATUS('VECTOR_USER','VIDX_HNSW');

Index objn: 74745
```

```
Stage: HNSW Index Loading Vectors
Percentage: 80%
```

211.1.8 INDEX_ACCURACY_QUERY

Use the `DBMS_VECTOR.INDEX_ACCURACY_QUERY` function to verify the accuracy of a vector index for a given query vector, top-K, and target accuracy.

Syntax

```
DBMS_VECTOR.INDEX_ACCURACY_QUERY (
    OWNER_NAME          IN VARCHAR2,
    INDEX_NAME          IN VARCHAR2,
    QV                  IN VECTOR,
    TOP_K               IN NUMBER,
    TARGET_ACCURACY     IN NUMBER
) return VARCHAR2;

DBMS_VECTOR.INDEX_ACCURACY_QUERY (
    OWNER_NAME          IN VARCHAR2,
    INDEX_NAME          IN VARCHAR2,
    QV                  IN VECTOR,
    TOP_K               IN NUMBER,
    QUERY_PARAM         IN JSON
) return VARCHAR2;
```

Parameters

Table 211-5 INDEX_ACCURACY_QUERY (IN) Parameters of DBMS_VECTOR

Parameter	Description
`owner_name`	The name of the vector index owner.
`index_name`	The name of the vector index.
`qv`	Specifies the query vector.
`top_k`	The `top_k` value for accuracy computation.
`target_accuracy`	The target accuracy value for the vector index.

For information about determining the accuracy of your vector indexes, see Index Accuracy Report in *Oracle Database AI Vector Search User's Guide*.

211.1.9 INDEX_ACCURACY_REPORT

Use the `DBMS_VECTOR.INDEX_ACCURACY_REPORT` function to capture from your past workloads, accuracy values achieved by approximate searches using a particular vector index for a certain period of time.

Syntax

```
DBMS_VECTOR.INDEX_ACCURACY_REPORT (
    OWNER_NAME          IN VARCHAR2,
    INDEX_NAME          IN VARCHAR2,
    START_TIME          IN TIMESTAMP WITH TIME ZONE,
    END_TIME            IN TIMESTAMP WITH TIME ZONE
) return NUMBER;
```

Parameters

Table 211-6 INDEX_ACCURACY_REPORT (IN) Parameters of DBMS_VECTOR

Parameter	Description
`owner_name`	The name of the vector index owner.
`index_name`	The name of the vector index.
`start_time`	Specifies from what time to capture query vectors to consider for the accuracy computation. A `NULL` `start_time` uses query vectors captured in the last 24 hours.
`end_time`	Specifies an end point up until which query vectors are considered for accuracy computation. A `NULL` `end_time` uses query vectors captured from `start_time` until the current time.

For information about determining the accuracy of your vector indexes, see Index Accuracy Report in *Oracle Database AI Vector Search User's Guide*.

211.1.10 INDEX_VECTOR_MEMORY_ADVISOR

Use the `INDEX_VECTOR_MEMORY_ADVISOR` procedure to determine the vector memory size needed for a particular vector index. This helps you evaluate the number of indexes that can fit for each simulated vector memory size.

Syntax

- Using the number and type of vector dimensions that you want to store in your vector index.

```
DBMS_VECTOR.INDEX_VECTOR_MEMORY_ADVISOR(
    INDEX_TYPE,
    NUM_VECTORS,
    DIM_COUNT,
    DIM_TYPE,
    PARAMETER_JSON,
    RESPONSE_JSON);
```

- Using the table and vector column on which you want to create your vector index:

```
DBMS_VECTOR.INDEX_VECTOR_MEMORY_ADVISOR(
    TABLE_OWNER,
    TABLE_NAME,
    COLUMN_NAME,
    INDEX_TYPE,
    PARAMETER_JSON,
    RESPONSE_JSON);
```

Table 211-7 Syntax Details: INDEX_VECTOR_MEMORY_ADVISOR

Parameter	Description
INDEX_TYPE	Type of vector index: • IVF for Inverted File Flat (IVF) vector indexes • HNSW for Hierarchical Navigable Small World (HNSW) vector indexes
NUM_VECTORS	Number of vectors that you plan to create the vector index with.
DIM_COUNT	Number of dimensions of a vector as a NUMBER.
DIM_TYPE	Type of dimensions of a vector. Possible values are: • FLOAT16 • FLOAT32 • FLOAT64 • INT8
TABLE_OWNER	Owner name of the table on which to create the vector index.
TABLE_NAME	Table name on which to create the vector index.
COLUMN_NAME	Name of the vector column on which to create the vector index.
PARAMETER_JSON	Input parameter in JSON format. You can specify only one of the following form: • PARAMTER_JSON=>{"accuracy":value} • INDEX_TYPE=>IVF, parameter_json=>{"neighbor_partitions":value} • INDEX_TYPE=>HNSW, parameter_json=>{"neighbors":value} **Note**: You cannot specify values for accuracy along with neighbor_partitions or neighbors.
RESPONSE_JSON	JSON-formatted response string.

Examples

- Using neighbors in the parameters list:

```
exec DBMS_VECTOR.INDEX_VECTOR_MEMORY_ADVISOR(
    INDEX_TYPE=>'HNSW',
    NUM_VECTORS=>10000,
    DIM_COUNT=>768,
    DIM_TYPE=>'FLOAT32',
    PARAMETER_JSON=>'{"neighbors":128}',
    RESPONSE_JSON=>:response_json);
```

```
Suggested vector memory pool size: 59918628 Bytes
```

- Using accuracy in the parameters list:

```
exec DBMS_VECTOR.INDEX_VECTOR_MEMORY_ADVISOR(
    INDEX_TYPE=>'HNSW',
    NUM_VECTORS=>10000,
    DIM_COUNT=>768,
    DIM_TYPE=>'FLOAT32',
    PARAMETER_JSON=>'{"accuracy":90}',
    RESPONSE_JSON=>:response_json);

Suggested vector memory pool size: 53926765 Bytes
```

- Using the table and vector column on which you want to create the vector index:

```
exec DBMS_VECTOR.INDEX_VECTOR_MEMORY_ADVISOR(
    'VECTOR_USER',
    'VECTAB',
    'DATA_VECTOR',
    'HNSW',
    RESPONSE_JSON=>:response_json);

Using default accuracy: 90%
Suggested vector memory pool size: 76396251 Bytes
```

Related Topics

- *Oracle Database AI Vector Search User's Guide*

211.1.11 LOAD_ONNX_MODEL

This procedure enables you to load an ONNX model into the Database.

Syntax

```
DBMS_VECTOR.LOAD_ONNX_MODEL (
     directory  VARCHAR2,
     file_name   VARCHAR2,
     model_name  VARCHAR2,
     metadata   JSON);


DBMS_VECTOR.LOAD_ONNX_MODEL(
model_name  IN  VARCHAR2,
model_data  IN  BLOB,
metadata    IN  JSON);
```

Parameters

Table 211-8 LOAD_ONNX_MODEL Procedure Parameters

Parameter	Description
`directory`	The directory name of the data dump. For example, `DM_DUMP`.

Table 211-8 (Cont.) LOAD_ONNX_MODEL Procedure Parameters

Parameter	Description
file_name	A VARCHAR2 type parameter that specifies the name of the ONNX model.
model_name	Name of the model in the form [schema_name.]model_name. If you do not specify a schema, then your own schema is used.
model_data	It is a BLOB holding the ONNX representation of the model. The BLOB contains the identical byte sequence as the one stored in an ONNX file.
metadata	A JSON description of the metadata describing the model. The metadata at minimum must describe the machine learning function supported by the model. The model's metadata parameters are described in JSON Metadata Parameters for ONNX Models.

Examples

The following examples illustrates a code snippet of using the DBMS_VECTOR.LOAD_ONNX_MODEL procedure. The complete step-by-step example is illustrated in Import ONNX Models and Generate Embeddings.

```
EXECUTE DBMS_VECTOR.LOAD_ONNX_MODEL(
    'DM_DUMP',
    'my_embedding_model.onnx',
    'doc_model',
    JSON('{"function" : "embedding", "embeddingOutput" : "embedding", "input": {"input":
["DATA"]}}'));


DBMS_VECTOR.LOAD_ONNX_MODEL('my_embedding_model.onnx',
                                                :blob_bind_variable,
                                                 JSON('{"function" :
"embedding",
                                                        "embeddingOutput" :
"embedding" ,
                                                         "input":{"input":
["DATA"]}}'));
```

For a complete example to illustrate how you can define a BLOB variable and use it in the LOAD_ONNX_MODEL procedure, you can have the following:

```
CREATE OR REPLACE MY_LOAD_EMBEDDING_MODEL(embedding_model_name VARCHAR2,
onnx_blob BLOB) IS
BEGIN
DBMS_VECTOR.LOAD_ONNX_MODEL(embedding_model_name,
                            onnx_blob,
                            JSON('{"function" : "embedding",
                                   "embeddingOutput" : "embedding" ,
                                   "input":{"input": ["DATA"]}}'));
END;
/
```

Usage Notes

The name of the model follows the same restrictions as those used for other machine learning models, namely:

- The schema name, if provided, is limited to 128 characters.
- The model name is limited to 123 characters and must follow the rules of unquoted identifiers: they contain only alphanumeric characters, the underscore (_), dollar sign ($), and pound sign (#). The initial character must be alphabetic.
- The model size is limited to 1 gigabyte.
- The model must not depend on external initializers. To know more about initializers and other ONNX concepts, see https://onnx.ai/onnx/intro/concepts.html.
- There are default input and output names for input and output attributes for models that are prepared by the Python utility. You can load those models without the JSON parameters. For example:

```
EXECUTE DBMS_VECTOR.LOAD_ONNX_MODEL('DM_DUMP', 'my_embedding_model.onnx',
'doc_model'));
```

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for examples of using ONNX models for machine learning tasks

211.1.11.1 JSON Metadata Parameters for ONNX Models

When importing models using the `IMPORT_ONNX_MODEL` (`DBMS_DATA_MINING`), `LOAD_ONNX_MODEL` (`DBMS_VECTOR`), or `LOAD_ONNX_MODEL_CLOUD` (`DBMS_VECTOR`) procedures, you supply metadata as JSON parameters.

Parameters

Field	Value Type	Description
`function`	String	Specify regression, classification, clustering, or embedding. This is a mandatory setting. **NOTE**: The only JSON parameter required when importing the model is the machine learning function.
`input`	NA	Describes the model input mapping. See "Input" in Usage Notes.
`regressionOutput`	String	The name of the regression model output that stores the regression results. The output is expected to be a tensor of supported shape of any supported regression output type. See "Output" in Usage Notes.
`classificationProbOutput`	String	The name of the classification model output storing probabilities. The output is expected to be a tensor value of type float (width 32/64) of supported shape. See "Automatic normalization of output probabilities" in Usage Notes.

Field	Value Type	Description
clusteringDistanceOutput	String	The name of the clustering model output storing distances. The output is of type float (width 16/32/64) of supported shape.
clusteringProbOutput	String	The name of the clustering model output storing probabilities. The output is of type float (width 16/32/64) of supported shape.
classificationLabelOutput	String	The name of the model output holding label information. You have the following metadata parameters to specify the labels for classification: • labels: specify the labels directly in the JSON metadata • classificationLabelOutput: specify the model output that provides labels If you do not specify any value for this parameter or the function of the model is not classification, you will receive an error. The user can specify to use labels from the model directly by setting classificationLabelOutput to the model output holding the label information. The tensor output holding the label information must be the same size as the number of classes and must be of integer or string type. If the tensor that holds the labels is of string type, the returned type of the PREDICTION operator is VARCHAR2. If the tensor that holds the labels is of integer type, the returned type of the PREDICTION operator is NUMBER.
normalizeProb	String	Describes automatic normalization of output probabilities. See "Automatic normalization of output probabilities" in Usage Notes.
labels	NA	The labels used for classification. If you want to use custom labels, specify the labels using the labels field in the JSON metadata. The field can be set to an array of length equal to the number of classes. The labels for the class i must be stored at index i of the label array. If an array of strings is used, the returned type of the PREDICTION operator is VARCHAR2. The size of the string labels specified by the user cannot exceed *4000* bytes. If an array of numbers is used, the returned type of the PREDICTION operator is NUMBER. If you do not specify labels or classificationLabelOutput, classes are identified by integers in the range 1 to N where N is the number of classes. In this case, the returned type of the PREDICTION operator is NUMBER.
embeddingOutput	String	The model output that holds the generated embeddings.

Field	Value Type	Description
`suitableDistanceMetrics`	String	An array of names of suitable distance metrics for the model. The names must be the names of the distance metrics used for the Oracle `VECTOR_DISTANCE` operator. To know the supported distance metrics, see Vector Distance Metrics.
`normalization`	Boolean	A boolean value indicates if normalization is applied to the output vector. The value 1 means normalization is applied. Normalization is process of converting an embedding vector so that it's norm or length equals 1. A normalized vector maintains its direction but its length becomes 1. The resulting vector is often called a unit vector.
`maxSequenceLength`	Number	The maximum length of the token (input) sequence that is meaningful for the model. This parameter sets a limit on the number of tokens, words, or elements in each input sequence that the model will process. This ensures uniform input size for the model. For example, the value could be 128, or 512 to 4096 depending on the task for which the parameter is used. A machine translation model might have a `maxSequenceLength` of 512, accommodating sentences or paragraphs up to 512 tokens for translation tasks.
`pooling`	String	Indicates the pooling function performed on the output vector.
`modelDescription`	Object	A JSON object that allows users to add additional descriptions to the models complementing the existing ONNX metadata for model description.
`languages`	String	A comma-separated list of language name or abbreviation, as described in "A.1 Languages" of *Oracle Database Globalization Support Guide*. If you import multi-lingual embedding model, specify the language or the language abbreviation as the metadata.
`tokenizer`	String	Tokenizers help in transforming text into words. There are several tokenizers available, including: bert, gpt2, bpe, wordpiece, sentencepiece, and clip.
`embeddingLayer`	String	An identifier for the embedding layer. An embedding layer, serving as a hidden layer in neural networks, transforms input data from high to lower dimensions, enhancing the network's understanding of input relationships and data processing efficiency. Embedding layer helps in processing and analyzing categorical or discrete data. It achieves this by transforming categories into continuous embeddings, capturing the essential semantic relationships and similarities between them. For example the last hidden state in some transformer, or a layer in a resnet network.

Field	Value Type	Description
defaultOnNull	NA	Specify the replacement of missing values in the JSON using the defaultOnNull field. If defaultOnNull is not specified, the replacement of missing values is not performed. The defaultOnNull sets the missing values to NULL by default. You can override the default value of NULL by providing meaningful default values to substitute for NULL. The field must be a JSON object literal, whose fields are the input attribute names and whose values are the default values for the input. Note that the default value is of type string and must be a valid Oracle PL/SQL NVL value for the given datatype.

Note: The parameters are case-sensitive. A number of default conventions for output parameter names and default values allows to minimize the information that you may have to provide. The parameters such as suitableDistanceMetrics are informational only and you are not expected to provide this information while importing the model. The JSON descriptor may specify only one input attribute. If more are specified, you will receive an error. You will receive an error if the normalizeProb field is specified as the JSON metadata parameter.

Usage Notes

The name of the model follows the same restrictions as those used for other machine learning models, namely:

- **Input**

  When importing a model from an ONNX representation, you must specify the name of the attribute used for scoring and how it maps to actual ONNX inputs. A scoring operator uses these attribute names to identify the columns to be used. (For example, PREDICTION). Follow these conventions to specify the attribute names using the input field:

  not specified: When the field input is not specified, attribute names are mapped directly to model inputs by name. That is, if the attribute name is not specified in the JSON metadata, then the name of the input tensor is used as an attribute name. Each model input must have dimension [batch_size, value]. If you do not specify input in the JSON metatdata, the value must be 1. You don't have to specify extra metadata if the input of the model already conforms to the format. For an embedding model, a single input is provided that may be used in batches. Here, if the input parameter is not specified in the JSON metadata, the valid model will have [batch_size, 1].

  You must ensure that all attribute names, whether implied by the model or explicitly set by you through the input field, are valid Oracle Database identifiers for column names. Each attribute name within a model must be unique, ensuring no duplicates exist.

  You can explicitly specify attribute name for model that use input tensors that have a dimension larger than 1 (for example, (batch_size, 2)). In this case, you must specify a name for each of these values for them to be interpreted as independent attribute name. This can be done for regression, classification, clustering which are models whose scoring operation can take multiple input attributes.

- **Output**

  As models might have multiple outputs, you can specify which output is of interest for a specific machine learning technique. You have the following ways to specify model outputs:

- Specify the output name of interest in the JSON during model import. If the specified name is not a valid model output (see the table with valid outputs for a given machine learning function), you will receive an error.
- If the model produces an output that matches the expected output name for the given machine learning technique (for example, `classificationProbOutput`) and you didn't explicitly specify it, the output is automatically assumed.
- If you do not specify any output name and the model has a single output, the system assumes that the single output corresponds to a default specific to the machine learning technique. For an embedding machine learning function, the default value is `embeddingOutput`.
  The system reports an error if you do not specify model outputs or if you supply outputs that the specified machine learning function does not support. The following table displays supported outputs for a specific machine learning function:

Machine learning function	Output
regression	`regressionOutput`
classification	`classificationProbOutput`
clustering	`clusteringDistanceOutput`
embedding	`embeddingOutput`

  If none of the mentioned model outputs are specified, or if you supply outputs that are not supported by the specified machine learning function, you will receive an error.

- **Automatic Normalization of Output Probabilities**

  Many users widely employ the softmax function to normalize the output of multi-class classification models, as it enables to easily interpret the results of these models. The **softmax function** is a mathematical function that converts a vector of real numbers into a probability distribution. It is also known as the softargmax, or normalized exponential function. This function is available to you to specify at the model import-time that a softmax normalization must be applied to the tensor holding output probabilities such as `classificationProbOutput` and `clusteringProbOutput`. Specify `normalizeProb` to define the normalization that must be applied for softmax normalization. The default setting is `none`, indicating that no normalization is applied. You can choose `softmax` to apply a softmax function to the probability output. Specifying any other value for this field will result in an error during import. Additionally, specifying this field for models other than classification and clustering will also lead to an error.

Example: Specifying JSON Metadata Parameters for Embedding Models

The following example illustrates a simple case of how you can specify JSON metadata parameters while importing an ONNX embedding model into the Database using the `DBMS_VECTOR.IMPORT_ONNX_MODEL` procedure.

```
DBMS_VECTOR.IMPORT_ONNX_MODEL('my_embedding_model.onnx', 'doc_model',
                JSON('{"function" : "embedding",
                       "embeddingOutput" : "embedding" ,
                        "input":{"input": ["DATA"]}}'));
```

Example: Specifying Complete JSON Metadata Parameters for Embedding Models

The following example illustrates how to provide a complete JSON metadata parameters, with an exception of `embeddingLayer`, for importing embedding models.

```
DECLARE
  metadata JSON;
  mdtxt varchar2(4000);
BEGIN
  metadata := JSON(q'#
           {
             "function"                : "embedding",
             "embeddingOutput"         : "embedding",
             "input"                   : { "input" : ["txt"]},
             "maxSequenceLength"       : 512,
             "tokenizer"               : "bert",
             "suitableDistanceMetrics" : [ "DOT", "COSINE", "EUCLIDEAN"],
             "pooling"                 : "Mean Pooling",
             "normalization"           : true,
             "languages"               : ["US"],
             "modelDescription"        : {
                 "description" : "This model was tuned for semantic search:
Given a query/question, if can find relevant passages. It was trained on a
large and diverse set of (question, a
nswer) pairs.",
                 "url" : "https://example.co/sentence-transformers/
my_embedding_model"}
           }
           #');
  -- load the onnx model
    DBMS_VECTOR.IMPORT_ONNX_MODEL('my_embedding_model.onnx', 'doc_model',
metadata);
END;
/
```

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for examples of using ONNX models for machine learning tasks

211.1.12 LOAD_ONNX_MODEL_CLOUD

This procedure enables you to load an ONNX model from object storage into the Database.

Syntax

```
DBMS_VECTOR.LOAD_ONNX_MODEL_CLOUD (
     model_name  IN  VARCHAR2,
     credential  IN  VARCHAR2,
     uri         IN  VARCHAR2,
     metadata    IN  JSON DEFAULT JSON('{"function" : "embedding", '||
                          '"embeddingOutput" : "embedding", "input": {"input":
["DATA"]}}')
);
```

Parameters

Table 211-9 LOAD_ONNX_MODEL_CLOUD Procedure Parameters

Parameter	Description
`model_name`	The name of the model in the form `[schema_name.]model_name`. If you do not specify a schema, then your own schema is used.
`credential`	The name of the credential to be used to access Object Store.
`uri`	The URI of the ONNX model.
`metadata`	A JSON description of the metadata describing the model. The metadata at minimum must describe the machine learning function supported by the model. The model's metadata parameters are described in JSON Metadata Parameters for ONNX Models.

Examples

The following example includes a code snippet of using the `DBMS_VECTOR.LOAD_ONNX_MODEL_CLOUD` procedure.

```
EXECUTE DBMS_VECTOR.LOAD_ONNX_MODEL_CLOUD(
    model_name => 'database',
    credential => 'MYCRED',
    uri => 'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-string/b/
bucketname/o/all-MiniLM-L6-v2.onnx',
    metadata => JSON('{"function" : "embedding", "embeddingOutput" : "embedding" , "input":
{"input": ["DATA"]}}')
);
```

Usage Notes

The name of the model follows the same restrictions as those used for other machine learning models, namely:

- The schema name, if provided, is limited to 128 characters.
- The model name is limited to 123 characters and must follow the rules of unquoted identifiers: they contain only alphanumeric characters, the underscore (_), dollar sign ($), and pound sign (#). The initial character must be alphabetic.
- The model size is limited to 1 gigabyte.

- The model must not depend on external initializers. To know more about initializers and other ONNX concepts, see https://onnx.ai/onnx/intro/concepts.html.
- There are default input and output names for input and output attributes for models that are prepared by the Python utility. You can load those models without the JSON parameters. For example:

```
EXECUTE DBMS_VECTOR.LOAD_ONNX_MODEL_CLOUD(
    'database',
    'MYCRED',
    'https://objectstorage.us-phoenix-1.oraclecloud.com/n/namespace-
string/b/bucketname/o/all-MiniLM-L6-v2.onnx'
);
```

> **See Also:**
>
> *Oracle Machine Learning for SQL User's Guide* for examples of using ONNX models for machine learning tasks

211.1.13 QUERY

Use the `DBMS_VECTOR.QUERY` function to perform a similarity search operation which returns the top-k results as a JSON array.

Syntax

Query is overloaded and supports a version with `query_vector` passed in as a `VECTOR` type in addition to `CLOB`.

```
DBMS_VECTOR.QUERY (
    TAB_NAME             IN VARCHAR2,
    VEC_COL_NAME         IN VARCHAR2,
    QUERY_VECTOR         IN CLOB,
    TOP_K                IN NUMBER,
    VEC_PROJ_COLS        IN JSON_ARRAY_T DEFAULT NULL,
    IDX_NAME             IN VARCHAR2 DEFAULT NULL,
    DISTANCE_METRIC      IN VARCHAR2 DEFAULT 'COSINE',
    USE_INDEX            IN BOOLEAN DEFAULT TRUE,
    ACCURACY             IN NUMBER DEFAULT '90',
    IDX_PARAMETERS       IN CLOB DEFAULT NULL
) return JSON_ARRAY_T;

DBMS_VECTOR.QUERY (
    TAB_NAME             IN VARCHAR2,
    VEC_COL_NAME         IN VARCHAR2,
    QUERY_VECTOR         IN VECTOR,
    TOP_K                IN NUMBER,
    VEC_PROJ_COLS        IN JSON_ARRAY_T DEFAULT NULL,
    IDX_NAME             IN VARCHAR2 DEFAULT NULL,
    DISTANCE_METRIC      IN VARCHAR2 DEFAULT 'COSINE',
    USE_INDEX            IN BOOLEAN DEFAULT TRUE,
    ACCURACY             IN NUMBER DEFAULT '90',
```

```
    IDX_PARAMETERS       IN CLOB DEFAULT NULL
) return JSON_ARRAY_T;
```

Parameters

Specify the input parameters in JSON format.

Table 211-10 DBMS_VECTOR.QUERY Parameters

Parameter	Description
tab_name	Table name to query
vec_col_name	Vector column name
query_vector	Query vector passed in as CLOB or VECTOR.
top_k	Number of results to be returned.
vec_proj_cols	Columns to be projected as part of the result.
idx_name	Name of the index queried.
distance_metric	Distance computation metric. Defaults to COSINE. Can also be MANHATTAN, HAMMING, DOT, EUCLIDEAN, L2_SQUARED, EUCLIDEAN_SQUARED. .
use_index	Specifies whether the search is an approximate search or exact search. Defaults to TRUE (that is, approximate).
accuracy	Specifies the minimum desired query accuracy.
idx_parameters	Specifies values of efsearch and neighbor partition probes passed in, formatted as JSON

DATA

This function accepts the input data type as VARCHAR2, NUMBER, JSON, BOOLEAN or CLOB.

211.1.14 REBUILD_INDEX

Use the DBMS_VECTOR.REBUILD_INDEX function to rebuild an index.

Syntax

This API rebuilds an index. In case only the idx_name is provided, it rebuilds the index using get_ddl. When all params are provided, it performs a drop index followed by a call to dbms_vector.create_index().

```
DBMS_VECTOR.REBUILD_INDEX (
    IDX_NAME                    IN VARCHAR2,
    TABLE_NAME                  IN VARCHAR2 DEFAULT NULL,
    IDX_VECTOR_COL              IN VARCHAR2 DEFAULT NULL,
    IDX_INCLUDE_COLS            IN VARCHAR2 DEFAULT NULL,
    IDX_PARTITIONING_SCHEME     IN VARCHAR2 DEFAULT NULL,
    IDX_ORGANIZATION            IN VARCHAR2 DEFAULT NULL,
    IDX_DISTANCE_METRIC         IN VARCHAR2 DEFAULT 'COSINE',
    IDX_ACCURACY                IN NUMBER DEFAULT 90,
    IDX_PARAMETERS              IN CLOB DEFAULT NULL,
```

```
    IDX_PARALLEL_CREATION      IN NUMBER DEFAULT 1,
);
```

Parameters

Specify the input parameters in JSON format.

Table 211-11 REBUILD_INDEX (IN) Parameters of DBMS_VECTOR

Parameter	Description
table_name	Table on which to create the index.
idx_name	Name of the index.
idx_vector_col	Vector column on which to rebuild the index.
idx_partitioning_scheme	Partitioning scheme. LOCAL or GLOBAL.
idx_organization	Index organization. Either NEIGHBOR PARTITIONS or INMEMORY NEIGHBOR GRAPH.
idx_distance_metric	Distance computation metric. Defaults to COSINE. Can also be MANHATTAN, HAMMING, DOT, EUCLIDEAN, L2_SQUARED, EUCLIDEAN_SQUARED.
idx_accuracy	Target accuracy.
idx_parameters	Takes in CLOB with type of index and associated parameters. Type, neighbor partitions for IVF, neighbors/efConstruction for HNSW) .
idx_parallel_creation	Number of parallel threads used for index construction.

DATA

This function accepts the input data type as VARCHAR2, NUMBER, or CLOB.

If you specify only the idx_name value, then the function rebuilds the index using get_ddl. If you specify all the parameter values, then the function first performs a drop index operation and then calls the DBMS_VECTOR.CREATE_INDEX() function.

211.1.15 RERANK

Use the DBMS_VECTOR.RERANK function to reassess and reorder an initial set of results to retrieve more relevant search output.

Purpose

To improve the relevance and quality of search results in both similarity search and Retrieval Augmented Generation (RAG) scenarios.

Reranking improves the quality of information ingested into an LLM by ensuring that the most relevant documents or chunks are prioritized. This helps to reduce hallucinations and improves the accuracy of generated outputs.

For this operation, Oracle AI Vector Search supports reranking models provided by Cohere and Vertex AI.

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

```
DBMS_VECTOR.RERANK(
                QUERY      IN CLOB,
                DOCUMENTS  IN JSON,
                PARAMS     IN JSON default NULL
) return JSON;
```

This function accepts the input containing a query as `CLOB` and a list of documents in `JSON` format. It then processes this information to generate a `JSON` object containing a reranked list of documents, sorted by score.

For example, a reranked output includes:

```
{
    "index"   : "1",
    "score"   : "0.99",
    "content" : "Jupiter boasts an impressive system of 95 known moons."
}
```

Where,

- `index` specifies the position of the document in the list of input text.
- `score` specifies the relevance score.
- `content` specifies the input text corresponding to the index.

QUERY

Specify the search query (typically from an initial search) as `CLOB`.

DOCUMENTS

Specify a JSON array of strings (list of potentially relevant documents to rerank) in the following format:

```
{
  "documents": [
  "string1",
  "string2",
    ...
```

```
  ]
}
```

PARAMS

Specify the following list of parameters in JSON format. All these parameters are mandatory.

```
{
  "provider"         : "<service provider>",
  "credential_name"  : "<credential name>",
  "url"              : "<REST endpoint URL for reranking>",
  "model"            : "<reranking model name>",
  ...
}
```

Table 211-12 RERANK Parameter Details

Parameter	Description
`provider`	Supported REST provider to access for reranking: • `cohere` • `vertexai`
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the reranking model in the form: *`schema.model_name`* If the model name is not schema-qualified, then the schema of the procedure invoker is used.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters for reranking.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"         : "cohere",
  "credential_name"  : "COHERE_CRED",
  "url"              : "https://api.cohere.example.com/rerank",
  "model"            : "rerank-model",
  "return_documents": false,
  "top_n"            : 3
}
```

Vertex AI example:

```
{
  "provider"          : "vertexai",
  "credential_name"   : "VERTEXAI_CRED",
  "url"               : "https://googleapis.example.com/
default_ranking_config:rank",
  "model"             : "rerank-model",
  "ignoreRecordDetailsInResponse" : true,
  "topN"              : 3
  }
```

Table 211-13 Additional REST Provider Parameter Details

Parameter	Description
`return_documents`	Whether to return search results with original documents or input text (`content`): • `false` (default, also recommended) to not return any input text (return only index and score) • `true` to return input text along with index and score **Note**: With Cohere as the provider, Oracle recommends that you keep this option disabled for better performance. You may choose to enable it for debugging purposes when you need to view the original text.
`ignoreRecordDetailsInResponse`	Whether to return search results with original record details or input text (`content`): • `false` (default) to return input text along with index and score • `true` (recommended) to not return any input text (return only index and score) **Note**: With Vertex AI as the provider, Oracle recommends that you keep this option enabled for better performance. You may choose to disable it for debugging purposes when you need to view the original text.
`top_n` or `topN`	The number of most relevant documents to return.

Examples

- Using Cohere:

```
declare
  params clob;
  reranked_output json;
begin
  params := '
```

```
{
  "provider": "cohere",
  "credential_name": "COHERE_CRED",
  "url": "https://api.cohere.com/v1/rerank",
  "model": "rerank-english-v3.0",
  "return_documents": true,
  "top_n": 3
}';

  reranked_output := dbms_vector.rerank(:query,
json(:initial_retrieval_docs), json(params));
  dbms_output.put_line(json_serialize(reranked_output));
end;
/
```

- Using Vertex AI:

```
declare
  params clob;
  reranked_output json;
begin
  params := '
{
  "provider": "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url": "https://discoveryengine.googleapis.com/v1/projects/1085581009881/
locations/global/rankingConfigs/default_ranking_config:rank",
  "model": "semantic-ranker-512@latest",
  "ignoreRecordDetailsInResponse": false,
  "topN": 3
}';

  reranked_output := dbms_vector.rerank(:query,
json(:initial_retrieval_docs), json(params));
  dbms_output.put_line(json_serialize(reranked_output));
end;
/
```

End-to-end example:

To run an end-to-end example scenario using this function, see Use Reranking for Better RAG Results.

211.1.16 UTL_TO_CHUNKS

Use the `DBMS_VECTOR.UTL_TO_CHUNKS` chainable utility function to split a large plain text document into smaller chunks of text.

Purpose

To perform a text-to-chunks transformation. This chainable utility function internally calls the `VECTOR_CHUNKS` SQL function for the operation.

To embed a large document, you may first need to split it into multiple appropriate-sized segments or chunks through a splitting process known as chunking (as explained in Understand the Stages of Data Transformations). A chunk can be words (to capture specific

words or word pieces), sentences (to capture a specific meaning), or paragraphs (to capture broader themes). A single document may be split into multiple chunks, each transformed into a vector.

Syntax

```
DBMS_VECTOR.UTL_TO_CHUNKS (
    DATA         IN CLOB | VARCHAR2
    PARAMS       IN JSON default NULL
) return VECTOR_ARRAY_T;
```

DATA

This function accepts the input data type as `CLOB` or `VARCHAR2`.

It returns an array of `CLOB`s, where each `CLOB` contains a chunk along with its metadata in JSON format, as follows:

```
{
    "chunk_id"     : NUMBER,
    "chunk_offset" : NUMBER,
    "chunk_length" : NUMBER,
    "chunk_data"   : "VARCHAR2(4000)"
}
```

For example:

```
{"chunk_id":1,"chunk_offset":1,"chunk_length":6,"chunk_data":"sample"}
```

Where,

- `chunk_id` specifies the chunk ID for each chunk.
- `chunk_offset` specifies the original position of each chunk in the source document, relative to the start of document which has a position of `1`.
- `chunk_length` specifies the character length of each chunk.
- `chunk_data` displays text pieces from each chunk.

PARAMS

Specify input parameters in JSON format:

```
{
    "by"           :     mode,
    "max"          :     max,
    "overlap"      :     overlap,
    "split"        :     split_condition,
    "custom_list"  :     [ split_chars1, ... ],
    "vocabulary"   :     vocabulary_name,
    "language"     :     nls_language,
    "normalize"    :     normalize_mode,
    "norm_options" :     [ normalize_option1, ... ],
    "extended"     :     boolean
}
```

For example:

```
JSON('
  { "by"           :    "vocabulary",
    "vocabulary"   :    "myvocab",
    "max"          :    "100",
    "overlap"      :    "0",
    "split"        :    "custom",
    "custom_list"  :    [ "<p>" , "<s>" ],
    "language"     :    "american",
    "normalize"    :    "options",
    "norm_options" :    [ "WHITESPACE" ]
  }')
```

Here is a complete description of these parameters:

Parameter	Description and Acceptable Values
`by`	Specify a mode for splitting your data, that is, to split by counting the number of characters, words, or vocabulary tokens. **Valid values**: • `characters` (or `chars`): Splits by counting the number of characters. • `words`: Splits by counting the number of words. Words are defined as sequences of alphabetic characters, sequences of digits, individual punctuation marks, or symbols. For segmented languages without whitespace word boundaries (such as Chinese, Japanese, or Thai), each native character is considered a word (that is, unigram). • `vocabulary`: Splits by counting the number of vocabulary tokens. Vocabulary tokens are words or word pieces, recognized by the vocabulary of the tokenizer that your embedding model uses. You can load your vocabulary file using the chunker helper API `DBMS_VECTOR_CHAIN.CREATE_VOCABULARY`. **Note**: For accurate results, ensure that the chosen model matches the vocabulary file used for chunking. If you are not using a vocabulary file, then ensure that the input length is defined within the token limits of your model. **Default value**: `words`
`max`	Specify a limit on the maximum size of each chunk. This setting splits the input text at a fixed point where the maximum limit occurs in the larger text. The units of `max` correspond to the `by` mode, that is, to split data when it reaches the maximum size limit of a certain number of characters, words, numbers, punctuation marks, or vocabulary tokens. **Valid values**: • `by characters`: `50` to `4000` characters • `by words`: `10` to `1000` words • `by vocabulary`: `10` to `1000` tokens **Default value**: `100`

Parameter	Description and Acceptable Values

split [by]	Specify where to split the input text when it reaches the maximum size limit. This helps to keep related data together by defining appropriate boundaries for chunks.
	Valid values:
	• none:
	Splits at the max limit of characters, words, or vocabulary tokens.
	• newline, blankline, and space:
	These are single-split character conditions that split at the last split character before the max value.
	Use newline to split at the end of a line of text. Use blankline to split at the end of a blank line (sequence of characters, such as two newlines). Use space to split at the end of a blank space.
	• recursively:
	This is a multiple-split character condition that breaks the input text using an ordered list of characters (or sequences).
	recursively is predefined as BLANKLINE, newline, space, none in this order:
	1. If the input text is more than the max value, then split by the first split character.
	2. If that fails, then split by the second split character.
	3. And so on.
	4. If no split characters exist, then split by max wherever it appears in the text.
	• sentence:
	This is an end-of-sentence split condition that breaks the input text at a sentence boundary.
	This condition automatically determines sentence boundaries by using knowledge of the input language's sentence punctuation and contextual rules. This language-specific condition relies mostly on end-of-sentence (EOS) punctuations and common abbreviations.
	Contextual rules are based on word information, so this condition is only valid when splitting the text by words or vocabulary (not by characters).
	Note: This condition obeys the by word and max settings, and thus may not determine accurate sentence boundaries in some cases. For example, when a sentence is larger than the max value, it splits the sentence at max. Similarly, it includes multiple sentences in the text only when they fit within the max limit.
	• custom:
	Splits based on a custom split characters list. You can provide custom sequences up to a limit of 16 split character strings, with a maximum length of 10 each.
	Specify an array of valid text literals using the custom_list parameter.
	`{ "split" : "custom", "custom_list" : [ "split_chars1", ... ] }`
	For example:
	`{ "split" : "custom", "custom_list" : [ "<p>" , "<s>" ] }`
	Note: You can omit sequences only for tab (\t), newline (\n), and linefeed (\r).
	Default value: recursively

Parameter	Description and Acceptable Values
`overlap`	Specify the amount (as a positive integer literal or zero) of the preceding text that the chunk should contain, if any. This helps in logically splitting up related text (such as a sentence) by including some amount of the preceding chunk text. The amount of overlap depends on how the maximum size of the chunk is measured (in characters, words, or vocabulary tokens). The overlap begins at the specified `split` condition (for example, at `newline`). **Valid value**: `5%` to `20%` of `max` **Default value**: `0`
`language`	Specify the language of your input data. This clause is important, especially when your text contains certain characters (for example, punctuations or abbreviations) that may be interpreted differently in another language. **Valid values**: • NLS-supported language name or its abbreviation, as listed in *Oracle Database Globalization Support Guide*. • Custom language name or its abbreviation, as listed in Supported Languages and Data File Locations. You use the `DBMS_VECTOR_CHAIN.CREATE_LANG_DATA` chunker helper API to load language-specific data (abbreviation tokens) into the database, for your specified language. **Default value**: `NLS_LANGUAGE` from session

Parameter	Description and Acceptable Values

normalize	Automatically pre-processes or post-processes issues (such as multiple consecutive spaces and smart quotes) that may arise when documents are converted into text. Oracle recommends you to use a normalization mode to extract high-quality chunks.
	Valid values:
	• none:
	Applies no normalization.
	• all:
	Normalizes common multi-byte (unicode) punctuation to standard single-byte.
	• options:
	Specify an array of normalization options using the norm_options parameter.
	`{ "normalize" : "options", "norm_options" : [ "normalize_option1", ... ] }`
	– punctuation:
	Includes smart quotes, smart hyphens, and other multi-byte equivalents to simple single-byte punctuation.
	For example:
	* 2018u 'map to 0027'
	* 2019u 'map to 0027'
	* 201Bu 'map to 0027'
	– whitespace:
	Minimizes whitespace by eliminating unnecessary characters.
	For example, retain blanklines, but remove any extra newlines and interspersed spaces or tabs: " \n \n " => "\n\n"
	– widechar:
	Normalizes wide, multi-byte digits and (a-z) letters to single-byte.
	These are multi-byte equivalents for 0-9 and a-z A-Z, which can show up in ZH/JA formatted text.
	For example:
	`{ "normalize" : "options", "norm_options" : [ "whitespace" ] }`
	Default value: None

extended	Increases the output limit of a VARCHAR2 string to 32767 bytes, without requiring you to set the max_string_size parameter to extended. **Default value**: 4000 or 32767 (when max_string_size=extended)

Example

```
SELECT D.id doc,
    JSON_VALUE(C.column_value, '$.chunk_id' RETURNING NUMBER) AS id,
    JSON_VALUE(C.column_value, '$.chunk_offset' RETURNING NUMBER) AS pos,
    JSON_VALUE(C.column_value, '$.chunk_length' RETURNING NUMBER) AS siz,
    JSON_VALUE(C.column_value, '$.chunk_data') AS txt
```

```
FROM docs D,
   dbms_vector.utl_to_chunks(D.text,
   JSON('{ "by"       : "words",
           "max"      : "100",
           "overlap"  : "0",
           "split"    : "recursively",
           "language" : "american",
           "normalize": "all" }')) C;
```

End-to-end examples:

To run end-to-end example scenarios using this function, see Perform Chunking With Embedding and Configure Chunking Parameters.

Related Topics

- VECTOR_CHUNKS

211.1.17 UTL_TO_EMBEDDING and UTL_TO_EMBEDDINGS

Use the `DBMS_VECTOR.UTL_TO_EMBEDDING` and `DBMS_VECTOR.UTL_TO_EMBEDDINGS` chainable utility functions to generate one or more vector embeddings from textual documents and images.

Purpose

To automatically generate vector embeddings from a set of textual documents and images.

- Text to Vector:

  You can perform a text-to-embedding transformation by accessing:

  - Oracle Database as the service provider: Calls an ONNX format embedding model that you load into the database (default setting)
  - Third-party embedding model: Makes a REST API call to either a remote service provider (Cohere, Google AI, Hugging Face, Generative AI, OpenAI, or Vertex AI) or a local service provider (Ollama)

- Image to Vector:

  You can also perform an image-to-embedding transformation. For this operation, this API makes a REST call to your chosen image embedding model or multimodal embedding model by Vertex AI (which is the only supported third-party provider).

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

- Text to Vector:

```
DBMS_VECTOR.UTL_TO_EMBEDDING (
    DATA           IN CLOB,
    PARAMS         IN JSON default NULL
) return VECTOR;


DBMS_VECTOR.UTL_TO_EMBEDDINGS (
    DATA           IN VECTOR_ARRAY_T,
    PARAMS         IN JSON default NULL
) return VECTOR_ARRAY_T;
```

- Image to Vector:

```
DBMS_VECTOR.UTL_TO_EMBEDDING (
    DATA           IN BLOB,
    MODALITY       IN VARCHAR2,
    PARAMS         IN JSON default NULL
) return VECTOR;
```

DATA

`UTL_TO_EMBEDDING` accepts the input as `CLOB` containing textual data (text strings or small documents) or as `BLOB` containing media data (for media files such as images). It then converts the text or image input to a single embedding (`VECTOR`).

`UTL_TO_EMBEDDINGS` converts an array of chunks (`VECTOR_ARRAY_T`) to an array of embeddings (`VECTOR_ARRAY_T`).

A generated embedding output includes:

```
{
    "embed_id"    :  NUMBER,
    "embed_data"  : "VARCHAR2(4000)",
    "embed_vector": "CLOB"
}
```

Where,

- `embed_id` displays the ID number of each embedding.
- `embed_data` displays the input text that is transformed into embeddings.
- `embed_vector` displays the generated vector representations.

MODALITY

For `BLOB` inputs, specify the type of content to vectorize. The only supported value is `image`.

PARAMS

Specify input parameters in JSON format, depending on the service provider that you want to use.

If using Oracle Database as the provider:

```
{
  "provider" : "database",
  "model"    : "<in-database ONNX embedding model filename>"
}
```

Table 211-14 Database Provider Parameter Details

Parameter	Description
`provider`	Specify `database` (default setting) to use Oracle Database as the provider. With this setting, you must load an ONNX format embedding model into the database.
`model`	User-specified name under which the imported ONNX embedding model is stored in Oracle Database. If you do not have an embedding model in ONNX format, then perform the steps listed in Convert Pretrained Models to ONNX Format.

If using a third-party provider:

Set the following parameters along with additional embedding parameters specific to your provider:

- For `UTL_TO_EMBEDDING`:

```
{
  "provider"        : "<AI service provider>",
  "credential_name" : "<credential name>",
  "url"             : "<REST endpoint URL for embedding service>",
  "model"           : "<REST provider embedding model name>",
  "transfer_timeout": <maximum wait time for the request to complete>,
  "<additional REST provider parameter>": "<REST provider parameter
value>"
}
```

- For `UTL_TO_EMBEDDINGS`:

```
{
  "provider"        : "<AI service provider>",
  "credential_name" : "<credential name>",
  "url"             : "<REST endpoint URL for embedding service>",
  "model"           : "<REST provider embedding model name>",
  "transfer_timeout": <maximum wait time for the request to complete>,
  "batch size"      : "<number of vectors to request at a time>",
  "<additional REST provider parameter>": "<REST provider parameter
value>"
}
```

Table 211-15 Third-Party Provider Parameter Details

Parameter	Description
`provider`	Third-party service provider that you want to access for this operation. A REST call is made to the specified provider to access its embedding model. For image input, specify `vertexai`. For text input, specify one of the following values: • `cohere` • `googleai` • `huggingface` • `ocigenai` • `openai` • `vertexai`
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the third-party embedding model in the form: *`schema.model_name`* If you do not specify a schema, then the schema of the procedure invoker is used. **Note**: • For accurate results, ensure that the chosen text embedding model matches the vocabulary file used for chunking. If you are not using a vocabulary file, then ensure that the input length is defined within the token limits of your model. • To get image embeddings, you can use any image embedding model or multimodal embedding model supported by Vertex AI. Multimodal embedding is a technique that vectorizes data from different modalities such as text and images. When using a multimodal embedding model to generate embeddings, ensure that you use the same model to vectorize both types of content (text and images). By doing so, the resulting embeddings are compatible and situated in the same vector space, which allows for effective comparison between the two modalities during similarity searches.
`transfer_timeout`	Maximum time to wait for the request to complete. The default value is `60` seconds. You can increase this value for busy web servers.
`batch_size`	Maximum number of vectors to request at a time. For example, for a batch size of `50`, if `100` chunks are passed, then this API sends two requests with an array of `50` strings each. If `30` chunks are passed (which is lesser than the defined batch size), then the API sends those in a single request. For REST calls, it is more efficient to send a batch of inputs at a time rather than requesting a single input per call. Increasing the batch size can provide better performance, whereas reducing the batch size may reduce memory and data usage, especially if your provider has a rate limit. The default or maximum allowed value depends on the third-party provider settings.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"       : "cohere",
  "credential_name": "COHERE_CRED",
  "url"            : "https://api.cohere.example.com/embed",
  "model"          : "embed-model",
  "input_type"     : "search_query"
}
```

Google AI example:

```
{
  "provider"       : "googleai",
  "credential_name": "GOOGLEAI_CRED",
  "url"            : "https://googleapis.example.com/models/",
  "model"          : "embed-model"
}
```

Hugging Face example:

```
{
  "provider"       : "huggingface",
  "credential_name": "HF_CRED",
  "url"            : "https://api.huggingface.example.com/",
  "model"          : "embed-model"
}
```

Generative AI example:

```
{
  "provider"       : "ocigenai",
  "credential_name": "OCI_CRED",
  "url"            : "https://generativeai.oci.example.com/embedText",
  "model"          : "embed-model",
  "batch_size"     : 10
}
```

Ollama example:

```
{
  "provider"       : "ollama",
  "host"           : "local",
```

```
  "url"            : "http://localhost:11434/api/embeddings",
  "model"          : "llama3"
}
```

OpenAI example:

```
{
  "provider"       : "openai",
  "credential_name": "OPENAI_CRED",
  "url"            : "https://api.openai.example.com/embeddings",
  "model"          : "embed-model"
}
```

Vertex AI example:

```
{
  "provider"       : "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url"            : "https://googleapis.example.com/models/",
  "model"          : "embed-model"
}
```

Table 211-16 Additional REST Provider Parameter Details

Parameter	Description
`input_type`	Type of input to vectorize.

> **Note:**
>
> The generated embedding results may be different between requests for the same input and configuration, depending on your embedding model or floating point precision. However, this does not affect your queries (and provides semantically correct results) because the vector distance will be similar.

Examples

You can use `UTL_TO_EMBEDDING` in a `SELECT` clause and `UTL_TO_EMBEDDINGS` in a `FROM` clause, as follows:

- The following examples use `UTL_TO_EMBEDDING` to generate an embedding with `Hello world` as input, by accessing the Generative AI embedding model:

```
-- declare embedding parameters

var params clob;

begin
 :params := '
{
  "provider": "ocigenai",
  "credential_name": "OCI_CRED",
```

```
  "url": "https://generativeai.oci.example.com/embedText",
  "model": "embed.modelname",
  "batch_size": 10
}';
end;
/

-- get text embedding: PL/SQL example

declare
  input clob;
  v vector;
begin
  input := 'Hello world';

  v := dbms_vector.utl_to_embedding(input, json(params));
  dbms_output.put_line(vector_serialize(v));
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/

-- get text embedding: select example

select dbms_vector.utl_to_embedding('Hello world', json(:params)) from
dual;
```

The following examples use `UTL_TO_EMBEDDING` to generate an embedding with `parrots.jpg` as input and modality as `image`, by accessing the Vertex AI's multimodal embedding model.

```
-- declare embedding parameters

var params clob;

begin
  :params := '
{
  "provider": "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url": "https://LOCATION-aiplatform.googleapis.com/v1/projects/PROJECT/
locations/LOCATION/publishers/google/models/",
  "model": "multimodalembedding:predict"
}';
end;
/

-- get image embedding: PL/SQL example

declare
  v vector;
  output clob;
begin
```

```
  v := dbms_vector.utl_to_embedding(
    to_blob(bfilename('VEC_DUMP', 'parrots.jpg')), 'image', json(:params));
  output := vector_serialize(v);
  dbms_output.put_line('vector data=' || dbms_lob.substr(output, 100) ||
'...');
end;
/

-- get image embedding: select example

select dbms_vector.utl_to_embedding(
  to_blob(bfilename('VEC_DUMP', 'parrots.jpg')), 'image', json(:params));
```

End-to-end examples:

To run end-to-end example scenarios using `UTL_TO_EMBEDDING`, see Generate Embedding.

- The following example uses `UTL_TO_EMBEDDINGS` to generate embeddings with a PDF document (stored in the `documentation_tab` table) as input, by calling an ONNX format model loaded into the database:

```
SELECT et.* from documentation_tab dt,dbms_vector.utl_to_embeddings(
  dbms_vector.utl_to_chunks(dbms_vector.utl_to_text(dt.data)),
  json(:embed_params)) et;
```

End-to-end examples:

To run end-to-end example scenarios using `UTL_TO_EMBEDDINGS`, see Perform Chunking With Embedding.

211.1.18 UTL_TO_GENERATE_TEXT

Use the `DBMS_VECTOR.UTL_TO_GENERATE_TEXT` chainable utility function to generate a text response for a given prompt or an image, by accessing third-party text generation models.

Purpose

To communicate with Large Language Models (LLMs) through natural language conversations. You can generate a textual answer, description, or summary for prompts and images, given as input to LLM-powered chat interfaces.

- Prompt to Text:

  A prompt can be an input text string, such as a question that you ask an LLM. For example, "`What is Oracle Text?`". A prompt can also be a command, such as "`Summarize the following ...`", "`Draft an email asking for ...`", or "`Rewrite the following ...`", and can include results from a search. The LLM responds with a textual answer or description based on the specified task in the prompt.

  For this operation, this API makes a REST call to either a remote third-party provider (Cohere, Google AI, Hugging Face, Generative AI, OpenAI, or Vertex AI) or a local third-party provider (Ollama).

- Image to Text:

  You can also prompt with a media file, such as an image, to extract text from pictures or photos. You supply a text question as the prompt (such as "`What is this image about?`"

or "`How many birds are there in this painting?`") along with the image. The LLM responds with a textual analysis or description of the contents of the image.

For this operation, this API makes a REST call to either a remote third-party provider (Google AI, Hugging Face, OpenAI, or Vertex AI) or a local third-party provider (Ollama).

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

This function accepts the input as `CLOB` containing text data (for textual prompts) or as `BLOB` containing media data (for media files such as images). It then processes this information to generate a new `CLOB` containing the generated text.

- Prompt to Text:

```
DBMS_VECTOR.UTL_TO_GENERATE_TEXT (
        DATA          IN CLOB,
        PARAMS        IN JSON default NULL
) return CLOB;
```

- Image to Text:

```
DBMS_VECTOR.UTL_TO_GENERATE_TEXT(
          TEXT_DATA      IN CLOB,
          MEDIA_DATA     IN BLOB,
          MEDIA_TYPE     IN VARCHAR2 default 'image/jpeg',
          PARAMS         IN JSON default NULL
) return CLOB;
```

DATA and TEXT_DATA

Specify the textual prompt as `CLOB` for the `DATA` or `TEXT_DATA` clause.

> **Note:**
>
> Hugging Face uses an image captioning model that does not require a prompt, when giving an image as input. If you input a prompt along with an image, then the prompt will be ignored.

MEDIA_DATA

Specify the `BLOB` file, such as an image or a visual PDF file.

MEDIA_TYPE

Specify the image format for the given image or visual PDF file (`BLOB` file) in one of the supported image data MIME types. For example:

- For PNG: `image/png`
- For JPEG: `image/jpeg`
- For PDF: `application/pdf`

> **Note:**
>
> For a complete list of the supported image formats, refer to your third-party provider's documentation.

PARAMS

Specify the following input parameters in JSON format, depending on the service provider that you want to access for text generation:

```
{
  "provider"         : "<AI service provider>",
  "credential_name"  : "<credential name>",
  "url"              : "<REST endpoint URL for text generation service>",
  "model"            : "<text generation model name>",
  "transfer_timeout" : <maximum wait time for the request to complete>,
  "<additional REST provider parameter>": "<REST provider parameter value>"
}
```

Table 211-17 UTL_TO_GENERATE_TEXT Parameter Details

Parameter	Description
`provider`	Supported REST provider that you want to access to generate text. Specify one of the following values: For `CLOB` input: • `cohere` • `googleai` • `huggingface` • `ocigenai` • `openai` • `vertexai` For `BLOB` input: • `googleai` • `huggingface` • `openai` • `vertexai`

Table 211-17 (Cont.) UTL_TO_GENERATE_TEXT Parameter Details

Parameter	Description
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the third-party text generation model in the form: *`schema.model_name`* If the model name is not schema-qualified, then the schema of the procedure invoker is used.
`transfer_timeout`	Maximum time to wait for the request to complete. The default value is `60` seconds. You can increase this value for busy web servers.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"        : "cohere",
  "credential_name" : "COHERE_CRED",
  "url"             : "https://api.cohere.example.com/chat",
  "model"           : "generate-text-model"
}
```

Google AI example:

```
{
  "provider"       : "googleai",
  "credential_name": "GOOGLEAI_CRED",
  "url"            : "https://googleapis.example.com/models/",
  "model"          : "generate-text-model"
}
```

Hugging Face example:

```
{
  "provider"       : "huggingface",
  "credential_name": "HF_CRED",
  "url"            : "https://api.huggingface.example.com/models/",
  "model"          : "generate-text-model"
}
```

Generative AI example:

```
{
  "provider"        : "ocigenai",
  "credential_name" : "GENAI_CRED",
  "url"             : "https://generativeai.oci.example.com/generateText",
  "model"           : "generate-text-model",
  "inferenceRequest": {
    "maxTokens"     : 300,
    "temperature"   : 1
  }
}
```

or more current chat models may implement:

```
{
  "provider"        : "ocigenai",
  "credential_name" : "OCI_CRED",
  "url"             : "https://inference.generativeai.oci.example.com/actions/
chat",
  "model"           : "generate-text-model",
  "chatRequest"     : {
    "maxTokens"     : 256
  }
}
```

Ollama example:

```
{
  "provider"       : "ollama",
  "host"           : "local",
  "url"            : "http://localhost:11434/api/generate",
  "model"          : "llama3"
}
```

OpenAI example:

```
{
  "provider"        : "openai",
  "credential_name" : "OPENAI_CRED",
  "url"             : "https://api.openai.example.com",
  "model"           : "generate-text-model",
  "max_tokens"      : 60,
```

```
  "temperature"     : 1.0
}
```

Vertex AI example:

```
{
  "provider"         : "vertexai",
  "credential_name"  : "VERTEXAI_CRED",
  "url"              : "https://googleapis.example.com/models/",
  "model"            : "generate-text-model",
  "generation_config": {
    "temperature"    : 0.9,
    "topP"           : 1,
    "candidateCount" : 1,
    "maxOutputTokens": 256
  }
}
```

Table 211-18 Additional REST Provider Parameter Details

Parameter	Description
`max_tokens`	Maximum number of tokens in the output text.
`temperature`	Degree of randomness used when generating the output text, in the range of `0.0-5.0`. To generate the same output for a prompt, use `0`. To generate a random new text for that prompt, increase the temperature. **Note**: Start with the temperature set to `0`. If you do not require random results, a recommended temperature value is between `0` and `1`. A higher value is not recommended because a high temperature may produce creative text, which might also include hallucinations.
`topP`	Probability of tokens in the output, in the range of `0.0-1.0`. A lower value provides less random responses and a higher value provides more random responses.
`candidateCount`	Number of response variations to return, in the range of `1-4`.
`maxOutputTokens`	Maximum number of tokens to generate for each response.

Examples

- Prompt to Text:

  The following statements generate text by making a REST call to Generative AI. Here, the prompt is "`What is Oracle Text?`".

```
-- select example

var params clob;
exec :params := '
{
  "provider": "ocigenai",
  "credential_name": "OCI_CRED",
  "url": "https://generativeai.oci.example.com/generateText",
  "model": "generate.modelname"
```

```
}';

select dbms_vector.utl_to_generate_text(
 'What is Oracle Text?',
 json(:params)) from dual;

-- PL/SQL example

declare
  input clob;
  params clob;
  output clob;
begin
  input := 'What is Oracle Text?';

  params := '
{
  "provider": "ocigenai",
  "credential_name": "OCI_CRED",
  "url": "https://generativeai.oci.example.com/generateText",
  "model": "generate.modelname"
}';

  output := dbms_vector.utl_to_generate_text(input, json(params));
  dbms_output.put_line(output);
  if output is not null then
    dbms_lob.freetemporary(output);
  end if;
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/
```

End-to-end examples:

To run end-to-end example scenarios, see Generate Text Response.

- Image to Text:

  The following statements generate text by making a REST call to OpenAI. Here, the input is an image (`sample_image.jpeg`) along with the prompt "`Describe this image?`".

```
-- select example

var input clob;
var media_data blob;
var media_type clob;
var params clob;

begin
  :input := 'Describe this image';
  :media_data := load_blob_from_file('DEMO_DIR', 'sample_image.jpeg');
  :media_type := 'image/jpeg';
  :params := '
{
```

```
  "provider": "openai",
  "credential_name": "OPENAI_CRED",
  "url": "https://api.openai.example.com/chat/completions",
  "model": "generate.modelname",
  "max_tokens": 256
}';
end;
/

select dbms_vector.utl_to_generate_text(:input, :media_data, :media_type,
json(:params));

-- PL/SQL example

declare
  input clob;
  media_data blob;
  media_type varchar2(32);
  params clob;
  output clob;

begin
  input := 'Describe this image';
  media_data := load_blob_from_file('DEMO_DIR', 'image_file');
  media_type := 'image/jpeg';
  params := '
{
  "provider": "openai",
  "credential_name": "OPENAI_CRED",
  "url": "https://api.openai.example.com/chat/completions",
  "model": "generate.modelname",
  "max_tokens": 256
}';

  output := dbms_vector.utl_to_generate_text(
    input, media_data, media_type, json(params));
  dbms_output.put_line(output);

  if output is not null then
    dbms_lob.freetemporary(output);
  end if;
  if media_data is not null then
    dbms_lob.freetemporary(media_data);
  end if;
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/
```

End-to-end examples:

To run end-to-end example scenarios, see Describe Image Content.

212
DBMS_VECTOR_CHAIN

The DBMS_VECTOR_CHAIN package provides APIs to support advanced operations with Oracle AI Vector Search, such as chunking and embedding data along with text generation and summarization capabilities. It is more suitable for text processing with similarity search and hybrid search, using functionality that can be pipelined together for an end-to-end search.

These functions accept their respective input parameters in JSON format.

Related Topics

- *Oracle Database AI Vector Search User's Guide*

212.1 Summary of DBMS_VECTOR_CHAIN Subprograms

This table lists the DBMS_VECTOR_CHAIN subprograms and briefly describes them.

Table 212-1 ***DBMS_VECTOR_CHAIN Package Subprograms***

Subprogram	Description
Chainable Utility (UTL) Functions:	
These functions are a set of modular and flexible functions within vector utility PL/SQL packages. You can chain these together to automate end-to-end data transformation and similarity search operations.	
UTL_TO_TEXT	Extracts plain text data from documents
UTL_TO_CHUNKS	Splits data into smaller pieces or chunks
UTL_TO_EMBEDDING and UTL_TO_EMBEDDINGS	Converts data to one or more vector embeddings
UTL_TO_SUMMARY	Extracts a summary from documents
UTL_TO_GENERATE_TEXT	Generates text for a prompt (input string) or an image
Credential Helper Procedures:	
These procedures enable you to securely manage authentication credentials in the database. You require these credentials to enable access to third-party service providers for making REST calls.	
CREATE_CREDENTIAL	Creates a credential name
DROP_CREDENTIAL	Drops an existing credential name
Preference Helper Procedures:	
These procedures enable you to manage vectorizer preferences, to be used with the CREATE_HYBRID_VECTOR_INDEX and ALTER_INDEX SQL statements when creating or managing hybrid vector indexes.	
CREATE_PREFERENCE	Creates a vectorizer preference
DROP_PREFERENCE	Drops an existing vectorizer preference
Chunker Helper Procedures:	
These procedures enable you to configure vocabulary and language data (abbreviations), to be used with the VECTOR_CHUNKS SQL function or UTL_TO_CHUNKS PL/SQL function.	
CREATE_VOCABULARY	Loads your token vocabulary file into the database
DROP_VOCABULARY	Removes existing vocabulary data

Table 212-1 (Cont.) *DBMS_VECTOR_CHAIN Package Subprograms*

Subprogram	Description
CREATE_LANG_DATA	Loads your language data file into the database
DROP_LANG_DATA	Removes existing abbreviation data
Data Access Function:	
This function enables you to enhance search operations.	
RERANK	Reorders search results for more relevant output

> **Note:**
>
> The `DBMS_VECTOR_CHAIN` package requires you to install the `CONTEXT` component of Oracle Text, an Oracle Database technology that provides indexing, term extraction, text analysis, text summarization, word and theme searching, and other utilities.
>
> Due to underlying dependance on the text processing capabilities of Oracle Text, note that both the `UTL_TO_TEXT` and `UTL_TO_SUMMARY` chainable utility functions and all the chunker helper procedures are available only in this package through Oracle Text.

212.1.1 CREATE_CREDENTIAL

Use the `DBMS_VECTOR_CHAIN.CREATE_CREDENTIAL` credential helper procedure to create a credential name for storing user authentication details in Oracle Database.

Purpose

To securely manage authentication credentials in the database. You require these credentials to enable access during REST API calls to your chosen third-party service provider, such as Cohere, Google AI, Hugging Face, Oracle Cloud Infrastructure (OCI) Generative AI, OpenAI, or Vertex AI.

A credential name holds authentication parameters, such as user name, password, access token, private key, or fingerprint.

Note that if you are using Oracle Database as the service provider, then you do not need to create a credential.

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

```
DBMS_VECTOR_CHAIN.CREATE_CREDENTIAL (
    CREDENTIAL_NAME     IN VARCHAR2,
    PARAMS              IN JSON DEFAULT NULL
);
```

CREDENTIAL_NAME

Specify a name of the credential that you want to create for holding authentication parameters.

PARAMS

Specify authentication parameters in JSON format, based on your chosen service provider.

Generative AI requires the following authentication parameters:

```
{
"user_ocid"       : "<user ocid>",
"tenancy_ocid"    : "<tenancy ocid>",
"compartment_ocid": "<compartment ocid>",
"private_key"     : "<private key>",
"fingerprint"     : "<fingerprint>"
}
```

Cohere, Google AI, Hugging Face, OpenAI, and Vertex AI require the following authentication parameter:

```
{ "access_token": "<access token>" }
```

Table 212-2 Parameter Details

Parameter	Description
`user_ocid`	Oracle Cloud Identifier (OCID) of the user, as listed on the User Details page in the OCI console.
`tenancy_ocid`	OCID of your tenancy, as listed on the Tenancy Details page in the OCI console.
`compartment_ocid`	OCID of your compartment, as listed on the Compartments information page in the OCI console.

Table 212-2 (Cont.) Parameter Details

Parameter	Description

`private_key`	OCI private key.
	Note: The generated private key may appear as:
	
	`-----BEGIN RSA PRIVATE KEY-----`
	`<private key string>`
	`-----END RSA PRIVATE KEY-----`
	
	You pass the *`<private key string>`* value (excluding the `BEGIN` and `END` lines), either as a single line or as multiple lines.

`fingerprint`	Fingerprint of the OCI profile key, as listed on the User Details page under API Keys in the OCI console.
`access_token`	Access token obtained from your third-party service provider.

Required Privilege

You need the `CREATE CREDENTIAL` privilege to call this API.

Examples

- For Generative AI:

```
declare
  jo json_object_t;
begin
  jo := json_object_t();

jo.put('user_ocid','ocid1.user.oc1..aabbalbbaa1112233aabbaabb1111222aa1111b
b');

jo.put('tenancy_ocid','ocid1.tenancy.oc1..aaaaalbbbb1112233aaaabbaa1111222a
aa111a');

jo.put('compartment_ocid','ocid1.compartment.oc1..ababalabab1112233abababab
1111222aba11ab');
  jo.put('private_key','AAAaaaBBB11112222333...AAA111AAABBB222aaa1a/+');
  jo.put('fingerprint','01:1a:a1:aa:12:a1:12:1a:ab:12:01:ab:a1:12:ab:1a');
  dbms_vector_chain.create_credential(
    credential_name   => 'OCI_CRED',
    params            => json(jo.to_string));
end;
/
```

- For Cohere:

```
declare
  jo json_object_t;
begin
  jo := json_object_t();
  jo.put('access_token', 'A1Aa0abA1AB1a1Abc123ab1A123ab123AbcA12a');
  dbms_vector_chain.create_credential(
```

```
        credential_name   => 'COHERE_CRED',
        params            => json(jo.to_string));
    end;
    /
```

End-to-end examples:

To run end-to-end example scenarios using this procedure, see Use LLM-Powered APIs to Generate Summary and Text.

212.1.2 CREATE_LANG_DATA

Use the `DBMS_VECTOR_CHAIN.CREATE_LANG_DATA` chunker helper procedure to load your own language data file into the database.

Purpose

To create custom language data for your chosen language (specified using the `language` chunking parameter).

A language data file contains language-specific abbreviation tokens. You can supply this data to the chunker to help in accurately determining sentence boundaries of chunks, by using knowledge of the input language's end-of-sentence (EOS) punctuations, abbreviations, and contextual rules.

Usage Notes

- All supported languages are distributed with the default language-specific abbreviation dictionaries. You can create a language data based on the abbreviation tokens loaded in the `schema.table.column`, using a user-specified language data name (`PREFERENCE_NAME`).
- After loading your language data, you can use language-specific chunking by specifying the `language` chunking parameter with `VECTOR_CHUNKS` or `UTL_TO_CHUNKS`.
- You can query these data dictionary views to access existing language data:
  - `ALL_VECTOR_LANG` displays all available languages data.
  - `USER_VECTOR_LANG` displays languages data from the schema of the current user.
  - `ALL_VECTOR_ABBREV_TOKENS` displays abbreviation tokens from all available language data.
  - `USER_VECTOR_ABBREV_TOKENS` displays abbreviation tokens from the language data owned by the current user.

Syntax

```
DBMS_VECTOR_CHAIN.CREATE_LANG_DATA (
    PARAMS       IN JSON default NULL
);
```

PARAMS

Specify the input parameters in JSON format:

```
{
    table_name,
    column_name,
```

```
    language,
    preference_name
}
```

Table 212-3 Parameter Details

Parameter	Description	Required	Default Value
table_name	Name of the table (along with the optional table owner) in which you want to load the language data	Yes	No value
column_name	Column name in the language data table in which you want to load the language data	Yes	No value
language	Any supported language name, as listed in Supported Languages and Data File Locations.	Yes	No value
preference_name	User-specified preference name for this language data	Yes	No value

Example

```
declare
    params CLOB := '{"table_name"      : "eos_data_1",
                     "column_name"     : "token",
                     "language"        : "indonesian",
                     "preference_name" : "my_lang_1"}';
begin
    DBMS_VECTOR_CHAIN.CREATE_LANG_DATA(
        JSON (params));
end;
/
```

End-to-end example:

To run an end-to-end example scenario using this procedure, see Create and Use Custom Language Data.

Related Topics

- VECTOR_CHUNKS
- UTL_TO_CHUNKS
- Text Processing Views

212.1.3 CREATE_PREFERENCE

Use the DBMS_VECTOR_CHAIN.CREATE_PREFERENCE preference helper procedure to create a vectorizer preference, to be used when creating or updating hybrid vector indexes.

Purpose

To create a vectorizer preference. This allows you to customize vector search parameters of a hybrid vector indexing pipeline. The goal of a vectorizer preference is to provide you with a

straightforward way to configure how to chunk or embed your documents, without requiring a deep understanding of various chunking or embedding strategies.

A **vectorizer** preference is a JSON object that collectively holds user-specified values related to the following chunking, embedding, or vector index creation parameters:

- Chunking (`UTL_TO_CHUNKS` or `VECTOR_CHUNKS`)
- Embedding (`UTL_TO_EMBEDDING`, `UTL_TO_EMBEDDINGS`, or `VECTOR_EMBEDDING`)
- Vector index creation (`distance`, `accuracy`, and `vector_idxtype`)

After creating a vectorizer preference, you can use the `VECTORIZER` parameter to pass this preference name in the `paramstring` of the `PARAMETERS` clause for `CREATE_HYBRID_VECTOR_INDEX` and `ALTER_INDEX` SQL statements.

Usage Notes

- Creating a preference is optional. If you do not specify any optional preference, then the index is created with system defaults.
- All vector index preferences follow the same JSON syntax as defined for their corresponding `DBMS_VECTOR` and `DBMS_VECTOR_CHAIN` APIs.

Syntax

```
DBMS_VECTOR_CHAIN.CREATE_PREFERENCE (
    PREF_NAME      IN VARCHAR2,
    PREF_TYPE      IN VARCHAR2,
    PARAMS         IN JSON default NULL
);
```

PREF_NAME

Specify the name of the vectorizer preference to create.

PREF_TYPE

Type of preference. The only supported preference type is:

```
DBMS_VECTOR_CHAIN.VECTORIZER
```

PARAMS

Specify vector search-specific parameters in JSON format.

- **Embedding Parameter**:

  ```
  { "model" : <embedding_model_for_vector_generation> }
  ```

  For example:

  ```
  { "model" : MY_INDB_MODEL }
  ```

  `model` specifies the name under which your ONNX embedding model is stored in the database.

  If you do not have an in-database embedding model in ONNX format, then perform the steps listed in *Oracle Database AI Vector Search User's Guide*.

- **Chunking Parameters**:

```
{
    "by"           :  mode,
    "max"          :  max,
    "overlap"      :  overlap,
    "split"        :  split_condition,
    "custom_list"  :  split_chars1, ...,
    "vocabulary"   :  vocabulary_name,
    "language"     :  nls_language,
    "normalize"    :  normalize_mode,
    "norm_options" :  normalize_option1, ...,
    "extended"     :  boolean
}
```

For example:

```
JSON('
  { "by"           :    "vocabulary",
    "vocabulary"   :    "myvocab",
    "max"          :    "100",
    "overlap"      :    "0",
    "split"        :    "custom",
    "custom_list"  :    [ "<p>" , "<s>" ],
    "language"     :    "american",
    "normalize"    :    "options",
    "norm_options" :    [ "whitespace" ]
  }')
```

Parameter	Description and Acceptable Values
`by`	Specify a mode for splitting your data, that is, to split by counting the number of characters, words, or vocabulary tokens. **Valid values**: – `characters` (or `chars`): Splits by counting the number of characters. – `words`: Splits by counting the number of words. Words are defined as sequences of alphabetic characters, sequences of digits, individual punctuation marks, or symbols. For segmented languages without whitespace word boundaries (such as Chinese, Japanese, or Thai), each native character is considered a word (that is, unigram). – `vocabulary`: Splits by counting the number of vocabulary tokens. Vocabulary tokens are words or word pieces, recognized by the vocabulary of the tokenizer that your embedding model uses. You can load your vocabulary file using the chunker helper API `DBMS_VECTOR_CHAIN.CREATE_VOCABULARY`. **Note**: For accurate results, ensure that the chosen model matches the vocabulary file used for chunking. If you are not using a vocabulary file, then ensure that the input length is defined within the token limits of your model. **Default value**: `words`

Parameter	Description and Acceptable Values
`max`	Specify a limit on the maximum size of each chunk. This setting splits the input text at a fixed point where the maximum limit occurs in the larger text. The units of `max` correspond to the `by` mode, that is, to split data when it reaches the maximum size limit of a certain number of characters, words, numbers, punctuation marks, or vocabulary tokens. **Valid values**: – `by characters`: `50` to `4000` characters – `by words`: `10` to `1000` words – `by vocabulary`: `10` to `1000` tokens **Default value**: `100`

Parameter	Description and Acceptable Values

split [by]	Specify where to split the input text when it reaches the maximum size limit. This helps to keep related data together by defining appropriate boundaries for chunks.
	Valid values:
	– none:
	Splits at the max limit of characters, words, or vocabulary tokens.
	– newline, blankline, and space:
	These are single-split character conditions that split at the last split character before the max value.
	Use newline to split at the end of a line of text. Use blankline to split at the end of a blank line (sequence of characters, such as two newlines). Use space to split at the end of a blank space.
	– recursively:
	This is a multiple-split character condition that breaks the input text using an ordered list of characters (or sequences).
	recursively is predefined as BLANKLINE, newline, space, none in this order:
	1. If the input text is more than the max value, then split by the first split character.
	2. If that fails, then split by the second split character.
	3. And so on.
	4. If no split characters exist, then split by max wherever it appears in the text.
	– sentence:
	This is an end-of-sentence split condition that breaks the input text at a sentence boundary.
	This condition automatically determines sentence boundaries by using knowledge of the input language's sentence punctuation and contextual rules. This language-specific condition relies mostly on end-of-sentence (EOS) punctuations and common abbreviations.
	Contextual rules are based on word information, so this condition is only valid when splitting the text by words or vocabulary (not by characters).
	Note: This condition obeys the by word and max settings, and thus may not determine accurate sentence boundaries in some cases. For example, when a sentence is larger than the max value, it splits the sentence at max. Similarly, it includes multiple sentences in the text only when they fit within the max limit.
	– custom:
	Splits based on a custom split characters list. You can provide custom sequences up to a limit of 16 split character strings, with a maximum length of 10 each.
	Specify an array of valid text literals using the custom_list parameter.
	`{ "split" : "custom", "custom_list" : [ "split_chars1", ... ] }`
	For example:
	`{ "split" : "custom", "custom_list" : [ "<p>" , "<s>" ] }`
	Note: You can omit sequences only for tab (\t), newline (\n), and linefeed (\r).
	Default value: recursively

Parameter	Description and Acceptable Values
`overlap`	Specify the amount (as a positive integer literal or zero) of the preceding text that the chunk should contain, if any. This helps in logically splitting up related text (such as a sentence) by including some amount of the preceding chunk text. The amount of overlap depends on how the maximum size of the chunk is measured (in characters, words, or vocabulary tokens). The overlap begins at the specified `split` condition (for example, at `newline`). **Valid value**: `5%` to `20%` of `max` **Default value**: `0`
`language`	Specify the language of your input data. This clause is important, especially when your text contains certain characters (for example, punctuations or abbreviations) that may be interpreted differently in another language. **Valid values**: – NLS-supported language name or its abbreviation, as listed in *Oracle Database Globalization Support Guide*. – Custom language name or its abbreviation, as listed in Supported Languages and Data File Locations. You use the `DBMS_VECTOR_CHAIN.CREATE_LANG_DATA` chunker helper API to load language-specific data (abbreviation tokens) into the database, for your specified language. **Default value**: `NLS_LANGUAGE` from session

Parameter	Description and Acceptable Values

normalize	Automatically pre-processes or post-processes issues (such as multiple consecutive spaces and smart quotes) that may arise when documents are converted into text. Oracle recommends you to use a normalization mode to extract high-quality chunks.
	Valid values:
	– none:
	Applies no normalization.
	– all:
	Normalizes common multi-byte (unicode) punctuation to standard single-byte.
	– options:
	Specify an array of normalization options using the norm_options parameter.
	`{ "normalize" : "options", "norm_options" : [ "normalize_option1", ... ] }`
	* punctuation:
	Includes smart quotes, smart hyphens, and other multi-byte equivalents to simple single-byte punctuation.
	For example:
	* 2018u 'map to 0027'
	* 2019u 'map to 0027'
	* 201Bu 'map to 0027'
	* whitespace:
	Minimizes whitespace by eliminating unnecessary characters.
	For example, retain blanklines, but remove any extra newlines and interspersed spaces or tabs: " \n \n " => "\n\n"
	* widechar:
	Normalizes wide, multi-byte digits and (a-z) letters to single-byte.
	These are multi-byte equivalents for 0-9 and a-z A-Z, which can show up in ZH/JA formatted text.
	For example:
	`{ "normalize" : "options", "norm_options" : [ "whitespace" ] }`
	Default value: None

extended	Increases the output limit of a VARCHAR2 string to 32767 bytes, without requiring you to set the max_string_size parameter to extended. **Default value**: 4000 or 32767 (when max_string_size=extended)

- **Vector Index Parameters**:

```
{
  "distance"        :  <vector_distance>,
  "accuracy"        :  <vector_accuracy>,
  "vector_idxtype"  :  <vector_idxtype>
}
```

For example:

```
{
  "distance"        :  COSINE,
  "accuracy"        :  95,
  "vector_idxtype"  :  HNSW
}
```

Parameter	Description
distance	Distance computation metric as COSINE, MANHATTAN, DOT, EUCLIDEAN, L2_SQUARED, or EUCLIDEAN_SQUARED. **Note**: Currently, the HAMMING and JACCARD vector distance metrics are not supported with hybrid vector indexes. For detailed information on each of these metrics, see Vector Distance Functions and Operators. **Default value**: COSINE
accuracy	Target accuracy at which the approximate search should be performed when running an approximate search query using vector indexes. You can specify non-default target accuracy values either by specifying a percentage value or by specifying internal parameters values, depending on the index type you are using. – For an IVF approximate search: Specify a target accuracy percentage value to influence the number of partitions used to probe the search. Instead of specifying a target accuracy percentage value, you can specify the NEIGHBOR PARTITION PROBES parameter to impose a certain maximum number of partitions to be probed by the search. See Understand Inverted File Flat Vector Indexes. – For an HNSW approximate search: Specify a target accuracy percentage value to influence the number of candidates considered to probe the search. Instead of specifying a target accuracy percentage value, you can specify the EFSEARCH parameter to impose a certain maximum number of candidates to be considered while probing the index. See Understand Hierarchical Navigable Small World Indexes. Valid range for both IVF and HNSW vector indexes is: ACCURACY: > 0 and <= 100 **Default value**: None
vector_idxtype	Type of vector index to create: – IVF for the Inverted File Flat (IVF) vector index – HNSW for the Hierarchical Navigable Small World (HNSW) vector index For detailed information on each of these indexes, see Manage the Different Categories of Vector Indexes. **Default value**: IVF

Example

```
begin
  DBMS_VECTOR_CHAIN.CREATE_PREFERENCE(
    'my_vec_spec',
     DBMS_VECTOR_CHAIN.VECTORIZER,
    json('{"vector_idxtype" :  "hnsw",
           "model"          :  "my_doc_model",
           "by"             :  "words",
           "max"            :  100,
```

```
           "overlap"        :  10,
           "split"          :  "recursively"}'));
end;
/

CREATE HYBRID VECTOR INDEX my_hybrid_idx on
    doc_table(text_column)
    parameters('VECTORIZER my_vec_spec');
```

Related Topics

- CREATE HYBRID VECTOR INDEX

212.1.4 CREATE_VOCABULARY

Use the `DBMS_VECTOR_CHAIN.CREATE_VOCABULARY` chunker helper procedure to load your own token vocabulary file into the database.

Purpose

To create custom token vocabulary that is recognized by the tokenizer used by your vector embedding model.

A vocabulary contains a set of tokens (words and word pieces) that are collected during a model's statistical training process. You can supply this data to the chunker to help in accurately selecting the text size that approximates the maximum input limit imposed by the tokenizer of your embedding model.

Usage Notes

- Usually, the supported vocabulary files (containing recognized tokens) are included as part of a model's distribution. Oracle recommends to use the vocabulary files associated with your model.

  If a vocabulary file is not available, then you may download one of the following files depending on the tokenizer type:

  – **WordPiece**:

    Vocabulary file (`vocab.txt`) for the "bert-base-uncased" (English) or "bert-base-multilingual-cased" model

  – **Byte-Pair Encoding (BPE)**:

    Vocabulary file (`vocab.json`) for the "GPT2" model

    Use the following python script to extract the file:

    ```
    import json
    import sys

    with open(sys.argv[1], encoding="utf-8") as f:
      d = json.load(f)
      for term in d:
        print(term)
    ```

  – **SentencePiece**:

    Vocabulary file (`tokenizer.json`) for the "xlm-roberta-base" model

Use the following python script to extract the file:

```
import json
import sys

with open(sys.argv[1], encoding="utf-8") as f:
  d = json.load(f)
  for entry in d["model"]["vocab"]:
    print(entry[0])
```

Ensure to save your vocabulary files in `UTF-8` encoding.

- You can create a vocabulary based on the tokens loaded in the `schema.table.column`, using a user-specified vocabulary name (`vocabulary_name`).

  After loading your vocabulary data, you can use the `by vocabulary` chunking mode (with `VECTOR_CHUNKS` or `UTL_TO_CHUNKS`) to split input data by counting the number of tokens.

- You can query these data dictionary views to access existing vocabulary data:
  - `ALL_VECTOR_VOCAB` displays all available vocabularies.
  - `USER_VECTOR_VOCAB` displays vocabularies from the schema of the current user.
  - `ALL_VECTOR_VOCAB_TOKENS` displays a list of tokens from all available vocabularies.
  - `USER_VECTOR_VOCAB_TOKENS` displays a list of tokens from the vocabularies owned by the current user.

Syntax

```
DBMS_VECTOR_CHAIN.CREATE_VOCABULARY(
    PARAMS      IN JSON default NULL
);
```

PARAMS

Specify the input parameters in JSON format:

```
{
    table_name,
    column_name,
    vocabulary_name,
    format,
    cased
}
```

Table 212-4 Parameter Details

Parameter	Description	Required	Default Value
`table_name`	Name of the table (along with the optional table owner) in which you want to load the vocabulary file	Yes	No value
`column_name`	Column name in the vocabulary table in which you want to load the vocabulary file	Yes	No value

Table 212-4 (Cont.) Parameter Details

Parameter	Description	Required	Default Value
vocabulary_name	User-specified name of the vocabulary, along with the optional owner name (if other than the current owner)	Yes	No value
format	• xlm for SentencePiece tokenization • bert for WordPiece tokenization • gpt2 for BPE tokenization	Yes	No value
cased	Character-casing of the vocabulary, that is, vocabulary to be treated as cased or uncased	No	false

Example

```
DECLARE
  params clob := '{"table_name"       : "doc_vocabtab",
                   "column_name"      : "token",
                   "vocabulary_name"  : "doc_vocab",
                   "format"           : "bert",
                   "cased"            : false}';

BEGIN
  dbms_vector_chain.create_vocabulary(json(params));
END;
/
```

End-to-end example:

To run an end-to-end example scenario using this procedure, see Create and Use Custom Vocabulary.

Related Topics

- VECTOR_CHUNKS
- UTL_TO_CHUNKS
- Text Processing Views

212.1.5 DROP_CREDENTIAL

Use the DBMS_VECTOR_CHAIN.DROP_CREDENTIAL credential helper procedure to drop an existing credential name from the data dictionary.

Syntax

```
DBMS_VECTOR_CHAIN.DROP_CREDENTIAL (
    CREDENTIAL_NAME      IN VARCHAR2
);
```

CREDENTIAL_NAME

Specify the credential name that you want to drop.

Examples

- For Generative AI:

```
exec dbms_vector_chain.drop_credential('OCI_CRED');
```

- For Cohere:

```
exec dbms_vector_chain.drop_credential('COHERE_CRED');
```

212.1.6 DROP_LANG_DATA

Use the `DBMS_VECTOR_CHAIN.DROP_LANG_DATA` chunker helper procedure to remove abbreviation data from the data dictionary.

Syntax

```
DBMS_VECTOR_CHAIN.DROP_LANG_DATA(
    PREF_NAME     IN VARCHAR2
);
```

LANG

Specify the name of the language data that you want to drop for a given language.

Example

```
DBMS_VECTOR_CHAIN.DROP_LANG_DATA('indonesian');
```

212.1.7 DROP_PREFERENCE

Use the `DBMS_VECTOR_CHAIN.DROP_PREFERENCE` preference helper procedure to remove an existing Vectorizer preference.

Syntax

```
DBMS_VECTOR_CHAIN.DROP_PREFERENCE (PREF_NAME);
```

PREF_NAME

Name of the Vectorizer preference to drop.

Example

```
DBMS_VECTOR_CHAIN.DROP_PREFERENCE ('scott_vectorizer');
```

212.1.8 DROP_VOCABULARY

Use the `DBMS_VECTOR_CHAIN.DROP_VOCABULARY` chunker helper procedure to remove vocabulary data from the data dictionary.

Syntax

```
DBMS_VECTOR_CHAIN.DROP_VOCABULARY(
    VOCABULARY_NAME    IN VARCHAR2
);
```

VOCAB_NAME

Specify the name of the vocabulary that you want to drop, in the form:

vocabulary_name

or

owner.vocabulary_name

Example

```
DBMS_VECTOR_CHAIN.DROP_VOCABULARY('MY_VOCAB_1');
```

212.1.9 RERANK

Use the `DBMS_VECTOR_CHAIN.RERANK` function to reassess and reorder an initial set of results to retrieve more relevant search output.

Purpose

To improve the relevance and quality of search results in both similarity search and Retrieval Augmented Generation (RAG) scenarios.

Reranking improves the quality of information ingested into an LLM by ensuring that the most relevant documents or chunks are prioritized. This helps to reduce hallucinations and improves the accuracy of generated outputs.

For this operation, Oracle AI Vector Search supports reranking models provided by Cohere and Vertex AI.

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

```
DBMS_VECTOR_CHAIN.RERANK(
                QUERY      IN CLOB,
                DOCUMENTS  IN JSON,
                PARAMS     IN JSON default NULL
) return JSON;
```

This function accepts the input containing a query as `CLOB` and a list of documents in `JSON` format. It then processes this information to generate a `JSON` object containing a reranked list of documents, sorted by score.

For example, a reranked output includes:

```
{
    "index"   : "1",
    "score"   : "0.99",
    "content" : "Jupiter boasts an impressive system of 95 known moons."
}
```

Where,

- `index` specifies the position of the document in the list of input text.
- `score` specifies the relevance score.
- `content` specifies the input text corresponding to the index.

QUERY

Specify the search query (typically from an initial search) as `CLOB`.

DOCUMENTS

Specify a JSON array of strings (list of potentially relevant documents to rerank) in the following format:

```
{
  "documents": [
  "string1",
  "string2",
    ...
```

```
  ]
}
```

PARAMS

Specify the following list of parameters in JSON format. All these parameters are mandatory.

```
{
  "provider"         : "<service provider>",
  "credential_name"  : "<credential name>",
  "url"              : "<REST endpoint URL for reranking>",
  "model"            : "<reranking model name>",
  ...
}
```

Table 212-5 RERANK Parameter Details

Parameter	Description
`provider`	Supported REST provider to access for reranking: • `cohere` • `vertexai`
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR_CHAIN.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the reranking model in the form: *`schema.model_name`* If the model name is not schema-qualified, then the schema of the procedure invoker is used.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters for reranking.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"         : "cohere",
  "credential_name"  : "COHERE_CRED",
  "url"              : "https://api.cohere.example.com/rerank",
  "model"            : "rerank-model",
  "return_documents": false,
  "top_n"            : 3
}
```

Vertex AI example:

```
{
  "provider"          : "vertexai",
  "credential_name"   : "VERTEXAI_CRED",
  "url"               : "https://googleapis.example.com/
default_ranking_config:rank",
  "model"             : "rerank-model",
  "ignoreRecordDetailsInResponse" : true,
  "topN"              : 3
  }
```

Table 212-6 Additional REST Provider Parameter Details

Parameter	Description
`return_documents`	Whether to return search results with original documents or input text (`content`): • `false` (default, also recommended) to not return any input text (return only index and score) • `true` to return input text along with index and score **Note**: With Cohere as the provider, Oracle recommends that you keep this option disabled for better performance. You may choose to enable it for debugging purposes when you need to view the original text.
`ignoreRecordDetailsInResponse`	Whether to return search results with original record details or input text (`content`): • `false` (default) to return input text along with index and score • `true` (recommended) to not return any input text (return only index and score) **Note**: With Vertex AI as the provider, Oracle recommends that you keep this option enabled for better performance. You may choose to disable it for debugging purposes when you need to view the original text.
`top_n` or `topN`	The number of most relevant documents to return.

Examples

- Using Cohere:

```
declare
  params clob;
  reranked_output json;
begin
  params := '
```

```
{
  "provider": "cohere",
  "credential_name": "COHERE_CRED",
  "url": "https://api.cohere.com/v1/rerank",
  "model": "rerank-english-v3.0",
  "return_documents": true,
  "top_n": 3
}';

  reranked_output := dbms_vector_chain.rerank(:query,
json(:initial_retrieval_docs), json(params));
  dbms_output.put_line(json_serialize(reranked_output));
end;
/
```

- Using Vertex AI:

```
declare
  params clob;
  reranked_output json;
begin
  params := '
{
  "provider": "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url": "https://discoveryengine.googleapis.com/v1/projects/1085581009881/
locations/global/rankingConfigs/default_ranking_config:rank",
  "model": "semantic-ranker-512@latest",
  "ignoreRecordDetailsInResponse": false,
  "topN": 3
}';

  reranked_output := dbms_vector_chain.rerank(:query,
json(:initial_retrieval_docs), json(params));
  dbms_output.put_line(json_serialize(reranked_output));
end;
/
```

End-to-end example:

To run an end-to-end example scenario using this function, see Use Reranking for Better RAG Results.

212.1.10 UTL_TO_CHUNKS

Use the `DBMS_VECTOR_CHAIN.UTL_TO_CHUNKS` chainable utility function to split a large plain text document into smaller chunks of text.

Purpose

To perform a text-to-chunks transformation. This chainable utility function internally calls the `VECTOR_CHUNKS` SQL function for the operation.

To embed a large document, you may first need to split it into multiple appropriate-sized segments or chunks through a splitting process known as chunking (as explained in Understand the Stages of Data Transformations). A chunk can be words (to capture specific

words or word pieces), sentences (to capture a specific meaning), or paragraphs (to capture broader themes). A single document may be split into multiple chunks, each transformed into a vector.

Syntax

```
DBMS_VECTOR_CHAIN.UTL_TO_CHUNKS (
    DATA         IN CLOB | VARCHAR2
    PARAMS       IN JSON default NULL
) return VECTOR_ARRAY_T;
```

DATA

This function accepts the input data type as `CLOB` or `VARCHAR2`.

It returns an array of `CLOB`s, where each `CLOB` contains a chunk along with its metadata in JSON format, as follows:

```
{
    "chunk_id"     : NUMBER,
    "chunk_offset" : NUMBER,
    "chunk_length" : NUMBER,
    "chunk_data"   : "VARCHAR2(4000)"
}
```

For example:

```
{"chunk_id":1,"chunk_offset":1,"chunk_length":6,"chunk_data":"sample"}
```

Where,

- `chunk_id` specifies the chunk ID for each chunk.
- `chunk_offset` specifies the original position of each chunk in the source document, relative to the start of document which has a position of `1`.
- `chunk_length` specifies the character length of each chunk.
- `chunk_data` displays text pieces from each chunk.

PARAMS

Specify input parameters in JSON format.

```
{
    "by"           :     mode,
    "max"          :     max,
    "overlap"      :     overlap,
    "split"        :     split_condition,
    "custom_list"  :     [ split_chars1, ... ],
    "vocabulary"   :     vocabulary_name,
    "language"     :     nls_language,
    "normalize"    :     normalize_mode,
    "norm_options" :     [ normalize_option1, ... ],
    "extended"     :     boolean
}
```

For example:

```
JSON('
  { "by"           :    "vocabulary",
    "vocabulary"   :    "myvocab",
    "max"          :    "100",
    "overlap"      :    "0",
    "split"        :    "custom",
    "custom_list"  :    [ "<p>" , "<s>" ],
    "language"     :    "american",
    "normalize"    :    "options",
    "norm_options" :    [ "whitespace" ]
  }')
```

Here is a complete description of these parameters:

Parameter	Description and Acceptable Values
`by`	Specify a mode for splitting your data, that is, to split by counting the number of characters, words, or vocabulary tokens. **Valid values**: • `characters` (or `chars`): Splits by counting the number of characters. • `words`: Splits by counting the number of words. Words are defined as sequences of alphabetic characters, sequences of digits, individual punctuation marks, or symbols. For segmented languages without whitespace word boundaries (such as Chinese, Japanese, or Thai), each native character is considered a word (that is, unigram). • `vocabulary`: Splits by counting the number of vocabulary tokens. Vocabulary tokens are words or word pieces, recognized by the vocabulary of the tokenizer that your embedding model uses. You can load your vocabulary file using the chunker helper API `DBMS_VECTOR_CHAIN.CREATE_VOCABULARY`. **Note**: For accurate results, ensure that the chosen model matches the vocabulary file used for chunking. If you are not using a vocabulary file, then ensure that the input length is defined within the token limits of your model. **Default value**: `words`
`max`	Specify a limit on the maximum size of each chunk. This setting splits the input text at a fixed point where the maximum limit occurs in the larger text. The units of `max` correspond to the `by` mode, that is, to split data when it reaches the maximum size limit of a certain number of characters, words, numbers, punctuation marks, or vocabulary tokens. **Valid values**: • `by characters`: `50` to `4000` characters • `by words`: `10` to `1000` words • `by vocabulary`: `10` to `1000` tokens **Default value**: `100`

Parameter	Description and Acceptable Values

split [by]	Specify where to split the input text when it reaches the maximum size limit. This helps to keep related data together by defining appropriate boundaries for chunks.
	Valid values:
	• none:
	Splits at the max limit of characters, words, or vocabulary tokens.
	• newline, blankline, and space:
	These are single-split character conditions that split at the last split character before the max value.
	Use newline to split at the end of a line of text. Use blankline to split at the end of a blank line (sequence of characters, such as two newlines). Use space to split at the end of a blank space.
	• recursively:
	This is a multiple-split character condition that breaks the input text using an ordered list of characters (or sequences).
	recursively is predefined as BLANKLINE, newline, space, none in this order:
	1. If the input text is more than the max value, then split by the first split character.
	2. If that fails, then split by the second split character.
	3. And so on.
	4. If no split characters exist, then split by max wherever it appears in the text.
	• sentence:
	This is an end-of-sentence split condition that breaks the input text at a sentence boundary.
	This condition automatically determines sentence boundaries by using knowledge of the input language's sentence punctuation and contextual rules. This language-specific condition relies mostly on end-of-sentence (EOS) punctuations and common abbreviations.
	Contextual rules are based on word information, so this condition is only valid when splitting the text by words or vocabulary (not by characters).
	Note: This condition obeys the by word and max settings, and thus may not determine accurate sentence boundaries in some cases. For example, when a sentence is larger than the max value, it splits the sentence at max. Similarly, it includes multiple sentences in the text only when they fit within the max limit.
	• custom:
	Splits based on a custom split characters list. You can provide custom sequences up to a limit of 16 split character strings, with a maximum length of 10 each.
	Specify an array of valid text literals using the custom_list parameter.
	`{ "split" : "custom", "custom_list" : [ "split_chars1", ... ] }`
	For example:
	`{ "split" : "custom", "custom_list" : [ "<p>" , "<s>" ] }`
	Note: You can omit sequences only for tab (\t), newline (\n), and linefeed (\r).
	Default value: recursively

Parameter	Description and Acceptable Values
`overlap`	Specify the amount (as a positive integer literal or zero) of the preceding text that the chunk should contain, if any. This helps in logically splitting up related text (such as a sentence) by including some amount of the preceding chunk text. The amount of overlap depends on how the maximum size of the chunk is measured (in characters, words, or vocabulary tokens). The overlap begins at the specified `split` condition (for example, at `newline`). **Valid value**: `5%` to `20%` of `max` **Default value**: `0`
`language`	Specify the language of your input data. This clause is important, especially when your text contains certain characters (for example, punctuations or abbreviations) that may be interpreted differently in another language. **Valid values**: • NLS-supported language name or its abbreviation, as listed in *Oracle Database Globalization Support Guide*. • Custom language name or its abbreviation, as listed in Supported Languages and Data File Locations. You use the `DBMS_VECTOR_CHAIN.CREATE_LANG_DATA` chunker helper API to load language-specific data (abbreviation tokens) into the database, for your specified language. **Default value**: `NLS_LANGUAGE` from session

Parameter	Description and Acceptable Values

normalize	Automatically pre-processes or post-processes issues (such as multiple consecutive spaces and smart quotes) that may arise when documents are converted into text. Oracle recommends you to use a normalization mode to extract high-quality chunks.
	Valid values:
	• none:
	Applies no normalization.
	• all:
	Normalizes common multi-byte (unicode) punctuation to standard single-byte.
	• options:
	Specify an array of normalization options using the norm_options parameter.
	`{ "normalize" : "options", "norm_options" : [ "normalize_option1", ... ] }`
	– punctuation:
	Includes smart quotes, smart hyphens, and other multi-byte equivalents to simple single-byte punctuation.
	For example:
	* 2018u 'map to 0027'
	* 2019u 'map to 0027'
	* 201Bu 'map to 0027'
	– whitespace:
	Minimizes whitespace by eliminating unnecessary characters.
	For example, retain blanklines, but remove any extra newlines and interspersed spaces or tabs: " \n \n " => "\n\n"
	– widechar:
	Normalizes wide, multi-byte digits and (a-z) letters to single-byte.
	These are multi-byte equivalents for 0-9 and a-z A-Z, which can show up in ZH/JA formatted text.
	For example:
	`{ "normalize" : "options", "norm_options" : [ "whitespace" ] }`
	Default value: None

extended	Increases the output limit of a VARCHAR2 string to 32767 bytes, without requiring you to set the max_string_size parameter to extended. **Default value**: 4000 or 32767 (when max_string_size=extended)

Example

```
SELECT D.id doc,
    JSON_VALUE(C.column_value, '$.chunk_id' RETURNING NUMBER) AS id,
    JSON_VALUE(C.column_value, '$.chunk_offset' RETURNING NUMBER) AS pos,
    JSON_VALUE(C.column_value, '$.chunk_length' RETURNING NUMBER) AS siz,
    JSON_VALUE(C.column_value, '$.chunk_data') AS txt
```

```
FROM docs D,
   dbms_vector_chain.utl_to_chunks(D.text,
   JSON('{ "by"        : "words",
           "max"       : "100",
           "overlap"   : "0",
           "split"     : "recursively",
           "language"  : "american",
           "normalize": "all" }')) C;
```

End-to-end examples:

To run end-to-end example scenarios using this function, see Perform Chunking With Embedding and Configure Chunking Parameters.

Related Topics

- VECTOR_CHUNKS

212.1.11 UTL_TO_EMBEDDING and UTL_TO_EMBEDDINGS

Use the `DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDING` and `DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDINGS` chainable utility functions to generate one or more vector embeddings from textual documents and images.

Purpose

To automatically generate vector embeddings from a set of textual documents and images.

- Text to Vector:

  You can perform a text-to-embedding transformation by accessing:

  - Oracle Database as the service provider: Calls an ONNX format embedding model that you load into the database (default setting)
  - Third-party embedding model: Makes a REST API call to either a remote service provider (Cohere, Google AI, Hugging Face, Generative AI, OpenAI, or Vertex AI) or a local service provider (Ollama)

- Image to Vector:

  You can also perform an image-to-embedding transformation. For this operation, this API makes a REST call to your chosen image embedding model or multimodal embedding model by Vertex AI (which is the only supported third-party provider).

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

- Text to Vector:

```
DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDING (
    DATA           IN CLOB,
    PARAMS         IN JSON default NULL
) return VECTOR;


DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDINGS (
    DATA           IN VECTOR_ARRAY_T,
    PARAMS         IN JSON default NULL
) return VECTOR_ARRAY_T;
```

- Image to Vector:

```
DBMS_VECTOR_CHAIN.UTL_TO_EMBEDDING (
    DATA           IN BLOB,
    MODALITY       IN VARCHAR2,
    PARAMS         IN JSON default NULL
) return VECTOR;
```

DATA

`UTL_TO_EMBEDDING` accepts the input as `CLOB` containing textual data (text strings or small documents) or as `BLOB` containing media data (for media files such as images). It then converts the text or image input to a single embedding (`VECTOR`).

`UTL_TO_EMBEDDINGS` converts an array of chunks (`VECTOR_ARRAY_T`) to an array of embeddings (`VECTOR_ARRAY_T`).

A generated embedding output includes:

```
{
    "embed_id"    : NUMBER,
    "embed_data"  : "VARCHAR2(4000)",
    "embed_vector": "CLOB"
}
```

Where,

- `embed_id` displays the ID number of each embedding.
- `embed_data` displays the input text that is transformed into embeddings.
- `embed_vector` displays the generated vector representations.

MODALITY

For `BLOB` inputs, specify the type of content to vectorize. The only supported value is `image`.

PARAMS

Specify input parameters in JSON format, depending on the service provider that you want to use.

If using Oracle Database as the provider:

```
{
  "provider" : "database",
  "model"    : "<in-database ONNX embedding model filename>"
}
```

Table 212-7 Database Provider Parameter Details

Parameter	Description
`provider`	Specify `database` (default setting) to use Oracle Database as the provider. With this setting, you must load an ONNX format embedding model into the database.
`model`	User-specified name under which the imported ONNX embedding model is stored in Oracle Database. If you do not have an embedding model in ONNX format, then perform the steps listed in Convert Pretrained Models to ONNX Format.

If using a third-party provider:

Set the following parameters along with additional embedding parameters specific to your provider:

- For `UTL_TO_EMBEDDING`:

```
{
  "provider"        : "<AI service provider>",
  "credential_name" : "<credential name>",
  "url"             : "<REST endpoint URL for embedding service>",
  "model"           : "<REST provider embedding model name>",
  "transfer_timeout": <maximum wait time for the request to complete>,
  "<additional REST provider parameter>": "<REST provider parameter
value>"
}
```

- For `UTL_TO_EMBEDDINGS`:

```
{
  "provider"        : "<AI service provider>",
  "credential_name" : "<credential name>",
  "url"             : "<REST endpoint URL for embedding service>",
  "model"           : "<REST provider embedding model name>",
  "transfer_timeout": <maximum wait time for the request to complete>,
  "batch size"      : "<number of vectors to request at a time>",
  "<additional REST provider parameter>": "<REST provider parameter
value>"
}
```

Table 212-8 Third-Party Provider Parameter Details

Parameter	Description
`provider`	Third-party service provider that you want to access for this operation. A REST call is made to the specified provider to access its embedding model. For image input, specify `vertexai`. For text input, specify one of the following values: • `cohere` • `googleai` • `huggingface` • `ocigenai` • `openai` • `vertexai`
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR_CHAIN.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the third-party embedding model in the form: *`schema.model_name`* If you do not specify a schema, then the schema of the procedure invoker is used. **Note**: • For accurate results, ensure that the chosen text embedding model matches the vocabulary file used for chunking. If you are not using a vocabulary file, then ensure that the input length is defined within the token limits of your model. • To get image embeddings, you can use any image embedding model or multimodal embedding model supported by Vertex AI. Multimodal embedding is a technique that vectorizes data from different modalities such as text and images. When using a multimodal embedding model to generate embeddings, ensure that you use the same model to vectorize both types of content (text and images). By doing so, the resulting embeddings are compatible and situated in the same vector space, which allows for effective comparison between the two modalities during similarity searches.
`transfer_timeout`	Maximum time to wait for the request to complete. The default value is `60` seconds. You can increase this value for busy web servers.
`batch size`	Maximum number of vectors to request at a time. For example, for a batch size of `50`, if `100` chunks are passed, then this API sends two requests with an array of `50` strings each. If `30` chunks are passed (which is lesser than the defined batch size), then the API sends those in a single request. For REST calls, it is more efficient to send a batch of inputs at a time rather than requesting a single input per call. Increasing the batch size can provide better performance, whereas reducing the batch size may reduce memory and data usage, especially if your provider has a rate limit. The default or maximum allowed value depends on the third-party provider settings.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"       : "cohere",
  "credential_name": "COHERE_CRED",
  "url"            : "https://api.cohere.example.com/embed",
  "model"          : "embed-model",
  "input_type"     : "search_query"
}
```

Google AI example:

```
{
  "provider"       : "googleai",
  "credential_name": "GOOGLEAI_CRED",
  "url"            : "https://googleapis.example.com/models/",
  "model"          : "embed-model"
}
```

Hugging Face example:

```
{
  "provider"       : "huggingface",
  "credential_name": "HF_CRED",
  "url"            : "https://api.huggingface.example.com/",
  "model"          : "embed-model"
}
```

Generative AI example:

```
{
  "provider"       : "ocigenai",
  "credential_name": "OCI_CRED",
  "url"            : "https://generativeai.oci.example.com/embedText",
  "model"          : "embed-model",
  "batch_size"     : 10
}
```

Ollama example:

```
{
  "provider"       : "ollama",
  "host"           : "local",
```

```
  "url"            : "http://localhost:11434/api/embeddings",
  "model"          : "llama3"
}
```

OpenAI example:

```
{
  "provider"       : "openai",
  "credential_name": "OPENAI_CRED",
  "url"            : "https://api.openai.example.com/embeddings",
  "model"          : "embed-model"
}
```

Vertex AI example:

```
{
  "provider"       : "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url"            : "https://googleapis.example.com/models/",
  "model"          : "embed-model"
}
```

Table 212-9 Additional REST Provider Parameter Details

Parameter	Description
`input_type`	Type of input to vectorize.

> **Note:**
>
> The generated embedding results may be different between requests for the same input and configuration, depending on your embedding model or floating point precision. However, this does not affect your queries (and provides semantically correct results) because the vector distance will be similar.

Examples

You can use `UTL_TO_EMBEDDING` in a `SELECT` clause and `UTL_TO_EMBEDDINGS` in a `FROM` clause, as follows:

- The following examples use `UTL_TO_EMBEDDING` to generate an embedding with `Hello world` as input, by accessing the Generative AI embedding model:

```
-- declare embedding parameters

var params clob;

begin
 :params := '
{
  "provider": "ocigenai",
  "credential_name": "OCI_CRED",
```

```
  "url": "https://generativeai.oci.example.com/embedText",
  "model": "embed.modelname",
  "batch_size": 10
}';
end;
/

-- get text embedding: PL/SQL example

declare
  input clob;
  v vector;
begin
  input := 'Hello world';

  v := dbms_vector_chain.utl_to_embedding(input, json(params));
  dbms_output.put_line(vector_serialize(v));
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/

-- get text embedding: select example

select dbms_vector_chain.utl_to_embedding('Hello world', json(:params))
from dual;
```

The following examples use `UTL_TO_EMBEDDING` to generate an embedding with `parrots.jpg` as input and modality as `image`, by accessing the Vertex AI's multimodal embedding model.

```
-- declare embedding parameters

var params clob;

begin
  :params := '
{
  "provider": "vertexai",
  "credential_name": "VERTEXAI_CRED",
  "url": "https://LOCATION-aiplatform.googleapis.com/v1/projects/PROJECT/
locations/LOCATION/publishers/google/models/",
  "model": "multimodalembedding:predict"
}';
end;
/

-- get image embedding: PL/SQL example

declare
  v vector;
  output clob;
begin
```

```
  v := dbms_vector_chain.utl_to_embedding(
    to_blob(bfilename('VEC_DUMP', 'parrots.jpg')), 'image', json(:params));
  output := vector_serialize(v);
  dbms_output.put_line('vector data=' || dbms_lob.substr(output, 100) ||
'...');
end;
/

-- get image embedding: select example

select dbms_vector_chain.utl_to_embedding(
  to_blob(bfilename('VEC_DUMP', 'parrots.jpg')), 'image', json(:params));
```

End-to-end examples:

To run end-to-end example scenarios using `UTL_TO_EMBEDDING`, see Generate Embedding.

- The following example uses `UTL_TO_EMBEDDINGS` to generate embeddings with a PDF document (stored in the `documentation_tab` table) as input, by calling an ONNX format model loaded into the database:

```
SELECT et.* from documentation_tab dt,dbms_vector_chain.utl_to_embeddings(
  dbms_vector_chain.utl_to_chunks(dbms_vector_chain.utl_to_text(dt.data)),
  json(:embed_params)) et;
```

End-to-end examples:

To run end-to-end example scenarios using `UTL_TO_EMBEDDINGS`, see Perform Chunking With Embedding.

212.1.12 UTL_TO_GENERATE_TEXT

Use the `DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT` chainable utility function to generate a text response for a given prompt or an image, by accessing third-party text generation models.

Purpose

To communicate with Large Language Models (LLMs) through natural language conversations. You can generate a textual answer, description, or summary for prompts and images, given as input to LLM-powered chat interfaces.

- Prompt to Text:

  A prompt can be an input text string, such as a question that you ask an LLM. For example, "`What is Oracle Text?`". A prompt can also be a command, such as "`Summarize the following ...`", "`Draft an email asking for ...`", or "`Rewrite the following ...`", and can include results from a search. The LLM responds with a textual answer or description based on the specified task in the prompt.

  For this operation, this API makes a REST call to either a remote third-party provider (Cohere, Google AI, Hugging Face, Generative AI, OpenAI, or Vertex AI) or a local third-party provider (Ollama).

- Image to Text:

  You can also prompt with a media file, such as an image, to extract text from pictures or photos. You supply a text question as the prompt (such as "`What is this image about?`"

or "How many birds are there in this painting?") along with the image. The LLM responds with a textual analysis or description of the contents of the image.

For this operation, this API makes a REST call to either a remote third-party provider (Google AI, Hugging Face, OpenAI, or Vertex AI) or a local third-party provider (Ollama).

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

This function accepts the input as CLOB containing text data (for textual prompts) or as BLOB containing media data (for media files such as images). It then processes this information to generate a new CLOB containing the generated text.

- Prompt to Text:

```
DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT (
        DATA          IN CLOB,
        PARAMS        IN JSON default NULL
) return CLOB;
```

- Image to Text:

```
DBMS_VECTOR_CHAIN.UTL_TO_GENERATE_TEXT(
          TEXT_DATA      IN CLOB,
          MEDIA_DATA     IN BLOB,
          MEDIA_TYPE     IN VARCHAR2 default 'image/jpeg',
          PARAMS         IN JSON default NULL
) return CLOB;
```

DATA and TEXT_DATA

Specify the textual prompt as CLOB for the DATA or TEXT_DATA clause.

> **Note:**
>
> Hugging Face uses an image captioning model that does not require a prompt, when giving an image as input. If you input a prompt along with an image, then the prompt will be ignored.

MEDIA_DATA

Specify the `BLOB` file, such as an image or a visual PDF file.

MEDIA_TYPE

Specify the image format for the given image or visual PDF file (`BLOB` file) in one of the supported image data MIME types. For example:

- For PNG: `image/png`
- For JPEG: `image/jpeg`
- For PDF: `application/pdf`

> **Note:**
>
> For a complete list of the supported image formats, refer to your third-party provider's documentation.

PARAMS

Specify the following input parameters in JSON format, depending on the service provider that you want to access for text generation:

```
{
  "provider"        : "<AI service provider>",
  "credential_name" : "<credential name>",
  "url"             : "<REST endpoint URL for text generation service>",
  "model"           : "<text generation model name>",
  "transfer_timeout": <maximum wait time for the request to complete>,
  "<additional REST provider parameter>": "<REST provider parameter value>"
}
```

Table 212-10 UTL_TO_GENERATE_TEXT Parameter Details

Parameter	Description
`provider`	Supported REST provider that you want to access to generate text. Specify one of the following values: For `CLOB` input: • `cohere` • `googleai` • `huggingface` • `ocigenai` • `openai` • `vertexai` For `BLOB` input: • `googleai` • `huggingface` • `openai` • `vertexai`

Table 212-10 (Cont.) UTL_TO_GENERATE_TEXT Parameter Details

Parameter	Description
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR_CHAIN.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the third-party text generation model in the form: *`schema.model_name`* If the model name is not schema-qualified, then the schema of the procedure invoker is used.
`transfer_timeout`	Maximum time to wait for the request to complete. The default value is `60` seconds. You can increase this value for busy web servers.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"       : "cohere",
  "credential_name": "COHERE_CRED",
  "url"            : "https://api.cohere.example.com/chat",
  "model"          : "generate-text-model"
}
```

Google AI example:

```
{
  "provider"        : "googleai",
  "credential_name" : "GOOGLEAI_CRED",
  "url"             : "https://googleapis.example.com/models/",
  "model"           : "generate-text-model"
}
```

Hugging Face example:

```
{
  "provider"        : "huggingface",
  "credential_name" : "HF_CRED",
  "url"             : "https://api.huggingface.example.com/models/",
  "model"           : "generate-text-model"
}
```

Generative AI example:

```
{
  "provider"        : "ocigenai",
  "credential_name" : "GENAI_CRED",
  "url"             : "https://generativeai.oci.example.com/generateText",
  "model"           : "generate-text-model",
  "inferenceRequest": {
    "maxTokens"     : 300,
    "temperature"   : 1
  }
}
```

or more current chat models may implement:

```
{
  "provider"        : "ocigenai",
  "credential_name" : "OCI_CRED",
  "url"             : "https://inference.generativeai.oci.example.com/actions/
chat",
  "model"           : "generate-text-model",
  "chatRequest"     : {
    "maxTokens"     : 256
  }
}
```

Ollama example:

```
{
  "provider"       : "ollama",
  "host"           : "local",
  "url"            : "http://localhost:11434/api/generate",
  "model"          : "llama3"
}
```

OpenAI example:

```
{
  "provider"        : "openai",
  "credential_name" : "OPENAI_CRED",
  "url"             : "https://api.openai.example.com",
  "model"           : "generate-text-model",
  "max_tokens"      : 60,
```

```
  "temperature"     : 1.0
}
```

Vertex AI example:

```
{
  "provider"         : "vertexai",
  "credential_name"  : "VERTEXAI_CRED",
  "url"              : "https://googleapis.example.com/models/",
  "model"            : "generate-text-model",
  "generation_config": {
    "temperature"    : 0.9,
    "topP"           : 1,
    "candidateCount" : 1,
    "maxOutputTokens": 256
  }
}
```

Table 212-11 Additional REST Provider Parameter Details

Parameter	Description
`max_tokens`	Maximum number of tokens in the output text.
`temperature`	Degree of randomness used when generating the output text, in the range of `0.0-5.0`. To generate the same output for a prompt, use `0`. To generate a random new text for that prompt, increase the temperature. **Note**: Start with the temperature set to `0`. If you do not require random results, a recommended temperature value is between `0` and `1`. A higher value is not recommended because a high temperature may produce creative text, which might also include hallucinations.
`topP`	Probability of tokens in the output, in the range of `0.0-1.0`. A lower value provides less random responses and a higher value provides more random responses.
`candidateCount`	Number of response variations to return, in the range of `1-4`.
`maxOutputTokens`	Maximum number of tokens to generate for each response.

Examples

- Prompt to Text:

  The following statements generate text by making a REST call to Generative AI. Here, the prompt is "`What is Oracle Text?`".

```
-- select example

var params clob;
exec :params := '
{
  "provider": "ocigenai",
  "credential_name": "OCI_CRED",
  "url": "https://generativeai.oci.example.com/generateText",
  "model": "generate.modelname"
```

```
}';

select dbms_vector_chain.utl_to_generate_text(
 'What is Oracle Text?',
 json(:params)) from dual;

-- PL/SQL example

declare
  input clob;
  params clob;
  output clob;
begin
  input := 'What is Oracle Text?';

  params := '
{
  "provider": "ocigenai",
  "credential_name": "OCI_CRED",
  "url": "https://generativeai.oci.example.com/generateText",
  "model": "generate.modelname"
}';

  output := dbms_vector_chain.utl_to_generate_text(input, json(params));
  dbms_output.put_line(output);
  if output is not null then
    dbms_lob.freetemporary(output);
  end if;
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/
```

End-to-end examples:

To run end-to-end example scenarios, see Generate Text Response.

- Image to Text:

  The following statements generate text by making a REST call to OpenAI. Here, the input is an image (`sample_image.jpeg`) along with the prompt "`Describe this image?`".

```
-- select example

var input clob;
var media_data blob;
var media_type clob;
var params clob;

begin
  :input := 'Describe this image';
  :media_data := load_blob_from_file('DEMO_DIR', 'sample_image.jpeg');
  :media_type := 'image/jpeg';
  :params := '
{
```

```
  "provider": "openai",
  "credential_name": "OPENAI_CRED",
  "url": "https://api.openai.example.com/chat/completions",
  "model": "generate.modelname",
  "max_tokens": 256
}';
end;
/

select
dbms_vector_chain.utl_to_generate_text(:input, :media_data, :media_type,
json(:params));

-- PL/SQL example

declare
  input clob;
  media_data blob;
  media_type varchar2(32);
  params clob;
  output clob;

begin
  input := 'Describe this image';
  media_data := load_blob_from_file('DEMO_DIR', 'image_file');
  media_type := 'image/jpeg';
  params := '
{
  "provider": "openai",
  "credential_name": "OPENAI_CRED",
  "url": "https://api.openai.example.com/chat/completions",
  "model": "generate.modelname",
  "max_tokens": 256
}';

  output := dbms_vector_chain.utl_to_generate_text(
    input, media_data, media_type, json(params));
  dbms_output.put_line(output);

  if output is not null then
    dbms_lob.freetemporary(output);
  end if;
  if media_data is not null then
    dbms_lob.freetemporary(media_data);
  end if;
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
end;
/
```

End-to-end examples:

To run end-to-end example scenarios, see Describe Image Content.

212.1.13 UTL_TO_TEXT

Use the `DBMS_VECTOR_CHAIN.UTL_TO_TEXT` chainable utility function to convert an input document (for example, PDF, DOC, JSON, XML, or HTML) to plain text.

Purpose

To perform a file-to-text transformation by using the Oracle Text component (`CONTEXT`) of Oracle Database.

Syntax

```
DBMS_VECTOR_CHAIN.UTL_TO_TEXT (
    DATA           IN CLOB | BLOB,
    PARAMS         IN JSON default NULL
) return CLOB;
```

DATA

This function accepts the input data type as `CLOB` or `BLOB`. It can read documents from a remote location or from files stored locally in the database tables.

It returns a plain text version of the document as `CLOB`.

Oracle Text supports around 150 file types. For a complete list of all the supported document formats, see *Oracle Text Reference*.

PARAMS

Specify the following input parameter in JSON format:

```
{
    "plaintext" : "true or false",
    "charset"   : "UTF8"
}
```

Table 212-12 Parameter Details

Parameter	Description
`plaintext`	Plain text output. The default value for this parameter is `true`, that is, by default the output format is plain text. If you do not want to return the document as plain text, then set this parameter to `false`.
`charset`	Character set encoding. Currently, only `UTF8` is supported.

Example

```
select DBMS_VECTOR_CHAIN.UTL_TO_TEXT (
    t.blobdata,
     json('{
            "plaintext": "true",
```

```
            "charset"  : "UTF8"
           }')
) from tab t;
```

End-to-end example:

To run an end-to-end example scenario using this function, see Convert File to Text to Chunks to Embeddings Within Oracle Database.

212.1.14 UTL_TO_SUMMARY

Use the `DBMS_VECTOR_CHAIN.UTL_TO_SUMMARY` chainable utility function to generate a summary for textual documents.

Purpose

To perform a text-to-summary transformation by accessing:

- Oracle Database as the service provider: Uses the in-house implementation with Oracle Database, leveraging a document gist (or summary) generated by Oracle Text (default setting)
- Third-party summarization model: Makes a REST API call to either a remote service provider (Cohere, Google AI, Hugging Face, Generative AI, OpenAI, or Vertex AI) or a local service provider (Ollama)

A summary is a short and concise extract with key features of a document that best represents what the document is about as a whole. A summary can be free-form paragraphs or bullet points based on the format that you specify.

> **WARNING:**
>
> Certain features of the database may allow you to access services offered separately by third-parties, for example, through the use of JSON specifications that facilitate your access to REST APIs.
>
> Your use of these features is solely at your own risk, and you are solely responsible for complying with any terms and conditions related to use of any such third-party services. Notwithstanding any other terms and conditions related to the third-party services, your use of such database features constitutes your acceptance of that risk and express exclusion of Oracle's responsibility or liability for any damages resulting from such access.

Syntax

```
DBMS_VECTOR_CHAIN.UTL_TO_SUMMARY (
    DATA          IN CLOB,
    PARAMS        IN JSON default NULL
) return CLOB;
```

DATA

This function accepts the input data type in plain text as `CLOB`.

It returns a summary of the input document also as `CLOB`.

PARAMS

Specify summary parameters in JSON format, depending on the service provider that you want to use for document summarization.

If using Oracle Database as the provider:

```
{
  "provider"     : "database",
  "glevel"       : "<summary format>",
  "numParagraphs": <number in the range 1-16>,
  "maxPercent"   : <number in the range 1-100>,
  "num_themes"   : <number in the range 1-50>,
  "language"     : "<name of the language>"
}
```

Table 212-13 Database Provider Parameter Details

Parameter	Description
`provider`	Specify `database` (default setting) to access Oracle Database as the provider. Oracle Text is internally used to extract a document `gist` or summary from user data.
`glevel`	Format to display the summary: • `SENTENCE` | `S`: As a list of sentences • `PARAGRAPH` | `P`: In a free-form paragraph
`numParagraphs`	Maximum number of document paragraphs (or sentences) selected for the summary. The default value is `16`. The `numParagraphs` parameter is used only when this parameter yields a smaller summary size than the summary size yielded by the `maxPercent` parameter, because the function always returns the smallest size summary.
`maxPercent`	Maximum number of document paragraphs (or sentences) selected for the summary, as a percentage of the total paragraphs (or sentences) in the document. The default value is `10`. The `maxPercent` parameter is used only when this parameter yields a smaller summary size than the summary size yielded by the `numParagraphs` parameter, because the function always returns the smallest size summary.
`num_themes`	Number of theme summaries to produce. For example, if you specify `10`, then this function returns the top 10 theme summaries. If you specify `0` or `NULL`, then this function returns all themes in a document. The default value is `50`. If the document contains more than 50 themes, only the top 50 themes show conceptual hierarchy.
`language`	Language name of your summary text, as listed in Supported Languages and Data File Locations.

If using a third-party provider:

Set the following parameters along with additional summarization parameters specific to your provider:

```
{
  "provider"        : "<AI service provider>",
  "credential_name" : "<credential name>",
  "url"             : "<REST endpoint URL for summarization service>",
```

```
  "model"           : "<REST provider summarization model name>",
  "transfer_timeout": <maximum wait time for the request to complete>,
  "<additional REST provider parameter>": "<REST provider parameter value>"
}
```

Table 212-14 Third-Party Provider Parameter Details

Parameter	Description
`provider`	Third-party service provider that you want to access to get the summary. A REST call is made to the specified provider to access its text summarization model. Specify one of the following values: • `cohere` • `googleai` • `huggingface` • `ocigenai` • `openai` • `vertexai`
`credential_name`	Name of the credential in the form: *`schema.credential_name`* A credential name holds authentication credentials to enable access to your provider for making REST API calls. You need to first set up your credential by calling the `DBMS_VECTOR_CHAIN.CREATE_CREDENTIAL` helper function to create and store a credential, and then refer to the credential name here. See CREATE_CREDENTIAL.
`url`	URL of the API endpoint for each REST call. **Note**: For a list of all supported REST endpoints, see Supported Third-Party Provider Operations and Endpoints.
`model`	Name of the third-party text summarization model in the form: *`schema.model_name`* If the model name is not schema-qualified, then the schema of the procedure invoker is used. **Note**: For Generative AI, you must specify *`schema.model_name`*.
`transfer_timeout`	Maximum time to wait for the request to complete. The default value is `60` seconds. You can increase this value for busy web servers.

Additional REST provider parameters:

Optionally, specify additional provider-specific parameters.

> **Important:**
>
> The following examples are for illustration purposes. For accurate and up-to-date information on additional parameters to use, refer to your third-party provider's documentation.

Cohere example:

```
{
  "provider"        : "cohere",
```

```
  "credential_name" : "COHERE_CRED",
  "url"             : "https://api.cohere.example.com/summarize",
  "model"           : "summarize-model",
  "length"          : "medium",
  "format"          : "paragraph",
  "temperature"     : 1.0
}
```

Google AI example:

```
{
  "provider"         : "googleai",
  "credential_name"  : "GOOGLEAI_CRED",
  "url"              : "https://googleapis.example.com/models/",
  "model"            : "summarize-model",
  "generation_config": {
    "temperature"    : 0.9,
    "topP"           : 1,
    "candidateCount" : 1,
    "maxOutputTokens": 256
  }
}
```

Hugging Face example:

```
{
  "provider"          : "huggingface",
  "credential_name"   : "HF_CRED",
  "url"               : "https://api.huggingface.example.co/models/",
  "model"             : "summarize-model"
}
```

Generative AI example:

```
{
 "provider"           : "ocigenai",
 "credential_name"    : "OCI_CRED",
 "url"                : "https://generativeai.oci.example.com/summarizeText",
 "model"              : "summarize-model",
 "length"             : "MEDIUM",
 "format"             : "PARAGRAPH"
}
```

Ollama example:

```
{
  "provider"       : "ollama",
  "host"           : "local",
  "url"            : "http://localhost:11434/api/generate",
  "model"          : "llama3"
}
```

OpenAI example:

```
{
  "provider"          : "openai",
  "credential_name"   : "OPENAI_CRED",
  "url"               : "https://api.openai.example.com",
  "model"             : "summarize-model",
  "max_tokens"        : 256,
  "temperature"       : 1.0
}
```

Vertex AI example:

```
{
  "provider"          : "vertexai",
  "credential_name"   : "VERTEXAI_CRED",
  "url"               : "https://googleapis.example.com/models/",
  "model"             : "summarize-model",
  "generation_config" : {
    "temperature"     : 0.9,
    "topP"            : 1,
    "candidateCount"  : 1,
    "maxOutputTokens" : 256
  }
}
```

Table 212-15 Additional REST Provider Parameter Details

Parameter	Description
`length`	Approximate length of the summary text: • `SHORT`: Roughly up to 2 sentences • `MEDIUM`: Between 3 and 5 sentences • `LONG`: 6 or more sentences • `AUTO`: The model chooses a length based on the input size **Note**: For Generative AI, you must enter this value in uppercase.
`format`	Format to display the summary: • `PARAGRAPH`: In a free-form paragraph • `BULLETS`: In bullet points **Note**: For Generative AI, you must enter this value in uppercase.
`temperature`	Degree of randomness used when generating output text, in the range of `0.0-5.0`. To generate the same output for a prompt, use `0`. To generate a random new text for that prompt, increase the temperature. Default temperature is `1` and the maximum temperature is `5`. **Note**: To summarize a text, start with the temperature set to `0`. If you do not require random results, a recommended temperature value is `0.2` for Generative AI and between `0` and `1` for Cohere. Use a higher value if for example you plan to perform a selection of the various summaries afterward. Do not use a high temperature for summarization because a high temperature encourages the model to produce creative text, which might also include hallucinations.

Table 212-15 (Cont.) Additional REST Provider Parameter Details

Parameter	Description
`extractiveness`	How much to reuse the input in the summary: • `LOW`: Summaries with low extractiveness tend to paraphrase. • `HIGH`: Summaries with high extractiveness lean toward reusing sentences verbatim. **Note**: For Generative AI, you must enter this value in uppercase.
`max_tokens`	Maximum number of tokens in the output text.
`topP`	Probability of tokens in the output, in the range of `0.0-1.0`. A lower value provides less random responses and a higher value provides more random responses.
`candidateCount`	Number of response variations to return, in the range of `1-4`.
`maxOutputTokens`	Maximum number of tokens to generate for each response.

> **Note:**
>
> When specifying the `length`, `format`, and `extractiveness` parameters for Generative AI, ensure to enter the values in uppercase letters.

Examples

These statements generate a summary from an extract on "Transactions", by accessing the Generative AI summarization model.

```
-- select example

var params clob;
exec :params := '
{
 "provider": "ocigenai",
 "credential_name": "OCI_CRED",
 "url": "https://generativeai.oci.example.com/summarizeText",
 "model": "summarize.modelname",
 "temperature": "0.0",
 "extractiveness": "LOW"
}';

select dbms_vector_chain.utl_to_summary(
 'A transaction is a logical, atomic unit of work that contains one or more
SQL statements.  An RDBMS must be able to group SQL statements so that they
are either all committed, which means they are applied to the database, or
all rolled back, which means they are undone.  An illustration of the need
for transactions is a funds transfer from a savings account to a checking
account. The transfer consists of the following separate operations:
    1. Decrease the savings account.
    2. Increase the checking account.
    3. Record the transaction in the transaction journal.
    Oracle Database guarantees that all three operations succeed or fail as a
unit. For example, if a hardware failure prevents a statement in the
```

```
transaction from executing, then the other statements must be rolled back.
    Transactions set Oracle Database apart from a file system. If you perform
an atomic operation that updates several files, and if the system fails
halfway through, then the files will not be consistent. In contrast, a
transaction moves an Oracle database from one consistent state to another.
The basic principle of a transaction is "all or nothing": an atomic operation
succeeds or fails as a whole.',
 json(:params)) from dual;

-- PL/SQL example

declare
  input clob;
  params clob;
  output clob;
begin
  input := 'A transaction is a logical, atomic unit of work that contains one
or more SQL statements.  An RDBMS must be able to group SQL statements so
that they are either all committed, which means they are applied to the
database, or all rolled back, which means they are undone.  An illustration
of the need for transactions is a funds transfer from a savings account to a
checking account. The transfer consists of the following separate operations:
    1. Decrease the savings account.
    2. Increase the checking account.
    3. Record the transaction in the transaction journal.
    Oracle Database guarantees that all three operations succeed or fail as a
unit. For example, if a hardware failure prevents a statement in the
transaction from executing, then the other statements must be rolled back.
    Transactions set Oracle Database apart from a file system. If you perform
an atomic operation that updates several files, and if the system fails
halfway through, then the files will not be consistent. In contrast, a
transaction moves an Oracle database from one consistent state to another.
The basic principle of a transaction is "all or nothing": an atomic operation
succeeds or fails as a whole.';

  params := '
  {
   "provider": "ocigenai",
   "credential_name": "OCI_CRED",
   "url": "https://generativeai.oci.example.com/summarizeText",
   "model": "summarize.modelname",
   "length": "MEDIUM",
   "format": "PARAGRAPH",
   "temperature": 1.0
  }';

  output := dbms_vector_chain.utl_to_summary(input, json(params));
  dbms_output.put_line(output);
  if output is not null then
    dbms_lob.freetemporary(output);
  end if;
exception
  when OTHERS THEN
    DBMS_OUTPUT.PUT_LINE (SQLERRM);
    DBMS_OUTPUT.PUT_LINE (SQLCODE);
```

```
end;
/
```

End-to-end examples:

To run end-to-end example scenarios using this function, see Generate Summary.

212.2 Supported Languages and Data File Locations

These are the supported languages for which language data files are distributed by default in the specified directories.

Language Name	Abbreviation	Data File
AFRIKAANS	af	ctx/data/eos/dreosaf.txt
AMERICAN	us	ctx/data/eos/dreosus.txt
ARABIC	ar	ctx/data/eos/dreosar.txt
BASQUE	eu	ctx/data/eos/dreoseu.txt
BELARUSIAN	be	ctx/data/eos/dreosbe.txt
BRAZILIAN PORTUGUESE	ptb	ctx/data/eos/dreosptb.txt
BULGARIAN	bg	ctx/data/eos/dreosbg.txt
CANADIAN FRENCH	frc	ctx/data/eos/dreosfrc.txt
CATALAN	ca	ctx/data/eos/dreosca.txt
CROATIAN	hr	ctx/data/eos/dreoshr.txt
CYRILLIC SERBIAN	csr	ctx/data/eos/dreoscsr.txt
CZECH	cs	ctx/data/eos/dreoscs.txt
DANISH	dk	ctx/data/eos/dreosdk.txt
DARI	prs	ctx/data/eos/dreosprs.txt
DUTCH	nl	ctx/data/eos/dreosnl.txt
EGYPTIAN	eg	ctx/data/eos/dreoseg.txt
ENGLISH	gb	ctx/data/eos/dreosgb.txt
ESTONIAN	et	ctx/data/eos/dreoset.txt
FINNISH	sf	ctx/data/eos/dreossf.txt
FRENCH	f	ctx/data/eos/dreosf.txt
GALICIAN	ga	ctx/data/eos/dreosga.txt
GERMAN	d	ctx/data/eos/dreosd.txt
GERMAN DIN	din	ctx/data/eos/dreosdin.txt
GREEK	el	ctx/data/eos/dreosel.txt
HEBREW	iw	ctx/data/eos/dreosiw.txt
HINDI	hi	ctx/data/eos/dreoshi.txt
HUNGARIAN	hu	ctx/data/eos/dreoshu.txt
ICELANDIC	is	ctx/data/eos/dreosis.txt
INDONESIAN	in	ctx/data/eos/dreosin.txt

Language Name	Abbreviation	Data File
ITALIAN	i	ctx/data/eos/dreosi.txt
JAPANESE	ja	ctx/data/eos/dreosja.txt
KOREAN	ko	ctx/data/eos/dreosko.txt
LATIN AMERICAN SPANISH	esa	ctx/data/eos/dreosesa.txt
LATIN BOSNIAN	lbs	ctx/data/eos/dreoslbs.txt
LATIN SERBIAN	lsr	ctx/data/eos/dreoslsr.txt
LATVIAN	lv	ctx/data/eos/dreoslv.txt
LITHUANIAN	lt	ctx/data/eos/dreoslt.txt
MACEDONIAN	mk	ctx/data/eos/dreosmk.txt
MALAY	ms	ctx/data/eos/dreosms.txt
MEXICAN SPANISH	esm	ctx/data/eos/dreosesm.txt
NORWEGIAN	n	ctx/data/eos/dreosn.txt
NYNORSK	nn	ctx/data/eos/dreosnn.txt
PERSIAN	fa	ctx/data/eos/dreosfa.txt
POLISH	pl	ctx/data/eos/dreospl.txt
PORTUGUESE	pt	ctx/data/eos/dreospt.txt
ROMANIAN	ro	ctx/data/eos/dreosro.txt
RUSSIAN	ru	ctx/data/eos/dreosru.txt
SIMPLIFIED CHINESE	zhs	ctx/data/eos/dreoszhs.txt
SLOVAK	sk	ctx/data/eos/dreossk.txt
SLOVENIAN	sl	ctx/data/eos/dreossl.txt
SPANISH	e	ctx/data/eos/dreose.txt
SWEDISH	s	ctx/data/eos/dreoss.txt
THAI	th	ctx/data/eos/dreosth.txt
TRADITIONAL CHINESE	zht	ctx/data/eos/dreoszht.txt
TURKISH	tr	ctx/data/eos/dreostr.txt
UKRAINIAN	uk	ctx/data/eos/dreosuk.txt
URDU	ur	ctx/data/eos/dreosur.txt

Related Topics

- CREATE_LANG_DATA
  Use the `DBMS_VECTOR_CHAIN.CREATE_LANG_DATA` chunker helper procedure to load your own language data file into the database.

213
DBMS_WARNING

The DBMS_WARNING package provides a way to manipulate the behavior of PL/SQL warning messages, in particular by reading and changing the setting of the PLSQL_WARNINGS initialization parameter to control what kinds of warnings are suppressed, displayed, or treated as errors. This package provides the interface to query, modify and delete current system or session settings.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_WARNING Subprograms

213.1 DBMS_WARNING Security Model

Note that for all the following interfaces, if value of the scope parameter is SYSTEM, then the user must have ALTER SYSTEM privilege.

213.2 Summary of DBMS_WARNING Subprograms

This table lists the DBMS_WARNING subprograms and briefly describes them.

Table 213-1 ***DBMS_WARNING Package Subprograms***

Subprogram	Description
ADD_WARNING_SETTING_CAT Procedure	Modifies the current session or system warning settings of the warning_category previously supplied
ADD_WARNING_SETTING_NUM Procedure	Modifies the current session or system warning settings of the or warning_number previously supplied
GET_CATEGORY Function	Returns the category name, given the message number
GET_WARNING_SETTING_CAT Function	Returns the specific warning category in the session
GET_WARNING_SETTING_NUM Function	Returns the specific warning number in the session
GET_WARNING_SETTING_STRING Function	Returns the entire warning string for the current session
SET_WARNING_SETTING_STRING Procedure	Replaces previous settings with the new value

213.2.1 ADD_WARNING_SETTING_CAT Procedure

You can modify the current session's or system's warning settings with the value supplied in this procedure. The value will be added to the existing parameter setting if the value for the `warning_category` or `warning_value` has not been set, or override the existing value.

The effect of calling this function is same as adding the qualifier (`ENABLE`/`DISABLE`/`ERROR`) on the category specified to the end of the current session or system setting.

Syntax

```
DBMS_WARNING.ADD_WARNING_SETTING_CAT (
   warning_category    IN    VARCHAR2,
   warning_value       IN    VARCHAR2,
   scope               IN    VARCHAR2);
```

Parameters

Table 213-2 ADD_WARNING_SETTING_CAT Procedure Parameters

Parameter	Description
`warning_category`	Name of the category. Allowed values are `ALL`, `INFORMATIONAL`, `SEVERE` and `PERFORMANCE`.
`warning_value`	Value for the category. Allowed values are `ENABLE`, `DISABLE`, and `ERROR`.
`scope`	Specifies if the changes are being performed in the session context or the system context. Allowed values are `SESSION` or `SYSTEM`.

213.2.2 ADD_WARNING_SETTING_NUM Procedure

You can modify the current `session` or `system` warning settings with the value supplied in this procedure. If the value was already set, you will override the existing value.

The effect of calling this function is same as adding the qualifier (`ENABLE` / `DISABLE`/ `ERROR`) on the category specified to the end of the current session or system setting.

Syntax

```
DBMS_WARNING.ADD_WARNING_SETTING_NUM (
   warning_number      IN    NUMBER,
   warning_value       IN    VARCHAR2,
   scope               IN    VARCHAR2);
```

Parameters

Table 213-3 ADD_WARNING_SETTING_NUM Procedure Parameters

Parameter	Description
`warning_number`	The warning number. Allowed values are all valid warning numbers.
`warning_value`	Value for the category. Allowed values are `ENABLE`, `DISABLE`, and `ERROR`.

Table 213-3 (Cont.) ADD_WARNING_SETTING_NUM Procedure Parameters

Parameter	Description
scope	Specifies if the changes are being performed in the session context or the system context. Allowed values are SESSION or SYSTEM.

Example 213-1 Enabling the Deprecation Warnings

This example shows how to enable the DEPRECATE pragma warnings using the DBMS_WARNING package for the session.

```
BEGIN
  DBMS_WARNING.ADD_WARNING_SETTING_NUM (6019,
                                        'ENABLE',
                                        'SESSION');

  DBMS_WARNING.ADD_WARNING_SETTING_NUM (6020,
                                        'ENABLE',
                                        'SESSION');

  DBMS_WARNING.ADD_WARNING_SETTING_NUM (6021,
                                        'ENABLE',
                                        'SESSION');

  DBMS_WARNING.ADD_WARNING_SETTING_NUM (6022,
                                        'ENABLE',
                                        'SESSION');
END;
```

213.2.3 GET_CATEGORY Function

This function returns the category name, given the message number.

Syntax

```
DBMS_WARNING.GET_CATEGORY (
   warning_number  IN   pls_integer)
RETURN VARCHAR2;
```

Parameters

Table 213-4 GET_CATEGORY Function Parameters

Parameter	Description
warning_number	The warning message number.

213.2.4 GET_WARNING_SETTING_CAT Function

This function returns the specific warning category setting for the current session.

Syntax

```
DBMS_WARNING.GET_WARNING_SETTING_CAT (
   warning_category    IN    VARCHAR2)
RETURN warning_value;
```

Parameters

Table 213-5 GET_WARNING_SETTING_CAT Function Parameters

Parameter	Description
warning_category	Name of the category. Allowed values are all valid category names (ALL, INFORMATIONAL, SEVERE and PERFORMANCE).

213.2.5 GET_WARNING_SETTING_NUM Function

This function returns the specific warning number setting for the current session.

Syntax

```
DBMS_WARNING.GET_WARNING_SETTING_NUM (
   warning_number      IN    NUMBER)
RETURN warning_value;
```

Parameters

Table 213-6 GET_WARNING_SETTING_NUM Function Parameters

Parameter	Description
warning_number	Warning number. Allowed values are all valid warning numbers.

213.2.6 GET_WARNING_SETTING_STRING Function

This function returns the entire warning string for the current session.

Syntax

```
DBMS_WARNING.GET_WARNING_SETTING_STRING
 RETURN VARCHAR2;
```

Usage Notes

Use this function when you do not have SELECT or READ privilege on v$parameter or v$paramater2 fixed tables, or if you want to parse the warning string yourself and then modify and set the new value using SET_WARNING_SETTING_STRING.

213.2.7 SET_WARNING_SETTING_STRING Procedure

This procedure replaces previous settings with the new value.

The warning string may contain mix of category and warning numbers using the same syntax as used on the right hand side of '=' when issuing an `ALTER SESSION` or `SYSTEM SET PLSQL_WARNINGS` command. This will have same effect as `ALTER SESSION` or `ALTER SYSTEM` command.

Syntax

```
DBMS_WARNING.SET_WARNING_SETTING_STRING (
   warning_value   IN   VARCHAR2,
   scope           IN   VARCHAR2);
```

Parameters

Table 213-7 SET_WARNING_SETTING_STRING Procedure Parameters

Parameter	Description
`warning_value`	The new string that will constitute the new value.
`scope`	This will specify if the changes are being done in the session context, or system context. Allowed values are `SESSION` or `SYSTEM`.

214
DBMS_WM

The `DBMS_WM` package provides an interface to Oracle Database Workspace Manager (often referred to as Workspace Manager).

For a complete description of this package, see `DBMS_WM` in the *Oracle Database Workspace Manager Developer's Guide*.

215
DBMS_WORKLOAD_CAPTURE

The `DBMS_WORKLOAD_CAPTURE` package configures the Workload Capture system and produce the workload capture data.

Replay of this capture is implemented by way of the DBMS_WORKLOAD_REPLAY package.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_WORKLOAD_CAPTURE Subprograms

> **See Also:**
>
> *Oracle Database Testing Guide* for more information about database replay

215.1 DBMS_WORKLOAD_CAPTURE Overview

Since the capture infrastructure is instance wide (and also within an Oracle Real Application Clusters (Oracle RAC)), only one workload capture is being produced at any point in time. Thus capture interfaces do not need a state object passed in as a parameter since there is one single state at any point in time. This means that all subprograms cannot be methods of an object but are package wide PL/SQL subprograms.

215.2 DBMS_WORKLOAD_CAPTURE Security Model

The security model describes the privileges needed for using `DBMS_WORKLOAD_CAPTURE`.

The following code describes the minimal set of privileges required to:

- Create directory objects
- Operate the interface provided by the `DBMS_WORKLOAD_CAPTURE` and `DBMS_WORKLOAD_REPLAY` packages
- Act as a replay client user (wrc `someuser`/`somepassword` or wrc `USER=someuser PASSWORD=somepassword`)

```
DROP USER rom1 CASCADE;
CREATE USER rom1 IDENTIFIED BY rom1;

GRANT EXECUTE ON DBMS_WORKLOAD_CAPTURE TO rom1;
GRANT EXECUTE ON DBMS_WORKLOAD_REPLAY TO rom1;

GRANT CREATE SESSION TO rom1;
GRANT CREATE ANY DIRECTORY TO rom1;
GRANT SELECT_CATALOG_ROLE TO rom1;
GRANT BECOME USER TO rom1;
```

Appropriate OS permissions are required to access and manipulate files and directories on both the capture and replay system. This means that the Oracle process(es) and the OS user performing the capture or replay must be able to access and manipulate at least one common directory accessible from the host where the instance is running. Additionally, the OS user performing the replay should be able to execute wrc on hosts that are used for the replay clients and be able to access the file system appropriately to copy the capture to the replay clients' hosts if required.

215.3 Summary of DBMS_WORKLOAD_CAPTURE Subprograms

This table lists the `DBMS_WORKLOAD_CAPTURE` package subprograms in alphabetical order.

Table 215-1 DBMS_WORKLOAD_CAPTURE Package Subprograms

Subprogram	Description
ADD_FILTER Procedures	Adds a specified filter
DECRYPT_CAPTURE Procedure	This procedure decrypts sensitive data in workload capture that was encrypted using various advanced encryption standards such as `AES128`, `AES192`, or `AES256`.
DELETE_CAPTURE_INFO Procedure	Deletes the rows in the `DBA_WORKLOAD_CAPTURES` and `DBA_WORKLOAD_FILTERS` views that corresponds to the given workload capture ID
DELETE_FILTER Procedure	Deletes a specified filter
ENCRYPT_CAPTURE Procedure	This procedure encrypts sensitive data in workload capture using various advanced encryption standards such as `AES128`, `AES192`, or `AES256`.
EXPORT_AWR Procedure	Exports the AWR snapshots associated with a given capture ID
FINISH_CAPTURE Procedure	Finalizes the workload capture by signaling all connected sessions to stop capture, and stops future requests to the database from being captured
GET_CAPTURE_INFO Function	Retrieves all the information regarding a workload capture present in the stipulated directory, imports the information into the `DBA_WORKLOAD_CAPTURES` and `DBA_WORKLOAD_FILTERS` views, and returns the appropriate `DBA_WORKLOAD_CAPTURES.ID`
IMPORT_AWR Function	Imports the AWR snapshots associated with a given capture ID
REPORT Function	Returns a report on the workload capture under consideration using one or more different sources
START_CAPTURE Procedure	Initiates workload capture on all instances

215.3.1 ADD_FILTER Procedures

This procedure adds a filter to capture a subset of the workload.

Syntax

```
DBMS_WORKLOAD_CAPTURE.ADD_FILTER (
   fname           IN   VARCHAR2 NOT NULL,
   fattribute      IN   VARCHAR2 NOT NULL,
   fvalue          IN   VARCHAR2 NOT NULL);
```

```
DBMS_WORKLOAD_CAPTURE.ADD_FILTER (
   fname           IN   VARCHAR2 NOT NULL,
   fattribute      IN   VARCHAR2 NOT NULL,
   fvalue          IN   NUMBER NOT NULL);
```

Parameters

Table 215-2 ADD_FILTER Procedure Parameters

Parameter	Description
fname	Name for the filter to be added. Can be used to delete the filter later if it is not required. (Mandatory)
fattribute	Specifies the attribute on which the filter needs to be applied (Mandatory). The possible values are: • INSTANCE_NUMBER - type NUMBER • USER - type STRING • MODULE - type STRING • ACTION - type STRING • PROGRAM - type STRING • SERVICE - type STRING • PDB_NAME - type STRING
fvalue	Specifies the value to which the given attribute should be equal to for the filter to be considered active. Wildcards like '%' are acceptable for all attributes that are of type STRING. This means that the filter for a NUMBER attribute is parsed as "attribute = value", with the filter for a STRING attribute parsed as "attribute like value" (Mandatory).

Usage Notes

- The workload capture filters work in either the DEFAULT INCLUSION or the DEFAULT EXCLUSION mode as determined by the default_action input to the START_CAPTURE Procedure.
- ADD_FILTER adds a new filter that affects the next workload capture, and whether the filters are considered as INCLUSION filters or EXCLUSION filters depends on the value of the default_action input to START_CAPTURE Procedure.
- Filters once specified are valid only for the next workload capture. If the same set of filters need to be used for subsequent capture, they need to be specified each time before the START_CAPTURE Procedure is executed.
- All the filters are listed in the DBA_WORKLOAD_FILTERS view.
- You can capture the workload for a particular PDB by specifying a filter of PDB type.

Examples

- By default, a capture works in an INCLUSION mode, which records everything except for those requests that satisfy conditions of specified filters. For example, if you want to exclude all requests from SCOTT, you can add the following filter before starting a capture.

  ```
  EXEC DBMS_WORKLOAD_CAPTURE.ADD_FILTER ('filter user1', 'USER', 'SCOTT');
  ```

- Multiple filters are evaluated according to the logical disjunction operator OR. Therefore, if you want to record workload for both SCOTT and JOHN, you add an additional filter:

  ```
  EXEC DBMS_WORKLOAD_CAPTURE.ADD_FILTER ('filter user2', 'USER', 'JOHN');
  ```

- In a CDB, you exclude the workload of a particular PDB by the filter:

```
EXEC DBMS_WORKLOAD_CAPTURE.ADD_FILTER ('filter pdb workload', 'PDB_NAME',
'CDB1_PDB1');
```

- To use DBMS_APPLICATION_INFO to identify workload that is issued to the database:

```
DBMS_APPLICATION_INFO.SET_MODULE('ORDER_ENTRY', NULL);
-- run some SQL here
DBMS_APPLICATION_INFO.SET_ACTION('ORDER_ENTRY_LOG');
-- run logging SQL
```

- If having captured workload, you want to exclude the logging SQL from the captured, specify a filter for capture:

```
DBMS_WORKLOAD_CAPTURE.ADD_FILTER('filter logging operations', 'ACTION',
'ORDER_ENTRY_LOG');
```

- To filter out the full order entry transaction, define a filter:

```
DBMS_WORKLOAD_CAPTURE.ADD_FILTER('filter order entry', 'MODULE', 'ORDER_ENTRY');
```

215.3.2 DECRYPT_CAPTURE Procedure

This procedure decrypts sensitive data in workload capture that was encrypted using various advanced encryption standards such as AES128, AES192, or AES256.

Syntax

```
DBMS_WORKLOAD_CAPTURE.DECRYPT_CAPTURE (
   src_dir               IN   VARCHAR2,
   dst_dir               IN   VARCHAR2);
```

Parameters

Table 215-3 DECRYPT_CAPTURE Procedure Parameters

Parameter	Description
src_dir	A directory object pointing to the workload capture to be decrypted. This parameter is case sensitive.
dst_dir	A directory object pointing to an OS path that has write permissions. The decrypted capture files will be written to this directory This parameter is case sensitive.

Usage Notes

This procedure relies on a software keystore. The identifier is `oracle.rat.database_replay.encryption` (case-sensitive).

215.3.3 DELETE_CAPTURE_INFO Procedure

This procedure deletes the rows in the DBA_WORKLOAD_CAPTURES and DBA_WORKLOAD_FILTERS views that corresponds to the given workload capture ID.

Syntax

```
DBMS_WORKLOAD_CAPTURE.DELETE_CAPTURE_INFO
   capture_id     IN   NUMBER);
```

Parameters

Table 215-4 DELETE_CAPTURE_INFO Procedure Parameters

Parameter	Description
capture_id	ID of the workload capture that needs to be deleted. Corresponds to DBA_WORKLOAD_CAPTURES.ID. (Mandatory)

Usage Notes

Passing the ID of a capture that is in progress will first automatically stop that capture.

215.3.4 DELETE_FILTER Procedure

This procedure deletes a specified filter.

Syntax

```
DBMS_WORKLOAD_CAPTURE.DELETE_FILTER (
   filter_name           IN   VARCHAR2(40) NOT NULL);
```

Parameters

Table 215-5 DELETE_FILTER Procedure Parameters

Parameter	Description
filter_name	Filter to be deleted

Usage Notes

The DELETE_FILTER Procedure only affects filters that have not been used by any previous capture. Consequently, filters can be deleted only if they have been added using the ADD_FILTER Procedures after any capture has been completed. Filters that have been added using ADD_FILTER before a START_CAPTURE and FINISH_CAPTURE cannot be deleted anymore using this subprogram.

215.3.5 ENCRYPT_CAPTURE Procedure

This procedure encrypts sensitive data in workload capture using various advanced encryption standards such as AES128, AES192, or AES256. To encrypt capture on the fly, use encryption parameter in START_CAPTURE procedure.

Syntax

```
DBMS_WORKLOAD_CAPTURE.ENCRYPT_CAPTURE
   src_dir     IN VARCHAR2,
   dst_dir     IN VARCHAR2,
   encryption  IN VARCHAR2 DEFAULT 'AES256');
```

Parameters

Table 215-6 ENCRYPT_CAPTURE Procedure Parameters

Parameter	Description
src_dir	A directory object pointing to the workload capture to be encrypted. The parameter value is case sensitive.
dst_dir	A directory object pointing to an OS path that has write permissions. The encrypted capture files will be written to this directory. The parameter value is case sensitive.
encryption	Specifies if sensitive data in workload capture is encrypted or not. The possible values are: AES128, AES192, and AES256. The default value is AES256.

Usage Notes

For encrypted capture, the ENCRYPT_CAPTURE Procedure relies on a software keystore. The identifier is oracle.rat.database_replay.encryption (case-sensitive).

215.3.6 EXPORT_AWR Procedure

This procedure exports the AWR snapshots associated with a given capture ID.

Syntax

```
DBMS_WORKLOAD_CAPTURE.EXPORT_AWR (
   capture_id     IN NUMBER);
```

Parameters

Table 215-7 EXPORT_AWR Procedure Parameters

Parameter	Description
capture_id	ID of the capture whose AWR snapshots are to be exported. (Mandatory)

Usage Notes

This procedure works only if the corresponding workload capture was performed in the current database (meaning that the corresponding row in DBA_WORKLOAD_CAPTURES was not created by calling the GET_CAPTURE_INFO Function) and the AWR snapshots that correspond to the original capture time period are still available.

215.3.7 FINISH_CAPTURE Procedure

This procedure signals all connected sessions to stop the workload capture and stops future requests to the database from being captured.

Syntax

```
DBMS_WORKLOAD_CAPTURE.FINISH_CAPTURE
   timeout     IN   NUMBER  DEFAULT 30
   reason       IN   VARCHAR2 DEFAULT NULL);
```

Parameters

Table 215-8 FINISH_CAPTURE Procedure Parameters

Parameter	Description
timeout	Specifies in seconds for how long the procedure should wait before it times out. Pass 0 if you want to cancel the current workload capture and not wait for any sessions to flush it's capture buffers. Default value: 30 seconds
reason	Specifies a reason for calling the procedure. The reason appears in the column ERROR_MESSAGE of the view DBA_WORKLOAD_CAPTURES.

Usage Notes

- By default, FINISH_CAPTURE waits for 30 seconds to receive a successful acknowledgement from all sessions in the database cluster before timing out.
- All sessions that either were in the middle of executing a user request or received a new user request, while FINISH_CAPTURE was waiting for acknowledgements, flush their buffers and send back their acknowledgement to FINISH_CAPTURE.
- If a database session remains idle (waiting for the next user request) throughout the duration of FINISH_CAPTURE, the session might have unflushed capture buffers and does not send it's acknowledgement to FINISH_CAPTURE.

  To avoid this, do not have sessions that remain idle (waiting for the next user request) while invoking FINISH_CAPTURE. Either close the database session(s) before running FINISH_CAPTURE or send new database requests to those sessions during FINISH_CAPTURE.

215.3.8 GET_CAPTURE_INFO Function

This procedure retrieves all information regarding a workload capture present in the stipulated directory, imports the information into the DBA_WORKLOAD_CAPTURES and DBA_WORKLOAD_FILTERS views, and returns the appropriate DBA_WORKLOAD_CAPTURES.ID

Syntax

```
DBMS_WORKLOAD_CAPTURE.GET_CAPTURE_INFO
   dir      IN    VARCHAR2)
  RETURN NUMBER;
```

Parameters

Table 215-9 GET_CAPTURE_INFO Function Parameters

Parameter	Description
dir	Name of the DIRECTORY object (case sensitive) where all the workload capture files are located (Mandatory)

Usage Notes

If an appropriate row describing the capture in the stipulated directory already exists in DBA_WORKLOAD_CAPTURES, the GET_CAPTURE_INFO Function simply returns that row's DBA_WORKLOAD_CAPTURES.ID. If no existing row matches the capture present in the stipulated directory a new row is inserted to DBA_WORKLOAD_CAPTURES and that row's ID is returned.

215.3.9 IMPORT_AWR Function

This procedure imports the AWR snapshots associated with a given capture ID provided those AWR snapshots were exported earlier from the original capture system using the EXPORT_AWR procedure.

Syntax

```
DBMS_WORKLOAD_CAPTURE.IMPORT_AWR (
   capture_id       IN   NUMBER,
   staging_schema   IN   VARCHAR2,
   force_cleanup    IN   BOOLEAN DEFAULT FALSE)
 RETURN NUMBER;
```

Parameters

Table 215-10 IMPORT_AWR Function Parameters

Parameter	Description
capture_id	ID of the capture whose AWR snapshots should be imported. (Mandatory)
staging_schema	Name of a valid schema in the current database which can be used as a staging area while importing the AWR snapshots from the capture directory to the SYS AWR schema.The SYS schema is not a valid input. (Mandatory, Case sensitive).

Table 215-10 (Cont.) IMPORT_AWR Function Parameters

Parameter	Description

force_cleanup	Values:
	• TRUE - any AWR data present in the given staging_schema are removed before the actual import operation. All tables with names that match any of the tables in AWR are dropped before the actual import. This typically is equivalent to dropping all tables returned by the following SQL:
	`SELECT table_name FROM dba_tables`
	`WHERE owner = staging_schema`
	`AND table_name like 'WR_$%';`
	Use this option only if you are sure that there are no important data in any such tables in the staging_schema.
	• FALSE - (default) no tables dropped from the staging_schema prior to the import operation

Return Values

Returns the new randomly generated database ID that was used to import the AWR snapshots. The same value can be found in the AWR_DBID column in the DBA_WORKLOAD_CAPTURES view.

Usage Notes

IMPORT_AWR fails if the staging_schema provided as input contains any tables with the same name as any of the AWR tables, such as WRM$_SNAPSHOT or WRH$_PARAMETER. Please drop any such tables in the staging_schema before invoking IMPORT_AWR.

Related Topics

- EXPORT_AWR Procedure
  This procedure exports the AWR snapshots associated with a given capture ID.

215.3.10 REPORT Function

This function generates a report on the stipulated workload capture.

Syntax

```
DBMS_WORKLOAD_CAPTURE.REPORT (
   capture_id      IN   NUMBER,
   format          IN   VARCHAR2)
 RETURN CLOB;
```

Parameters

Table 215-11 REPORT Function Parameters

Parameter	Description
capture_id	ID of the workload capture whose capture report is required. (Mandatory) This relates to the directory that contains the workload capture on which the Report needs to be generated. Should be a valid DIRECTORY object that points to a valid directory in the host system that contains a workload capture.
format	Specifies the report format. Valid values are DBMS_WORKLOAD_CAPTURE.TYPE_TEXT and DBMS_WORKLOAD_CAPTURE.TYPE_HTML.(Mandatory)

Return Values

The report body in the desired format returned as a CLOB.

Table 215-12 Constants Used by Report Function

Constant	Type	Value	Description
TYPE_HTML	VARCHAR2(4)	'HTML'	Generates the HTML version of the report
TYPE_TEXT	VARCHAR2(4)	'TEXT'	Used as input to the format argument to generate the text version of the report

215.3.11 START_CAPTURE Procedure

This procedure initiates workload capture on all instances.

Syntax

```
DBMS_WORKLOAD_CAPTURE.START_CAPTURE (
   name              IN  VARCHAR2,
   dir               IN  VARCHAR2,
   duration          IN  NUMBER   DEFAULT NULL,
   default_action    IN  VARCHAR2 DEFAULT 'INCLUDE',
   auto_unrestrict   IN  BOOLEAN  DEFAULT TRUE,
   capture_sts       IN  BOOLEAN  DEFAULT FALSE,
   sts_cap_interval  IN  NUMBER   DEFAULT 300,
   plsql_mode        IN  VARCHAR2 DEFAULT 'TOP_LEVEL',
   encryption        IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 215-13 START_CAPTURE Procedure Parameters

Parameter	Description
name	Name of the workload capture. Allows the workload capture to be given a label, such as "Thanksgiving weekend" or "Christmas peak workload" for future reference. The workload capture's name is preserved along with the captured workload actions. (Mandatory)
dir	Name of the DIRECTORY object (case sensitive) where all the workload capture files are stored. Should contain enough space to hold all the workload capture files. (Mandatory)
duration	Optional input to specify the duration (in seconds) for which the workload needs to be captured. DEFAULT is NULL which means that workload capture continues until the user executes DBMS_WORKLOAD_CAPTURE.FINISH_CAPTURE.
default_action	Can be either INCLUDE or EXCLUDE. Determines whether, by default, every user request should be captured or not. Also determines whether the workload filters specified should be considered as INCLUSION filters or EXCLUSION filters. • If INCLUDE, by default all user requests to the database are captured, except for the part of the workload defined by the filters. In this case, all the filters specified using the ADD_FILTER Procedures are treated as EXCLUSION filters, determining the workload that is not captured. (DEFAULT, and so all the filters specified are assumed to be EXCLUSION filters.) • If EXCLUDE, by default no user request to the database is captured, except for the part of the workload defined by the filters. In this case, all the filters specified using the ADD_FILTER Procedures are treated as INCLUSION filters, determining the workload that is captured.
auto_unrestrict	Can be either TRUE or FALSE. • If TRUE, all instances started up in RESTRICTED mode using STARTUP RESTRICT are automatically unrestricted upon a successful START_CAPTURE. (DEFAULT) • If FALSE, no database instance is automatically unrestricted.
capture_sts	If this parameter is TRUE, a SQL tuning set capture is also started in parallel with workload capture. The resulting SQL tuning set can be exported using the EXPORT_AWR Procedure along with the AWR data. Currently, parallel STS capture is not supported in an Oracle RAC environment, so this parameter has no effect if used in that context. Capture filters defined using the DBMS_WORKLOAD_REPLAY interface do not apply to the SQL tuning set capture. The calling user must have the appropriate privileges ('ADMINISTER SQL TUNING SET'). If starting SQL set capture fails, workload capture is stopped. The reason is stored in DBA_WORKLOAD_CAPTURES.ERROR_MESSAGE. The default value is FALSE.
sts_cap_interval	Specifies the capture interval of the SQL set capture from the cursor cache in seconds. The default value is 300.

Table 215-13 (Cont.) START_CAPTURE Procedure Parameters

Parameter	Description
plsql_mode	Specifies the PL/SQL capture mode: • TOP_LEVEL — only top-level PL/SQL calls are captured • EXTENDED — both top-level PL/SQL calls and SQL called from PL/SQL are captured
encryption	Specify if sensitive data in workload capture is encrypted or not. The possible values are: • NULL — capture files are not encrypted • AES128 — capture files are encrypted using AES128 • AES192 — capture files are encrypted using AES192 • AES256 — capture files is encrypted using AES256 The default value is NULL.

Usage Notes

- All user requests sent to database after a successful invocation of START_CAPTURE are recorded in the given dir directory for the given duration provided that one was specified. If no duration was specified, the capture lasts indefinitely until the FINISH_CAPTURE Procedure is executed.
- A workload capture once started continues to record user requests across database instance shutdowns and startups for the specified duration, or until FINISH_CAPTURE is executed, whichever occurs first.
- One can use workload filters (as described with regard to the ADD_FILTER Procedures) to capture only a subset of the user requests sent to the database. By default, when no workload filters are defined, all user requests are captured.
- Workload that is initiated from Oracle Database background processes (such as SMON, PMON, MMON) and Oracle Database Scheduler Jobs (as detailed in the DBMS_SCHEDULER package) is not captured, no matter how the workload filters are defined. These activities should happen automatically on an appropriately configured replay system.
- By default, all database instances that were started up in RESTRICTED mode using STARTUP RESTRICT are UNRESTRICTED upon a successful invocation of START_CAPTURE Use FALSE for the auto_unrestrict input parameter, if you do not want this behavior.
- It is important to have a well-defined starting point for the workload so that the replay system can be restored to that point before initiating a replay of the captured workload. To have a well-defined starting point for the workload capture, it is preferable not to have any active user sessions when START_CAPTURE is executed. If ongoing sessions have ongoing transactions, those transactions are not replayed properly in subsequent database replays, since only that part of the transaction whose calls were executed after START_CAPTURE are replayed.
- For encrypted capture, the START_CAPTURE Procedure relies on a software keystore. The identifier is oracle.rat.database_replay.encryption (case-sensitive).
- You must configure a software keystore in auto-login mode. Otherwise, if the database is bounced during capture, the capture is automatically terminated.

216
DBMS_WORKLOAD_REPLAY

The `DBMS_WORKLOAD_REPLAY` package provides a interface to replay a workload capture.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_WORKLOAD_REPLAY Subprograms

> **See Also:**
>
> *Oracle Database Testing Guide* for more information about database replay

216.1 DBMS_WORKLOAD_REPLAY Overview

The `DBMS_WORKLOAD_REPLAY` package provides an interface to replay a workload capture that was originally created by way of the DBMS_WORKLOAD_CAPTURE package.

Typically, the `DBMS_WORKLOAD_CAPTURE` package is used in the production system to capture a production workload, and the `DBMS_WORKLOAD_REPLAY` package is subsequently used in a test system to replay the captured production workload for testing purposes.

Related Topics

- DBMS_WORKLOAD_CAPTURE
  The `DBMS_WORKLOAD_CAPTURE` package configures the Workload Capture system and produce the workload capture data.

216.2 DBMS_WORKLOAD_REPLAY Security Model

The security model describes the privileges needed for using `DBMS_WORKLOAD_REPLAY`.

The following code sample shows the minimum set of privileges required to:

- Create directory objects
- Operate the interface provided by the `DBMS_WORKLOAD_CAPTURE` and `DBMS_WORKLOAD_REPLAY` packages
- Act as a replay client user (wrc `someuser/somepassword` or wrc `USER=someuser PASSWORD=somepassword`)

```
DROP USER rom1 CASCADE;
CREATE USER rom1 IDENTIFIED BY rom1;

GRANT EXECUTE ON DBMS_WORKLOAD_CAPTURE TO rom1;
GRANT EXECUTE ON DBMS_WORKLOAD_REPLAY TO rom1;
```

```
GRANT CREATE SESSION TO rom1;
GRANT CREATE ANY DIRECTORY TO rom1;
GRANT SELECT_CATALOG_ROLE TO rom1;
GRANT BECOME USER TO rom1;
```

Appropriate OS permissions are required to access and manipulate files and directories on both the capture and replay system. The Oracle process(es) and the OS user performing the capture or replay must be able to access and manipulate at least one common directory accessible from the host where the instance is running.

The replay client is a multithreaded program (an executable named `wrc` located in the `$ORACLE_HOME/bin` directory) where each thread submits a workload from a captured session. The OS user performing the replay must be able to execute `wrc` on hosts that are used for the replay clients and be able to access the file system appropriately to be able to copy the capture to the replay clients' hosts if required.

216.3 Summary of DBMS_WORKLOAD_REPLAY Subprograms

This table lists the `DBMS_WORKLOAD_REPLAY` package subprograms in alphabetical order.

Table 216-1 DBMS_WORKLOAD_REPLAY Package Subprograms

Subprogram	Description
ADD_CAPTURE Function	Adds the given capture to the current schedule
ADD_FILTER Procedure	Adds a filter to replay only a subset of the captured workload
ADD_SCHEDULE_ORDERING Function	Adds a schedule order between two captures
ASSIGN_GROUP_TO_INSTANCE Procedure	Modifies the view `DBA_WORKLOAD_GROUP_ASSIGNMENTS`
BEGIN_REPLAY_SCHEDULE Procedure	Initiates the creation of a reusable replay schedule
CALIBRATE Function	Operates on a processed workload capture directory to estimate the number of hosts and workload replay clients needed to faithfully replay the given workload
CANCEL_REPLAY Procedure	Cancels the workload replay in progress
COMPARE_PERIOD_REPORT Procedure	Generates a report comparing a replay to its capture or to another replay of the same capture
COMPARE_SQLSET_REPORT Function	Generates a report comparing a sqlset captured during replay to one captured during workload capture or to one captured during another replay of the same capture
CREATE_FILTER_SET Procedure	Uses the replay filters added to create a set of filters to use against the replay in `replay_dir`
DELETE_FILTER Procedure	Deletes the named filter
DELETE_REPLAY_INFO Procedure	Deletes the rows in `DBA_WORKLOAD_REPLAYS` that corresponds to the given workload replay ID
END_REPLAY_SCHEDULE Procedure	Wraps up the creation of the current schedule
EXPORT_AWR Procedure	Exports the Automatic Workload Repository (AWR) snapshots associated with a given replay ID
GENERATE_CAPTURE_SUBSET Procedure	Creates a new capture from an existing workload capture

Table 216-1 (Cont.) DBMS_WORKLOAD_REPLAY Package Subprograms

Subprogram	Description
GET_DIVERGING_STATEMENT Function	Exports the Automatic Workload Repository (AWR) snapshots associated with a given replay ID
GET_REPLAY_DIRECTORY Function	Returns the current replay directory set by the SET_REPLAY_DIRECTORY Procedure.
GET_REPLAY_INFO Function	Retrieves information about the workload capture and the history of all the workload replay attempts from the related directory
GET_REPLAY_TIMEOUT Procedure	Retrieves the replay timeout setting
IMPORT_AWR Function	Imports the Automatic Workload Repository (AWR) snapshots associated with a given replay ID
INITIALIZE_CONSOLIDATED_REPLAY Procedure	Puts the database state in `INIT` for a multiple-capture replay
INITIALIZE_REPLAY Procedure	Initializes replay, and loads specific data produced during processing into the database
IS_REPLAY_PAUSED Function	Reports whether the replay is currently paused
LOAD_LONG_SQLTEXT Procedure	Loads the captured SQL statements that are longer than 1000 characters to the `DBA_WORKLOAD_LONG_SQLTEXT` view
PAUSE_REPLAY Procedure	Pauses the in-progress workload replay
POPULATE_DIVERGENCE Procedure	Precomputes the divergence information for the given call, stream, or the whole replay so that the GET_DIVERGING_STATEMENT Function returns as quickly as possible for the precomputed calls
PREPARE_CONSOLIDATED_REPLAY Procedure	Puts the database in a special "Prepare" mode for a multiple-capture replay
PREPARE_REPLAY Procedure	Puts the database in a special "Prepare" mode
PROCESS_CAPTURE Procedure	Processes the workload capture found in `capture_dir` in place
REMAP_CONNECTION Procedure	Remaps the captured connection to a new one so that the user sessions can connect to the database in a desired way during workload replay
REMOVE_CAPTURE Procedure	Removes the given capture from the current schedule
REMOVE_SCHEDULE_ORDERING Procedure	Removes an existing schedule order from the current replay schedule
REPORT Function	Generates a report on the given workload replay
RESUME_REPLAY Procedure	Resumes a paused workload replay
REUSE_REPLAY_FILTER_SET Procedure	Reuses filters in the specified filter set as if each were added using the ADD_SCHEDULE_ORDERING Function
SET_ADVANCED_PARAMETER Procedure	Sets an advanced parameter for replay besides the ones used with the PREPARE_REPLAY Procedure
SET_REPLAY_DIRECTORY Procedure	Sets a directory that contains multiple workload captures as the current replay directory
SET_REPLAY_TIMEOUT Procedure	Sets the replay timeout setting
SET_SQL_MAPPING Procedure	Specifies SQL statements to be skipped or replaced during a database replay operation

Table 216-1 (Cont.) DBMS_WORKLOAD_REPLAY Package Subprograms

Subprogram	Description
SET_USER_MAPPING Procedure	Sets a new schema or user name to be used during replay instead of the captured user
START_CONSOLIDATED_REPLAY Procedure	Starts the replay of a multiple-capture capture
START_REPLAY Procedure	Starts the workload replay
USE_FILTER_SET Procedure	Uses the given filter set that has been created by calling the CREATE_FILTER_SET Procedure to filter the current replay

216.3.1 ADD_CAPTURE Function

This function adds the given capture to the current schedule. The directory has to be a valid capture processed in the current database's version. It returns a unique ID that identifies this capture within this schedule.

Syntax

```
DBMS_WORKLOAD_REPLAY.ADD_CAPTURE (
   capture_dir_name      IN    VARCHAR2,
   start_delay_seconds   IN    NUMBER  DEFAULT 0,
   stop_replay           IN    BOOLEAN FALSE,
   take_begin_snapshot   IN    BOOLEAN TRUE,
   take_end_snapshot     IN    BOOLEAN TRUE,
   query_only            IN    BOOLEAN DEFAULT FALSE)
 RETURN NUMBER;

DBMS_WORKLOAD_REPLAY.ADD_CAPTURE (
   capture_dir_name      IN    VARCHAR2,
   start_delay_seconds   IN    NUMBER  DEFAULT 0,
   stop_replay           IN    BOOLEAN FALSE,
   take_begin_snapshot   IN    BOOLEAN TRUE,
   take_end_snapshot     IN    BOOLEAN TRUE,
   query_only            IN    VARCHAR2 DEFAULT 'N')
 RETURN NUMBER;
```

Parameters

Table 216-2 ADD_CAPTURE Function Parameters

Parameter	Description
capture_dir_name	Name of the OS directory containing the capture under the replay top-level directory
start_delay_seconds	Delay time in seconds before the replay of this capture starts
stop_replay	Stop the replay after it finishes
take_begin_snapshot	Take an AWR snapshot when the replay of this capture starts
take_end_snapshot	Take an AWR snapshot when the replay of this capture finishes
query_only	Replay only the read-only queries of this workload capture

Usage Notes

The SET_REPLAY_DIRECTORY Procedure must have already been called.

216.3.2 ADD_FILTER Procedure

This procedure adds a filter to replay only a subset of the captured workload.

The procedure adds a new filter that is used in the next replay filter set created using the CREATE_FILTER_SET Procedure. This filter will be considered an "INCLUSION" or "EXCLUSION" filter depending on the argument passed to CREATE_FILTER_SET when creating the filter set.

Syntax

```
DBMS_WORKLOAD_REPLAY.ADD_FILTER (
   fname           IN VARCHAR2,
   fattribute      IN VARCHAR2,
   fvalue          IN VARCHAR2);

DBMS_WORKLOAD_REPLAY.ADD_FILTER (
   fname           IN VARCHAR2,
   fattribute      IN VARCHAR2,
   fvalue          IN NUMBER);
```

Parameters

Table 216-3 ADD_FILTER Procedure Parameters

Parameter	Description
fname	(Mandatory) Name of the filter. Can be used to delete the filter later if it is not required.
fattribute	(Mandatory) Specifies the attribute on which the filter is defined as one of the following values of type STRING : • USER • MODULE • ACTION • PROGRAM • SERVICE • CONNECTION_STRING
fvalue	(Mandatory) Specifies the value to which the given 'attribute' must be equal to for the filter to be considered active. Wildcards such as '%' are acceptable for all attributes that are of type STRING. Currently all the listed values of fattribute are of type STRING. INSTANCE_NUMBER is a NUMBER attribute. It is currently only supported for capture.

216.3.3 ADD_SCHEDULE_ORDERING Function

This function adds a schedule order between two captures.

Together, schedule_capture_id and waitfor_capture_id form a schedule ordering that previously added by the ADD_SCHEDULE_ORDERING Function. The order is that replay of capture indicated by schedule_capture_id will not start unless the replay of capture indicated by waiting_for_capture_id finishes.

Syntax

```
DBMS_WORKLOAD_REPLAY.ADD_SCHEDULE_ORDERING (
   schedule_capture_id     IN VARCHAR2,
   waitfor_capture_id      IN VARCHAR2)
  RETURN NUMBER;
```

Parameters

Table 216-4 ADD_SCHEDULE_ORDERING Function Parameters

Parameter	Description
schedule_capture_id	Points to a capture that has been added to the current replay schedule. According to the new schedule ordering added by this subprogram, its replay will not start until the replay of another capture specified by waitfor_capture_id runs to completion.
waitfor_capture_id	Points to a capture that has been added to the current replay schedule. According to the new schedule ordering added by this subprogram, the replay of capture specified by schedule_capture_id will not start until the replay of this capture runs to completion.

Return Values

Returns a non-zero error code if the constraint cannot be added

Usage Notes

The two captures must have already been added to the replay schedule.

216.3.4 ASSIGN_GROUP_TO_INSTANCE Procedure

This procedure modifies the view DBA_WORKLOAD_GROUP_ASSIGNMENTS.

Syntax

```
DBMS_WORKLOAD_REPLAY.ASSIGN_GROUP_TO_INSTANCE (
   group_id          IN INTEGER,
   instance_number   IN INTEGER);
```

Parameters

Table 216-5 ASSIGN_GROUP_TO_INSTANCE Procedure Parameters

Parameter	Description
group_id	The identifier of the specified group of capture files
instance_number	The number used for instance registration. It is equivalent to the INSTANCE_NUMBER column in V$INSTANCE.

> **See Also:**
>
> - `DBA_WORKLOAD_GROUP_ASSIGNMENTS` in *Oracle Database Reference*
> - `V$INSTANCE` in *Oracle Database Reference*

216.3.5 BEGIN_REPLAY_SCHEDULE Procedure

This procedure initiates the creation of a reusable replay schedule.

Syntax

```
DBMS_WORKLOAD_REPLAY.BEGIN_REPLAY_SCHEDULE  (
   replay_dir_obj      IN     VARCHAR2,
   schedule_name       IN     VARCHAR2);
```

Parameters

Table 216-6 BEGIN_REPLAY_SCHEDULE Procedure Parameters

Parameter	Description
`replay_dir_obj`	Directory object that points to the replay directory that contains all the capture directories involved in the schedule
`schedule_name`	Name of the schedule to be replayed

Usage Notes

- Only one schedule can be in creation mode at a time. Calling the subprogram again before `end_replay_schedule` will raise an error.
- Prerequisites:
  - The workload capture was already processed using the PROCESS_CAPTURE Procedure in the same database version.
  - The user must have copied the capture directory appropriately.
  - The database is not in replay mode.
  - The SET_REPLAY_DIRECTORY Procedure has already been called.

216.3.6 CALIBRATE Function

This function operates on a processed workload capture directory to estimate the number of hosts and workload replay clients needed to faithfully replay the given workload. This function returns the results as an XML `CLOB`.

Syntax

```
DBMS_WORKLOAD_REPLAY.CALIBRATE (
   capture_dir          IN VARCHAR2,
   process_per_cpu      IN BINARY_INTEGER DEFAULT 4,
   threads_per_process  IN BINARY_INTEGER DEFAULT 50)
  RETURN CLOB;
```

Parameters

Table 216-7 CALIBRATE Function Parameters

Parameter	Description
capture_dir	Name of the directory object that points to the (case sensitive) OS directory that contains processed capture data
process_per_cpu	Maximum number of processes allowed for each CPU (default is 4)
threads_per_process	Maximum number of threads allowed for each process (default is 50)

Return Values

Returns a CLOB formatted as XML that contains:

- Information about the capture
- Current database version
- Input parameters to this function
- Number of CPUs and replay clients needed to replay the given workload
- Information about the sessions captured (total number and maximum concurrency)

Usage Notes

- Prerequisite: The input workload capture was already processed using the PROCESS_CAPTURE Procedure in the same database version.
- This procedure will return the same results as the workload replay client in calibrate mode, which can be run as follows.

```
$ wrc mode=calibrate replaydir=
```

216.3.7 CANCEL_REPLAY Procedure

This procedure cancels workload replay in progress. All the external replay clients (WRC) will automatically be notified to stop issuing the captured workload and exit.

Syntax

```
DBMS_WORKLOAD_REPLAY.CANCEL_REPLAY (
   error_msg    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 216-8 CANCEL_REPLAY Procedure Parameters

Parameter	Description
error_msg	An optional reason for cancelling the replay can be passed which is recorded into DBA_WORKLOAD_REPLAYS.ERROR_MESSAGE.

Usage Notes

Prerequisite: A call to the INITIALIZE_REPLAY Procedure, or PREPARE_REPLAY Procedure, or START_REPLAY Procedure was already issued.

216.3.8 COMPARE_PERIOD_REPORT Procedure

This procedure generates a report comparing a replay to its capture or to another replay of the same capture.

Syntax

```
DBMS_WORKLOAD_REPLAY.COMPARE_PERIOD_REPORT (
   replay_id1   IN   NUMBER,
   replay_id2   IN   NUMBER,
   format       IN   VARCHAR2,
   result       OUT  CLOB );
```

Parameters

Table 216-9 COMPARE_PERIOD_REPORT Procedure Parameters

Parameter	Description
replay_id1	First ID of the workload replay whose report is requested
replay_id2	Second ID of the workload replay whose report is requested. If this is NULL, then the comparison is done with the capture.
format	Specifies the report format. Valid values are DBMS_WORKLOAD_CAPTURE.TYPE_HTML and DBMS_WORKLOAD_CAPTURE.TYPE_XML.
result	Output of the report (CLOB)

216.3.9 COMPARE_SQLSET_REPORT Function

This procedure generates a report comparing a sqlset captured during replay to one captured during workload capture or to one captured during another replay of the same capture.

Syntax

```
DBMS_WORKLOAD_REPLAY.COMPARE_SQLSET_REPORT (
   replay_id1    IN NUMBER,
   replay_id2    IN NUMBER,
   format        IN VARCHAR2,
   r_level       IN VARCHAR2 DEFAULT 'ALL',
   r_sections    IN VARCHAR2 DEFAULT 'ALL',
   result        OUT CLOB )
  RETURN VARCHAR2;
```

Parameters

Table 216-10 COMPARE_SQLSET_REPORT Function Parameters

Parameter	Description
replay_id1	First ID of the workload replay after a change
replay_id2	Second ID of the workload replay before a change. If this is NULL, then the comparison is done with the capture.

Table 216-10 (Cont.) COMPARE_SQLSET_REPORT Function Parameters

Parameter	Description
format	Specifies the report format. Valid values are DBMS_WORKLOAD_CAPTURE.TYPE_HTML, DBMS_WORKLOAD_CAPTURE.TYPE_XML and DBMS_WORKLOAD_CAPTURE.TYPE_TEXT.
r_level	See level parameter in the REPORT_ANALYSIS_TASK Function in the DBMS_SQLPA package
r_sections	See section parameter in the REPORT_ANALYSIS_TASK Function in the DBMS_SQLPA package
result	Output of the report (CLOB)

216.3.10 CREATE_FILTER_SET Procedure

This procedure creates a new filter set for the replays at replay_dir.

It includes all the replay filters that have already been added by the ADD_FILTER Procedure. After the procedure has completed and replay initiated, the newly-created filter set can be used to filter the replay in replay_dir by calling the USE_FILTER_SET Procedure.

Syntax

```
DBMS_WORKLOAD_REPLAY.CREATE_FILTER_SET(
   replay_dir       IN  VARCHAR2,
   filter_set       IN  VARCHAR2,
   default_action   IN  VARCHAR2 DEFAULT 'INCLUDE');
```

Parameters

Table 216-11 CREATE_FILTER_SET Procedure Parameters

Parameter	Description
replay_dir	Object directory of the replay to be filtered
filter_set	Name of the filter set to create (to use in USE_FILTER_SET Procedure)
default_action	Can be either INCLUDE or EXCLUDE. Determines whether, by default, every captured call must be replayed or not. Also determines whether the workload filters specified must be considered as INCLUSION filters or EXCLUSION filters.) If it is INCLUDE, then by default all captured calls are replayed, except for the part of the workload defined by the filters. In this case, all the filters that were specified using the ADD_SCHEDULE_ORDERING Function are treated as EXCLUSION filters, and will determine the workload that will not be replayed. If it is EXCLUDE, then by default no captured call to the database is replayed, except for the part of the workload defined by the filters. In this case, all the filters that were specified using he ADD_SCHEDULE_ORDERING Function are treated as INCLUSION filters, and will determine the workload that is replayed. Default: INCLUDE and all the filters specified are assumed to be EXCLUSION filters

Usage Notes

This operation must be invoked when no replay is initialized, prepared, or in progress.

216.3.11 DELETE_FILTER Procedure

This procedure deletes the named filter.

Syntax

```
DBMS_WORKLOAD_REPLAY.DELETE_FILTER(
   fname    IN  VARCHAR2);
```

Parameters

Table 216-12 DELETE_FILTER Procedure Parameters

Parameter	Description
fname	(Mandatory) Name of the filter that must be deleted

Usage Notes

The DELETE_FILTER Procedure only affects filters that have not been used by any previous capture. Consequently, filters can be deleted only if they have been added using the ADD_FILTER Procedures after any capture has been completed. Filters that have been added using ADD_FILTER before a START_CAPTURE and FINISH_CAPTURE cannot be deleted anymore using this subprogram.

216.3.12 DELETE_REPLAY_INFO Procedure

This procedure deletes the rows in DBA_WORKLOAD_REPLAYS that correspond to the given workload replay ID.

Syntax

```
DBMS_WORKLOAD_REPLAY.DELETE_REPLAY_INFO (
   replay_id    IN  NUMBER);
```

Parameters

Table 216-13 DELETE_REPLAY_INFO Procedure Parameters

Parameter	Description
replay_id	(Mandatory) ID of the workload replay that must be deleted. Corresponds to DBA_WORKLOAD_REPLAYS.ID

216.3.13 END_REPLAY_SCHEDULE Procedure

This procedure wraps up the creation of the current schedule. The schedule is now saved and associated with the replay directory and can be used for a replay.

Syntax

```
DBMS_WORKLOAD_REPLAY.END_REPLAY_SCHEDULE;
```

Usage Notes

The BEGIN_REPLAY_SCHEDULE Procedure must have already been called.

216.3.14 EXPORT_AWR Procedure

This procedure exports the AWR snapshots associated with a stipulated replay ID.

Syntax

```
DBMS_WORKLOAD_REPLAY.EXPORT_AWR (
   replay_id    IN  NUMBER);
```

Parameters

Table 216-14 EXPORT_AWR Function Parameters

Parameter	Description
replay_id	(Mandatory) ID of the replay whose AWR snapshots are to be exported

Usage Notes

- At the end of each replay, the corresponding AWR snapshots are automatically exported. Consequently, there is no need to do this manually after a workload replay is complete, unless the automatic `EXPORT_AWR` invocation failed.
- This procedure will work only if the corresponding workload replay was performed in the current database (meaning that the corresponding row in `DBA_WORKLOAD_REPLAYS` was not created by calling the GET_REPLAY_INFO Function) and the AWR snapshots that correspond to that replay time period are still available.

216.3.15 GENERATE_CAPTURE_SUBSET Procedure

This procedure creates a new capture from an existing workload capture.

Syntax

```
DBMS_WORKLOAD_REPLAY.GENERATE_CAPTURE_SUBSET (
   input_capture_dir          IN   VARCHAR2,
   output_capture_dir         IN   VARCHAR2,
   new_capture_name           IN   VARCHAR2,
   begin_time                 IN   NUMBER,
   begin_include_incomplete   IN   BOOLEAN DEFAULT TRUE,
   end_time                   IN   NUMBER,
   end_include_incomplete     IN   BOOLEAN DEFAULT FALSE,
   parallel_level             IN   NUMBER DEFAULT NULL);
```

Parameters

Table 216-15 GENERATE_CAPTURE_SUBSET Procedure Parameters

Parameter	Description
input_capture_dir	(Mandatory) Name of the directory object that points to an existing workload capture
output_capture_dir	(Mandatory) Name of the directory object that points to the new capture
new_capture_name	(Mandatory) Name of new capture
begin_time	Start of the time range - time offset in seconds from the start of a workload capture
begin_include_incomplete	Column to include incomplete calls caused by begin_time
end_time	End of the time range - time offset in seconds from the start of a workload capture. If end_time is zero or end_time is less or equal than begin_time, the time range is invalid. The new capture will use the whole duration of the input capture.
end_include_incomplete	Column to include incomplete calls caused by end_time
parallel_level	Number of Oracle processes used to process the input captures in a parallel fashion. The NULL default value will auto-compute the parallelism level based on number of CPUs, whereas a value of 1 will enforce serial execution.

216.3.16 GET_DIVERGING_STATEMENT Function

This function retrieves information about a diverging call, including the statement text, the SQL ID, and the binds. If the replay of a recorded user call has data or error divergence, it is a diverging call.

Syntax

```
DBMS_WORKLOAD_REPLAY.GET_DIVERGING_STATEMENT (
   replay_id    IN NUMBER,
   stream_id    IN NUMBER,
   call_counter IN NUMBER)
  RETURN CLOB;
```

Parameters

Table 216-16 GET_DIVERGING_STATEMENT Function Parameters

Parameter	Description
replay_id	ID of the replay in which that call diverged
stream_id	Stream ID of the diverging call
call_counter	Call counter of the diverging call

Usage Notes

- Returns a CLOB formatted as XML that contains:
  - SQL ID

  - SQL Text
  - Bind information: position, name and value
- This function will silently invoke the POPULATE_DIVERGENCE Procedure to read the information from the capture files. Therefore, if divergence has not been populated, then the first call to this function for a particular diverging call might take longer, especially in very large captures.

216.3.17 GET_REPLAY_DIRECTORY Function

This function returns the current replay directory set by the SET_REPLAY_DIRECTORY Procedure. It returns `NULL` if no replay directory has been set.

Syntax

```
DBMS_WORKLOAD_REPLAY.GET_REPLAY_DIRECTORY
   RETURN VARCHAR2;
```

Related Topics

- SET_REPLAY_DIRECTORY Procedure
  This procedure sets a directory that contains multiple workload captures as the current replay directory.

216.3.18 GET_REPLAY_INFO Function

This function retrieves information about the workload capture and the history of all the workload replay attempts from the stipulated directory.

Syntax

```
DBMS_WORKLOAD_REPLAY.GET_REPLAY_INFO (
   replay_dir    IN VARCHAR2,
   load_details  IN BOOLEAN DFAULT FALSE)
 RETURN NUMBER;
```

Parameters

Table 216-17 GET_REPLAY_INFO Function Parameters

Parameter	Description
`replay_dir`	(Mandatory) Name of the workload replay directory object (case sensitive).
`load_details`	Load the divergence and tracked commits data. The default value is `FALSE`.

Return Values

The procedure returns the `CAPTURE_ID`, which can be associated with both `DBA_WORKLOAD_CAPTURES.ID` and `DBA_WORKLOAD_REPLAYS.CAPTURE_ID` to access the imported information.

Usage Notes

- The procedure first imports a row into `DBA_WORKLOAD_CAPTURES` which will contain information about the capture. It then imports a row for every replay attempt retrieved from the given replay directory into `DBA_WORKLOAD_REPLAYS`.
- The procedure will not insert new rows to `DBA_WORKLOAD_CAPTURES` and `DBA_WORKLOAD_REPLAYS` if these views already contain rows describing the capture and replay history present in the given directory.

216.3.19 GET_REPLAY_TIMEOUT Procedure

This procedure gets the replay timeout setting.

Syntax

```
DBMS_WORKLOAD_REPLAY.GET_REPLAY_TIMEOUT (
   enabled       OUT  BOOLEAN,
   min_delay     OUT  NUMBER,
   max_delay     OUT  NUMBER,
   delay_factor  OUT  NUMBER);
```

Parameters

Table 216-18 GET_REPLAY_TIMEOUT Procedure Parameters

Parameter	Description
`enabled`	`TRUE` if the timeout action is enabled, `FALSE` otherwise.
`min_delay`	Lower bound of call delay in minutes. The replay action is activated only when the delay is equal to or more than `min_delay`.
`max_delay`	Upper bound of call delay in minutes. The timeout action throws `ORA-15569` when the delay is more than `max_delay`.
`delay_factor`	Factor for the call delay that is between `min_delay` and `max_delay`. The timeout action throws `ORA-15569` when the current replay elapsed time is more than the product of capture elapsed time and `delay_factor`.

Usage Notes

This procedure can be called anytime during replay.

216.3.20 IMPORT_AWR Function

This procedure imports the AWR snapshots from a given replay.

Syntax

```
DBMS_WORKLOAD_REPLAY.IMPORT_AWR (
   replay_id       IN   NUMBER,
   staging_schema  IN   VARCHAR2)
  RETURN NUMBER;
```

Parameters

Table 216-19 IMPORT_AWR Function Parameters

Parameter	Description
replay_id	(Mandatory) ID of the replay whose AWR snapshots must be imported
staging_schema	(Mandatory) Name of a valid schema in the current database which can be used as a staging area while importing the AWR snapshots from the replay directory to the SYS AWR schema. The SYS schema is not a valid input.

Return Values

Returns the new randomly generated database ID that was used to import the AWR snapshots. The same value can be found in the AWR_DBID column in the DBA_WORKLOAD_REPLAYS view.

Usage Notes

- This procedure will work provided those AWR snapshots were exported earlier from the original replay system using the EXPORT_AWR Procedure.
- IMPORT_AWR will fail if the staging_schema provided as input contains any tables with the same name as any of the AWR tables, such as WRM$_SNAPSHOT or WRH$_PARAMETER. Drop any such tables in the staging_schema before invoking IMPORT_AWR.

216.3.21 INITIALIZE_CONSOLIDATED_REPLAY Procedure

This procedure puts the database state in INIT for a multiple-capture replay.

It uses the replay_dir which has already been defined by the SET_REPLAY_DIRECTORY Procedure, pointing to a directory that contains all the capture directories involved in the schedule. It reads data about schedule schedule_name from the directory, and loads required connection data into the replay system.

Syntax

```
DBMS_WORKLOAD_REPLAY.INITIALIZE_CONSOLIDATED_REPLAY (
   replay_name        IN    VARCHAR2,
   schedule_name      IN    VARCHAR2,
   plsql_mode         IN    VARCHAR2 DEFAULT 'TOP_LEVEL');
```

Parameters

Table 216-20 INITIALIZE_CONSOLIDATED_REPLAY Procedure Parameters

Parameter	Description
replay_name	(Mandatory) Name of the workload replay. Every replay of a processed workload capture can be given a name.
schedule_name	Name of the schedule to be replayed. It must have been created through the BEGIN_REPLAY_SCHEDULE Procedure for the replay directory replay_dir.

Table 216-20 (Cont.) INITIALIZE_CONSOLIDATED_REPLAY Procedure Parameters

Parameter	Description
plsql_mode	Specifies the replay options for PL/SQL calls: • TOP_LEVEL — only top-level PL/SQL calls are replayed • EXTENDED — SQL executed from PL/SQL or top-level SQL PL/SQL if there is no SQL recorded inside the PL/SQL are replayed. All captures must have been done in 'EXTENDED' PL/SQL mode.

Usage Notes

Prerequisites:

- Workload capture was already processed using the PROCESS_CAPTURE Procedure in the same database version.
- Database state has been logically restored to what it was at the beginning of the original workload capture.
- The SET_REPLAY_DIRECTORY Procedure has been called.

216.3.22 INITIALIZE_REPLAY Procedure

This procedure puts the database state in INIT for REPLAY mode, and loads data into the replay system that is required before preparing for the replay (by executing the PAUSE_REPLAY Procedure).

Syntax

```
DBMS_WORKLOAD_REPLAY.INITIALIZE_REPLAY (
   replay_name     IN  VARCHAR2,
   replay_dir      IN  VARCHAR2,
   plsql_mode      IN  VARCHAR2 DEFAULT 'TOP_LEVEL',
   rac_inst_list   IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 216-21 INITIALIZE_REPLAY Procedure Parameters

Parameter	Description
replay_name	(Mandatory) Name of the workload replay. Every replay of a processed workload capture can be given a name.
replay_dir	Name of the directory object that points to the OS directory (case sensitive) that contains processed capture data
plsql_mode	Specifies the replay options for PL/SQL calls: • TOP_LEVEL — only top-level PL/SQL calls are replayed • EXTENDED — SQL executed from PL/SQL or top-level SQL PL/SQL if there is no SQL recorded inside the PL/SQL are replayed. All captures must have been done in 'EXTENDED' PL/SQL mode.

Table 216-21 (Cont.) INITIALIZE_REPLAY Procedure Parameters

Parameter	Description
rac_inst_list	Specifies a list of Oracle Real Application Clusters (Oracle RAC) instances that will be used for replay. The parameter is a string of instance numbers that are separated by commas. For example: `rac_inst_list='1,3,5'`

Usage Notes

- Prerequisites:
  - Workload capture was already processed using the PROCESS_CAPTURE Procedure in the same database version.
  - Database state has been logically restored to what it was at the beginning of the original workload capture.
- The subprogram loads data into the replay system that is required before preparing for the replay by calling the PAUSE_REPLAY Procedure.

  For instance, during capture the user may record the connection string each session used to connect to the server. The INITIALIZE_REPLAY Procedure loads this data and allows the user to re-map the recorded connection string to new connection strings or service points.

  Elaborating on the example described in the PROCESS_CAPTURE Procedure, the user could invoke the following:

  ```
  DBMS_WORKLOAD_REPLAY.INITIALIZE_REPLAY('replay foo #1', 'rec_dir');
  ```

  This command will load up the connection map and by default will set all replay time connection strings to be equal to `NULL`. A `NULL` replay time connection string means that the workload replay clients (WRCs) will connect to the default host as determined by the replay client's runtime environment settings. The user can change a particular connection string to a new one (or a new service point) for replay by using the REMAP_CONNECTION Procedure.
- For encrypted capture, the `INITIALIZE_REPLAY` Procedure relies on Oracle wallet. The identifier is `oracle.rat.database_replay.encryption` (case-sensitive).

Related Topics

- PAUSE_REPLAY Procedure
  This procedure pauses the in-progress workload replay.
- INITIALIZE_REPLAY Procedure
  This procedure puts the database state in `INIT` for `REPLAY` mode, and loads data into the replay system that is required before preparing for the replay (by executing the PAUSE_REPLAY Procedure).
- PROCESS_CAPTURE Procedure
  This procedure processes the workload capture found in `capture_dir` in place.
- REMAP_CONNECTION Procedure
  This procedure remaps the captured connection to a new one so that the user sessions can connect to the database in a desired way during workload replay.

216.3.23 IS_REPLAY_PAUSED Function

This function reports whether the replay is currently paused.

Syntax

```
DBMS_WORKLOAD_REPLAY.IS_REPLAY_PAUSED
   RETURN BOOLEAN;
```

Return Values

Returns `TRUE` if the PAUSE_REPLAY Procedure has been called successfully and the RESUME_REPLAY Procedure has not been called yet.

Usage Notes

A call to the START_REPLAY Procedure must have already been issued as a prerequisite.

216.3.24 LOAD_LONG_SQLTEXT Procedure

This procedure loads the captured SQL statements that are longer than 1000 characters to the `DBA_WORKLOAD_LONG_SQLTEXT` view.

Syntax

```
DBMS_WORKLOAD_REPLAY.LOAD_LONG_SQLTEXT (
   capture_id     IN  NUMBER);
```

Parameters

Table 216-22 LOAD_LONG_SQLTEXT Procedure Parameters

Parameter	Description
capture_id	Internal key for the workload capture

> **Note:**
>
> This procedure is available starting with Oracle Database Release 18c.

> **See Also:**
>
> `DBA_WORKLOAD_LONG_SQLTEXT` in *Oracle Database Reference*

216.3.25 PAUSE_REPLAY Procedure

This procedure pauses the in-progress workload replay.

All subsequent user calls from the replay clients will be stalled until either a call to the RESUME_REPLAY Procedure is issued or the replay is cancelled.

Syntax

```
DBMS_WORKLOAD_REPLAY.PAUSE_REPLAY;
```

Usage Notes

- Prerequisite: A call to the START_REPLAY Procedure must have already been issued.
- User calls that were already in-progress when this procedure was invoked are allowed to run to completion. Only subsequent user calls, when issued, are paused.

216.3.26 POPULATE_DIVERGENCE Procedure

This procedure precomputes the divergence information for the given call, stream, or the whole replay so that the GET_DIVERGING_STATEMENT Function returns as quickly as possible for the precomputed calls.

Syntax

```
DBMS_WORKLOAD_REPLAY.POPULATE_DIVERGENCE (
   replay_id    IN   NUMBER,
   stream_id    IN   NUMBER DEFAULT NULL,
   call_counter IN   NUMBER DEFAULT NULL);
```

Parameters

Table 216-23 POPULATE_DIVERGENCE Procedure Parameters

Parameter	Description
replay_id	ID of the replay
stream_id	Stream ID of the diverging call. If NULL is provided, then divergence information is precomputed for all diverging calls in the given replay.
call_counter	Call counter of the diverging call. If NULL is provided, then divergence information is precomputed for all diverging calls in the given stream.

Related Topics

- GET_DIVERGING_STATEMENT Function
  This function retrieves information about a diverging call, including the statement text, the SQL ID, and the binds. If the replay of a recorded user call has data or error divergence, it is a diverging call.

216.3.27 PREPARE_CONSOLIDATED_REPLAY Procedure

Similar to the PREPARE_REPLAY Procedure, this procedure puts the database in a special "Prepare" mode for a multiple-capture replay. The difference is that this subprogram should be used only for consolidated replays.

Syntax

```
DBMS_WORKLOAD_REPLAY.PREPARE_CONSOLIDATED_REPLAY (
   synchronization          IN BOOLEAN,
   connect_time_scale       IN NUMBER    DEFAULT 100,
   think_time_scale         IN NUMBER    DEFAULT 100,
   think_time_auto_correct  IN BOOLEAN   DEFAULT TRUE,
   capture_sts              IN BOOLEAN   DEFAULT FALSE,
```

```
   sts_cap_interval          IN NUMBER    DEFAULT 300);

DBMS_WORKLOAD_REPLAY.PREPARE_CONSOLIDATED_REPLAY (
   synchronization           IN VARCHAR2  DEFAULT 'SCN',
   connect_time_scale        IN NUMBER    DEFAULT 100,
   think_time_scale          IN NUMBER    DEFAULT 100,
   think_time_auto_correct   IN BOOLEAN   DEFAULT TRUE,
   capture_sts               IN BOOLEAN   DEFAULT FALSE,
   sts_cap_interval          IN NUMBER    DEFAULT 300);
```

Parameters

Table 216-24 PREPARE_CONSOLIDATED_REPLAY Procedure Parameters

Parameter	Description
`synchronization`	Sets the synchronization mode for replay: • `'TIME'` — The synchronization will be based on the time the action took place during capture (clock-based time). • `'SCN'` — The synchronization will be based on the capture-time commits; the commit order will be preserved during replay. This is the default mode. • `'OBJECT_ID'` — Every replayed action will be executed only *after* the relevant commits have finished execution. The relevant commits are those that were issued before the given action in the captured workload and that modified at least one of the database objects the given action is referencing (either implicitly or explicitly). This synchronization mode makes sure that any replay action will see the same data that the action saw during capture, but allows greater concurrency for the actions that do not touch the same objects/tables. This value is deprecated. For compatibility, an overloaded version of this procedure uses BOOLEAN for this parameter: • TRUE means `'SCN'` • FALSE means `'TIME'`
`connect_time_scale`	Scales the time elapsed between the instant the workload capture was started and the session connects with the given value. The input is interpreted as a % value. Can potentially be used to increase or decrease the number of concurrent users during the workload replay. `DEFAULT VALUE` is 100. See "Example 216-1".
`think_time_scale`	Scales the time elapsed between two successive user calls "Example 216-1"from the same session. The input is interpreted as a % value. Can potentially be used to increase or decrease the number of concurrent users during the workload replay. `DEFAULT VALUE` is 100. See "Example 216-2".
`think_time_auto_correct`	Auto corrects the think time between calls appropriately when user calls takes longer to complete during replay than during the original capture. `DEFAULT` is `TRUE` which is to reduce think time if replay goes slower than capture. See "Example 216-3".

Table 216-24 (Cont.) PREPARE_CONSOLIDATED_REPLAY Procedure Parameters

Parameter	Description
capture_sts	If this parameter is TRUE, then a SQL tuning set capture is also started in parallel with workload replay. The resulting SQL tuning set can be exported using the EXPORT_AWR Procedure along with the Automatic Workload Repository (AWR) data. Currently, parallel SQL tuning set (STS) capture is not supported in an Oracle RAC environment. So, this parameter has no effect in that context. The calling user must have the appropriate privileges ('ADMINISTER SQL TUNING SET'). The default value is FALSE.
sts_cap_interval	Specifies the capture interval of the SQL set capture from the cursor cache in seconds. The default value is 300.

Usage Notes

A consolidated replay replays multiple captures in one replay. Each capture records different system change number (SCN) values. For this reason SCN-based sync is not supported for consolidated replays. Consolidated replays only support non-sync mode and the Object-ID based synchronization, and SCN-based synchronization is currently not supported.

Related Topics

- PREPARE_REPLAY Procedure
  This procedure puts the database state in PREPARE FOR REPLAY mode.

216.3.28 PREPARE_REPLAY Procedure

This procedure puts the database state in PREPARE FOR REPLAY mode.

Syntax

```
DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY (
   synchronization           IN BOOLEAN,
   connect_time_scale        IN NUMBER    DEFAULT 100,
   think_time_scale          IN NUMBER    DEFAULT 100,
   think_time_auto_correct   IN BOOLEAN   DEFAULT TRUE,
   scale_up_multiplier       IN NUMBER    DEFAULT 1,
   capture_sts               IN BOOLEAN   DEFAULT FALSE,
   sts_cap_interval          IN NUMBER    DEFAULT 300,
   rac_mode                  IN NUMBER    DEFAULT GLOBAL_SYNC,
   query_only                IN BOOLEAN   DEFAULT FALSE);

DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY (
   synchronization           IN VARCHAR2  DEFAULT 'SCN',
   connect_time_scale        IN NUMBER    DEFAULT 100,
   think_time_scale          IN NUMBER    DEFAULT 100,
   think_time_auto_correct   IN BOOLEAN   DEFAULT TRUE,
   scale_up_multiplier       IN NUMBER    DEFAULT 1,
   capture_sts               IN BOOLEAN   DEFAULT FALSE,
   sts_cap_interval          IN NUMBER    DEFAULT 300),
   rac_mode                  IN NUMBER    DEFAULT GLOBAL_SYNC,
   query_only                IN BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 216-25 PREPARE_REPLAY Procedure Parameters

Parameter	Description
synchronization	Sets the synchronization mode for replay: • 'TIME' — The synchronization will be based on the time the action took place during capture (clock-based time). • 'SCN' — The synchronization will be based on the capture-time commits; the commit order will be preserved during replay. This is the default mode. • 'OBJECT_ID' — Every replayed action will be executed only *after* the relevant commits have finished execution. The relevant commits are those that were issued before the given action in the captured workload and that modified at least one of the database objects the given action is referencing (either implicitly or explicitly). This synchronization mode makes sure that any replay action will see the same data that the action saw during capture, but allows greater concurrency for the actions that do not touch the same objects/tables. This value is deprecated. For compatibility, an overloaded version of this procedure uses BOOLEAN for this parameter: • TRUE means 'SCN' • FALSE means 'TIME'
connect_time_scale	Scales the time elapsed between the instant the workload capture was started and the session connects with the given value. The input is interpreted as a % value. Can potentially be used to increase or decrease the number of concurrent users during the workload replay. DEFAULT VALUE is 100. See "Example 216-1".
think_time_scale	Scales the time elapsed between two successive user calls from the same session. The input is interpreted as a % value. Can potentially be used to increase or decrease the number of concurrent users during the workload replay. DEFAULT VALUE is 100. See "Example 216-2".
think_time_auto_correct	Auto corrects the think time between calls appropriately when a user call takes longer to complete during replay than during the original capture. DEFAULT is TRUE which is to reduce think time if replay goes slower than capture. See "Example 216-3".
scale_up_multiplier	Defines the number of times the query workload is scaled up during replay. Each captured session is replayed concurrently as many times as the value of the scale_up_multiplier. However, only one of the sessions in each set of identical replay sessions executes both queries and updates. The remaining sessions only execute queries.
capture_sts	If this parameter is TRUE, then a SQL tuning set capture is also started in parallel with workload replay. The resulting SQL tuning set can be exported using the EXPORT_AWR Procedure along with the AWR data. Currently, parallel SQL tuning set (STS) capture is not supported in an Oracle RAC environment. So, this parameter has no effect in that context. The calling user must have the appropriate privileges ('ADMINISTER SQL TUNING SET'). The default value is FALSE.
sts_cap_interval	Specifies the capture interval of the SQL set capture from the cursor cache in seconds. The default value is 300.

Table 216-25 (Cont.) PREPARE_REPLAY Procedure Parameters

Parameter	Description
rac_mode	Specifies replay options in an Oracle Real Application Cluster (Oracle RAC) environment. This parameter accepts the following values: • GLOBAL_SYNC — Synchronization across all Oracle RAC instances. This is the default value. • PER_INSTANCE_CLIENT — Synchronization is global and each WRC client handles part of the workload that is sent to only one instance. • PER_INSTANCE_SYNC — Synchronization is local (within each instance only) and each WRC client handles part of the workload that is sent to only one instance
query_only	Replays only the read-only queries of the workload capture. The default value is FALSE.

Usage Notes

- Prerequisites:
  - The database has been initialized for replay using the INITIALIZE_REPLAY Procedure.
  - Any capture time connection strings that require remapping have been already done using the REMAP_CONNECTION Procedure.
- One or more external replay clients (WRC) can be started once the PREPARE_REPLAY procedure has been executed.
- With regard to scale_up_multiplier:
  - One replay session (base session) of each set of identical sessions will replay every call from the capture as usual.
  - The remaining sessions (scale-up sessions) will only replay calls that are read-only. Thus, DDL, DML, and PL/SQL calls that modified the database is skipped. SELECT FOR UPDATE statements are also skipped.
  - Read-only calls from the scale-up are synchronized appropriately and obey the timings defined by think_time_scale, connect_time_scale, and think_time_auto_correct. Also, the queries are made to wait for the appropriate commits.
  - No replay data or error divergence records are generated for the scale-up sessions.
  - All base or scale-up sessions that replay the same capture file will connect from the same workload replay client.

Example 216-1 Application of the connect_time_scale Parameter

If the following was observed during the original workload capture:

```
12:00 : Capture was started
12:10 : First session connect  (10m after)
12:30 : Second session connect (30m after)
12:42 : Third session connect  (42m after)
```

If the connect_time_scale is 50, then the session connects will happen as follows:

```
12:00 : Replay was started with 50% connect time scale
12:05 : First session connect  ( 5m after)
```

```
12:15 : Second session connect (15m after)
12:21 : Third session connect  (21m after)
```

If the connect_time_scale is 200, then the session connects will happen as follows:

```
12:00 : Replay was started with 200% connect time scale
12:20 : First session connect  (20m after)
13:00 : Second session connect (60m after)
13:24 : Third session connect  (84m after)
```

Example 216-2 Application of the think_time_scale Parameter

If the following was observed during the original workload capture:

```
12:00 : User SCOTT connects
12:10 : First user call issued (10m after completion of prevcall)
12:14 : First user call completes in 4mins
12:30 : Second user call issued (16m after completion of prevcall)
12:40 : Second user call completes in 10m
12:42 : Third user call issued ( 2m after completion of prevcall)
12:50 : Third user call completes in 8m
```

If the think_time_scale is 50 during the workload replay, then the user calls will look something like below:

```
12:00 : User SCOTT connects
12:05 : First user call issued 5 mins (50% of 10m) after the completion of
        previous call
12:10 : First user call completes in 5m (takes a minute longer)
12:18 : Second user call issued 8 mins (50% of 16m) after the completion of prev
        call
12:25 : Second user call completes in 7m (takes 3 minutes less)
12:26 : Third user call issued 1 min  (50% of 2m) after the completion of prev
        call
12:35 : Third user call completes in 9m (takes a minute longer)
```

Example 216-3 Application of the think_time_auto_correct Parameter

If the following was observed during the original workload capture:

```
12:00 : User SCOTT connects
12:10 : First user call issued (10m after completion of prevcall)
12:14 : First user call completes in 4m
12:30 : Second user call issued (16m after completion of prevcall)
12:40 : Second user call completes in 10m
12:42 : Third user call issued ( 2m after completion of prevcall)
12:50 : Third user call completes in 8m
```

If the think_time_scale is 100 and the think_time_auto_correct is TRUE during the workload replay, then the user calls will look something like below:

```
12:00 : User SCOTT connects
12:10 : First user call issued 10 mins after the completion of prev call
12:15 : First user call completes in 5m (takes 1 minute longer)
12:30 : Second user call issued 15 mins (16m minus the extra time of 1m the prev
        call took) after the completion of prev call
12:44 : Second user call completes in 14m (takes 4 minutes longer)
12:44 : Third user call issued immediately (2m minus the extra time of 4m the prev
        call took) after the completion of prev call
12:52 : Third user call completes in 8m
```

216.3.29 PROCESS_CAPTURE Procedure

This procedure processes the workload capture found in `capture_dir` in place.

Syntax

```
DBMS_WORKLOAD_REPLAY.PROCESS_CAPTURE (
   capture_dir            IN   VARCHAR2,
   parallel_level         IN   NUMBER DEFAULT NULL,
   synchronization        IN   VARCHAR2 DEFAULT 'SCN',
   plsql_mode             IN   VARCHAR2 DEFAULT 'TOP_LEVEL');
```

Parameters

Table 216-26 PROCESS_CAPTURE Procedure Parameters

Parameter	Description
`capture_dir`	(Mandatory) Name of the workload capture directory object (case sensitive). The directory object must point to a valid OS directory that has the appropriate permissions. New files are added to this directory.
`parallel_level`	Number of Oracle processes used to process the capture in parallel. The `NULL` default value will auto-compute the parallelism level, whereas a value of 1 will enforce serial execution.
`synchronization`	Determines the synchronization mode that the user will be able to use for replay: • `'TIME'` — When `'TIME'` is selected, the replay can use `'TIME'` synchronization mode only. When `'TIME'` synchronization mode is used for replay, the synchronization will be based on the time the action took place during capture (clock-based time). • `'SCN'` — When `'SCN'` is selected, the replay can use the `'TIME'` or `'SCN'` synchronization mode. This is the default. When `'SCN'` synchronization mode is used for replay, the synchronization will be based on the capture-time commits; the commit order will be preserved during replay. This is the default mode. • `'OBJECT_ID'` — When `'OBJECT_ID'` is selected, replay can use the `'TIME'`, `'SCN'`, or `'OBJECT_ID'` synchronization mode. When `'OBJECT_ID'` synchronization mode is used for replay, every replayed action will be executed only *after* the relevant commits have finished execution. The relevant commits are those that were issued before the given action in the captured workload and that modified at least one of the database objects the given action is referencing (either implicitly or explicitly). This synchronization mode makes sure that any replay action will see the same data that the action saw during capture, but allows greater concurrency for the actions that do not touch the same objects/tables. This synchronization mode is deprecated.

Table 216-26 (Cont.) PROCESS_CAPTURE Procedure Parameters

Parameter	Description
plsql_mode	Specifies the processing mode for PL/SQL: • 'TOP_LEVEL' — metadata is generated for top-level PL/SQL calls only; 'TOP_LEVEL' will be the only option for replay. • 'EXTENDED' — metadata is generated for both top-level PL/SQL calls and the SQL called from PL/SQL. A new directory ppe_X.X.X.X (where the Xs represent the current Oracle version) is created under the capture root directory. Capture must have been done with this same value for the plsql_mode parameter. Replay can use either 'TOP_LEVEL' or 'EXTENDED'.

Usage Notes

- This subprogram analyzes the workload capture found in the capture_dir and creates new workload replay specific metadata files that are required to replay the given workload capture. It only creates new files and does not modify any files that were originally created during the workload capture. Therefore, this procedure can be run multiple times on the same capture directory, such as when the procedure encounters unexpected errors or is cancelled by the user.
- Once this procedure runs successfully, the capture_dir can be used as input to the INITIALIZE_REPLAY Procedure in order to replay the captured workload present in capture_dir.
- Before a workload capture can be replayed in a particular database version, the capture must be processed using PROCESS_CAPTURE in the same database version. Once created, a processed workload capture can be used to replay the captured workload multiple times in the same database version.

  For example, suppose workload "foo" was captured in rec_dir in Oracle database version 10.2.0.5. In order to replay the workload "foo" in version 11.1.0.1 the workload must be processed in version 11.1.0.1. The following procedure must be executed in an 11.1.0.1 database in order to process the capture directory rec_dir:

  ```
  DBMS_WORKLOAD_REPLAY.PROCESS_CAPTURE('rec_dir');
  ```

  Now, rec_dir contains a valid 11.1.0.1 processed workload capture that can be used to replay the workload "foo" in 11.1.0.1 databases as many times as required.
- For encrypted capture, the PROCESS_CAPTURE procedure relies on Oracle wallet. The identifier is oracle.rat.database_replay.encryption (case-sensitive).

216.3.30 REMAP_CONNECTION Procedure

This procedure remaps the captured connection to a new one so that the user sessions can connect to the database in a desired way during workload replay.

Syntax

```
DBMS_WORKLOAD_REPLAY.REMAP_CONNECTION (
   connection_id         IN  NUMBER,
   replay_connection     IN  VARCHAR2);

DBMS_WORKLOAD_REPLAY.REMAP_CONNECTION (
   capture_number        IN  VARCHAR2,
```

```
   connection_id         IN  NUMBER,
   replay_connection     IN  VARCHAR2);
```

Parameters

Table 216-27 REMAP_CONNECTION Procedure Parameters

Parameter	Description
capture_number	Pointing to a capture of the current replay schedule
connection_id	ID of the connection to be remapped. Corresponds to DBA_WORKLOAD_CONNECTION_MAP.CONN_ID.
replay_connection	New connection string to be used during replay

Usage Notes

- Prior to calling REMAP_CONNECTION all replay connection strings are set to NULL by default. If a replay_connection is NULL, then the replay sessions will connect as determined by the replay client's runtime environment. For example, if the environment variable TNS_ADMIN is defined and the user does not call the REMAP_CONNECTION Procedure, then the wrc executable will connect to the server specified in the tnsnames.ora file pointed to by TNS_ADMIN.
- A valid replay_connection must specify a connect identifier or a service point. See the *Oracle Database Net Services Reference* for ways to specify connect identifiers (such as net service names, database service names, and net service aliases) and naming methods that can be used to resolve a connect identifier to a connect descriptor.
- An error is returned if no row matches the given connection_id.
- Use the DBA_WORKLOAD_CONNECTION_MAP view to review all the connection strings that are used by the subsequent workload replay, and also to examine connection string remappings used for previous workload replays.

216.3.31 REMOVE_CAPTURE Procedure

This procedure removes the given capture from the current schedule.

Syntax

```
DBMS_WORKLOAD_REPLAY.REMOVE_CAPTURE (
   schedule_capture_number    IN     NUMBER);
```

Parameters

Table 216-28 REMOVE_CAPTURE Procedure Parameters

Parameter	Description
schedule_capture_number	Unique ID that identifies this capture within this schedule

216.3.32 REMOVE_SCHEDULE_ORDERING Procedure

This procedure removes an existing schedule order from the current replay schedule.

Together, `schedule_capture_id` and `waitfor_capture_id` form a schedule ordering that previously added by the ADD_SCHEDULE_ORDERING Function (`schedule_capture_id`, `waitfor_capture_id`). The order is that replay of capture indicated by `schedule_capture_id` will not start unless the replay of capture indicated by `waiting_for_capture_id` finishes.

Syntax

```
DBMS_WORKLOAD_REPLAY.REMOVE_SCHEDULE_ORDERING (
   schedule_capture_id     IN         NUMBER,
   waitfor_capture_id      IN         NUMBER);
```

Parameters

Table 216-29 REMOVE_SCHEDULE_ORDERING Procedure Parameters

Parameter	Description
`schedule_capture_id`	Points to a capture that has been added to the current replay schedule (see procedure description).
`waitfor_capture_id`	Points to a capture that has been added to the current replay schedule.

Usage Notes

Prerequisites:

- The BEGIN_REPLAY_SCHEDULE Procedure must have been called.
- The replay schedule order should have already been added using the ADD_SCHEDULE_ORDERING Function.

216.3.33 REPORT Function

This function generates a report on the stipulated workload replay.

Syntax

```
DBMS_WORKLOAD_REPLAY.REPORT (
   replay_id          IN NUMBER,
   format             IN VARCHAR2)
  RETURN CLOB;
```

Parameters

Table 216-30 REPORT Function Parameters

Parameter	Description
`replay_id`	(Mandatory) Specifies the ID of the workload replay whose report is requested.

Table 216-30 (Cont.) REPORT Function Parameters

Parameter	Description
format	(Mandatory) Specifies the report format. Valid values: • HTML - Generates the HTML version of the report • XML - Generates the XML version of the report • TEXT - Generates the text version of the report

Return Values

The report body in the desired format returned as a CLOB

216.3.34 RESUME_REPLAY Procedure

This procedure resumes a paused workload replay.

Syntax

```
DBMS_WORKLOAD_REPLAY.RESUME_REPLAY;
```

Usage Notes

Prerequisite: A call to the PAUSE_REPLAY Procedure must have already been issued.

216.3.35 REUSE_REPLAY_FILTER_SET Procedure

This procedure reuses filters in the specified filter set as if each were added using the ADD_SCHEDULE_ORDERING Function.

Each call adds one filter set, which is a collection of individual filters on various attributes. Also, a new filter rule can be added, and an existing filter can be deleted before invoking the CREATE_FILTER_SET Procedure to create a new filter set.

Syntax

```
DBMS_WORKLOAD_REPLAY.REUSE_REPLAY_FILTER_SET(
   replay_dir  IN VARCHAR2,
   filter_set  IN VARCHAR2);
```

Parameters

Table 216-31 REUSE_REPLAY_FILTER_SET Procedure Parameters

Parameter	Description
replay_dir	Capture ID of the existing filter set with which it is associated
filter_set	Name of the filter set to be reused

Related Topics

- ADD_SCHEDULE_ORDERING Function
  This function adds a schedule order between two captures.

216.3.36 SET_ADVANCED_PARAMETER Procedure

This procedure sets an advanced parameter for replay besides the ones used with the PREPARE_REPLAY Procedure.

The advanced parameters control aspects of the replay that are more specialized. The advanced parameters are reset to their default values after the replay has finished.

Syntax

```
DBMS_WORKLOAD_REPLAY.SET_ADVANCED_PARAMETER(
   pname    IN   VARCHAR2,
   pvalue   IN   VARCHAR2);

DBMS_WORKLOAD_REPLAY.SET_ADVANCED_PARAMETER(
   pname    IN   VARCHAR2,
   pvalue   IN   NUMBER);

DBMS_WORKLOAD_REPLAY.SET_ADVANCED_PARAMETER(
   pname    IN   VARCHAR2,
   pvalue   IN   BOOLEAN);
```

Parameters

Table 216-32 SET_ADVANCED_PARAMETER Procedure Parameters

Parameter	Description
pname	Name of the parameter (case insensitive)
pvalue	Value of the parameter

Usage Notes

The current parameters and values that can be used are:

```
'DO_NO_WAIT_COMMITS': (default: FALSE)
```

This parameter controls whether the `COMMIT` issued by replay sessions is `NOWAIT`. The default value for this parameter is `FALSE`. In this case all the `COMMIT`s are issued with the mode they were captured (`wait`, `no-wait`, `batch`, `no-batch`). If the parameter is set to `TRUE`, then all `COMMIT`s are issued in `no-wait` mode. This is useful in cases where the replay is becoming noticeably slow because of a high volume of concurrent `COMMIT`s. Setting the parameter to `TRUE` will significantly decrease the waits on the `'log file sync'` event during the replay with respect to capture.

Related Topics

- PREPARE_REPLAY Procedure
  This procedure puts the database state in `PREPARE FOR REPLAY` mode.

216.3.37 SET_REPLAY_DIRECTORY Procedure

This procedure sets a directory that contains multiple workload captures as the current replay directory.

Syntax

```
DBMS_WORKLOAD_REPLAY.SET_REPLAY_DIRECTORY (
   replay_dir    IN    VARCHAR2);
```

Parameters

Table 216-33 SET_REPLAY_DIRECTORY Procedure Parameters

Parameter	Description
replay_dir	Name of the OS directory containing the captures for a workload consolidation

216.3.38 SET_REPLAY_TIMEOUT Procedure

This procedure sets the replay timeout setting. The purpose is to abort user calls that might make the replay much slower or even cause a replay hang.

Syntax

```
DBMS_WORKLOAD_REPLAY.SET_REPLAY_TIMEOUT (
   enabled       OUT  BOOLEAN DEFAULT TRUE,
   min_delay     OUT  NUMBER DEFAULT 10,
   max_delay     OUT  NUMBER DEFAULT 120,
   delay_factor  OUT  NUMBER DEFAULT 8);
```

Parameters

Table 216-34 SET_REPLAY_TIMEOUT Procedure Parameters

Parameter	Description
enabled	TRUE to enable the timeout action, and FALSE to disable.
min_delay	Lower bound of call delay in minutes. The replay action is activated only when the delay is equal to or more than min_delay. Default = 10.
max_delay	Upper bound of call delay in minutes. The timeout action throws ORA-15569 when the delay is more than max_delay. Default = 120.
delay_factor	Factor for the call delay that is between min_delay and max_delay. The timeout action throws ORA-15569 when the current replay elapsed time is more than the product of capture elapsed time and delay_factor. Default = 8.

Usage Notes

- This procedure can be called anytime during replay.
- Call delay is defined as the difference between replay and capture if replay elapsed time is longer than call elapsed time.

- Once a replay timeout action is enabled, a user call will exit with `ORA-15569` if it has been delayed more than the condition specified by the replay action. The call and its error are reported as error divergence.
- Replay timeout operates as follows:
  - The timeout action has no effect if it is not enabled.
  - If the call delay in minutes is less than a lower bound specified by parameter `min_delay`, then the timeout action is non-operational.
  - If the delay in minutes is more than a upper bound specified by parameter `max_delay`, the timeout action will abort the user call and throw `ORA-15569`.
  - For delay that is between the lower bound and upper bound, the user call will abort with `ORA-15569` only when the current replay elapsed time is more than the product of capture elapsed time and parameter `delay_factor`.

216.3.39 SET_SQL_MAPPING Procedure

This procedure specifies SQL statements to be skipped or replaced during a database replay operation.

Syntax

```
PROCEDURE SET_SQL_MAPPING (
   schedule_cap_id      IN NUMBER,
   sql_id               IN VARCHAR2,
   operation            IN VARCHAR2,
   replacement_sql_text IN VARCHAR2 DEFAULT NULL);

PROCEDURE SET_SQL_MAPPING (
   sql_id               IN VARCHAR2,
   operation            IN VARCHAR2,
   replacement_sql_text IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 216-35 SET_SQL_MAPPING Procedure Parameters

Parameter	Description
`schedule_cap_id`	ID of a capture in the schedule
`sql_id`	SQL identifier of the SQL statement at the time of capture
`operation`	Directs that one of the following actions be performed for the specified statement during database replay: • `'SKIP'` — Skip the SQL statement identified by `sql_id` • `'REPLACE'` — Replace this SQL statement identified by `sql_id` with the SQL statement in the `replacement_sql_text` parameter

Usage Notes

- `replacement_sql_text`: When `'SKIP'` is specified for the `operation` parameter, this parameter is NULL. When `'REPLACE'` is specified for the `operation` parameter, this parameter's value is the SQL statement to be used.
- `schedule_cap_id` is used for consolidated replay.

216.3.40 SET_USER_MAPPING Procedure

This procedure sets a new schema or user name to be used during replay instead of the captured user.

Syntax

```
DBMS_WORKLOAD_REPLAY.SET_USER_MAPPING (
   schedule_cap_id      IN NUMBER,
   capture_user         IN VARCHAR2,
   replay_user          IN VARCHAR2);

DBMS_WORKLOAD_REPLAY.SET_USER_MAPPING (
   capture_user         IN VARCHAR2,
   replay_user          IN VARCHAR2);
```

Parameters

Table 216-36 SET_USER_MAPPING Procedure Parameters

Parameter	Description
schedule_cap_id	ID of the a capture in the schedule
capture_user	User name during the time of the workload capture
replay_user	User name to which captured user is remapped during replay.

Usage Notes

- A schedule_cap_id of NULL is used for regular non-consolidate replay.
- The replay must be initialized but not prepared in order to use this subprogram.
- If replay_user is set to NULL, then the mapping is disabled.
- After multiple calls with the same capture_user, the last call always takes effect.
- To list all the mappings that will be in effect during the subsequent replay execute the following:

  ```
  SELECT * FROM DBA_WORKLOAD_ACTIVE_USER_MAP
  ```

- The overloaded version without the schedule_cap_id calls the one with the schedule_cap_id argument by passing in NULL.
- Mappings are stored in a table made public through the view DBA_WORKLOAD_USER_MAP. To remove old mappings execute

  ```
  DELETE * FROM DBA_WORKLOAD_USER_MAP
  ```

216.3.41 START_CONSOLIDATED_REPLAY Procedure

This procedure starts the replay of a multiple-capture capture. It should be used only for consolidated replays.

Syntax

```
DBMS_WORKLOAD_REPLAY.START_CONSOLIDATED_REPLAY;
```

Usage Notes

Prerequisites:

- The call to the PREPARE_REPLAY Procedure was already issued.
- A sufficient number of external replay clients (WRC) that can faithfully replay the captured workload already started. The status of such external replay clients can be monitored using `V$WORKLOAD_REPLAY_CLIENTS`.

216.3.42 START_REPLAY Procedure

This procedure starts the workload replay.

All the external replay clients (WRC) that are currently connected to the replay database will automatically be notified, and those replay clients (WRC) will begin issuing the captured workload. It should only be used for consolidated replays.

Syntax

```
DBMS_WORKLOAD_REPLAY.START_REPLAY;
```

Usage Notes

- Prerequisites:
  - The call to the PREPARE_REPLAY Procedure was already issued.
  - A sufficient number of external replay clients (WRC) that can faithfully replay the captured workload already started. The status of such external replay clients can be monitored using `V$WORKLOAD_REPLAY_CLIENTS`.
- Use the WRC's `CALIBRATE` mode to determine the number of replay clients that might be required to faithfully replay the captured workload. For example:

  ```
  $ wrc mode=calibrate replaydir=.
  ```

216.3.43 USE_FILTER_SET Procedure

This procedure applies a filter set to a capture in the current replay schedule.

The filter set must have been created by calling the CREATE_FILTER_SET Procedure.

Syntax

```
DBMS_WORKLOAD_REPLAY.USE_FILTER_SET(
   capture_number   IN    VARCHAR2,
   filter_set       IN    VARCHAR2);

DBMS_WORKLOAD_REPLAY.USE_FILTER_SET(
   filter_set       IN    VARCHAR2);
```

Parameters

Table 216-37 USE_FILTER_SET Procedure Parameters

Parameter	Description
`capture_number`	Pointing to a capture of the current replay schedule

Table 216-37 (Cont.) USE_FILTER_SET Procedure Parameters

Parameter	Description
`filter_set`	Name of the filter set

Usage Notes

The filter set must have been created by calling the CREATE_FILTER_SET Procedure.

217
DBMS_WORKLOAD_REPOSITORY

The `DBMS_WORKLOAD_REPOSITORY` package lets you manage the Automatic Workload Repository (AWR) by performing operations, such as, managing snapshots and baselines.

The chapter contains the following topics:

- Examples
- Data Structures
- Summary of DBMS_WORKLOAD_REPOSITORY Subprograms

> **See Also:**
>
> *Oracle Database Performance Tuning Guide* for more information about the Automatic Workload Repository

217.1 DBMS_WORKLOAD_REPOSITORY Examples

This example shows how to generate an AWR text report with the `DBMS_WORKLOAD_REPOSITORY` package for database ID 1557521192, instance ID 1, snapshot IDs 5390 and 5391, and with default options.

```
-- make sure to set line size appropriately
-- set linesize 152
SELECT output FROM TABLE(
   DBMS_WORKLOAD_REPOSITORY.AWR_REPORT_TEXT(
     1557521192,  1, 5390, 5392) ) ;
```

You can call the `DBMS_WORKLOAD_REPOSITORY` packaged functions directly as in the example, but Oracle recommends you use the corresponding supplied SQL script (`awrrpt.sql` in this case) for the packaged function, which prompts the user for required information.

217.2 DBMS_WORKLOAD_REPOSITORY Data Structures

The `DBMS_WORKLOAD_REPOSITORY` package defines an object and associated table types.

OBJECT Types

- AWR_BASELINE_METRIC_TYPE Object Type

TABLE Types

- AWR_BASELINE_METRIC_TYPE_TABLE Table Type
- AWRRPT_INSTANCE_LIST_TYPE Table Type

217.2.1 DBMS_WORKLOAD_REPOSITORY AWR_BASELINE_METRIC_TYPE Object Type

This type shows the values of the metrics corresponding to a baseline.

Syntax

```
TYPE awr_baseline_metric_type AS OBJECT (
   baseline_name      VARCHAR2(64),
   dbid               NUMBER NOT NULL,
   instance_number    NUMBER NOT NULL,
   beg_time           DATE NOT NULL,
   end_time           DATE NOT NULL,
   metric_id          NUMBER NOT NULL,
   metric_name        VARCHAR2(64) NOT NULL,
   metric_unit        VARCHAR2(64) NOT NULL,
   num_interval       NUMBER NOT NULL,
   interval_size      NUMBER NOT NULL,
   average            NUMBER NOT NULL,
   minimum            NUMBER NOT NULL,
   maximum            NUMBER NOT NULL);
```

Fields

Table 217-1 AWR_BASELINE_METRIC_TYPE Fields

Field	Description
baseline_name	Name of the Baseline
dbid	Database ID for the snapshot
instance_number	Instance number for the snapshot
beg_time	Begin time of the interval
end_time	End time of the interval
metric_id	Metric ID
metric_name	Metric name
metric_unit	Unit of measurement
num_interval	Number of intervals observed
interval_size	Interval size (in hundredths of a second)
average	Average over the period
minimum	Minimum value observed
maximum	Maximum value observed

217.2.2 DBMS_WORKLOAD_REPOSITORY AWR_BASELINE_METRIC_TYPE_TABLE Table Type

This type is used by the SELECT_BASELINE_METRIC Function.

Syntax

```
CREATE TYPE awr_baseline_metric_type_table AS TABLE OF awr_baseline_metric_type;
```

Related Topics

- SELECT_BASELINE_METRIC Function
  This table function shows the values of the metrics corresponding to a baseline for all the snapshots.

217.2.3 DBMS_WORKLOAD_REPOSITORY AWRRPT_INSTANCE_LIST_TYPE Table Type

This type provides an alternative to a comma-separated list.

Syntax

```
CREATE TYPE awrrpt_instance_list_type AS TABLE OF NUMBER;
```

217.3 Summary of DBMS_WORKLOAD_REPOSITORY Subprograms

This table lists the `DBMS_WORKLOAD_REPOSITORY` subprograms and briefly describes them.

Table 217-2 DBMS_WORKLOAD_REPOSITORY Package Subprograms

Subprogram	Description
ADD_COLORED_SQL Procedure	Adds a colored SQL ID
ASH_GLOBAL_REPORT_HTML Function	Displays a global or Oracle Real Application Clusters (Oracle RAC) ASH Spot report in HTML format.
ASH_GLOBAL_REPORT_TEXT Function	Displays a global or Oracle Real Application Clusters (Oracle RAC) ASH Spot report in Text format.
ASH_REPORT_ANALYTICS Function	Displays the ASH Analytics active report
ASH_REPORT_HTML Function	Displays the ASH report in HTML
ASH_REPORT_TEXT Function	Displays the ASH report in text
AWR_DIFF_REPORT_HTML Function	Displays the AWR Diff-Diff report in HTML
AWR_DIFF_REPORT_TEXT Function	Displays the AWR Diff-Diff report in text
AWR_GLOBAL_DIFF_REPORT_HTML Functions	Displays the Global AWR Compare Periods Report in HTML
AWR_GLOBAL_DIFF_REPORT_TEXT Functions	Displays the Global AWR Compare Periods Report in text
AWR_EXP Procedure	Extracts AWR data from the AWR schema into a dump file.
AWR_GLOBAL_REPORT_HTML Functions	Displays the Global AWR report in HTML
AWR_GLOBAL_REPORT_TEXT Functions	Displays the Global AWR report in text
AWR_IMP Procedure	Loads the AWR data from a dump file into the `SYS` schema.
AWR_REPORT_HTML Function	Displays the AWR report in HTML
AWR_REPORT_TEXT Function	Displays the AWR report in text

Table 217-2 (Cont.) DBMS_WORKLOAD_REPOSITORY Package Subprograms

Subprogram	Description
AWR_SET_REPORT_THRESHOLDS Procedure	Configures specified report thresholds, including the number of rows in the report
AWR_SQL_REPORT_HTML Function	Displays the AWR SQL Report in HTML format
AWR_SQL_REPORT_TEXT Function	Displays the AWR SQL Report in text format
CONTROL_RESTRICTED_SNAPSHOT Procedure	Controls the AWR snapshot creation for a database in the *restricted session* mode.
CREATE_BASELINE Functions & Procedures	Creates a single baseline
CREATE_BASELINE_TEMPLATE Procedures	Creates a baseline template
CREATE_REMOTE_SNAPSHOT Function and Procedure	Creates a manual *remote* snapshot immediately using the Remote Management Framework (RMF)
CREATE_SNAPSHOT Function and Procedure	Creates a manual *local* snapshot immediately
DROP_BASELINE Procedure	Drops a previously-defined baseline
DROP_BASELINE_TEMPLATE Procedure	Removes a baseline template that is no longer needed
DROP_SNAPSHOT_RANGE Procedure	Drops a range of snapshots
LOCAL_AWR_DBID Function	Returns the database identifier for the local AWR database
MODIFY_BASELINE_WINDOW_SIZE Procedure	Modifies the window size for the Default Moving Window Baseline
MODIFY_SNAPSHOT_SETTINGS Procedures	Modifies the snapshot settings
PURGE_SQL_DETAILS Procedure	Purges SQL details, specifically rows from `WRH$_SQLTEXT` and `WRH$_SQL_PLAN` that do not have corresponding rows (DBID, SQL_ID) in `WRH$_SQLSTAT`.
REGISTER_REMOTE_DATABASE Procedures	Registers a remote database in the AWR using the Remote Management Framework (RMF)
REMOVE_COLORED_SQL Procedure	Removes a colored SQL ID
RENAME_BASELINE Procedure	Renames a baseline
SELECT_BASELINE_DETAILS Function	Shows the values of the metrics corresponding to a baseline for a range of snapshots
SELECT_BASELINE_METRIC Function	Shows the values of the metrics corresponding to a baseline for all the snapshots
UNREGISTER_REMOTE_DATABASE Procedures	Removes all the statistics, metadata, and partitions of a remote database from the AWR using the Remote Management Framework (RMF)
UPDATE_DATAFILE_INFO Procedure	Updates the data file and tablespace information stored in the AWR with the current information in the database
UPDATE_OBJECT_INFO Procedure	Updates rows of `WRH$_SEG_STAT_OBJ` table that represent objects in the local database

217.3.1 ADD_COLORED_SQL Procedure

This procedure adds a colored SQL ID.

If an SQL ID is colored, it will be captured in every snapshot, independent of its level of activities (so that it does not have to be a `TOP SQL`). Capture occurs if the SQL is found in the cursor cache at snapshot time.To uncolor the SQL, invoke the REMOVE_COLORED_SQL Procedure.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.ADD_COLORED_SQL(
   sql_id          IN VARCHAR2,
   dbid            IN NUMBER DEFAULT NULL);
```

Parameters

Table 217-3 *ADD_COLORED_SQL Procedure Parameters*

Parameter	Description
`sql_id`	13-character external SQL ID
`dbid`	Optional DBID, defaults to Local `DBID`

217.3.2 ASH_GLOBAL_REPORT_HTML Function

This table function displays a global or Oracle Real Application Clusters (Oracle RAC) ASH Spot report in HTML format.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.ASH_GLOBAL_REPORT_HTML(
   l_dbid           IN NUMBER,
   l_inst_num       IN VARCHAR2((1023),
   l_btime          IN DATE,
   l_etime          IN DATE,
   l_options        IN NUMBER    DEFAULT 0,
   l_slot_width     IN NUMBER    DEFAULT 0,
   l_sid            IN NUMBER    DEFAULT NULL,
   l_sql_id         IN VARCHAR2  DEFAULT NULL,
   l_wait_class     IN VARCHAR2  DEFAULT NULL,
   l_service_hash   IN NUMBER    DEFAULT NULL,
   l_module         IN VARCHAR2  DEFAULT NULL,
   l_action         IN VARCHAR2  DEFAULT NULL,
   l_client_id      IN VARCHAR2  DEFAULT NULL,
   l_plsql_entry    IN VARCHAR2  DEFAULT NULL,
   l_data_src       IN NUMBER    DEFAULT 0,
   l_container      IN VARCHAR2  DEFAULT NULL)
 RETURN awrrpt_html_type_table PIPELINED;
```

Parameters

Table 217-4 *ASH_GLOBAL_REPORT_HTML Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	List of instances (such as '1,2,3'), or NULL to report on all instances in the database
l_btime	The 'begin time'
l_etime	The 'end time'
l_options	Report level (currently not used)
l_slot_width	Specifies (in seconds) how wide the slots used in the "Top Activity" section of the report should be. This argument is optional, and if it is not specified the time interval between l_btime and l_etime is appropriately split into not more than 10 slots.
l_sid	Session ID (see Usage Notes)
l_sql_id	SQL ID (see Usage Notes)
l_wait_class	Wait class name (see Usage Notes)
l_service_hash	Service name hash (see Usage Notes)
l_module	Module name (see Usage Notes)
l_action	Action name (see Usage Notes)
l_client_id	Client ID for end-to-end backtracing (see Usage Notes)
l_plsql_entry	PL/SQL entry point (see Usage Notes)
l_data_src	Ignored since the report works off of data on disk only
l_container	Name of the container for which report activity is limited. Valid values other than NULL (default) should be taken from container names in V$CONTAINERS. Behavior is as follows: • If NULL: When connected to a root container the report is on all containers. When connected to a PDB the report is on only that PDB. • If not NULL: When connected to a root container the report is on activity from the specified container. When connected to a PDB the report is the same as NULL value for l_container regarding the connected PDB. Note: If while connected to a PDB you request information from another PDB this produces an empty report.

Return Values

The output will be one column of VARCHAR2(1500).

Usage Notes

- You can call the function directly but Oracle recommends you use the ashrpti.sql script which prompts users for the required information.
- The unspecified optional arguments are used to generate an ASH Reports that specify 'report targets' such as a SQL statement, or a session, or a particular Service/Module combination. These arguments are specified to restrict the ASH rows that would be used to generate the report. For example, to generate an ASH report on a particular SQL

statement, such as SQL_ID 'abcdefghij123 ' pass that sql_id value to the l_sql_id argument:

```
l_sql_id =>  'abcdefghij123'
```

Any combination of those optional arguments can be passed in, and only rows in ASH that satisfy all of those 'report targets' will be used. If multiple 'report targets' are specified, AND conditional logic is used to connect them. For example, to generate an ASH report on MODULE "PAYROLL" and ACTION "PROCESS", use the following predicate:

```
l_module =>  'PAYROLL', l_action =>  'PROCESS'
```

Valid SQL wildcards can be used in all the arguments that are of type VARCHAR2.

Table 217-5 ASH_REPORT_HTML: Wildcards Allowed (or Not) in Arguments

Argument Name	Comment	Wildcard Allowed
l_sid	Session ID (for example, V$SESSION.SID)	No
l_sql_id	SQL ID (for example, V$SQL.SQL_ID)	Yes
l_wait_class	Wait class name (for example, V$EVENT_NAME.WAIT_CLASS)	Yes
l_service_hash	Service name hash (for example, V$ACTIVE_SERVICES.NAME_HASH)	No
l_module	Module name (for example, V$SESSION.MODULE)	Yes
l_action	Action name (for example, V$SESSION.ACTION)	Yes
l_client_id	Client ID for end-to-end backtracing (for example, V$SESSION.CLIENT_IDENTIFIER)	Yes
l_data_src	Wildcards are not allowed for l_data_src as it is of numeric datatype	No

217.3.3 ASH_GLOBAL_REPORT_TEXT Function

This table function displays a global or Oracle Real Application Clusters (Oracle RAC) ASH Spot report in text format.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.ASH_GLOBAL_REPORT_TEXT(
   l_dbid          IN VARCHAR2(1023),
   l_inst_num      IN NUMBER,
   l_btime         IN DATE,
   l_etime         IN DATE,
   l_options       IN NUMBER    DEFAULT 0,
   l_slot_width    IN NUMBER    DEFAULT 0,
   l_sid           IN NUMBER    DEFAULT NULL,
   l_sql_id        IN VARCHAR2  DEFAULT NULL,
   l_wait_class    IN VARCHAR2  DEFAULT NULL,
   l_service_hash  IN NUMBER    DEFAULT NULL,
   l_module        IN VARCHAR2  DEFAULT NULL,
   l_action        IN VARCHAR2  DEFAULT NULL,
   l_client_id     IN VARCHAR2  DEFAULT NULL,
   l_plsql_entry   IN VARCHAR2  DEFAULT NULL,
   l_data_src      IN NUMBER    DEFAULT 0,
   l_container     IN VARCHAR2  DEFAULT NULL)
 RETURN awrrpt_text_type_table PIPELINED;
```

Parameters

Table 217-6 *ASH_GLOBAL_REPORT_TEXT Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	List of instances (such as '1,2,3'), or NULL to report on all instances in the database
l_btime	The 'begin time'
l_etime	The 'end time'
l_options	Report level (currently not used)
l_slot_width	Specifies (in seconds) how wide the slots used in the "Top Activity" section of the report should be. This argument is optional, and if it is not specified the time interval between l_btime and l_etime is appropriately split into not more than 10 slots.
l_sid	Session ID (see Usage Notes)
l_sql_id	SQL ID (see Usage Notes)
l_wait_class	Wait class name (see Usage Notes)
l_service_hash	Service name hash (see Usage Notes)
l_module	Module name (see Usage Notes)
l_action	Action name (see Usage Notes)
l_client_id	Client ID for end-to-end backtracing (see Usage Notes)
l_plsql_entry	PL/SQL entry point (see Usage Notes)
l_data_src	Ignored since the report works off of data on disk only
l_container	Name of the container for which report activity is limited. Valid values other than NULL (default) should be taken from container names in V$CONTAINERS. Behavior is as follows: • If NULL: When connected to a root container the report is on all containers. When connected to a PDB the report is on only that PDB. • If not NULL: When connected to a root container the report is on activity from the specified container. When connected to a PDB the report is the same as NULL value for l_container regarding the connected PDB. Note: If while connected to a PDB you request information from another PDB this produces an empty report.

Return Values

The output will be one column of VARCHAR2(320).

Usage Notes

- You can call the function directly but Oracle recommends you use the ashrpti.sql script which prompts users for the required information.
- The unspecified optional arguments are used to generate an ASH Reports that specify 'report targets' such as a SQL statement, or a session, or a particular Service/Module combination. These arguments are specified to restrict the ASH rows that would be used to generate the report. For example, to generate an ASH report on a particular SQL

statement, such as SQL_ID 'abcdefghij123 ' pass that SQL_ID value to the l_sql_id argument:

```
l_sql_id =>  'abcdefghij123'
```

Table 217-7 ASH_GLOBAL_REPORT_TEXT: Wildcards Allowed (or Not) in Arguments

Argument Name	Comment	Wildcard Allowed
l_sid	Session ID (for example, V$SESSION.SID)	No
l_sql_id	SQL ID (for example, V$SQL.SQL_ID)	Yes
l_wait_class	Wait class name (for example, V$EVENT_NAME.WAIT_CLASS)	Yes
l_service_hash	Service name hash (for example, V$ACTIVE_SERVICES.NAME_HASH)	No
l_module	Module name (for example, V$SESSION.MODULE)	Yes
l_action	Action name (for example, V$SESSION.ACTION)	Yes
l_client_id	Client ID for end-to-end backtracing (for example, V$SESSION.CLIENT_IDENTIFIER)	Yes
l_plsql_entry	PL/SQL entry point (for example, "SYS.DBMS_LOB.*")	Yes
l_data_src	Wildcards are not allowed for l_data_src as it is of numeric datatype	No

- Any combination of those optional arguments can be passed in, and only rows in ASH that satisfy all of those 'report targets' will be used. If multiple 'report targets' are specified, AND conditional logic is used to connect them. For example, to generate an ASH report on MODULE "PAYROLL" and ACTION "PROCESS", use the following predicate:

  ```
  l_module =>  'PAYROLL', l_action =>  'PROCESS'
  ```

  Valid SQL wildcards can be used in all the arguments that are of type VARCHAR2.

217.3.4 ASH_REPORT_ANALYTICS Function

This function returns the ASH Analytics active report.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.ASH_REPORT_ANALYTICS(
   dbid           IN NUMBER DEFAULT NULL,
   inst_id        IN NUMBER DEFAULT NULL,
   begin_time     IN DATE,
   end_time       IN DATE,
   report_level   IN VARCHAR2 DEFAULT NULL,
   filter_list    IN VARCHAR2 DEFAULT NULL)
 RETURN CLOB;
```

Parameters

Table 217-8 ASH_REPORT_ANALYTICS Parameters

Parameter	Description
dbid	Database identifier. If its value is set to NULL, then the database identifier for the local database is used. Its default value is NULL.
inst_id	Instance number of the database for which the statistics are required. If its value is set to NULL, then the statistics for the local database are returned. Its default value is NULL.
begin_time	The start time of the interval for which the ASH report is required.
end_time	The end time of the interval for which the ASH report is required.
report_level	Describes the list of components to build.
filter_list	Describes the list of filters to apply. Its default value is NULL (no filters to apply).

Return Values

Returns the ASH Analytics active report.

217.3.5 ASH_REPORT_HTML Function

This table function displays the ASH Spot report in HTML.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.ASH_REPORT_HTML(
   l_dbid          IN NUMBER,
   l_inst_num      IN NUMBER,
   l_btime         IN DATE,
   l_etime         IN DATE,
   l_options       IN NUMBER    DEFAULT 0,
   l_slot_width    IN NUMBER    DEFAULT 0,
   l_sid           IN NUMBER    DEFAULT NULL,
   l_sql_id        IN VARCHAR2  DEFAULT NULL,
   l_wait_class    IN VARCHAR2  DEFAULT NULL,
   l_service_hash  IN NUMBER    DEFAULT NULL,
   l_module        IN VARCHAR2  DEFAULT NULL,
   l_action        IN VARCHAR2  DEFAULT NULL,
   l_client_id     IN VARCHAR2  DEFAULT NULL,
   l_plsql_entry   IN VARCHAR2  DEFAULT NULL,
   l_data_src      IN NUMBER    DEFAULT 0,
   l_container     IN VARCHAR2  DEFAULT NULL)
 RETURN awrrpt_html_type_table PIPELINED;
```

Parameters

Table 217-9 *ASH_REPORT_HTML Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	Instance number

Table 217-9 (Cont.) *ASH_REPORT_HTML Parameters*

Parameter	Description
l_btime	The 'begin time'
l_etime	The 'end time'
l_options	Report level (currently not used)
l_slot_width	Specifies (in seconds) how wide the slots used in the "Top Activity" section of the report should be. This argument is optional, and if it is not specified the time interval between l_btime and l_etime is appropriately split into not more than 10 slots.
l_sid	Session ID (see Usage Notes)
l_sql_id	SQL ID (see Usage Notes)
l_wait_class	Wait class name (see Usage Notes)
l_service_hash	Service name hash (see Usage Notes)
l_module	Module name (see Usage Notes)
l_action	Action name (see Usage Notes)
l_client_id	Client ID for end-to-end backtracing (see Usage Notes)
l_plsql_entry	PL/SQL entry point (see Usage Notes)
l_data_src	Can be used to specify a data source (see Usage Notes) • 1 => memory (V$ACTIVE_SESION_HISTORY) • 2 => disk (DBA_HIST_ACTIVE_SESS_HISTORY) • 0 => both. This is the default value. Here, the begin and end time parameters are used to get the samples from the appropriate data source, which can be memory, disk, or both.
l_container	Name of the container for which report activity is limited. Valid values other than NULL (default) should be taken from container names in V$CONTAINERS. Behavior is as follows: • If NULL: When connected to a root container the report is on all containers. When connected to a PDB the report is on only that PDB. • If not NULL: When connected to a root container the report is on activity from the specified container. When connected to a PDB the report is the same as NULL value for l_container regarding the connected PDB. Note: If while connected to a PDB you request information from another PDB this produces an empty report.

Return Values

The output will be one column of VARCHAR2(500).

Usage Notes

- You can call the function directly but Oracle recommends you use the ashrpti.sql script which prompts users for the required information.
- By default, the report uses the begin and end time parameters (l_btime and l_etime, respectively) to find all rows in that time range either from memory, or disk, or both. However, using l_data_src, one can explicitly specify one of those data sources. For

example, to generate an ASH report on all rows between `l_btime` and `l_time` found in memory, use

```
l_data_src =>  1
```

Similarly, to generate a report on samples found only on disk, use

```
l_data_src => 2
```

- The unspecified optional arguments are used to generate an ASH Reports that specify 'report targets' such as a SQL statement, or a session, or a particular Service/Module combination. These arguments are specified to restrict the ASH rows that would be used to generate the report. For example, to generate an ASH report on a particular SQL statement, such as SQL_ID '`abcdefghij123` ' pass that `sql_id` value to the `l_sql_id` argument:

```
l_sql_id =>  'abcdefghij123'
```

Any combination of those optional arguments can be passed in, and only rows in ASH that satisfy all of those 'report targets' will be used. If multiple 'report targets' are specified, `AND` conditional logic is used to connect them. For example, to generate an ASH report on `MODULE` "`PAYROLL`" and `ACTION` "`PROCESS`", use the following predicate:

```
l_module =>  'PAYROLL', l_action =>  'PROCESS'
```

Valid SQL wildcards can be used in all the arguments that are of type `VARCHAR2`.

Table 217-10 ASH_REPORT_HTML: Wildcards Allowed (or Not) in Arguments

Argument Name	Comment	Wildcard Allowed
`l_sid`	Session ID (for example, `V$SESSION.SID`)	No
`l_sql_id`	SQL ID (for example, `V$SQL.SQL_ID`)	Yes
`l_wait_class`	Wait class name (for example, `V$EVENT_NAME.WAIT_CLASS`)	Yes
`l_service_hash`	Service name hash (for example, `V$ACTIVE_SERVICES.NAME_HASH`)	No
`l_module`	Module name (for example, `V$SESSION.MODULE`)	Yes
`l_action`	Action name (for example, `V$SESSION.ACTION`)	Yes
`l_client_id`	Client ID for end-to-end backtracing (for example, `V$SESSION.CLIENT_IDENTIFIER`)	Yes

217.3.6 ASH_REPORT_TEXT Function

This table function displays the ASH Spot report in text.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.ASH_REPORT_TEXT(
   l_dbid          IN NUMBER,
   l_inst_num      IN NUMBER,
   l_btime         IN DATE,
   l_etime         IN DATE,
   l_options       IN NUMBER    DEFAULT 0,
   l_slot_width    IN NUMBER    DEFAULT 0,
   l_sid           IN NUMBER    DEFAULT NULL,
   l_sql_id        IN VARCHAR2  DEFAULT NULL,
```

```
   l_wait_class    IN VARCHAR2  DEFAULT NULL,
   l_service_hash  IN NUMBER    DEFAULT NULL,
   l_module        IN VARCHAR2  DEFAULT NULL,
   l_action        IN VARCHAR2  DEFAULT NULL,
   l_client_id     IN VARCHAR2  DEFAULT NULL,
   l_plsql_entry   IN VARCHAR2  DEFAULT NULL,
   l_data_src      IN NUMBER    DEFAULT 0,
   l_container     IN VARCHAR2  DEFAULT NULL)
 RETURN awrrpt_text_type_table PIPELINED;
```

Parameters

Table 217-11 *ASH_REPORT_TEXT Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	Instance number
l_btime	The 'begin time'
l_etime	The 'end time'
l_options	Report level (currently not used)
l_slot_width	Specifies (in seconds) how wide the slots used in the "Top Activity" section of the report should be. This argument is optional, and if it is not specified the time interval between l_btime and l_etime is appropriately split into not more than 10 slots.
l_sid	Session ID (see Usage Notes)
l_sql_id	SQL ID (see Usage Notes)
l_wait_class	Wait class name (see Usage Notes)
l_service_hash	Service name hash (see Usage Notes)
l_module	Module name (see Usage Notes)
l_action	Action name (see Usage Notes)
l_client_id	Client ID for end-to-end backtracing (see Usage Notes)
l_plsql_entry	PL/SQL entry point (see Usage Notes)
l_data_src	Can be used to specify a data source (see Usage Notes) • 1 => memory (V$ACTIVE_SESION_HISTORY) • 2 => disk (DBA_HIST_ACTIVE_SESS_HISTORY) • 0 => both. This is the default value. Here, the begin and end time parameters are used to get the samples from the appropriate data source, which can be memory, disk, or both.
l_container	Name of the container for which report activity is limited. Valid values other than NULL (default) should be taken from container names in V$CONTAINERS. Behavior is as follows: • If NULL: When connected to a root container the report is on all containers. When connected to a PDB the report is on only that PDB. • If not NULL: When connected to a root container the report is on activity from the specified container. When connected to a PDB the report is the same as NULL value for l_container regarding the connected PDB. Note: If while connected to a PDB you request information from another PDB this produces an empty report.

Return Values

The output will be one column of `VARCHAR2(80)`.

Usage Notes

- You can call the function directly but Oracle recommends you use the `ashrpti.sql` script which prompts users for the required information.
- By default, the report uses the begin and end time parameters (`l_btime` and `l_etime`, respectively) to find all rows in that time range either from memory, or disk, or both. However, using `l_data_src`, one can explicitly specify one of those data sources. For example, to generate an ASH report on all rows between `l_btime` and `l_time` found in memory, use

  ```
  l_data_src =>  1
  ```

  Similarly, to generate a report on samples found only on disk, use

  ```
  l_data_src => 2
  ```

- The unspecified optional arguments are used to generate an ASH Reports that specify 'report targets' such as a SQL statement, or a session, or a particular Service/Module combination. These arguments are specified to restrict the ASH rows that would be used to generate the report. For example, to generate an ASH report on a particular SQL statement, such as SQL_ID '`abcdefghij123` ' pass that `SQL_ID` value to the `l_sql_id` argument:

  ```
  l_sql_id =>  'abcdefghij123'
  ```

Table 217-12 ASH_REPORT_TEXT: Wildcards Allowed (or Not) in Arguments

Argument Name	Comment	Wildcard Allowed
`l_sid`	Session ID (for example, `V$SESSION.SID`)	No
`l_sql_id`	SQL ID (for example, `V$SQL.SQL_ID`)	Yes
`l_wait_class`	Wait class name (for example, `V$EVENT_NAME.WAIT_CLASS`)	Yes
`l_service_hash`	Service name hash (for example, `V$ACTIVE_SERVICES.NAME_HASH`)	No
`l_module`	Module name (for example, `V$SESSION.MODULE`)	Yes
`l_action`	Action name (for example, `V$SESSION.ACTION`)	Yes
`l_client_id`	Client ID for end-to-end backtracing (for example, `V$SESSION.CLIENT_IDENTIFIER`)	Yes
`l_plsql_entry`	PL/SQL entry point (for example, "`SYS.DBMS_LOB.*`")	Yes
`l_data_src`	Wildcards are not allowed for `l_data_src` as it is of numeric datatype	No

- Any combination of those optional arguments can be passed in, and only rows in ASH that satisfy all of those 'report targets' will be used. If multiple 'report targets' are specified, `AND` conditional logic is used to connect them. For example, to generate an ASH report on `MODULE` "`PAYROLL`" and `ACTION` "`PROCESS`", use the following predicate:

  ```
  l_module =>  'PAYROLL', l_action =>  'PROCESS'
  ```

Valid SQL wildcards can be used in all the arguments that are of type VARCHAR2.

217.3.7 AWR_DIFF_REPORT_HTML Function

This table function displays the AWR Compare Periods report in HTML.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_DIFF_REPORT_HTML(
   dbid1     IN NUMBER,
   inst_num1 IN NUMBER,
   bid1      IN NUMBER,
   eid1      IN NUMBER,
   dbid2     IN NUMBER,
   inst_num2 IN NUMBER,
   bid2      IN NUMBER,
   eid2      IN NUMBER)
  RETURN awrdrpt_text_type_table PIPELINED;
```

Parameters

Table 217-13 *AWR_DIFF_REPORT_HTML Parameters*

Parameter	Description
dbid1	1st database identifier
inst_num1	1st instance number
bid1	1st beginning snapshot ID
eid1	1st ending snapshot ID
dbid2	2nd database identifier
inst_num2	2nd instance number
bid2	2nd beginning snapshot ID
eid2	2nd ending snapshot ID

Return Values

The output will be one column of VARCHAR2(500).

Usage Notes

You can call the function directly but Oracle recommends you use the awrddrpt.sql script which prompts users for the required information.

217.3.8 AWR_DIFF_REPORT_TEXT Function

This table function displays the AWR Compare Periods report in text.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_DIFF_REPORT_TEXT(
   dbid1     IN NUMBER,
   inst_num1 IN NUMBER,
   bid1      IN NUMBER,
   eid1      IN NUMBER,
   dbid2     IN NUMBER,
```

```
   inst_num2 IN NUMBER,
   bid2      IN NUMBER,
   eid2      IN NUMBER)
  RETURN awrdrpt_text_type_table PIPELINED;
```

Parameters

Table 217-14 *AWR_DIFF_REPORT_TEXT Parameters*

Parameter	Description
dbid1	1st database identifier
inst_num1	1st instance number
bid1	1st beginning snapshot ID
eid1	1st ending snapshot ID
dbid2	2nd database identifier
inst_num2	2nd instance number
bid2	2nd beginning snapshot ID
eid2	2nd ending snapshot ID

Return Values

The output will be one column of `VARCHAR2(500)`.

Usage Notes

You can call the function directly but Oracle recommends you use the `awrddrpt.sql` script which prompts users for the required information.

217.3.9 AWR_EXP Procedure

This procedure extracts AWR data from the AWR schema and dump the information into a file. You can specify the snapshot range for the data that you want to extract.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_EXP(
   dmpfile              IN   VARCHAR2   DEFAULT 'awrdat',
   dmpdir               IN   VARCHAR2   DEFAULT 'DATA_PUMP_DIR',
   dbid                 IN   NUMBER     DEFAULT NULL,
   bid                  IN   NUMBER     DEFAULT 1,
   eid                  IN   NUMBER     DEFAULT 1000000);
```

Parameters

Table 217-15 *AWR_EXP Procedure Parameters*

Parameter	Description
dmpfile	The prefix for the name of the extract dump file and log file. The name of the dump file where all the data from the AWR table will be written is dmpfile.dmp. The `dmpfile.log` log file shows the status of the export job. The default value for the prefix is `awrdat`.

Table 217-15 (Cont.) *AWR_EXP Procedure Parameters*

Parameter	Description
dmpdir	Name of the Directory Object for the file system directory where the extract dump file will be written. The list of Directory Objects can be queried using the DBA_DIRECTORIES view, and a new directory object can be created using the following command: create directory dmpdir as '/directory/path'. The default value is DATA_PUMP_DIR.
dbid	The database ID for the snapshots that you want to export. The default value is NULL, for the local database ID.
bid	The begin snapshot ID for snapshots to be exported. The default value is 1.
eid	The end Snapshot Id for snapshots to be exported. The default value is 10000000.

217.3.10 AWR_GLOBAL_DIFF_REPORT_HTML Functions

This table function displays Global AWR Compare Periods Report in HTML format.

The first overload accepts comma-separated lists of instance numbers for inst_num1 and inst_num2. No leading zeroes are allowed and there is a limit of 1023 characters.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_DIFF_REPORT_HTML (
   dbid1        IN    NUMBER,
   inst_num1    IN    AWRRPT_INSTANCE_LIST_TYPE,
   bid1         IN    NUMBER,
   eid1         IN    NUMBER,
   dbid2        IN    NUMBER,
   inst_num2    IN    AWRRPT_INSTANCE_LIST_TYPE,
   bid2         IN    NUMBER,
   eid2         IN    NUMBER)
  RETURN awrrpt_html_type_table PIPELINED;

DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_DIFF_REPORT_HTML (
   dbid1        IN    NUMBER,
   inst_num1    IN    VARCHAR2,
   bid1         IN    NUMBER,
   eid1         IN    NUMBER,
   dbid2        IN    NUMBER,
   inst_num2    IN    VARCHAR2,
   bid2         IN    NUMBER,
   eid2         IN    NUMBER)
  RETURN awrrpt_html_type_table PIPELINED;
```

Parameters

Table 217-16 *AWR_GLOBAL_DIFF_REPORT_HTML Function Parameters*

Parameter	Description
dbid1	1st database identifier
inst_num1	1st list of instance numbers. If set to NULL, all instances for which begin and end snapshots are available, and which have not been restarted between snapshots, will be included in the report.
bid1	1st beginning snapshot ID
eid1	1st ending snapshot ID
dbid2	2nd database identifier
inst_num2	2nd list of instance numbers to be included in report. If set to NULL, all instances for which begin and end snapshots are available, and which have not been restarted between snapshots, will be included in the report.
bid2	2nd beginning snapshot ID
eid2	2nd ending snapshot ID

Return Values

The output will be one column of VARCHAR2(1500).

217.3.11 AWR_GLOBAL_DIFF_REPORT_TEXT Functions

This table function displays Global AWR Compare Periods Report in text format.

The first overload accepts comma-separated lists of instance numbers for inst_num1 and inst_num2. No leading zeroes are allowed and there is a limit of 1023 characters.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_DIFF_REPORT_TEXT (
   dbid1         IN    NUMBER,
   inst_num1     IN    AWRRPT_INSTANCE_LIST_TYPE,
   bid1          IN    NUMBER,
   eid1          IN    NUMBER,
   dbid2         IN    NUMBER,
   inst_num2     IN    AWRRPT_INSTANCE_LIST_TYPE,
   bid2          IN    NUMBER,
   eid2          IN    NUMBER)
  RETURN awrdrpt_text_type_table PIPELINED;

DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_DIFF_REPORT_TEXT (
   dbid1         IN    NUMBER,
   inst_num1     IN    VARCHAR2,
   bid1          IN    NUMBER,
   eid1          IN    NUMBER,
   dbid2         IN    NUMBER,
   inst_num2     IN    VARCHAR2,
   bid2          IN    NUMBER,
   eid2          IN    NUMBER)
  RETURN awrdrpt_text_type_table PIPELINED;
```

Parameters

Table 217-17 *AWR_GLOBAL_DIFF_REPORT_TEXT Functions Parameters*

Parameter	Description
dbid1	1st database identifier
inst_num1	1st list of instance numbers. If set to NULL, all instances for which begin and end snapshots are available, and which have not been restarted between snapshots, will be included in the report.
bid1	1st beginning snapshot ID
eid1	1st ending snapshot ID
dbid2	2nd database identifier
inst_num2	2nd list of instance numbers to be included in report. If set to NULL, all instances for which begin and end snapshots are available, and which have not been restarted between snapshots, will be included in the report.
bid2	2nd beginning snapshot ID
eid2	2nd ending snapshot ID

Return Values

The output will be one column of VARCHAR2(320).

217.3.12 AWR_GLOBAL_REPORT_HTML Functions

This table function displays the Global AWR report in HTML.

The first overload accepts a comma-separated list of instance numbers. No leading zeroes are allowed and there is a limit of 1023 characters.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_REPORT_HTML (
   l_dbid        IN    NUMBER,
   l_inst_num    IN    AWRRPT_INSTANCE_LIST_TYPE,
   l_bid         IN    NUMBER,
   l_eid         IN    NUMBER,
   l_options     IN    NUMBER DEFAULT 0)
  RETURN awrrpt_html_type_table PIPELINED;

DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_REPORT_HTML (
   l_dbid       IN    NUMBER,
   l_inst_num   IN    VARCHAR2,
   l_bid        IN    NUMBER,
   l_eid        IN    NUMBER,
   l_options    IN    NUMBER DEFAULT 0)
  RETURN awrrpt_html_type_table PIPELINED;
```

Parameters

Table 217-18 *AWR_GLOBAL_REPORT_HTML Function Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	List of instance numbers to be included in report. If set to NULL, all instances for which begin and end snapshots are available, and which have not been restarted between snapshots, will be included in the report.
l_bid	Beginning snapshot ID
l_eid	Ending snapshot ID
l_options	Report level (currently not used)

Return Values

The output will be one column of VARCHAR2(1500).

217.3.13 AWR_GLOBAL_REPORT_TEXT Functions

This table function displays the Global AWR report in text.

The first overload accepts a comma-separated list of instance numbers. No leading zeroes are allowed and there is a limit of 1023 characters

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_REPORT_TEXT(
   l_dbid        IN    NUMBER,
   l_inst_num    IN    AWRRPT_INSTANCE_LIST_TYPE,
   l_bid         IN    NUMBER,
   l_eid         IN    NUMBER,
   l_options     IN    NUMBER DEFAULT 0)
  RETURN awrdrpt_text_type_table PIPELINED;

DBMS_WORKLOAD_REPOSITORY.AWR_GLOBAL_REPORT_TEXT(
   l_dbid       IN    NUMBER,
   l_inst_num   IN    VARCHAR2,
   l_bid        IN    NUMBER,
   l_eid        IN    NUMBER,
   l_options    IN    NUMBER DEFAULT 0)
  RETURN awrdrpt_text_type_table PIPELINED;
```

Parameters

Table 217-19 *AWR_GLOBAL_REPORT_TEXT Function Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	List of instance numbers to be included in report. If set to NULL, all instances for which begin and end snapshots are available, and which have not been restarted between snapshots, will be included in the report.

Table 217-19 (Cont.) *AWR_GLOBAL_REPORT_TEXT Function Parameters*

Parameter	Description
l_bid	Beginning snapshot ID
l_eid	Ending snapshot ID
l_options	A flag to specify to control the output of the report. Currently, not used.

Return Values

The output will be one column of VARCHAR2(320).

217.3.14 AWR_IMP Procedure

This procedure loads the AWR data from a dump file into the SYS schema.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_IMP(
   dmpfile              IN   VARCHAR2   DEFAULT 'awrdat',
   dmpdir               IN   VARCHAR2   DEFAULT 'DATA_PUMP_DIR',
   new_dbid             IN   NUMBER     DEFAULT NULL);
```

Parameters

Table 217-20 *AWR_IMP Procedure Parameters*

Parameter	Description
dmpfile	The prefix for the name of the dump file and log file. This file will be the source of the imported AWR data. The default value is awrdat.
dmpdir	The name of the Directory Object for the file system directory where the load dump file is located. The default value is DATA_PUMP_DIR.
new_dbid	The database ID to be used instead of existing database ID.

217.3.15 AWR_REPORT_HTML Function

This table function displays the AWR report in HTML.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_REPORT_HTML(
   l_dbid       IN    NUMBER,
   l_inst_num   IN    NUMBER,
   l_bid        IN    NUMBER,
   l_eid        IN    NUMBER,
   l_options    IN    NUMBER DEFAULT 0)
 RETURN awrrpt_text_type_table PIPELINED;
```

Parameters

Table 217-21 *AWR_REPORT_HTML Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	Instance number
l_bid	Beginning snapshot ID
l_eid	Ending snapshot ID
l_options	A flag to specify to control the output of the report. Currently, Oracle supports one value: • l_options - 8. Displays the ADDM specific portions of the report. These sections include the Buffer Pool Advice, Shared Pool Advice, and PGA Target Advice.

Return Values

The output will be one column of VARCHAR2(1500).

Usage Notes

You can call the function directly but Oracle recommends you use the awrrpt.sql script which prompts users for the required information.

217.3.16 AWR_REPORT_TEXT Function

This table function displays the AWR report in text.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_REPORT_TEXT(
   l_dbid       IN    NUMBER,
   l_inst_num   IN    NUMBER,
   l_bid        IN    NUMBER,
   l_eid        IN    NUMBER,
   l_options    IN    NUMBER DEFAULT 0)
 RETURN awrrpt_text_type_table PIPELINED;
```

Parameters

Table 217-22 AWR_REPORT_TEXT Parameters

Parameter	Description
l_dbid	Database identifier
l_insT_num	Instance number
l_bid	Beginning snapshot ID
l_eid	Ending snapshot ID

Table 217-22 (Cont.) AWR_REPORT_TEXT Parameters

Parameter	Description
l_options	A flag to specify to control the output of the report. Currently, Oracle supports one value: • l_options - 8. Displays the ADDM specific portions of the report. These sections include the Buffer Pool Advice, Shared Pool Advice, and PGA Target Advice.

Return Values

The output will be one column of VARCHAR2(80).

Usage Notes

You can call the function directly but Oracle recommends you use the awrrpt.sql script which prompts users for the required information.

217.3.17 AWR_SET_REPORT_THRESHOLDS Procedure

This procedure configure specified report thresholds, including the number of rows in the report.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_SET_REPORT_THRESHOLDS(
   top_n_events         IN   NUMBER DEFAULT NULL,
   top_n_files          IN   NUMBER DEFAULT NULL,
   top_n_segments       IN   NUMBER DEFAULT NULL,
   top_n_services       IN   NUMBER DEFAULT NULL,
   top_n_sql            IN   NUMBER DEFAULT NULL,
   top_n_sql_max        IN   NUMBER DEFAULT NULL,
   top_sql_pct          IN   NUMBER DEFAULT NULL,
   shmem_threshold      IN   NUMBER DEFAULT NULL,
   versions_threshold   IN   NUMBER DEFAULT NULL,
   top_n_disks          IN   NUMBER DEFAULT NULL,
   outlier_pct          IN   NUMBER DEFAULT NULL,
   outlier_cpu_pct      IN   NUMBER DEFAULT NULL);
```

Parameters

Table 217-23 *AWR_SET_REPORT_THRESHOLDS Procedure Parameters*

Parameter	Description
top_n_events	Number of most significant wait events to be included
top_n_files	Number of most active files to be included
top_n_segments	Number of most active segments to be included
top_n_services	Number of most active services to be included
top_n_sql	Number of most significant SQL statements to be included
top_n_sql_max	Number of SQL statements to be included if their activity is greater than that specified by top_sql_pct

Table 217-23 (Cont.) *AWR_SET_REPORT_THRESHOLDS Procedure Parameters*

Parameter	Description
top_sql_pct	Significance threshold for SQL statements between top_n_sql and top_n_sql_max
shmem_threshold	Shared memory low threshold
versions_threshold	Plan version count low threshold
top_n_disks	Number of cell disks with most I/O
outlier_pct	Percentage of maximum capacity before displaying outliers for Exadata sections
outlier_cpu_pct	Threshold for mean percentage CPU to display outliers

User Notes

- The effect of each setting depends on the type of report being generated as well as on the underlying AWR data. Not all settings are meaningful for each report type. Invalid settings (such as negative numbers) are ignored.
- Settings are effective only in the context of the session that executes the AWR_SET_REPORT_THRESHOLDS procedure. For example, to get a report that lists top 12 segments as compared to the default, one can invoke as follows:

```
DBMS_WORKLOAD_REPOSITORY.AWR_SET_REPORT_THRESHOLDS (top_n_segments=>12);
```

217.3.18 AWR_SQL_REPORT_HTML Function

This table function displays the AWR SQL Report in HTML format.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_SQL_REPORT_HTML(
   l_dbid       IN    NUMBER,
   l_inst_num   IN    NUMBER,
   l_bid        IN    NUMBER,
   l_eid        IN    NUMBER,
   l_sqlid      IN    VARCHAR2,
   l_options    IN    NUMBER DEFAULT 0,
   l_con_dbid  IN    NUMBER DEFAULT NULL)
 RETURN awrrpt_html_type_table PIPELINED;
```

Parameters

Table 217-24 *AWR_SQL_REPORT_HTML Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	Instance number
l_bid	Beginning snapshot ID
l_eid	Ending snapshot ID

Table 217-24 (Cont.) *AWR_SQL_REPORT_HTML Parameters*

Parameter	Description
l_sqlid	SQL ID of statement to be analyzed
l_options	A flag to specify to control the output of the report. Currently, not used.
l_con_dbid	Container database identifier. If NULL (default), then l_con_dbid is set equal to l_dbid

Return Values

The output will be one column of VARCHAR2(500).

Usage Notes

You can call the function directly but Oracle recommends you use the awrsqrpt.sql script which prompts users for the required information.

217.3.19 AWR_SQL_REPORT_TEXT Function

This table function displays the AWR SQL Report in text format.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.AWR_SQL_REPORT_TEXT(
   l_dbid       IN    NUMBER,
   l_inst_num   IN    NUMBER,
   l_bid        IN    NUMBER,
   l_eid        IN    NUMBER,
   l_sqlid      IN    VARCHAR2,
   l_options    IN    NUMBER DEFAULT 0,
   l_con_dbid  IN    NUMBER DEFAULT NULL)
 RETURN awrrpt_text_type_table PIPELINED;
```

Parameters

Table 217-25 *AWR_SQL_REPORT_TEXT Parameters*

Parameter	Description
l_dbid	Database identifier
l_inst_num	Instance number
l_bid	Beginning snapshot ID
l_eid	Ending snapshot ID
l_sqlid	SQL ID of statement to be analyzed
l_options	Flag to specify to control the output of the report. Currently, not used.
l_con_dbid	Container database identifier. If NULL (default), then l_con_dbid is set equal to l_dbid

Return Values

The output will be one column of VARCHAR2(120).

Usage Notes

You can call the function directly but Oracle recommends you use the `awrsqrpt.sql` script which prompts users for the required information.

217.3.20 CONTROL_RESTRICTED_SNAPSHOT Procedure

This procedure controls the AWR snapshot creation for a database in the *restricted session* mode. If this procedure is not used, then by default, the AWR snapshots cannot be created for a database in the *restricted session* mode. This procedure affects the local database on which it is executed.

> **Note:**
>
> To enable AWR snapshot creation for an Oracle RAC in the *restricted session* mode, this procedure must be executed on every database instance in the Oracle RAC.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.CONTROL_RESTRICTED_SNAPSHOT(
   allow IN BOOLEAN);
```

Parameters

Table 217-26 CONTROL_RESTRICTED_SNAPSHOT Parameters

Parameter	Description
`allow`	This parameter can have one of the following values: • `TRUE`: AWR snapshots can be created for the database in the *restricted session* mode. • `FALSE`: AWR snapshots cannot be created for the database in the *restricted session* mode.

217.3.21 CREATE_BASELINE Functions and Procedures

This function and procedure creates a baseline.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE(
   start_snap_id    IN  NUMBER,
   end_snap_id      IN  NUMBER,
   baseline_name    IN  VARCHAR2,
   dbid             IN  NUMBER DEFAULT NULL,
   expiration       IN  NUMBER DEFAULT NULL);

DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE(
   start_snap_id    IN  NUMBER,
   end_snap_id      IN  NUMBER,
   baseline_name    IN  VARCHAR2,
   dbid             IN  NUMBER DEFAULT NULL,
   expiration       IN  NUMBER DEFAULT NULL)
 RETURN NUMBER;
```

```
DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE(
   start_time       IN  DATE,
   end_time         IN  DATE,
   baseline_name    IN  VARCHAR2,
   dbid             IN  NUMBER DEFAULT NULL,
   expiration       IN  NUMBER DEFAULT NULL);

DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE(
   start_time       IN  DATE,
   end_time         IN  DATE,
   baseline_name    IN  VARCHAR2,
   dbid             IN  NUMBER DEFAULT NULL,
   expiration       IN  NUMBER DEFAULT NULL);
 RETURN NUMBER;
```

Parameters

Table 217-27 CREATE_BASELINE Function & Procedure Parameters

Parameter	Description
start_snap_id	Start snapshot sequence number for the baseline.
end_snap_id	End snapshot sequence number for the baseline.
start_time	Start time for the baseline.
end_time	End time for the baseline.
baseline_name	Name of baseline.
dbid	Database Identifier for baseline. If NULL, this takes the database identifier for the local database. Defaults to NULL.
expiration	Expiration in number of days for the baseline. If NULL, then expiration is infinite, meaning do not drop baseline ever. Defaults to NULL.

Exceptions

- An error will be returned if this baseline name already exists in the system.
- The snapshot range that is specified for this interface must be an existing pair of snapshots in the database. An error will be returned if the inputted snapshots do not exist in the system.

Examples

This example creates a baseline (named 'oltp_peakload_bl') between snapshots 105 and 107 for the local database:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE (start_snap_id => 105,
                                                  end_snap_id   => 107,
                                                  baseline_name => 'oltp_peakload_bl');
```

If you query the DBA_HIST_BASELINE view after the CREATE BASELINE action, you will see the newly created baseline in the Workload Repository.

217.3.22 CREATE_BASELINE_TEMPLATE Procedures

This procedure specifies a template for how they would like baselines to be created for future time periods.

Syntax

Specifies a template for generating a baseline for a single time period in the future.

```
DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE_TEMPLATE(
   start_time              IN DATE,
   end_time                IN DATE,
   baseline_name           IN VARCHAR2,
   template_name           IN VARCHAR2,
   expiration              IN NUMBER,
   dbid                    IN NUMBER   DEFAULT NULL);
```

Specifies a template for creating and dropping baseline based on repeating time periods:

```
DBMS_WORKLOAD_REPOSITORY.CREATE_BASELINE_TEMPLATE(
   day_of_week             IN VARCHAR2,
   hour_in_day             IN NUMBER,
   duration                IN NUMBER,
   start_time              IN DATE,
   end_time                IN DATE,
   baseline_name_prefix    IN VARCHAR2,
   template_name           IN VARCHAR2,
   expiration              IN NUMBER,
   dbid                    IN NUMBER   DEFAULT NULL);
```

Parameters

Table 217-28 CREATE_BASELINE_TEMPLATE Procedure Parameters

Parameter	Description
start_time	Start Time for the baseline to be created'
end_time	End Time for the baseline to be created
baseline_name	Name of baseline to be created
template_name	Name for the template
expiration	Expiration in number of days for the baseline. If NULL, then expiration is infinite, meaning do not drop baseline ever. Defaults to NULL
dbid	Database ID for which the baseline template needs to be used. If NULL, this takes the database identifier of the local database. Defaults to NULL.
day_of_week	Day of week that the baseline should repeat on. Specify one of the following values: SUNDAY, MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY.
hour_in_day	Value of 0-23 to specify the Hour in the Day the baseline should start
duration	Duration (in number of hours) after hour in the day that the baseline should last
baseline_name_prefix	Name for baseline prefix. When creating the baseline, the name of the baseline will be the prefix appended with the date information.

217.3.23 CREATE_REMOTE_SNAPSHOT Function and Procedure

This function and procedure create a remote snapshot using the Remote Management Framework (RMF). The function returns the snapshot ID.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.CREATE_REMOTE_SNAPSHOT(
   node_id     IN NUMBER,
   flush_level IN VARCHAR2 DEFAULT 'BESTFIT');

DBMS_WORKLOAD_REPOSITORY.CREATE_REMOTE_SNAPSHOT(
   node_id     IN NUMBER,
   flush_level IN VARCHAR2 DEFAULT 'BESTFIT')
 RETURN NUMBER;

DBMS_WORKLOAD_REPOSITORY.CREATE_REMOTE_SNAPSHOT(
   node_name      IN VARCHAR2,
   topology_name  IN VARCHAR2 DEFAULT NULL,
   flush_level    IN VARCHAR2 DEFAULT 'BESTFIT');

DBMS_WORKLOAD_REPOSITORY.CREATE_REMOTE_SNAPSHOT(
   node_name      IN VARCHAR2,
   topology_name  IN VARCHAR2 DEFAULT NULL,
   flush_level    IN VARCHAR2 DEFAULT 'BESTFIT')
 RETURN NUMBER;
```

Parameters

Table 217-29 CREATE_REMOTE_SNAPSHOT Parameters

Parameter	Description
node_id	RMF node identifier of the database for which the snapshot needs to be created.
node_name	RMF node name of the database for which the snapshot needs to be created.
topology_name	RMF topology name of the database for which the snapshot needs to be created.

Table 217-29 (Cont.) CREATE_REMOTE_SNAPSHOT Parameters

Parameter	Description
flush_level	The flush level can be one of the following: • BESTFIT: Uses the default value depending on the type of snapshot being taken. • LITE: Lightweight snapshot. Only the most important statistics are collected. This is default for a pluggable database (PDB) and application container. • TYPICAL: Regular snapshot. Most of the statistics are collected. This is default for a container database root (CDB root) and non-CDB database. • ALL: Heavyweight snapshot. All the possible statistics are collected. This consumes a considerable amount of disk space and takes a long time to create.

Examples

This example creates a remote snapshot of the database having the RMF node identifier of 10:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.CREATE_REMOTE_SNAPSHOT(10);
```

If you query the DBA_HIST_SNAPSHOT view after executing the above procedure, you will see a new snapshot record added to the Automatic Workload Repository (AWR).

217.3.24 CREATE_SNAPSHOT Function and Procedure

This function and procedure create a snapshot. The function returns the snapshot ID. If both, the database ID and the database name are not specified in this subprogram, then the snapshot is created for the local database on which the subprogram is executed.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT(
   flush_level IN VARCHAR2 DEFAULT 'BESTFIT',
   dbid        IN NUMBER   DEFAULT NULL,
   source_name IN VARCHAR2 DEFAULT NULL);

DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT(
   flush_level IN VARCHAR2 DEFAULT 'BESTFIT',
```

```
   dbid        IN NUMBER   DEFAULT NULL,
   source_name IN VARCHAR2 DEFAULT NULL)
 RETURN NUMBER;
```

Parameters

Table 217-30 CREATE_SNAPSHOT Parameters

Parameter	Description
flush_level	The flush level can be one of the following: • BESTFIT: Uses the default value depending on the type of snapshot being taken. • LITE: Lightweight snapshot. Only the most important statistics are collected. This is default for a pluggable database (PDB) and application container. • TYPICAL: Regular snapshot. Most of the statistics are collected. This is default for a container database root (CDB root) and non-CDB database. • ALL: Heavyweight snapshot. All the possible statistics are collected. This consumes a considerable amount of disk space and takes a long time to create.
dbid	Database ID of the database for which the snapshot needs to be created.
source_name	Name of the database for which the snapshot needs to be created.

Examples

This example creates a snapshot of the local database with the flush level of ALL:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT('ALL');
```

If you query the DBA_HIST_SNAPSHOT view after executing the above procedure, you will see a new snapshot record added to the Automatic Workload Repository (AWR).

217.3.25 DROP_BASELINE Procedure

This procedure drops a previously-defined baseline.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.DROP_BASELINE(
   baseline_name  IN  VARCHAR2,
   cascade        IN  BOOLEAN  DEFAULT FALSE,
   dbid           IN  NUMBER   DEFAULT NULL);
```

Parameters

Table 217-31 DROP_BASELINE Parameters

Parameter	Description
baseline_name	Name of baseline to drop from the system
cascade	If TRUE, the pair of snapshots associated with the baseline will also be dropped. Otherwise, only the baseline is removed.
dbid	Database ID for which the baseline needs to be dropped (defaults to local DBID).

Examples

This example drops the baseline 'oltp_peakload_bl' without dropping the underlying snapshots:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.DROP_BASELINE (
       baseline_name => 'oltp_peakload_bl');
```

If you query the DBA_HIST_BASELINE view after the DROP_BASELINE action, you will see the specified baseline definition is removed. You can query the DBA_HIST_SNAPSHOT view to find that the underlying snapshots are left intact.

217.3.26 DROP_BASELINE_TEMPLATE Procedure

This procedure removes a template that is no longer needed.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.DROP_BASELINE_TEMPLATE(
   template_name    IN VARCHAR2,
   dbid             IN NUMBER   DEFAULT NULL);
```

Parameters

Table 217-32 DROP_BASELINE_TEMPLATE Procedure Parameters

Parameter	Description
template_name	Name of the template to remove
dbid	Database ID for which the baseline template needs to be dropped. If NULL, this takes the database identifier of the local database. Defaults to NULL.

217.3.27 DROP_SNAPSHOT_RANGE Procedure

This procedure drops a range of snapshots.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.DROP_SNAPSHOT_RANGE(
   low_snap_id    IN  NUMBER,
   high_snap_id   IN  NUMBER,
   dbid           IN  NUMBER   DEFAULT NULL);
```

Parameters

Table 217-33 DROP_SNAPSHOT_RANGE Procedure Parameters

Parameter	Description
low_snap_id	Low snapshot id of snapshots to drop.
high_snap_id	High snapshot id of snapshots to drop.
dbid	Database id (defaults to local DBID).

Examples

This example drops the range of snapshots between snapshot id 102 to 105 for the local database:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.DROP_SNAPSHOT_RANGE(102, 105);
```

If you query the `dba_hist_snapshot` view after the `Drop Snapshot` action, you will see that snapshots 102 to 105 are removed from the Workload Repository.

217.3.28 LOCAL_AWR_DBID Function

This function returns the database identifier for the local AWR database.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.LOCAL_AWR_DBID()
 RETURN NUMBER;
```

217.3.29 MODIFY_BASELINE_WINDOW_SIZE Procedure

This procedure modifies the window size for the Default Moving Window Baseline.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.MODIFY_BASELINE_WINDOW_SIZE(
   window_size     IN   NUMBER,
   dbid            IN   NUMBER   DEFAULT NULL);
```

Parameters

Table 217-34 MODIFY_BASELINE_WINDOW_SIZE Procedure Parameters

Parameter	Description
window_size	New Window size for the default Moving Window Baseline, in number of days.
dbid	Database ID (defaults to local DBID).

Usage Notes

The window size must be less than or equal to the AWR retention setting. If the window size needs to be greater than the retention setting, the MODIFY_SNAPSHOT_SETTINGS Procedures can be used to adjust the retention setting. A moving window can be set to a maximum of 13 weeks.

217.3.30 MODIFY_SNAPSHOT_SETTINGS Procedures

This procedure controls three aspects of snapshot generation.

- The INTERVAL setting affects how often snapshots are automatically captured.
- The RETENTION setting affects how long snapshots are retained in the Workload Repository.
- The number of SQL captured for each Top criteria. If the user manually specifies a value for Top N SQL, the AWR SQL collection will use the user-specified number for both automatic and manual snapshots.

There are two overloads. The first takes a NUMBER and the second takes a VARCHAR2 for the topnsql argument. The differences are described under the Parameters description.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.MODIFY_SNAPSHOT_SETTINGS(
   retention       IN  NUMBER    DEFAULT  NULL,
   interval        IN  NUMBER    DEFAULT  NULL,
   topnsql         IN  NUMBER    DEFAULT  NULL,
   dbid            IN  NUMBER    DEFAULT  NULL,
   tablespace_name     IN  VARCHAR2  DEFAULT  NULL);

DBMS_WORKLOAD_REPOSITORY.MODIFY_SNAPSHOT_SETTINGS(
   retention       IN  NUMBER    DEFAULT NULL,
   interval        IN  NUMBER    DEFAULT NULL,
   topnsql         IN  VARCHAR2,
   dbid            IN  NUMBER    DEFAULT  NULL,
   tablespace_name     IN  VARCHAR2  DEFAULT  NULL);
```

Parameters

Table 217-35 MODIFY_SNAPSHOT_SETTINGS Procedure Parameters

Parameter	Description
retention	New retention time (in minutes). The specified value must be in the range of MIN_RETENTION (1 day) to MAX_RETENTION (100 years). If ZERO is specified, snapshots will be retained forever. A large system-defined value will be used as the interval setting. If NULL is specified, the old value for retention is preserved. NOTE: The retention setting must be greater than or equal to the window size of the 'SYSTEM_MOVING_WINDOW' baseline. If the retention needs to be less than the window size, the MODIFY_BASELINE_WINDOW_SIZE Procedure can be used to adjust the window size.
interval	New interval setting between each snapshot, in units of minutes. The specified value must be in the range MIN_INTERVAL (10 minutes) to MAX_INTERVAL (1 year). If ZERO is specified, automatic and manual snapshots will be disabled. A large system-defined value will be used as the retention setting. If NULL is specified, the current value is preserved.

Table 217-35 (Cont.) MODIFY_SNAPSHOT_SETTINGS Procedure Parameters

Parameter	Description
topnsql	• If NUMBER: Top N SQL size. The number of Top SQL to flush for each SQL criteria (Elapsed Time, CPU Time, Parse Calls, Shareable Memory, Version Count). The value for this setting will not be affected by the statistics/flush level and will override the system default behavior for the AWR SQL collection. The setting will have a minimum value of 30 and a maximum value of 50,000. Specifying NULL will keep the current setting. • If VARCHAR2: Users are allowed to specify the following values: (DEFAULT, MAXIMUM, N), where N is the number of Top SQL to flush for each SQL criteria. Specifying DEFAULT will revert the system back to the default behavior of Top 30 for statistics level TYPICAL and Top 100 for statistics level ALL. Specifying MAXIMUM will cause the system to capture the complete set of SQL in the cursor cache. Specifying the number N is equivalent to setting the Top N SQL with the NUMBER type. Specifying NULL for this argument will keep the current setting.
dbid	Database identifier in AWR for which to modify the snapshot settings. If NULL is specified, the local dbid will be used. Defaults to NULL.
tablespace_name	Specify a user-defined tablespace for storing AWR data (snapshot data). If this parameter is not used, then AWR data is stored in the SYSAUX tablespace by default.

Examples

This example changes the interval setting to one hour and the retention setting to two weeks for the local database:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.MODIFY_SNAPSHOT_SETTINGS(
  interval  =>  60,
  retention =>  20160);
```

If you query the DBA_HIST_WR_CONTROL table after this procedure is executed, you will see the changes to these settings.

217.3.31 PURGE_SQL_DETAILS Procedure

This procedure purges SQL details, specifically rows from WRH$_SQLTEXT, WRH$_SQL_PLAN, and WRH$_SQL_BIND_METADATA that do not have corresponding rows (DBID, SQL_ID) in WRH$_SQLSTAT.

The subprogram calls for the DBID for which to run the purge. If the DBID is not specified, the database DBID is used. You can constrain runtime by specifying the maximum number of rows to purge per table. If no maximum is specified, the subprograms tries to purge all applicable rows.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.PURGE_SQL_DETAILS(
   numrows IN NUMBER DEFAULT NULL,
   dbid    IN NUMBER DEFAULT NULL);
```

Parameters

Table 217-36 PURGE_SQL_DETAILS Procedure Parameters

Parameter	Description
numrows	Number of rows
dbid	Database ID (default to local DBID)

217.3.32 REGISTER_REMOTE_DATABASE Procedures

This procedure registers a remote database in the Automatic Workload Repository (AWR) using the Remote Management Framework (RMF).

Syntax

```
DBMS_WORKLOAD_REPOSITORY.REGISTER_REMOTE_DATABASE(
   node_id  IN NUMBER);

DBMS_WORKLOAD_REPOSITORY.REGISTER_REMOTE_DATABASE(
   node_name      IN VARCHAR2,
   topology_name  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 217-37 REGISTER_REMOTE_DATABASE Parameters

Parameter	Description
node_id	RMF node identifier of the database that needs to be registered in the AWR.
node_name	RMF node name of the database that needs to be registered in the AWR.
topology_name	RMF topology name of the database that needs to be registered in the AWR.

Examples

This example registers the remote database having the RMF node identifier of 10 in the AWR:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.REGISTER_REMOTE_DATABASE(10);
```

217.3.33 REMOVE_COLORED_SQL Procedure

This procedure removes a colored SQL ID. After a SQL is uncolored, it will no longer be captured in a snapshot automatically, unless it makes the TOP list.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.REMOVE_COLORED_SQL(
   sql_id         IN VARCHAR2,
   dbid           IN NUMBER DEFAULT NULL);
```

Parameters

Table 217-38 *REMOVE_COLORED_SQL Procedure Parameters*

Parameter	Description
sql_id	13-character external SQL ID
dbid	Optional dbid, defaults to Local DBID

217.3.34 RENAME_BASELINE Procedure

This procedure renames a baseline.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.RENAME_BASELINE(
   old_baseline_name     IN   VARCHAR2,
   new_baseline_name     IN VARCHAR2,
   dbid                  IN NUMBER   DEFAULT NULL);
```

Parameters

Table 217-39 RENAME_BASELINE Procedure Parameters

Parameter	Description
old_baseline_name	Old baseline name.
new_baseline_name	New baseline name.
dbid	Database ID for which the baseline needs to be renamed (defaults to local DBID).

217.3.35 SELECT_BASELINE_DETAILS Function

This table function shows the values of the metrics corresponding to a baseline for a range of snapshots.

This table function returns an object of AWR_BASELINE_METRIC_TYPE Object Type.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.SELECT_BASELINE_DETAILS(
   l_baseline_id    IN  NUMBER,
   l_begin_snap     IN  NUMBER DEFAULT NULL,
   l_end_snap       IN  NUMBER DEFAULT NULL,
   l_dbid           IN  NUMBER DEFAULT NULL)
 RETURN awrbl_details_type_table PIPELINED;
```

Parameters

Table 217-40 SELECT_BASELINE_DETAILS Function Parameters

Parameter	Description
l_baseline_id	ID of the baseline for which the statistics need to be retrieved. Specifying the value 0 returns the statistics for the moving window baseline.
l_begin_snap	Start snapshot sequence number for the baseline.
l_end_snap	End snapshot sequence number for the baseline.
l_dbid	Database identifier for the baseline. If its value is set to NULL, then the database identifier for the local database is used. Its default value is NULL.

217.3.36 SELECT_BASELINE_METRIC Function

This table function shows the values of the metrics corresponding to a baseline for all the snapshots.

This table function returns an object of AWR_BASELINE_METRIC_TYPE Object Type.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.SELECT_BASELINE_METRIC(
   l_baseline_name     IN VARCHAR2,
   l_dbid              IN NUMBER DEFAULT NULL,
   l_instance_num      IN NUMBER DEFAULT NULL)
 RETURN awr_metric_type_table PIPELINED;
```

Parameters

Table 217-41 SELECT_BASELINE_METRIC Function Parameters

Parameter	Description
l_baseline_name	Name of the baseline for which the metrics need to be viewed.
l_dbid	Database identifier for the baseline. If set to NULL, the database identifier for the local database is used. Default is NULL.
l_instance_num	The instance number for which the metrics need to be viewed. If set to NULL, metrics for the local database instance are shown. Default is NULL.

217.3.37 UNREGISTER_REMOTE_DATABASE Procedures

This procedure removes all the statistics, metadata, partitions, and so on of a remote database from the Automatic Workload Repository (AWR). After executing this procedure, the remote database cannot be used for any AWR operations, such as creating remote snapshots.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.UNREGISTER_REMOTE_DATABASE(
   node_id        IN NUMBER,
   remote_check   IN BOOLEAN  DEFAULT TRUE);
```

```
DBMS_WORKLOAD_REPOSITORY.UNREGISTER_REMOTE_DATABASE(
   node_name      IN VARCHAR2,
   topology_name  IN VARCHAR2 DEFAULT NULL,
   remote_check   IN BOOLEAN  DEFAULT TRUE);
```

Parameters

Table 217-42 UNREGISTER_REMOTE_DATABASE Parameters

Parameter	Description
node_id	Identifier of the remote database whose data needs to be removed from the AWR.
node_name	Name of the remote database whose data needs to be removed from the AWR.
topology_name	RMF topology name of the remote database.
remote_check	If set to TRUE, the remote database's metadata is validated before removing its data from the AWR. This option requires the remote database to be available. If set to FALSE, the remote database's data is removed from the AWR without validating its metadata. This option must be selected to unregister a remote database that is not available (it is offline or there is a network outage).

Examples

This example removes the AWR data related to the remote database having the database identifier of 10:

```
EXECUTE DBMS_WORKLOAD_REPOSITORY.UNREGISTER_REMOTE_DATABASE(10);
```

217.3.38 UPDATE_DATAFILE_INFO Procedure

This procedure updates the data file and tablespace information stored in the Automatic Workload Repository (AWR) with the current information in the database. This procedure is useful when a data file or a tablespace has been moved or renamed. As this change is not always captured in the next snapshot, AWR report may not show the most current information.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.UPDATE_DATAFILE_INFO();
```

217.3.39 UPDATE_OBJECT_INFO Procedure

This procedure updates rows of WRH$_SEG_STAT_OBJ table that represent objects in the local database. It attempts to determine the current names for all object belonging to the local database, except those with 'MISSING' and/or 'TRANSIENT' values in the name columns.

The amount of work performed at each invocation of this routine may be controlled by setting the input parameter.

Syntax

```
DBMS_WORKLOAD_REPOSITORY.UPDATE_OBJECT_INFO(
   maxrows    IN  NUMBER  DEFAULT 0);
```

Parameters

Table 217-43 UPDATE_OBJECT_INFO Procedure Parameters

Parameter	Description
`maxrows`	Maximum number of rows to be updated. Default= `0`, meaning there is no limit.

218
DBMS_XDB

The `DBMS_XDB` package supports the following features:

> **Note:**
>
> The Oracle XML DB Repository is deprecated with Oracle Database 23ai.
>
> Oracle recommends that you replace any functionality used in XML DB Repository with alternative technologies.

The `DBMS_XDB` package supports the following features:

- Resource Management subprograms which complement Resource Views
- The Access Control List (ACL)-based Security Mechanism
- Configuration Session Management
- Creation of the XDB username

This chapter contains the following topics:

- DBMS_XDB Overview
- DBMS_XDB Security Model
- DBMS_XDB Constants
- Summary of DBMS_XDB Subprograms

> **See Also:**
>
> - *Oracle XML DB Developer's Guide*

218.1 DBMS_XDB Overview

`DBMS_XDB` provides a range of management functions.

The `DBMS_XDB` package supports the following features:

- The Resource Management functionality provides LINK Procedures, EXISTSRESOURCE Function, LOCKRESOURCE Function, GETLOCKTOKEN Procedure, UNLOCKRESOURCE Function, CREATERESOURCE Functions, RENAMERESOURCE Procedure, DELETERESOURCE Procedure, GETRESOID Function, CREATEOIDPATH Function, and CREATEFOLDER Function subprograms which complement Resource Views.
- The Access Control List (ACL)-based Security Mechanism can be used with in-hierarchy ACLs stored by the database or in-memory ACLs that may be stored outside the database.

Some of these methods can be used for both Oracle resources and arbitrary database objects. Use CHECKPRIVILEGES Function, GETACLDOCUMENT Function, and CHANGEPRIVILEGES Function for Oracle Resources. ACLCHECKPRIVILEGES Function provides access to Oracle's ACL-based Security mechanism without storing objects in the Hierarchy.

- The XDB username is created during XDB installation. This user owns a set of default tables and packages. GETXDB_TABLESPACE Function and the `DBMS_XDB_ADMIN.MOVEXDB_TABLESPACE` Procedure enable movement of schemas to a specified tablespace, and support the default `SYSAUX` tablespace introduction

218.2 DBMS_XDB Security Model

Owned by `XDB`, the `DBMS_XDB` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`. Subprograms in this package are executed using the privileges of the current user. Subprograms that operate on the XDB Configuration will succeed only if the current user is `SYS` or `XDB`, or the current user has the `XDBADMIN` or `DBA` role.

218.3 DBMS_XDB Constants

All constants described in the following table are deprecated in Oracle Release 12*c*.

They are relocated to either the DBMS_XDB_CONFIG package or the DBMS_XDB_REPOS package. The specifics of transference in each case are detailed in the **Relocated** column.

Oracle recommends that you do not use constants in their `DBMS_XDB` context in new applications. Support for deprecated features is for backward compatibility only and may be terminated in future releases.

Table 218-1 DBMS_XDB Constants

Constant	Type	Value	Description	Relocated
DELETE_RESOURCE	NUMBER	1	Deletes a resource; fails if the resource has children.	DBMS_XDB_REPOS
DELETE_RECURSIVE	NUMBER	2	Deletes a resource and its children, if any.	DBMS_XDB_REPOS
DELETE_FORCE	NUMBER	3	Deletes the resource, even if the object it contains is invalid.	DBMS_XDB_REPOS
DELETE_RECURSIVE_FORCE	NUMBER	4	Deletes a resource and its children, if any, even if the object it contains is invalid.	DBMS_XDB_REPOS
DELETE_RES_METADATA_CASCADE	NUMBER	1	Deletes the corresponding row in the metadata table	DBMS_XDB_REPOS
DELETE_RES_METADATA_NOCASCADE	NUMBER	2	Does not delete the row in the metadata table	DBMS_XDB_REPOS
DEFAULT_LOCK_TIMEOUT CONSTANT	PLS_INTEGER	(60*60)	Default time (in seconds) after which lock will expire	DBMS_XDB_REPOS
LINK_TYPE_HARD	NUMBER	1	Type of link to be created (default)	DBMS_XDB_REPOS
LINK_TYPE_WEAK	NUMBER	2	Type of link to be created	DBMS_XDB_REPOS

Table 218-1 (Cont.) DBMS_XDB Constants

Constant	Type	Value	Description	Relocated
LINK_TYPE_SYMBOLIC	NUMBER	3	Type of link to be created	DBMS_XDB_REPOS
ON_DENY_NEXT_CUSTOM	NUMBER	1	If access denied, the next custom authorization is tried	DBMS_XDB_CONFIG
ON_DENY_BASIC	NUMBER	2	If access denied, basic authentication is used	DBMS_XDB_CONFIG

218.4 Summary of DBMS_XDB Subprograms

This table lists the DBMS_XDB subprograms and briefly describes them.

Table 218-2 DBMS_XDB Package Subprograms

Subprogram	Description
ACLCHECKPRIVILEGES Function	Checks access privileges granted to the current user by specified ACL document on a resource whose owner is specified by the 'owner' parameter.
APPENDRESOURCEMETADATA Procedure	Takes in user-defined metadata either as a REF to XMLTYPE or an XMLTYPE and adds it to the desired resource
CHANGEOWNER Procedure	Changes the owner of the resource/s to the specified owner.
CHANGEPRIVILEGES Function	Adds a specified ACE to a specified resource's ACL
CHECKPRIVILEGES Function	Checks access privileges granted to the current user on the specified resource
CREATEFOLDER Function	Creates a new folder resource in the hierarchy
CREATEOIDPATH Function	Creates a virtual path to the resource based on object ID
CREATERESOURCE Functions	Creates a new resource
DELETERESOURCE Procedure	Deletes a resource from the hierarchy
DELETERESOURCEMETADATA Procedures	Deletes metadata from a resource (can be used for schema-based or nonschema-based metadata)
ENABLEDIGESTAUTHENTICATION Procedure	Enables digest authentication
EXISTSRESOURCE Function	Determines if a resource is the hierarchy, based on its absolute path
GETACLDOCUMENT Function	Retrieves ACL document that protects resource given its path name
GETCONTENTBLOB Function	Retrieves the contents of a resource returned as a BLOB
GETCONTENTCLOB Function	Retrieves the contents of a resource returned as a CLOB
GETCONTENTVARCHAR2 Function	Retrieves the contents of a resource returned as a string
GETCONTENTXMLREF Function	Retrieves the contents of a resource returned as a a REF to an XMLTYPE

Table 218-2 (Cont.) DBMS_XDB Package Subprograms

Subprogram	Description
GETCONTENTXMLTYPE Function	Retrieves the contents of a resource returned as an `XMLTYPE`
GETHTTPREQUESTHEADER Function	Gets the values of the passed header
GETLOCKTOKEN Procedure	Returns that resource's lock token for the current user given a path to a resource
GETPRIVILEGES Function	Gets all privileges granted to the current user on a specified resource
GETRESOID Function	Returns the object ID of the resource from its absolute path
GETXDB_TABLESPACE Function	Returns the current tablespace of the XDB (user)
HASBLOBCONTENT Function	Returns `TRUE` if the resource has BLOB content
HASCHARCONTENT Function	Returns `TRUE` if the resource has character content
HASXMLCONTENT Function	Returns `TRUE` if the resource has XML content
HASXMLREFERENCE Function	Returns `TRUE` if the resource has `REF` to XML content
ISFOLDER Function	Returns `TRUE` if the resource is a folder or container
LINK Procedures	Creates a link to an existing resource
LOCKRESOURCE Function	Gets a WebDAV-style lock on that resource given a path to that resource
PROCESSLINKS Procedure	Processes document links in the specified resource
PURGERESOURCEMETADATA Procedure	Deletes all user metadata from a resource
RENAMERESOURCE Procedure	Renames the XDB resource
SETACL Procedure	Sets the ACL on a specified resource
SPLITPATH Procedure	Splits the path into a parentpath and childpath
TOUCHRESOURCE Procedure	Changes the modification time of the resource to the current time
UNLOCKRESOURCE Function	Unlocks the resource given a lock token and resource path
UPDATERESOURCEMETADATA Procedures	Updates metadata for a resource

218.4.1 ACLCHECKPRIVILEGES Function

This function checks access privileges granted to the current user by specified ACL document by the `OWNER` of the resource. Returns positive integer if all privileges are granted.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the ACLCHECKPRIVILEGES Function.

Syntax

```
DBMS_XDB.ACLCHECKPRIVILEGES(
   acl_path  IN  VARCHAR2,
   owner     IN  VARCHAR2,
   privs     IN  xmltype)
 RETURN PLS_INTEGER;
```

Parameters

Table 218-3 *ACLCHECKPRIVILEGES Function Parameters*

Parameter	Description
`acl_path`	Absolute path in the Hierarchy for ACL document
`owner`	Resource owner name; the pseudo user "DAV:owner" is replaced by this user during ACL privilege resolution
`privs`	An `XMLType` instance of the privilege element specifying the requested set of access privileges. See description for CHECKPRIVILEGES Function.

218.4.2 APPENDRESOURCEMETADATA Procedure

This procedure takes in user-defined metadata either as a `REF` to `XMLTYPE` or an `XMLTYPE` and adds it to the desired resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the APPENDRESOURCEMETADATA Procedure.

Syntax

```
DBMS_XDB.APPENDRESOURCEMETADATA (
 abspath   IN  VARCHAR2,
 metadata  IN  XMLTYPE);

DBMS_XDB.APPENDRESOURCEMETADATA (
 abspath   IN  VARCHAR2,
 metadata  IN  REF SYS.XMLTYPE);
```

Parameters

Table 218-4 APPENDRESOURCEMETADATA Procedure

Parameter	Description
`abspath`	Absolute path of the resource
`metadata`	Metadata can be schema based or nonschema-based. Schema-based metadata is stored in its own table.

Usage Notes

- In the case in which a `REF` is passed in, the procedure stores the `REF` in the resource, and the metadata is stored in a separate table. In this case you are responsible for populating the `RESID` column for the metadata table. Note that the`REF` passed in must be unique. In other words, there must not be a`REF` with the same value in the resource metadata, as this would violate uniqueness of properties. An error is thrown if users attempt to add a `REF` that already exists.
- In the case where the XMLTYPE is passed in, the data is parsed to determine if it is schema-based or not and stored accordingly.

218.4.3 CHANGEOWNER Procedure

This procedure changes the owner of the resource/s to the specified owner. This procedure is deprecated in Release 12*c*.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the CHANGEOWNER Procedure.

Syntax

```
DBMS_XDB.CHANGEOWNER(
     abspath    IN   VARCHAR2,
     owner      IN   VARCHAR2,
     recurse    IN   BOOLEAN := FALSE);
```

Parameters

Table 218-5 CHANGEOWNER Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource
`owner`	New owner for the resource
`recurse`	If `TRUE`, recursively change owner of all resources in the folder tree

218.4.4 CHANGEPRIVILEGES Function

This function adds a specified ACE to a specified resource's ACL. This procedure is deprecated in Release 12*c*.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the CHANGEPRIVILEGES Function.

Syntax

```
DBMS_XDB.CHANGEPRIVILEGES(
   res_path    IN    VARCHAR2,
   ace         IN    xmltype)
 RETURN PLS_INTEGER;
```

Parameters

Table 218-6 CHANGEPRIVILEGES Function Parameters

Parameter	Description
res_path	Path name of the resource for which privileges need to be changed
ace	An XMLType instance of the <ace> element which specifies the <principal>, the operation <grant> and the list of privileges

Return Values

A positive integer if the ACL was successfully modified.

Usage Notes

If no ACE with the same principal and the same operation (grant/deny) already exists in the ACL, the new ACE is added at the end of the ACL.

218.4.5 CHECKPRIVILEGES Function

This function checks access privileges granted to the current user on the specified resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the CHECKPRIVILEGES Function.

Syntax

```
DBMS_XDB.CHECKPRIVILEGES(
   res_path    IN  VARCHAR2,
   privs       IN  xmltype)
 RETURN PLS_INTEGER;
```

Parameters

Table 218-7 CHECKPRIVILEGES Function Parameters

Parameter	Description
`res_path`	Absolute path in the Hierarchy for resource
`privs`	An `XMLType` instance of the privilege element specifying the requested set of access privileges

Return Values

A positive integer if all requested privileges granted.

218.4.6 CREATEFOLDER Function

This deprecated function creates a new folder resource in the hierarchy.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the CREATEFOLDER Function.

Syntax

```
DBMS_XDB.CREATEFOLDER(
   path    IN  VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-8 CREATEFOLDER Function Parameters

Parameter	Description
`path`	Path name for the new folder

Return Values

`TRUE` if operation successful; `FALSE`, otherwise.

Usage Notes

The given path name's parent folder must already exist in the hierarchy: if `'/folder1/folder2'` is passed as the path parameter, then `'/folder1'` must already exist.

218.4.7 CREATEOIDPATH Function

This deprecated function creates a virtual path to the resource based on object ID.

> **Note:**
>
> This funtion is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the CREATEOIDPATH Function.

Syntax

```
DBMS_XDB.CREATEOIDPATH(
   oid    IN   RAW)
 RETURN VARCHAR2;
```

Parameters

Table 218-9 CREATEOIDPATH Function Parameters

Parameter	Description
`oid`	Object ID of the resource

218.4.8 CREATERESOURCE Functions

The deprecated function creates a new resource. The description of the overload options precede each version of the syntax

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the CREATERESOURCE Functions.

Syntax

Given a `REF` to an existing `XMLType` row, creates a resource whose contents point to that row. That row should not already exist inside another resource:

```
DBMS_XDB.CREATERESOURCE(
     abspath        IN  VARCHAR2,
     datarow        IN  REF SYS.XMLTYPE,
     createfolders  IN  BOOLEAN := FALSE)
 RETURN BOOLEAN;
```

Creates a resource with a specified `BLOB` as its contents, and specifies character set of the source `BLOB`:

```
DBMS_XDB.CREATERESOURCE(
     abspath        IN  VARCHAR2,
```

```
    data           IN  BLOB,
    csid           IN  NUMBER :=0,
    createfolders  IN  BOOLEAN := FALSE)
 RETURN BOOLEAN;
```

Creates a resource with a specified `BFILE` as its contents, and specifies character set of the source `BFILE`:

```
DBMS_XDB.CREATERESOURCE (
     abspath        IN  VARCHAR2,
     data           IN  BFILE,
     csid           IN  NUMBER :=0,
     createfolders  IN  BOOLEAN := FALSE)
 RETURN BOOLEAN;
```

Creates a resource with a specified `CLOB` as its contents:

```
DBMS_XDB.CREATERESOURCE (
     abspath        IN  VARCHAR2,
     data           IN  CLOB,
     createfolders  IN  BOOLEAN := FALSE)
 RETURN BOOLEAN;
```

Given a string, inserts a new resource into the hierarchy with the string as the contents:

```
DBMS_XDB.CREATERESOURCE (
   abspath        IN  VARCHAR2,
   data           IN  VARCHAR2,
   schemaurl      IN  VARCHAR2 := NULL,
   elem           IN  VARCHAR2 := NULL)
 RETURN BOOLEAN;
```

Given an `XMLTYPE` and a schema URL, inserts a new resource into the hierarchy with the `XMLTYPE` as the contents:

```
DBMS_XDB.CREATERESOURCE (
   abspath        IN  VARCHAR2,
   data           IN  SYS.XMLTYPE,
   schemaurl      IN  VARCHAR2 := NULL,
   elem           IN  VARCHAR2 := NULL)
 RETURN BOOLEAN;
```

Parameters

Table 218-10 CREATERESOURCE Function Parameters

Parameter	Description
`abspath`	Absolute path of the resource to create. The path name's parent folder must already exist in the hierarchy. In other words, if `/foo/bar.txt` is passed in, then folder `/foo` must already exist.
`data`	String buffer containing new resource's contents. The data is parsed to check if it contains a schema-based XML document, and the contents are stored as schema-based in the schema's default table. Otherwise, it is saved as binary data.
`datarow`	`REF` to an `XMLType` row to be used as the contents

Table 218-10 (Cont.) CREATERESOURCE Function Parameters

Parameter	Description
csid	Character set id of the document. Must be a valid Oracle ID; otherwise returns an error. If CSID is not specified, or if a zero CSID is specified, then the character set id of the document is determined as follows: • From the abspath extension, determine the resource's MIME type. • If the MIME type is */xml, then the encoding is detected based on Appendix F of the W3C XML 1.0 Reference at http://www.w3.org/TR/2000/REC-xml-20001006; • Otherwise, it is defaulted to the database character set.
createfolders	If TRUE, create the parent folders if they do not exist
schemaurl	For XML data, schema URL data conforms to (default NULL)
elem	Element name (default NULL)

Return Values

TRUE if operation successful; FALSE, otherwise.

218.4.9 DELETERESOURCE Procedure

This **deprecated procedure** deletes a resource from the hierarchy.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the DELETERESOURCE Procedure.

Syntax

```
DBMS_XDB.DELETERESOURCE(
   path          IN      VARCHAR2,
   delete_option IN      PLS_INTEGER);
```

Parameters

Table 218-11 DELETERESOURCE Procedure Parameters

Parameter	Description
path	Path name of the resource to delete
delete_option	The option that controls how a a resource is deleted; defined in Table 218-1: • DELETE_RESOURCE • DELETE_RECURSIVE • DELETE_FORCE • DELETE_RECURSIVE_FORCE

218.4.10 DELETERESOURCEMETADATA Procedures

This **deprecated procedure** takes in a resource by absolute path and removes either the schema-based metadata identified by the REF, or the metadata identified by the namespace and name combination, which can be either schema-based or non-schema based. It also takes an additional (optional) parameter that specifies how to delete it. This parameter is only relevant for schema-based resource metadata that needs to be deleted. For non-schema based metadata, this parameter is ignored.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the DELETERESOURCEMETADATA Procedures.

Syntax

Can be used only for schema-based metadata:

```
DBMS_XDB.DELETERESOURCEMETADATA (
   abspath        IN  VARCHAR2,
   metadata       IN  REF SYS.XMLTYPE,
   delete_option  IN  pls_integer := dbms_xdb.DELETE_RESOURCE_METADATA_CASCADE);
```

Can be used for schema-based or nonschema-based metadata:

```
DBMS_XDB.DELETERESOURCEMETADATA (
   abspath        IN VARCHAR2,
   metadatans     IN VARCHAR2,
   metadataname   IN VARCHAR2,
   delete_option  IN pls_integer := dbms_xdb.DELETE_RESOURCE_METADATA_CASCADE);
```

Parameters

Table 218-12 DELETERESOURCEMETADATA Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource
`metadata`	`REF` to the piece of metadata (schema based) to be deleted
`mettadatans`	Namespace of the metadata fragment to be removed
`mettadataname`	Local name of the metadata fragment to be removed
`delete_option`	Only applicable for schema-based metadata, this can be one of the following: • `DELETE_RES_METADATA_CASCADE` - deletes the corresponding row in the metadata table • `DELETE_RES_METADATA_NOCASCADE` - does not delete the row in the metadata table

218.4.11 ENABLEDIGESTAUTHENTICATION Procedure

This deprecated procedure enabling digest authentication. It will list digest as the first authentication mechanism to be used by the XML DB HTTP server.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_CONFIG package - the ENABLEDIGESTAUTHENTICATION Procedure.

Syntax

```
DBMS_XDB.ENABLEDIGESTAUTHENTICATION;
```

218.4.12 EXISTSRESOURCE Function

This deprecated function indicates if a resource is in the hierarchy. Matches resource by a string that represents its absolute path.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the EXISTSRESOURCE Function.

Syntax

```
DBMS_XDB.EXISTSRESOURCE(
   abspath    IN    VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-13 EXISTSRESOURCE Function Parameters

Parameter	Description
abspath	Path name of the resource whose ACL document is required

Return Values

TRUE if the resource is found.

218.4.13 GETACLDOCUMENT Function

This deprecated function retrieves ACL document that protects resource given its path name.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETACLDOCUMENT Function.

Syntax

```
DBMS_XDB.GETACLDOCUMENT(
   abspath  IN  VARCHAR2)
  RETURN sys.xmltype;
```

Parameters

Table 218-14 GETACLDOCUMENT Function Parameters

Parameter	Description
`abspath`	Path name of the resource whose ACL document is required

Return Values

The `XMLType` for ACL document.

218.4.14 GETCONTENTBLOB Function

This deprecated function retrieves the contents of a resource returned as a `BLOB`.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETCONTENTBLOB Function.

Syntax

```
DBMS_XDB.GETCONTENTBLOB(
     abspath     IN      VARCHAR2,
     csid        OUT     PLS_INTEGER,
     locksrc     IN      BOOLEAN := FALSE)
  RETURN BLOB;
```

Parameters

Table 218-15 GETCONTENTBLOB Function Parameters

Parameter	Description
abspath	Absolute path of the resource
csid	If TRUE, lock and return the source LOB. If FALSE, return a temp LOB copy.
locksrc	Contents of the resource as a BLOB

Return Values

The contents of the resource as a BLOB.

218.4.15 GETCONTENTCLOB Function

This deprecated function gets the contents of a resource returned as a CLOB.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETCONTENTCLOB Function.

Syntax

```
DBMS_XDB.GETCONTENTCLOB(
     abspath    IN     VARCHAR2,
  RETURN CLOB;
```

Parameters

Table 218-16 GETCONTENTCLOB Function Parameters

Parameter	Description
abspath	Absolute path of the resource

Return Values

The contents of the resource as a CLOB.

218.4.16 GETCONTENTVARCHAR2 Function

This deprecated function gets the contents of a resource returned as a string.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETCONTENTVARCHAR2 Function.

Syntax

```
DBMS_XDB.GETCONTENTVARCHAR2(
     abspath     IN      VARCHAR2,
  RETURN BLOB;
```

Parameters

Table 218-17 GETCONTENTVARCHAR2 Function Parameters

Parameter	Description
`abspath`	Absolute path of the resource

Return Values

The contents of the resource as a string.

218.4.17 GETCONTENTXMLREF Function

This deprecated function retrieves the contents of a resource returned as a a `REF` to an `XMLTYPE`.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETCONTENTXMLREF Function.

Syntax

```
DBMS_XDB.GETCONTENTXMLREF(
     abspath     IN      VARCHAR2,
  RETURN SYS.XMLTYPE;
```

Parameters

Table 218-18 GETCONTENTXMLREF Function Parameters

Parameter	Description
abspath	Absolute path of the resource

Return Values

The contents of the resource as a REF to an XMLTYPE.

218.4.18 GETCONTENTXMLTYPE Function

This deprecated function retrieves the contents of a resource returned as an XMLTYPE.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETCONTENTXMLTYPE Function.

Syntax

```
DBMS_XDB.GETCONTENTXMLTYPE(
     abspath    IN     VARCHAR2,
  RETURN SYS.XMLTYPE;
```

Parameters

Table 218-19 GETCONTENTXMLTYPE Function Parameters

Parameter	Description
abspath	Absolute path of the resource

Return Values

The contents of the resource as an XMLTYPE.

218.4.19 GETHTTPREQUESTHEADER Function

This deprecated function, if called during an HTTP request serviced by XDB, returns the values of the passed header. It is used by routines that implement custom authentication.

> **Note:**
>
> This procedure is deprecated in Release 12*c*.

Syntax

```
DBMS_XDB.GETHTTPREQUESTHEADER(
    header_name    IN   VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 218-20 GETHTTPREQUESTHEADER Function Parameters

Parameter	Description
header_name	Passed header

Return Values

Returns NULL in case the header is not present in the request, or for AUTHENTICATION, for security reasons.

218.4.20 GETLOCKTOKEN Procedure

Given a path to a resource, this deprecated procedure returns that resource's lock token for the current user.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETLOCKTOKEN Procedure.

Syntax

```
DBMS_XDB.GETLOCKTOKEN(
   path         IN      VARCHAR2,
   locktoken    OUT     VARCHAR2);
```

Parameters

Table 218-21 GETLOCKTOKEN Procedure Parameters

Parameter	Description
path	Path name to the resource
locktoken	Logged-in user's lock token for the resource

Usage Notes

The user must have READPROPERTIES privilege on the resource.

218.4.21 GETPRIVILEGES Function

This deprecated function gets all privileges granted to the current user on a specified resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETPRIVILEGES Function.

Syntax

```
DBMS_XDB.GETPRIVILEGES(
   res_path    IN     VARCHAR2)
 RETURN sys.xmltype;
```

Parameters

Table 218-22 GETPRIVILEGES Function Parameters

Parameter	Description
res_path	Absolute path in the hierarchy of the resource

Return Values

An XMLType instance of <privilege> element, which contains the list of all leaf privileges granted on this resource to the current user.

218.4.22 GETRESOID Function

This deprecated procedure returns the object ID of the resource from its absolute path.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETRESOID Function.

Syntax

```
DBMS_XDB.GETRESOID(
   abspath IN VARCHAR2)
RETURN RAW;
```

Parameters

Table 218-23 GETRESOID Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

NULL if the resource is not present.

218.4.23 GETXDB_TABLESPACE Function

This deprecated function returns the current tablespace of the XDB (user).

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the GETXDB_TABLESPACE Function.

Syntax

```
DBMS_XDB.GETXDB_TABLESPACE
 RETURN VARCHAR2;
```

218.4.24 HASBLOBCONTENT Function

This deprecated function returns TRUE if the resource has BLOB content.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the HASBLOBCONTENT Function.

Syntax

```
DBMS_XDB.HASBLOBCONTENT
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-24 HASBLOBCONTENT Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

`TRUE` if the resource has BOB content.

218.4.25 HASCHARCONTENT Function

This deprecated function returns `TRUE` if the resource has character content.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the HASCHARCONTENT Function.

Syntax

```
DBMS_XDB.HASCHARCONTENT
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-25 HASCHARCONTENT Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

`TRUE` if the resource has character content.

218.4.26 HASXMLCONTENT Function

This deprecated function returns `TRUE` if the resource has XML content.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the HASXMLCONTENT Function.

Syntax

```
DBMS_XDB.HASXMLCONTENT
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-26 HASXMLCONTENT Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

TRUE if the resource has XML content.

218.4.27 HASXMLREFERENCE Function

This deprecated function returns TRUE if the resource has a REF to XML content.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the HASXMLREFERENCE Function.

Syntax

```
DBMS_XDB.HASXMLREFERENCE
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-27 HASXMLREFERENCE Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

TRUE resource has a REF to XML content.

218.4.28 ISFOLDER Function

This deprecated function returns TRUE if the resource is a folder or container.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the ISFOLDER Function.

Syntax

```
DBMS_XDB.ISFOLDER
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-28 DBMS_XDB.ISFOLDER Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

TRUE if the resource is a folder or container.

218.4.29 LINK Procedures

This deprecated procedure creates a link from a specified folder to a specified resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the LINK Procedures.

Syntax

```
DBMS_XDB.LINK(
   srcpath      IN   VARCHAR2,
   linkfolder   IN   VARCHAR2,
   linkname     IN   VARCHAR2);

DBMS_XDB.LINK(
   srcpath      IN   VARCHAR2,
   linkfolder   IN   VARCHAR2,
   linkname     IN   VARCHAR2,
   linktype     IN   PLS_INTEGER := DBMS_XDB.LINK_TYPE_HARD);
```

Parameters

Table 218-29 LINK Procedure Parameters

Parameter	Description
srcpath	Path name of the resource to which a link is created
linkfolder	Folder in which the new link is placed
linkname	Name of the new link

Table 218-29 (Cont.) LINK Procedure Parameters

Parameter	Description
`linktype`	Type of link to be created: • `DBMS_XDB.LINK_TYPE_HARD` (default) • `DBMS_XDB.LINK_TYPE_WEAK` • `DBMS_XDB.LINK_TYPE_SYMBOLIC`

218.4.30 LOCKRESOURCE Function

Given a path to a resource, this deprecated function gets a WebDAV-style lock on that resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the LOCKRESOURCE Function.

Syntax

```
DBMS_XDB.LOCKRESOURCE(
   path      IN  VARCHAR2,
   depthzero IN  BOOLEAN,
   shared    IN  boolean)
RETURN BOOLEAN;
```

Parameters

Table 218-30 LOCKRESOURCE Function Parameters

Parameter	Description
`path`	Path name of the resource to lock.
`depthzero`	Currently not supported
`shared`	Passing `TRUE` obtains a shared write lock

Return Values

`TRUE` if successful.

Usage Notes

The user must have `UPDATE` privileges on the resource.

218.4.31 PROCESSLINKS Procedure

This deprecated procedure processes document links in the specified resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the PROCESSLINKS Procedure.

Syntax

```
DBMS_XDB.PURGERESOURCEMETADATA(
 abspath  IN  VARCHAR2,
 recurse  IN  BOOLEAN := FALSE);
```

Parameters

Table 218-31 PROCESSLINKS Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource. If the path is a folder, use the `recurse` flag.
`recurse`	Used only if `abspath` specifies a folder. If `TRUE`, process links of all resources in the folder hierarchy rooted at the specified resource. If `FALSE`, process links of all documents in this folder only.

218.4.32 PURGERESOURCEMETADATA Procedure

This deprecated procedure deletes all user metadata from a resource. Schema-based metadata is removed in cascade mode, rows being deleted from the corresponding metadata tables.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the PURGERESOURCEMETADATA Procedure.

Syntax

```
DBMS_XDB.PURGERESOURCEMETADATA(
 abspath  IN  VARCHAR2);
```

Parameters

Table 218-32 PURGERESOURCEMETADATA Procedure Parameters

Parameter	Description
abspath	Absolute path of the resource

218.4.33 RENAMERESOURCE Procedure

This deprecated procedure renames the XDB resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the RENAMERESOURCE Procedure.

Syntax

```
DBMS_XDB.RENAMERESOURCE(
   srcpath    IN  VARCHAR2,
   destfolder IN  CARCHAR2,
   newname    IN  VARCHAR2);
```

Parameters

Table 218-33 RENAMERESOURCE Procedure Parameters

Parameter	Description
srcpath	Absolute path in the Hierarchy for the source resource destination folder
destfolder	Absolute path in the Hierarchy for the destination folder
newname	Name of the child in the destination folder

218.4.34 SETACL Procedure

This deprecated procedure set the ACL on a specified resource to be the ACL specified by path.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the SETACL Procedure.

Syntax

```
DBMS_XDB.SETACL(
   res_path   IN  VARCHAR2,
   acl_path   IN  VARCHAR2);
```

Parameters

Table 218-34 SETACL Procedure Parameters

Parameter	Description
res_path	Absolute path in the Hierarchy for resource
acl_path	Absolute path in the Hierarchy for ACL

Usage Notes

The user must have <write-acl> privileges on the resource.

218.4.35 SPLITPATH Procedure

This deprecated procedure splits the path into a parentpath and childpath.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the SPLITPATH Procedure.

Syntax

```
DBMS_XDB.SPLITPATH(
      abspath     IN  VARCHAR2,
      parentpath  OUT VARCHAR2,
     childpath    OUT VARCHAR2);
```

Parameters

Table 218-35 SPLITPATH Procedure Parameters

Parameter	Description
abspath	Absolute path to be split
parentpath	Parentpath
childpath	Childpath

218.4.36 TOUCHRESOURCE Procedure

This **deprecated procedure** changes the modification time of the resource to the current time.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the TOUCHRESOURCE Procedure.

Syntax

```
DBMS_XDB.TOUCHRESOURCE
     abspath    IN     VARCHAR2);
```

Parameters

Table 218-36 DBMS_XDB.TOUCHRESOURCE Procedure Parameters

Parameter	Description
abspath_path	Absolute path of the resource

218.4.37 UNLOCKRESOURCE Function

This deprecated function unlocks the resource given a lock token and a path to the resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the UNLOCKRESOURCE Function.

Syntax

```
DBMS_XDB.UNLOCKRESOURCE(
   path     IN  VARCHAR2,
   deltoken IN  VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 218-37 UNLOCKRESOURCE Function Parameters

Parameter	Description
path	Path name to the resource
deltoken	Lock token to be removed

Return Values

`TRUE` if operation successful.

Usage Notes

The user must have `UPDATE` privileges on the resource.

218.4.38 UPDATERESOURCEMETADATA Procedures

This deprecated procedure updates metadata for a resource.

> **Note:**
>
> This procedure is deprecated in Release 12*c*. This functionality is replaced by a subprogram of the same name in the DBMS_XDB_REPOS package - the UPDATERESOURCEMETADATA Procedures.

The procedure takes in a resource identified by absolute path and the metadata in it to replace identified by its `REF`. It replaces that piece of metadata with user-defined metadata which is either in the form of a `REF` to `XMLTYPE` or an `XMLTYPE`.

Syntax

Can be used to update schema-based metadata only. The new metadata must be schema-based:

```
DBMS_XDB.UPDATERESOURCEMETADATA(
   abspath  IN VARCHAR2,
   oldmetadata    IN   REF SYS.XMLTYPE,
   newmetadata    IN   REF SYS.XMLTYPE)
```

Can be used to update schema-based metadata only. The new metadata must be schema-based or nonschema-based:

```
DBMS_XDB.UPDATERESOURCEMETADATA(
   abspath        IN   VARCHAR2,
   oldmetadata    IN   REF SYS.XMLTYPE,
   newmetadata    IN   XMLTYPE);
```

Can be used for both schema-based and nonschema-based metadata:

```
DBMS_XDB.UPDATERESOURCEMETADATA(
   abspath        IN   VARCHAR2,
   oldns          IN   VARCHAR2,
   oldname IN VARCHAR,
   newmetadata IN XMLTYPE);
```

Can be used for both schema-based or nonschema-based metadata. New metadata must be schema-based:

```
DBMS_XDB.UPDATERESOURCEMETADATA(
   abspath       IN   VARCHAR2,
```

```
   oldns          IN   VARCHAR2,
   oldname        IN   VARCHAR,
   newmetadata    IN   REF SYS.XMLTYPE);
```

Parameters

Table 218-38 UPDATERESOURCEMETADATA Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource
`oldmetadata`	`REF` to the old of metadata
`newmetadata`	`REF` to the new, replacement metadata (can be either schema-based or nonschema-based depending on the overload)
`oldns`	Namespace identifying old metadata
`oldname`	Local name identifying old metadata

Usage Notes

In the case of `REF`, it stores the `REF` in the resource and the metadata is stored in a separate table. Uniqueness of `REFs` is enforced. In the case where the `XMLTYPE` is passed in, data is parsed to determine if it is schema-based or not and is stored accordingly.

219
DBMS_XA

The `DBMS_XA` package contains the XA/Open interface for applications to call XA interface in PL/SQL. Using this package, application developers can switch or share transactions across SQL*Plus sessions or processes using PL/SQL.

The chapter contains the following topics:

- Overview
- Security Model
- Constants
- Operational Notes
- Data Structures
- Summary of DBMS_XA Subprograms

> **See Also:**
>
> *Oracle Database Advanced Application Developer's Guide* for more information about "Developing Applications with Oracle XA"

219.1 DBMS_XA Overview

These subprograms allow a PL/SQL application to define a global transaction branch ID (`XID`) and associate or disassociate the current session with the transaction branch.

Subsequently, these transaction branches may be prepared and committed by following the two-phase commit protocol. A single-phase commit protocol is also supported if only one resource manager is involved.

Interfaces are also provided for a PL/SQL application to set the timeout values for any new global transaction branches that may start with the current session.

219.2 DBMS_XA Security Model

This package is created under `SYS`. Operations provided by this package are performed under the current calling user, not under the package owner `SYS`. Any `DBMS_XA` subprogram called from an anonymous PL/SQL block is executed using the privileges of the current user. Any `DBMS_XA` subprogram called from a stored procedure is executed using the privileges of the owner of the stored procedure.

`SELECT` or `READ` privilege on `SYS.DBA_PENDING_TRANSACTIONS` is required for users who need to execute `XA_RECOVER` subprogram.

`FORCE ANY TRANSACTION` privilege is required for users who need to manipulate XA transactions created by other users.

219.3 DBMS_XA Constants

The DBMS_XA package defines several constants that can be used for specifying parameter values.

The package uses the constants shown in Table 219-1 for use in the flag field of the XA_START Function and the XA_END Function.

Table 219-1 DBMS_XA Constants for Flag Field of XA_START & XA_END Functions

Name	Type	Value	Description
TMNOFLAGS	PLS_INTEGER	00000000	Indicates no flag value is selected.
TMSUCCESS	PLS_INTEGER	UTL_RAW.CAST_TO_BINARY_INTEGER ('04000000')	Dissociates caller from transaction branch
TMJOIN	PLS_INTEGER	UTL_RAW.CAST_TO_BINARY_INTEGER ('00200000')	Caller is joining existing transaction branch.
TMSUSPEND	PLS_INTEGER	UTL_RAW.CAST_TO_BINARY_INTEGER ('02000000')	Caller is suspending, not ending, association
TMRESUME	PLS_INTEGER	UTL_RAW.CAST_TO_BINARY_INTEGER ('08000000')	Caller is resuming association with suspended transaction branch.

The DBMS_XA package uses the constants shown in Table 219-2 for Possible Return Values

Table 219-2 DBMS_XA Constants for Possible Return Values

Name	Type	Value	Description
XA_RBBASE	PLS_INTEGER	100	Inclusive lower bound of the rollback codes
XA_RBROLLBACK	PLS_INTEGER	XA_RBBASE	Rollback was caused by an unspecified reason
XA_RBCOMMFAIL	PLS_INTEGER	XA_RBBASE+1	Rollback was caused by a communication failure
XA_RBDEADLOCK	PLS_INTEGER	XA_RBBASE+2	Deadlock was detected
XA_RBINTEGRITY	PLS_INTEGER	XA_RBBASE+3	Condition that violates the integrity of the resources was detected
XA_RBOTHER	PLS_INTEGER	XA_RBBASE+4	Resource manager rolled back the transaction for an unlisted reason
XA_RBPROTO	PLS_INTEGER	XA_RBBASE+5	Protocol error occurred in the resource manager
XA_RBTIMEOUT	PLS_INTEGER	XA_RBBASE+6	transaction branch took long
XA_RBTRANSIENT	PLS_INTEGER	XA_RBBASE+7	May retry the transaction branch

Table 219-2 (Cont.) DBMS_XA Constants for Possible Return Values

Name	Type	Value	Description
XA_RBEND	PLS_INTEGER	XA_RBTRANSIENT	Inclusive upper bound of the rollback codes
XA_NOMIGRATE	PLS_INTEGER	9	Transaction branch may have been heuristically completed
XA_HEURHAZ	PLS_INTEGER	8	Transaction branch may have been heuristically completed
XA_HEURCOM	PLS_INTEGER	7	Transaction branch has been heuristically committed
XA_HEURRB	PLS_INTEGER	6	Transaction branch has been heuristically rolled back
XA_HEURMIX	PLS_INTEGER	5	Some of the transaction branches have been heuristically committed, others rolled back
XA_RETRY	PLS_INTEGER	4	Routine returned with no effect and may be re-issued
XA_RDONLY	PLS_INTEGER	3	Transaction was read-only and has been committed
XA_OK	PLS_INTEGER	0	Normal execution
XAER_ASYNC	PLS_INTEGER	-2	Asynchronous operation already outstanding
XAER_RMERR	PLS_INTEGER	-3	Resource manager error occurred in the transaction branch
XAER_NOTA	PLS_INTEGER	-4	XID is not valid
XAER_INVAL	PLS_INTEGER	-5	Invalid arguments were given
XAER_PROTO	PLS_INTEGER	-6	Routine invoked in an improper context
XAER_RMFAIL	PLS_INTEGER	-7	Resource manager unavailable
XAER_DUPID	PLS_INTEGER	-8	XID already exists
XAER_OUTSIDE	PLS_INTEGER	-9	Resource manager doing work outside global transaction

219.4 DBMS_XA Operational Notes

In compliance with the XA specification of the X/Open CAE Standard for Distributed Transaction Processing, XA_PREPARE/COMMIT/ ROLLBACK/FORGET may not be called when the transaction is still associated with the current session. Only after XA_END has been called so that there is not any transaction associated with the current session, the application may call XA_PREPARE/COMMIT/ ROLLBACK/FORGET.

XAER_PROTO error is returned from XA_PREPARE/COMMIT/ROLLBACK/FORGET if a transaction is being associated with the current session.

Prior to calling any of the package subprograms, a connection/session must have already been established to the Oracle database server backend, or a resource manager. Resource manager identifiers are not supported. If multiple resource managers are involved, multiple

connections/sessions must be pre-established to each resource manager before calling any the package subprograms. If multiple connections/sessions are established during the course of global transaction processing, the caller must ensure that all of those connections/sessions associated with a specific global transaction branch identifier (XID) are established to the same resource manager.

219.5 DBMS_XA Data Structures

The `DBMS_XA` package uses this `OBJECT` type and associated `TABLE` type.

OBJECT Types

- DBMS_XA_XID Object Type

TABLE Types

- DBMS_XA_XID_ARRAY Table Type

219.5.1 DBMS_XA DBMS_XA_XID Object Type

The PL/SQL XA interface allows the PL/SQL application to define a global transaction branch id (XID) and associate/disassociate the current session with the transaction branch. XID is defined as a PL/SQL object type.

> **Note:**
>
> For more information, see "Distributed Transaction Processing: The XA Specification" in the public XA Standard.

Syntax

```
TYPE DBMS_XA_XID IS OBJECT(
  formatid      NUMBER,
  gtrid         RAW(64),
  bqual         RAW(64),
  constructor function DBMS_XA_XID(
       gtrid     IN   NUMBER)
    RETURN SELF AS RESULT,
  constructor function DBMS_XA_XID (
       gtrid     IN   RAW,
       bqual     IN   RAW)
    RETURN SELF AS RESULT,
  constructor function DBMS_XA_XID(
       formatid  IN   NUMBER,
       gtrid     IN   RAW,
       bqual     IN   RAW DEFAULT HEXTORAW('00000000000000000000000000000001'))
    RETURN SELF AS RESULT)
```

Attributes

Table 219-3 DBMS_XA_XID Object Type

Attribute	Description
formatid	Format identifier, a number identifying different transaction managers (TM)
gtrid	Global transaction identifier uniquely identifying a global transaction, of which the maximum size is 64 bytes
bqual	Branch qualifier, of which the maximum size is 64 bytes

219.5.2 DBMS_XA DBMS_XA_XID_ARRAY Table Type

This type is used to define an array of xid that represent a list of global transaction branches.

Syntax

```
TYPE DBMS_XA_XID_ARRAY as TABLE of DBMS_XA_XID
```

219.6 Summary of DBMS_XA Subprograms

This table lists the DBMS_XA subprograms and briefly describes them.

Table 219-4 DBMS_XA Package Subprograms

Subprogram	Description
DIST_TXN_SYNC Procedure	Used in recovery of synchronization when utilizing Oracle Real Application Clusters (Oracle RAC)
XA_COMMIT Function	Commits the global transaction specified by xid
XA_END Function	Disassociates the current session from the transaction branch specified by xid
XA_FORGET Function	Informs the resource manager to forget about a heuristically committed or rolled back transaction branch.
XA_GETLASTOER Function	Obtains the last Oracle error code, in case of failure of previous XA calls.
XA_PREPARE Function	Prepares the transaction branch specified in xid for committing the transaction subsequently if possible
XA_RECOVER Function	Obtains a list of prepared or heuristically completed transaction branches from a resource manager
XA_ROLLBACK Function	Informs the resource manager to roll back work done on behalf of a transaction branch
XA_SETTIMEOUT Function	Sets the transaction timeout in seconds for the current session
XA_START Function	Associates the current session with the transaction branch specified by xid

219.6.1 DIST_TXN_SYNC Procedure

This procedure can be used to synchronize in-doubt transactions when one of the Oracle Real Application Clusters (Oracle RAC) instances fails.

Syntax

```
DBMS_XA.DIST_TXN_SYNC;
```

219.6.2 XA_COMMIT Function

This function commits the global transaction specified by `xid`.

Syntax

```
DBMS_XA.XA_COMMIT (
   xid       IN  DBMS_XA_XID,
   onePhase  IN  BOOLEAN)
RETURN PLS_INTEGER;
```

Parameters

Table 219-5 *XA_COMMIT Function Parameters*

Parameter	Description
`xid`	See DBMS_XA_XID Object Type
`onePhase`	If `TRUE`, apply single phase commit

Return Values

See Table 219-2. Possible return values indicating error are: `XAER_RMERR`, `XAER_RMFAIL`, `XAER_NOTA`, `XAER_INVAL`, or `XAER_PROTO`. Other possible return values include: `XA_OK`, `XA_RB*`, `XA_HEURHAZ`, `XA_HEURCOM`, `XA_HEURRB`, and `XA_HEURMIX`.

Usage Notes

- An application must not call `COMMIT`, but instead must call `XA_COMMIT` to commit the global transaction specified by `xid`. If a user needs to commit a transaction branch that is created by other users, `FORCE ANY TRANSACTION` must be granted to the user.
- If `onePhase` is `TRUE`, the resource manager should use a one-phase commit protocol to commit the work done on behalf of `xid`. Otherwise, only if all branches of the global transaction have been prepared successfully and the preceding `XA_PREPARE` call has returned `XA_OK`, should `XA_COMMIT` be called.
- The application must make a separate `XA_COMMIT` call for each of the transaction branches of the global transaction for which `XA_PREPARE` has returned `XA_OK`.
- If the resource manager did not commit the transaction and the parameter `onePhase` is set to `TRUE`, the resource manager may return one of the `XA_RB*` code. Upon return, the resource manager has rolled back the branch's work and has released all held resources.

219.6.3 XA_END Function

This function disassociates the current session from the transaction branch specified by `xid`.

A transaction manager calls `XA_END` when a thread of control finishes, or needs to suspend work on, a transaction branch. This occurs when the application completes a portion of its work, either partially or in its entirety (for example, before blocking on some event in order to let other threads of control work on the branch). When `XA_END` successfully returns, the calling thread of control is no longer actively associated with the branch but the branch still exists

Syntax

```
DBMS_XA.XA_END (
   xid   IN  DBMS_XA_XID,
   flag  IN  PLS_INTEGER)
RETURN PLS_INTEGER;
```

Parameters

Table 219-6 *XA_END Function Parameters*

Parameter	Description
`xid`	See DBMS_XA_XID Object Type
`flag`	See Table 219-1.

Return Values

See Table 219-2. Possible return values in error are `XAER_RMERR`, `XAER_RMFAILED`, `XAER_NOTA`, `XAER_INVAL`, `XAER_PROTO`, or `XA_RB*`.

Usage Notes

- `TMSUCCESS` or `TMSUSPEND` may be specified in flag, and the transaction branch is disassociated with the current session in detached state if the return value is `XA_OK`. `TMFAIL` is not supported. `XA_END` may be called with either `TMSUCCESS` or `TMSUSPEND` to disassociate the transaction branch identified by xid from the current session.
- `XA_OK` is returned if `XA_END` succeeds. An application must check the return value and handle error cases. Only when `XA_OK` is returned, the application should proceed for other normal operations.
- Executing a `ROLLBACK` statement without calling `XA_END` first will rollback the changes made by the current transaction. However, the transaction context is still associated with the current session until `XA_END` is called.
- Executing a `COMMIT` statement without calling `XA_END` first will result in `ORA-02089: COMMIT` is not allowed in a subordinate session.
- Executing a `COMMIT` or a `ROLLBACK` statement after `XA_END` has no effect on the transaction identified by `xid`, since this transaction is no longer associated with the current session.

219.6.4 XA_FORGET Function

This function informs the resource manager to forget about a heuristically committed or rolled back transaction branch.

Syntax

```
DBMS_XA.XA_FORGET (
   xid       IN  DBMS_XA_XID)
RETURN PLS_INTEGER;
```

Parameters

Table 219-7 *XA_FORGET Function Parameters*

Parameter	Description
xid	See DBMS_XA_XID Object Type

Return Values

See Table 219-2. Possible return values are XA_OK, XAER_RMERR, XAER_RMFAIL, XAER_NOTA, XAER_INVAL, or XAER_PROTO.

219.6.5 XA_GETLASTOER Function

This function obtains the last Oracle error code, in case of failure of previous XA calls.

Syntax

```
DBMS_XA.XA_GETLASTOER
 RETURN PLS_INTEGER;
```

Return Values

The return value carries the last Oracle error code.

219.6.6 XA_PREPARE Function

This function prepares the transaction branch specified in xid for committing the transaction subsequently if possible.

Syntax

```
DBMS_XA.XA_PREPARE (
   xid   IN  DBMS_XA_XID)
RETURN PLS_INTEGER;
```

Parameters

Table 219-8 *XA_PREPARE Function Parameters*

Parameter	Description
xid	See DBMS_XA_XID Object Type

Return Values

See Table 219-2. Possible return codes include: `XA_OK`, `XA_RDONLY`, `XA_RB*`, `XAER_RMERR`, `XAER_RMFAIL`, `XAER_NOTA`, `XAER_INVAL`, or `XAER_PROTO`.

Usage Notes

- If a user needs to prepare a transaction branch that is created by other users, `FORCE ANY TRANSACTION` must be granted to the user.
- An application must keep track of all the branches of one global transaction, and prepare each transaction branch. Only if all branches of the global transaction have been prepared successfully and `XA_PREPARE` has returned `XA_OK`, the application may proceed to call `XA_COMMIT`.

219.6.7 XA_RECOVER Function

This function obtains a list of prepared or heuristically completed transaction branches from a resource manager.

Syntax

```
DBMS_XA.XA_RECOVER
 RETURN DBMS_XA_XID_ARRAY;
```

Return Values

See DBMS_XA_XID_ARRAY Table Type

Usage Notes

- The flags `TMSTARTSCAN`, `TMENDSCAN`, `TMNOFLAGS` are not supported.
- The privilege `SELECT ON DBA_PENDING_TRANSACTIONS` must be granted to the user who needs to call `XA_RECOVER`.

219.6.8 XA_ROLLBACK Function

This function informs the resource manager to roll back work done on behalf of a transaction branch.

Syntax

```
DBMS_XA.XA_ROLLBACK (
   xid       IN  DBMS_XA_XID)
 RETURN PLS_INTEGER;
```

Parameters

Table 219-9 *XA_ROLLBACK Function Parameters*

Parameter	Description
`xid`	See DBMS_XA_XID Object Type

Return Values

See Table 219-2. Possible return values are: `XA_OK`, `XA_RB*`, `XA_HEURHAZ`, `XA_HEURCOM`, `XA_HEURRB`, or `XA_HEURMIX`.

Usage Notes

If a user needs to rollback a transaction branch that created by other users, the privilege `FORCE ANY TRANSACTION` must be granted to the user.

219.6.9 XA_SETTIMEOUT Function

This function sets the transaction timeout in seconds for the current session.

Syntax

```
DBMS_XA.XA_SETTIMEOUT (
   seconds  IN  PLS_INTEGER)
RETURN PLS_INTEGER;
```

Parameters

Table 219-10 *XA_SETTIMEOUT Function Parameters*

Parameter	Description
`seconds`	The timeout value indicates the maximum time in seconds that a transaction branch may be disassociated from the session before the system automatically aborts the transaction. The default value is 60 seconds.

Return Values

See Table 219-2. Possible return values are `XA_OK`, `XAER_RMERR`, `XAER_RMFAIL`, or `XAER_INVAL`.

Usage Notes

Only if return value is `XA_OK`, is the timeout value successfully set.

219.6.10 XA_START Function

This function associates the current session with a transaction branch specified by the `xid`.

Syntax

```
DBMS_XA.XA_START (
   xid   IN  DBMS_XA_XID,    flag  IN  PLS_INTEGER)  RETURN PLS_INTEGER;
```

Parameters

Table 219-11 *XA_START Function Parameters*

Parameter	Description
`xid`	See DBMS_XA_XID Object Type
`flag`	See Table 219-1.

Return Values

See Table 219-2

Usage Notes

- If TMJOIN or TMRESUME is specified in flag, the start is for joining an existing transaction branch identified by the xid. TMJOIN flag should be used when the transaction is detached with TMSUCCESS flag. TMRESUME should be used when the transaction branch is detached with TMSUSPEND flag. XA_START may be called with either flag to join an existing transaction branch.
- If TMNOFLAGS is specified in flag, and neither TMJOIN nor TMRESUME is specified, a new transaction branch is to be started. If the transaction branch specified in xid already exists, XA_START returns an XAER_DUPID error code.
- Possible return values in error include: XAER_RMERR, XAER_RMFAIL, XAER_DUPID, XAER_OUTSIDE, XAER_NOTA, XAER_INVAL, and XAER_PROTO.
- XA_OK is returned if XA_START succeeds. An application must check the return value and handle error cases. Only when XA_OK is returned, the PL/SQL application should proceed for other normal operations. Transaction stacking is not supported. If there is an active transaction associated with the current session, may not be called to start or join another transaction. XAER_PROTO will be returned if XA_START is called with an active global transaction branch associated with the session. XAER_OUTSIDE will be returned if XA_START is called with a local transaction associated with the current session.

220
DBMS_XDB_ADMIN

The `DBMS_XDB_ADMIN` package provides an interface to manage the Oracle XML DB repository.

This chapter contains the following topics:

- DBMS_XDB_ADMIN Security Model
- Summary of DBMS_XDB_ADMIN Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for information about Oracle XML DB Repository

220.1 DBMS_XDB_ADMIN Security Model

Owned by `XDB`, the `DBMS_XDB_ADMIN` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `SYS` or `XDB` or `DBA`. Subprograms in this package are executed using the privileges of the current user.

220.2 Summary of DBMS_XDB_ADMIN Subprograms

This table lists the package subprograms in alphabetical order.

Table 220-1 DBMS_XDB_ADM Package Subprograms

Subprogram	Description
CREATENONCEKEY Procedure	Generates a nonce value for use in digest authentication
INSTALLDEFAULTWALLET Procedure	Installs the default XDB wallet in the default XDB wallet directory
MOVEXDB_TABLESPACE Procedure	Moves the XDB (user) to the specified tablespace
REBUILDHIERARCHICALINDEX Procedure	Rebuilds the hierarchical index after import or export operations

220.2.1 CREATENONCEKEY Procedure

This procedure generates a nonce value for use in digest authentication.

Syntax

```
DBMS_XDB_ADMIN.CREATENONCEKEY;
```

220.2.2 INSTALLDEFAULTWALLET Procedure

This procedure installs the default XDB wallet in the default XDB wallet directory.

The directory name where the XDB wallet is stored is prefixed either by `ORACLE_BASE` when it is defined, or `ORACLE_HOME`. It is then followed by `/admin/`*`db_name`*`/xdb_wallet` where *db_name* is the unique database name.

Syntax

```
DBMS_XDB_ADMIN.INSTALLDEFAULTWALLET;
```

Usage Notes

Only `SYS` can install or replace the default wallet.

220.2.3 MOVEXDB_TABLESPACE Procedure

This procedure moves the XDB (user) to the specified tablespace.

Syntax

```
DBMS_XDB_ADMIN.MOVEXDB_TABLESPACE(
   new_tablespace   IN  VARCHAR2);
```

Parameters

Table 220-2 MOVEXDB_TABLESPACE Procedure Parameters

Parameter	Description
`new_tablespace`	Name of the tablespace to where the XDB is moved

Usage Notes

- This operation waits for all concurrent XDB sessions to exit.
- If `MOVEXDB_TABLESPACE` fails, the user should restart the database before issuing any further command.Failure to do so will result into unexpected behavior from the database.
- The XDB repository by default resides in the `SYSAUX` tablespace. Using this procedure it can be moved to another tablespace. As a best practice we recommend to create a dedicated tablespace for the XDB repository only and not share it with other objects (such as tables). The tablespace containing the XDB repository should never be set to `READ ONLY` because this might affect various XML operations being executed.

220.2.4 REBUILDHIERARCHICALINDEX Procedure

This procedure rebuilds the hierarchical index after import or export operations. This is necessary because data cannot be exported from index tables.

Syntax

```
DBMS_XDB_ADMIN.REBUILDHIERARCHICALINDEX;
```

221
DBMS_XDB_CONFIG

The `DBMS_XDB_CONFIG` package provides an interface for configuring Oracle XML DB and its repository.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_XDB_CONFIG Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

221.1 DBMS_XDB_CONFIG Overview

PL/SQL package `DBMS_XDB_CONFIG` is the Oracle XML DB resource application program interface (API) for PL/SQL for DBAs to configure their system.
This API provides functions and procedures to access and manage Oracle XML DB Repository resources using PL/SQL. It includes methods for managing resource security and Oracle XML DB configuration.

> **Note:**
>
> The Oracle XML DB Repository is deprecated with Oracle Database 23ai.
>
> Oracle recommends that you replace any functionality used in XML DB Repository with alternative technologies.

Oracle XML DB Repository is modeled on XML, and provides a database file system for any data. The repository maps path names (or URLs) onto database objects of XMLType and provides management facilities for these objects.

PL/SQL package DBMS_XDB_CONFIG is an API that you can use to manage all of the following:

- Oracle XML DB resources
- Oracle XML DB security based on access control lists (ACLs). An ACL is a list of access control entries (ACEs) that determines which principals (users and roles) have access to which resources.
- Oracle XML DB configuration

221.2 DBMS_XDB_CONFIG Security Model

Owned by XDB, the DBMS_XDB_CONFIG package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

221.3 DBMS_XDB_CONFIG Constants

The DBMS_XDB_CONFIG package defines several enumerated constants that should be used for specifying parameter values.

These constants are shown in the following table.

Table 221-1 DBMS_XDB_CONFIG Constants

Constant	Type	Value	Description
XDB_ENDPOINT_HTTP	NUMBER	1	Defining listener for first HTTP port
XDB_ENDPOINT_HTTP2	NUMBER	2	Defining listener for second HTTP port
XDB_ENDPOINT_RHTTP	NUMBER	3	Defining listener for a remote HTTP port
XDB_ENDPOINT_RHTTPS	NUMBER	4	Defining listener for a remote HTTPS port
XDB_PROTOCOL_TCP	NUMBER	1	Defining listener for HTTP protocol
XDB_PROTOCOL_TCPS	NUMBER	2	Defining listener for HTTPS protocol

221.4 Summary of DBMS_XDB_CONFIG Subprograms

The DBMS_XDB_CONFIG package uses subprograms for configuring Oracle XML DB and its repository.

These subprograms are listed in the following table.

Table 221-2 DBMS_XDB_CONFIG Package Subprograms

Subprogram	Description
ADDHTTPEXPIREMAPPING Procedure	Adds to XDB$CONFIG a mapping of the URL pattern to an expiration date. This will control the Expire headers for URLs matching the pattern.
ADDMIMEMAPPING Procedure	Adds a mime mapping to XDB configuration
ADDSCHEMALOCMAPPING Procedure	Adds a schema location mapping to the XDB configuration
ADDSERVLET Procedure	Adds a servlet to XDB configuration
ADDSERVLETMAPPING Procedure	Adds a servlet mapping to XDB configuration
ADDSERVLETSECROLE Procedure	Adds a security role REF to a specified servlet in XDB configuration
ADDXMLEXTENSION Procedure	Adds adds the following XML extension to the XDB configuration under <xml-extensions>: <extension>extension</extension>

Table 221-2 (Cont.) DBMS_XDB_CONFIG Package Subprograms

Subprogram	Description
CFG_GET Function	Retrieves the session's configuration information as an `XMLType` instance
CFG_REFRESH Procedure	Refreshes the session's configuration information to the latest configuration
CFG_UPDATE Procedure	Updates the configuration information and commits the change.
DELETEHTTPEXPIREMAPPING Procedure	Deletes from `XDB$CONFIG` all mappings of the URL pattern to an expiration date
DELETEMIMEMAPPING Procedure	Deletes the mime mapping for a specified extension from the XDB configuration
DELETESCHEMALOCMAPPING Procedure	Deletes the schema location mapping for a specified schema URL from the XDB configuration
DELETESERVLET Procedure	Deletes a servlet from the XDB configuration
DELETESERVLETMAPPING Procedure	Deletes the servlet mapping for a specified servlet name from the XDB configuration
DELETESERVLETSECROLE Procedure	Deletes the specified role from a servlet in the XDB configuration
DELETEXMLEXTENSION Procedure	Deletes the specified XML extension from the XDB configuration
ENABLEDIGESTAUTHENTICATION Procedure	Enables enables digest authentication
GETFTPPORT Function	Gets the value of the current FTP port
GETHTTPPORT Function	Gets the value of the current HTTP port
GETHTTPCONFIGREALM Function	Gets the realm name
GETHTTPSPORT Function	Gets the value of the current HTTPS port
GETLISTENERENDPOINT Procedure	Gets the parameters of a listener end point corresponding to the XML DB HTTP server
GETREMOTEHTTPPORT Function	Gets the value of a remote HTTP port
GETREMOTEHTTPSPORT Function	Gets the value of a remote HTTPS port
ISGLOBALPORTENABLED Function	Returns the value of the `DBMS_XDB_CONFIG.SETGLOBALPORTENABLED` procedure setting
SETFTPPORT Procedure	Sets the FTP port to a new value
SETHTTPPORT Procedure	Sets the HTTP port to a new value
SETHTTPCONFIGREALM Procedure	Sets the realm to a new value
SETHTTPSPORT Procedure	Sets the HTTPS port to a new value
SETLISTENERENDPOINT Procedure	Sets the parameters of a listener end point corresponding to the XML DB HTTP server
SETLISTENERLOCALACCESS Procedure	Restricts all listener end points of the XML DB HTTP server to listen either only on the localhost interface (when `l_access` is set to `TRUE`) or to listen on both localhost and non-localhost interfaces (when `l_access` is set to `FALSE`)

Table 221-2 (Cont.) DBMS_XDB_CONFIG Package Subprograms

Subprogram	Description
SETREMOTEHTTPPORT Procedure	Sets the remote HTTP port to a new value
SETREMOTEHTTPSPORT Procedure	Sets the remote HTTPS port to a new value
USEDPORT Procedure	Obtains the port numbers used by other pluggable databases in the same consolidation database

221.4.1 ADDHTTPEXPIREMAPPING Procedure

This procedure adds to `XDB$CONFIG` a mapping of the URL pattern to an expiration date. This will control the Expire headers for URLs matching the pattern.

Syntax

```
DBMS_XDB_REPOS.ADDHTTPEXPIREMAPPING (
     pattern    IN    VARCHAR2,
     expire     IN    VARCHAR2);
```

Parameters

Table 221-3 *ADDHTTPEXPIREMAPPING Procedure Parameters*

Parameter	Description
`pattern`	URL pattern (only * accepted as wildcards)

`expire`	Expiration directive, follows the `ExpireDefault` in Apache's `mod_expires`:
	`base [plus] (num type)*`
	`-- base: now | modification`
	`-- type: year|years|month|months|week|weeks|day|days|minute|minutess|second|seconds`

Examples

```
DBMS_XDB_REPOS.ADDHTTPEXPIREMAPPING ('/public/test1/*', 'now plus 4 weeks');
DBMS_XDB_REPOS.ADDHTTPEXPIREMAPPING (
             '/public/test2/*', 'modification plus 1 day 30 seconds');
```

221.4.2 ADDMIMEMAPPING Procedure

This procedure adds the following mime mapping to XDB configuration:

```
<mime-mapping>
<extension>extension</extension>
<mime-type>mimetype</mime-type>
</mime-mapping>
```

Syntax

```
DBMS_XDB_CONFIG.ADDMIMEMAPPING(
     extension     IN   VARCHAR2,
     mimetype      IN   VARCHAR2);
```

Parameters

Table 221-4 *ADDMIMEMAPPING Procedure Parameters*

Parameter	Description
extension	Extension for which a mime type is being added
mimetype	Mime type

221.4.3 ADDSCHEMALOCMAPPING Procedure

This procedure adds the following schema location mapping to the XDB configuration:

```
<schemaLocation-mapping>
     <namespace>namespace</namespace>
     <element>element</element>
      <schemaURL>schemaURL</schemaURL>
</schemaLocation-mapping>
```

Syntax

```
DBMS_XDB_CONFIG.ADDSCHEMALOCMAPPING(
     namespace     IN   VARCHAR2,
     element       IN   VARCHAR2,
     schemaURL     IN   VARCHAR2);
```

Parameters

Table 221-5 *ADDSCHEMALOCMAPPING Procedure Parameters*

Parameter	Description
namespace	Namespace
element	Element
schemaURL	Schema URL

221.4.4 ADDSERVLET Procedure

This procedure adds a servlet to XDB configuration.

IT adds the following servlet:

```
<servlet>
     <servlet-name>name</servlet-name>     <servlet-language>language</servlet-language>
     <display-name>dispname</display-name>
     <description>descript</description>
     <servlet-class>class</servlet-class>
     <servlet-schema>schema</servlet-schema>
</servlet>
```

Syntax

```
DBMS_XDB_CONFIG.ADDSERVLET(
    name         IN   VARCHAR2,
    language     IN   VARCHAR2,
    dispname     IN   VARCHAR2,
```

```
    icon        IN   VARCHAR2 := NULL,
    descript    IN   VARCHAR2 := NULL,
    class       IN   VARCHAR2 := NULL,
    jspfile     IN   VARCHAR2 := NULL,
    plsql       IN   VARCHAR2 := NULL,
    schema      IN   VARCHAR2 := NULL);
```

Parameters

Table 221-6 ADDSERVLET Procedure Parameters

Parameter	Description
`name`	Servlet name
`language`	Must be one of "C", "Java", "PL/SQL"
`dispname`	Display name
`icon`	Icon
`descript`	Description
`class`	The `class` / `jspfile` / `plsql` function corresponding to this servlet. The first non-`NULL` argument amongst these three is chosen, and the others are treated as `NULL`.
`jspfile`	This parameter is not supported. Always provide `NULL` value for this parameter. System throws an error if the value is not `NULL`.
`plsql`	The `class` / `jspfile` / `plsql` function corresponding to this servlet. The first non-`NULL` argument amongst these three is chosen, and the others are treated as `NULL`.
`schema`	Schema is used to specify `servlet-schema` xml element name in `xdbconfig.xml`. It indicates the Oracle schema in which the Java class is loaded. If you do not specify any value, then the schema is searched using the default resolver specification.

221.4.5 ADDSERVLETMAPPING Procedure

This procedure adds a servlet mapping to XDB configuration.

It adds the following servlet mapping:

```
<servlet-mapping>
     <servlet-pattern>pattern</servlet-pattern>
     <servlet-name>name</servlet-name>
</servlet-mapping>
```

Syntax

```
DBMS_XDB_CONFIG.ADDSERVLETMAPPING(
     pattern   IN   VARCHAR2,     name      IN   VARCHAR2);
```

Parameters

Table 221-7 *ADDSERVLETMAPPING Procedure Parameters*

Parameter	Description
`pattern`	Sservlet pattern
`name`	Servlet name

221.4.6 ADDSERVLETSECROLE Procedure

This procedure adds the security role `REF` to a specified servlet in XDB configuration.

It adds the following security role as shown in the following:

```
<security-role-ref>
     <role-name>rolename</role-name>
     <role-link>rolelink</role-link>
     <description>descript</description>
</security-role-ref>
```

Syntax

```
DBMS_XDB_CONFIG.ADDSERVLETSECROLE(
     servname    IN   VARCHAR2,     rolename    IN   VARCHAR2,     rolelink    IN
VARCHAR2,     descript    IN   VARCHAR2 := NULL);
```

Parameters

Table 221-8 *ADDSERVLETSECROLE Procedure Parameters*

Parameter	Description
`servname`	Sservlet name
`rolename`	Role name
`rolelink`	Role link
`descript`	Description

221.4.7 ADDXMLEXTENSION Procedure

This procedure adds an XML extension to the XDB configuration under `<xml-extensions>`.

It adds the following XML extension to the XDB configuration, under `<xml-extensions>`:

```
<extension>extension</extension>
```

Syntax

```
DBMS_XDB_CONFIG.ADDXMLEXTENSION(
     extension    IN   VARCHAR2);
```

Parameters

Table 221-9 *ADDXMLEXTENSION Procedure Parameters*

Parameter	Description
`extension`	XML extension to be added

221.4.8 CFG_GET Function

This function retrieves the session's configuration information as an `XMLType` instance.

Syntax

```
DBMS_XDB_CONFIG.CFG_GET
 RETURN SYS.XMLType;
```

221.4.9 CFG_REFRESH Procedure

This procedure refreshes the session's configuration information to the latest configuration.

Syntax

```
DBMS_XDB_CONFIG.CFG_REFRESH;
```

221.4.10 CFG_UPDATE Procedure

This procedure updates the configuration information and commits the change.

Syntax

```
DBMS_XDB_CONFIG.CFG_UPDATE(
   xdbconfig    IN   SYS.XMLTYPE);
```

Parameters

Table 221-10 CFG_UPDATE Procedure Parameters

Parameter	Description
`xdbconfig`	The new configuration data

221.4.11 DELETEHTTPEXPIREMAPPING Procedure

This procedure deletes from `XDB$CONFIG` all mappings of the URL pattern to an expiration date.

Syntax

```
DBMS_XDB_REPOS.DELETEHTTPEXPIREMAPPING(
     pattern   IN    VARCHAR2);
```

Parameters

Table 221-11 DELETEHTTPEXPIREMAPPING Procedure Parameters

Parameter	Description
`pattern`	URL pattern (only * accepted as wildcards)

221.4.12 DELETEMIMEMAPPING Procedure

This procedure deletes the mime mapping for a specified extension from the XDB configuration.

Syntax

```
DBMS_XDB_CONFIG.DELETEMIMEMAPPING(
     extension    IN   VARCHAR2);
```

Parameters

Table 221-12 *DELETEMIMEMAPPING Procedure Parameters*

Parameter	Description
extension	Extension for which a mime type is to be deleted

221.4.13 DELETESCHEMALOCMAPPING Procedure

This procedure deletes the schema location mapping for a specified schema URL from the XDB configuration.

Syntax

```
DBMS_XDB_CONFIG.DELETESCHEMALOCMAPPING(
     schemaURL    IN   VARCHAR2);
```

Parameters

Table 221-13 *DELETESCHEMALOCMAPPING Procedure Parameters*

Parameter	Description
schemaURL	Schema URL

221.4.14 DELETESERVLET Procedure

This procedure deletes a servlet from the XDB configuration.

Syntax

```
DBMS_XDB_CONFIG.DELETESERVLET(
     name         IN   VARCHAR2);
```

Parameters

Table 221-14 *DELETESERVLET Procedure Parameters*

Parameter	Description
name	Servlet name

221.4.15 DELETESERVLETMAPPING Procedure

This procedure deletes the servlet mapping for a specified servlet name from the XDB configuration.

Syntax

```
DBMS_XDB_CONFIG.DELETESERVLETMAPPING(
     name       IN    VARCHAR2);
```

Parameters

Table 221-15 ***DELETESERVLETMAPPING Procedure Parameters***

Parameter	Description
name	Servlet name

221.4.16 DELETESERVLETSECROLE Procedure

This procedure deletes the specified role from a servlet in the XDB configuration.

Syntax

```
DBMS_XDB_CONFIG.DELETESERVLETSECROLE(
     servname     IN    VARCHAR2,     rolename     IN    VARCHAR2);
```

Parameters

Table 221-16 ***DELETESERVLETSECROLE Procedure Parameters***

Parameter	Description
servname	Servlet name
rolename	Name of the role to be deleted

221.4.17 DELETEXMLEXTENSION Procedure

This procedure deletes the specified XML extension from the XDB configuration.

Syntax

```
DBMS_XDB_CONFIG.DELETEXMLEXTENSION(
     extension     IN    VARCHAR2);
```

Parameters

Table 221-17 ***DELETEXMLEXTENSION Procedure Parameters***

Parameter	Description
extension	XML extension to be deleted

221.4.18 ENABLEDIGESTAUTHENTICATION Procedure

This procedure enables digest authentication. It will list digest as the first authentication mechanism to be used by the XML DB HTTP server.

Syntax

```
DBMS_XDB_CONFIG.ENABLEDIGESTAUTHENTICATION;
```

221.4.19 GETFTPPORT Function

This procedure gets the value of the current FTP port.

Syntax

```
DBMS_XDB_CONFIG.GETFTPPORT
 RETURN NUMBER;
```

221.4.20 GETHTTPCONFIGREALM Function

This function gets the realm name. Definition of a realm is referenced in IETF's RFC2617.

Syntax

```
DBMS_XDB_CONFIG.GETHTTPCONFIGREALM
 RETURN VARCHAR2;
```

221.4.21 GETHTTPPORT Function

This function gets the value of the current HTTP port.

Syntax

```
DBMS_XDB_CONFIG.GETHTTPPORT
 RETURN NUMBER;
```

221.4.22 GETHTTPSPORT Function

This procedure gets the value of the current HTTPS port.

Syntax

```
DBMS_XDB_CONFIG.GETHTTPSPORT
 RETURN NUMBER;
```

Usage Notes

Returns `NULL` if no port has been configured

221.4.23 GETLISTENERENDPOINT Procedure

This procedure retrieves the parameters of a listener end point corresponding to the XML DB HTTP server. The parameters of both HTTP and HTTP2 end points can be retrieved by invoking this procedure.

Syntax

```
DBMS_XDB_CONFIG.GETLISTENERENDPOINT (
   endpoint  IN   NUMBER,
   host      OUT  VARCHAR2,    port      OUT  NUMBER,
   protocol  OUT  NUMBER);
```

Parameters

Table 221-18 GETLISTENERENDPOINT Procedure Parameters

Parameter	Description
`endpoint`	End point to be retrieved. Its value can be `XDB_ENDPOINT_HTTP` or `XDB_ENDPOINT_HTTP2`.
`host`	Interface on which the listener end point listens
`port`	Port on which the listener end point listens
`protocol`	Transport protocol accepted by the listener end point

221.4.24 GETREMOTEHTTPPORT Function

This function gets the value of a remote HTTP port.

Syntax

```
DBMS_XDB_CONFIG.GETREMOTEHTTPPORT
 RETURN NUMBER;
```

221.4.25 GETREMOTEHTTPSPORT Function

This function gets the value of a remote HTTPS port.

Syntax

```
DBMS_XDB_CONFIG.GETREMOTEHTTPSPORT
 RETURN NUMBER;
```

221.4.26 ISGLOBALPORTENABLED Function

This procedure returns the value of the `GlobalPortEnabled` setting.

Syntax

```
DBMS_XDB_CONFIG.ISGLOBALPORTENABLED RETURN BOOLEAN;
```

Usage Notes

- This procedure returns `TRUE` if `GlobalPortEnabled` has been set to `TRUE`; otherwise it returns `FALSE`.
- In a multitenant environment, you can execute this function in both the CDB root and PDBs.

221.4.27 SETFTPPORT Procedure

This procedure sets the FTP port to a new value.

Syntax

```
DBMS_XDB_CONFIG.SETFTPPORT(
   new_port  IN  NUMBER);
```

Parameters

Table 221-19 SETFTPPORT Procedure Parameters

Parameter	Description
`new_port`	Value to which the FTP port is set

221.4.28 SETHTTPPORT Procedure

This procedure sets the HTTP port to a new value.

Syntax

```
DBMS_XDB_CONFIG.SETHTTPPORT(
   new_port  IN  NUMBER);
```

Parameters

Table 221-20 SETHTTPPORT Procedure Parameters

Parameter	Description
`new_port`	Value to which the HTTP port is set

221.4.29 SETHTTPCONFIGREALM Procedure

This procedure modifies the realm value.

Syntax

```
DBMS_XDB_CONFIG.SETHTTPCONFIGREALM(
   realm IN VARCHAR2);
```

Parameters

Table 221-21 SETHTTPPORT Procedure Parameters

Parameter	Description
realm	Realm as defined in IETF's RFC2617

221.4.30 SETHTTPSPORT Procedure

This procedure sets the HTTPS port to a new value.

Syntax

```
DBMS_XDB_CONFIG.SETHTTPSPORT(
   new_port  IN  NUMBER);
```

Parameters

Table 221-22 SETHTTPSPORT Procedure Parameters

Parameter	Description
new_port	Value to which the HTTPS port is set

221.4.31 SETLISTENERENDPOINT Procedure

This procedure sets the parameters of a listener end point corresponding to the XML DB HTTP server.

Both HTTP and HTTP2 end points can be set by invoking this procedure.

Syntax

```
DBMS_XDB_CONFIG.SETLISTENERENDPOINT (
   endpoint  IN  NUMBER,
   host      IN  VARCHAR2,    port      IN  NUMBER,
   protocol  IN  NUMBER);
```

Parameters

Table 221-23 SETLISTENERENDPOINT Procedure Parameters

Parameter	Description
endpoint	End point to be set. Its value can be XDB_ENDPOINT_HTTP or XDB_ENDPOINT_HTTP2.
host	Interface on which the listener end point is to listen. Its value can be 'LOCALHOST,' NULL, or a hostname. If its value is 'LOCALHOST' the listener end point is permitted to only listen on the localhost interface. If its value is NULL or hostname, the listener end point is permitted to listen on both localhost and non-localhost interfaces.
port	Port on which the listener end point is to listen

Table 221-23 (Cont.) SETLISTENERENDPOINT Procedure Parameters

Parameter	Description
protocol	Transport protocol that the listener end point is to accept. Its value can be XDB_PROTOCOL_TCP or XDB_PROTOCOL_TCPS

221.4.32 SETLISTENERLOCALACCESS Procedure

This procedure restricts all listener end points of the XML DB HTTP server to listen either only on the localhost interface (when l_access is set to TRUE) or to listen on both localhost and non-localhost interfaces (when l_access is set to FALSE).

Syntax

```
DBMS_XDB_CONFIG.SETLISTENERLOCALACCESS (
   l_access   BOOLEAN);
```

Parameters

Table 221-24 SETLISTENERLOCALACCESS Procedure Parameters

Parameter	Description
l_access.	TRUE or FALSE.

221.4.33 SETREMOTEHTTPPORT Procedure

This procedure sets a remote HTTP port to a new value.

Syntax

```
DBMS_XDB_CONFIG.SETREMOTEHTTPPORT(
   new_port  IN  NUMBER);
```

Parameters

Table 221-25 SETREMOTEHTTPPORT Procedure Parameters

Parameter	Description
new_port	Value to which the remote HTTP port is set

221.4.34 SETREMOTEHTTPSPORT Procedure

This procedure sets a remote HTTPS port to a new value.

Syntax

```
DBMS_XDB_CONFIG.SETREMOTEHTTPSPORT (
   new_port  IN  NUMBER);
```

Parameters

Table 221-26 SETREMOTEHTTPSPORT Procedure Parameters

Parameter	Description
new_port	Value to which the remote HTTPS port is set

221.4.35 USEDPORT Function

This function obtains the port numbers used by other pluggable databases in the same consolidation database.

Syntax

```
DBMS_XDB_CONFIG.USEDPORT RETURN XMLTYPE;
```

222
DBMS_XDB_CONSTANTS

The DBMS_XDB_CONSTANTS package provides an interface to commonly used constants. Users should use constants instead of dynamic strings to avoid typographical errors.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_XDB_CONSTANTS Subprograms

222.1 DBMS_XDB_CONSTANTS Security Model

Owned by XDB, the DBMS_XDB_CONSTANTS package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

222.2 Summary of DBMS_XDB_CONSTANTS Subprograms

This table lists the DBMS_XDB_CONSTANTS subprograms and briefly describes them.

Table 222-1 DBMS_XDB_CONSTANTS Package Subprograms

Subprograms	Description
ENCODING_DEFAULT Function	Returns 'AL32UTF8'
ENCODING_ISOLATIN1 Function	Returns 'WE8ISO8859P1'
ENCODING_UTF8 Function	Returns 'AL32UTF8'
ENCODING_WIN1252 Function	Returns 'WE8MSWIN1252'
NAMESPACE_ACL Function	Returns 'http://xmlns.oracle.com/xdb/acl.xsd'
NAMESPACE_METADATA Function	Returns 'http://xmlns.oracle.com/xdbuserMetaData'
NAMESPACE_ORACLE Function	Returns 'http://xmlns.oracle.com'
NAMESPACE_ORACLE_XDB Function	Returns 'http://xmlns.oracle.com/xdb'
NAMESPACE_RESOURCE Function	Returns 'http://xmlns.oracle.com/xdb/XDBResource.xsd'
NAMESPACE_RESOURCE_EVENT Function	Returns 'http://xmlns.oracle.com/resourceEvent'
NAMESPACE_RESOURCE_CONFIG Function	Returns 'http://xmlns.oracle.com/XDBResConfig.xsd'
NAMESPACE_XDBSCHEMA Function	Returns 'http://xmlns.oracle.com/xdb/XDBSchema.xsd'
NAMESPACE_XMLDIFF Function	Returns 'http://xmlns.oracle.com/xdb/xdiff.xsd'
NAMESPACE_XMLINSTANCE Function	Returns 'http://www.w3.org/2001/XMLSchema-instance'

Table 222-1 (Cont.) DBMS_XDB_CONSTANTS Package Subprograms

Subprograms	Description
NAMESPACE_XMLSCHEMA Function	Returns 'http://www.w3.org/2001/XMLSchema'
NSPREFIX_ACL_ACL Function	Returns 'xmlns:acl= 'http://xmlns.oracle.com/acs.xsd'
NSPREFIX_RESCONFIG_RC Function	Returns 'xmlns:rc="http://xmlns.oracle.com/xdb/XDBResConfig.xsd"'
NSPREFIX_RESOURCE_R Function	Returns 'xmlns:r="http://xmlns.oracle.com/XDBResource.xsd'
NSPREFIX_XDB_XDB Function	Returns 'http://xmlns.oracle.com/xdb'
NSPREFIX_XMLINSTANCE_XSI Function	Returns 'xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"'
NSPREFIX_XMLDIFF_XD Function	Returns 'xmlns:xd="http://xmlns.oracle.com/xdb/xdiff.xsd"'
NSPREFIX_XMLSCHEMA_XSD Function	Returns 'xmlns:xsd="http://www.w3.org/2001/XMLSchema"'
SCHEMAURL_ACL Function	Returns 'http://xmlns.oracle.com/xdb/acl.xsd'
SCHEMAELEM_RES_ACL Function	Returns 'http://xmlns.oracle.com/xdb/acl.xsd#acl'
SCHEMAELEM_RESCONTENT_BINARY Function	Returns 'http://xmlns.oracle.com/xdb/XDBSchema.xsd#binary'
SCHEMAELEM_RESCONTENT_TEXT Function	Returns 'http://xmlns.oracle.com/xdb/XDBSchema.xsd#text'
SCHEMAURL_RESOURCE Function	Returns 'http://xmlns.oracle.com/xdb/XDBResource.xsd'
SCHEMAURL_XDBSCHEMA Function	Returns 'http://xmlns.oracle.com/xdb/XDBSchema.xsd'
XDBSCHEMA_PREFIXES Function	Returns 'xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xdb="http://xmlns.oracle.com/xdb'
XSD_ATTRIBUTE Function	Returns 'attribute'
XSD_COMPLEX_TYPE Function	Returns 'complexType'
XSD_ELEMENT Function	Returns 'element'
XSD_GROUP Function	Returns 'group'

222.2.1 ENCODING_DEFAULT Function

This function returns 'AL32UTF8'.

Syntax

```
DBMS_XDB_CONSTANTS.ENCODING_DEFAULT
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

'AL32UTF8'

222.2.2 ENCODING_ISOLATIN1 Function

This function returns `'WE8ISO8859P1'`.

Syntax

```
DBMS_XDB_CONSTANTS.ENCODING_ISOLATIN1
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'WE8ISO8859P1'
```

222.2.3 ENCODING_UTF8 Function

This function returns `'AL32UTF8'`.

Syntax

```
DBMS_XDB_CONSTANTS.ENCODING_UTF8
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'AL32UTF8'
```

222.2.4 ENCODING_WIN1252 Function

This function returns `'WE8MSWIN1252'`.

Syntax

```
DBMS_XDB_CONSTANTS.ENCODING_WIN1252
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'WE8MSWIN1252'
```

222.2.5 NAMESPACE_ACL Function

This function returns `'http://xmlns.oracle.com/xdb/acl.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_ACL
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'http://xmlns.oracle.com/xdb/acl.xsd'
```

222.2.6 NAMESPACE_METADATA Function

This function returns `'http://xmlns.oracle.com/xdbuserMetaData'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_METADATA
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdbuserMetaData'`

222.2.7 NAMESPACE_ORACLE Function

This function returns `'http://xmlns.oracle.com'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_ORACLE
  RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com'`

222.2.8 NAMESPACE_ORACLE_XDB Function

This function returns `'http://xmlns.oracle.com/xdb'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_ORACLE_XDB
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb'`

222.2.9 NAMESPACE_RESOURCE Function

This function returns `'/XDBResource.xsd'`

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_RESOURCE
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb/XDBResource.xsd`

222.2.10 NAMESPACE_RESOURCE_EVENT Function

This function returns `'http://xmlns.oracle.com/resourceEvent'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_RESOURCE_EVENT
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/resourceEvent'`

222.2.11 NAMESPACE_RESOURCE_CONFIG Function

This function returns `'http://xmlns.oracle.com/XDBResConfig.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_RESOURCE_CONFIG
  RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/XDBResConfig.xsd'`

222.2.12 NAMESPACE_XDBSCHEMA Function

This function returns `'http://xmlns.oracle.com/xdb/XDBSchema.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_XDBSCHEMA
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb/XDBSchema.xsd'`

222.2.13 NAMESPACE_XMLDIFF Function

This function returns `'http://xmlns.oracle.com/xdb/xdiff.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_XMLDIFF
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb/xdiff.xsd'`

222.2.14 NAMESPACE_XMLINSTANCE Function

This function returns `'http://www.w3.org/2001/XMLSchema-instance'`.

Syntax

```
DBMS_XDB_CONSTANTS.NAMESPACE_XMLINSTANCE
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://www.w3.org/2001/XMLSchema-instance'`

222.2.15 NAMESPACE_XMLSCHEMA Function

This function returns `'http://www.w3.org/2001/XMLSchema'`.

Syntax

```
DBMS_XDB_CONSTANTS. NAMESPACE_XMLSCHEMA
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://www.w3.org/2001/XMLSchema'`

222.2.16 NSPREFIX_ACL_ACL Function

This function returns `'xmlns:acl= 'http://xmlns.oracle.com/acs.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_ACL_ACL
  RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'xmlns:acl= 'http://xmlns.oracle.com/acs.xsd'`

222.2.17 NSPREFIX_RESCONFIG_RC Function

This function returns `''xmlns:rc="http://xmlns.oracle.com/xdb/XDBResConfig.xsd"'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_RESCONFIG_RC
  RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`Returns ''xmlns:rc="http://xmlns.oracle.com/xdb/XDBResConfig.xsd"'`

222.2.18 NSPREFIX_RESOURCE_R Function

This function returns `'xmlns:r="http://xmlns.oracle.com/XDBResource.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_RESOURCE_R
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'xmlns:r="http://xmlns.oracle.com/XDBResource.xsd'
```

222.2.19 NSPREFIX_XDB_XDB Function

This function returns `'http://xmlns.oracle.com/xdb'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_XDB_XDB
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'http://xmlns.oracle.com/xdb'
```

222.2.20 NSPREFIX_XMLINSTANCE_XSI Function

This function returns `'xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_XMLINSTANCE_XSI
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"'
```

222.2.21 NSPREFIX_XMLDIFF_XD Function

This function returns `'xmlns:xd="http://xmlns.oracle.com/xdb/xdiff.xsd"'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_XMLDIFF_XD
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'xmlns:xd="http://xmlns.oracle.com/xdb/xdiff.xsd"'
```

222.2.22 NSPREFIX_XMLSCHEMA_XSD Function

This function returns `'xmlns:xsd="http://www.w3.org/2001/XMLSchema"'`.

Syntax

```
DBMS_XDB_CONSTANTS.NSPREFIX_XMLSCHEMA_XSD
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'xmlns:xsd="http://www.w3.org/2001/XMLSchema"'
```

222.2.23 SCHEMAURL_ACL Function

This function returns `'http://xmlns.oracle.com/xdb/acl.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.SCHEMAURL_ACL
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'http://xmlns.oracle.com/xdb/acl.xsd'
```

222.2.24 SCHEMAELEM_RES_ACL Function

This function returns `'http://xmlns.oracle.com/xdb/acl.xsd#acl'`.

Syntax

```
DBMS_XDB_CONSTANTS.SCHEMAELEM_RES_ACL
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'http://xmlns.oracle.com/xdb/acl.xsd#acl'
```

222.2.25 SCHEMAELEM_RESCONTENT_BINARY Function

This function returns `'http://xmlns.oracle.com/xdb/XDBSchema.xsd#binary'`.

Syntax

```
DBMS_XDB_CONSTANTS.SCHEMAELEM_RESCONTENT_BINARY
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

```
'http://xmlns.oracle.com/xdb/XDBSchema.xsd#binary'
```

222.2.26 SCHEMAELEM_RESCONTENT_TEXT Function

This function returns `'http://xmlns.oracle.com/xdb/XDBSchema.xsd#text'`.

Syntax

```
DBMS_XDB_CONSTANTS.SCHEMAELEM_RESCONTENT_TEXT
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb/XDBSchema.xsd#text'`

222.2.27 SCHEMAURL_RESOURCE Function

This function returns `'http://xmlns.oracle.com/xdb/XDBResource.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.SCHEMAURL_RESOURCE
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb/XDBResource.xsd'`

222.2.28 SCHEMAURL_XDBSCHEMA Function

This function returns `'http://xmlns.oracle.com/xdb/XDBSchema.xsd'`.

Syntax

```
DBMS_XDB_CONSTANTS.SCHEMAURL_XDBSCHEMA
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'http://xmlns.oracle.com/xdb/XDBSchema.xsd'`

222.2.29 XDBSCHEMA_PREFIXES Function

This function returns `'xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xdb="http://xmlns.oracle.com/xdb'`.

Syntax

```
DBMS_XDB_CONSTANTS.XDBSCHEMA_PREFIXES
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xdb="http://xmlns.oracle.com/xdb'`

222.2.30 XSD_ATTRIBUTE Function

This function returns `'attribute'`.

Syntax

```
DBMS_XDB_CONSTANTS.XSD_ATTRIBUTE
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'attribute'`

222.2.31 XSD_COMPLEX_TYPE Function

This function returns `'complexType'`.

Syntax

```
DBMS_XDB_CONSTANTS.XSD_COMPLEX_TYPE
  RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'complexType'`

222.2.32 XSD_ELEMENT Function

This function returns `'element'`.

Syntax

```
DBMS_XDB_CONSTANTS. XSD_ELEMENT
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'element'`

222.2.33 XSD_GROUP Function

This function returns `'group'`

Syntax

```
DBMS_XDB_CONSTANTS.XSD_GROUP
 RETURN VARCHAR2 DETERMINISTIC;
```

Return Value

`'group'`

223
DBMS_XDB_REPOS

The `DBMS_XDB_REPOS` package provides an interface to operate on the Oracle XML database Repository.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Summary of DBMS_XDB_REPOS Subprograms

> **See Also:**
>
> Oracle XML DB Developer's Guide for more information regarding:
>
> - Using and managing repository resources
> - ACL-based security management (controlling access to repository resources)
> - Managing XLink and XInclude links
> - Loading documents into the repository
> - Creating, deleting, and managing resource metadata

223.1 DBMS_XDB_REPOS Overview

The `DBMS_XDB_REPOS` package lets you operate on the Oracle XML DB Repository to create, modify and delete resources, including managing security based on access control lists (ACLs). The interface provides both query and DML functions.

Using a combination of PL/SQL packages - `DBMS_XDB_REPOS`, DBMS_XDBZ, and DBMS_XDB_VERSION - you can create, delete, and rename documents and folders, move a file or folder within the folder hierarchy, set and change the access permissions on a file or folder, and initiate and manage versioning.

223.2 DBMS_XDB_REPOS Security Model

Owned by XDB, the `DBMS_XDB_REPOS` package must be created by `SYS` or XDB. The `EXECUTE` privilege is granted to `PUBLIC`. Subprograms in this package are executed using the privileges of the current user. Subprograms that operate on the XDB Configuration will succeed only if the current user is `SYS` or XDB, or the current user has the `XDBADMIN` or `DBA` role.

223.3 DBMS_XDB_REPOS Constants

The DBMS_XDB_REPOS package defines several constants that can be used for specifying parameter values.

These constants are shown in the following table.

Table 223-1 DBMS_XDB_REPOS Constants

Constant	Type	Value	Description
DELETE_RESOURCE	NUMBER	1	Deletes a resource; fails if the resource has children.
DELETE_RECURSIVE	NUMBER	2	Deletes a resource and its children, if any.
DELETE_FORCE	NUMBER	3	Deletes the resource, even if the object it contains is invalid
DELETE_RECURSIVE_FORCE	NUMBER	4	Deletes a resource and its children, if any, even if the object it contains is invalid
DELETE_RES_METADATA_CASCADE	NUMBER	1	Deletes the row in the metadata
DELETE_RES_METADATA_NOCASCADE	NUMBER	2	Does not delete the row
DEFAULT_LOCK_TIMEOUT	PLS_INTEGER	(60*60)	Timeout value (in seconds) of the webdav lock
LINK_TYPE_HARD	NUMBER	1	Hard link of a folder to a resource
LINK_TYPE_WEAK	NUMBER	2	Weak link of a folder to a resource
LINK_TYPE_SYMBOLIC	NUMBER	3	Symbolic link of a folder to a resource

223.4 Summary of DBMS_XDB_REPOS Subprograms

This table lists the DBMS_XDB_REPOS subprograms and briefly describes them.

Table 223-2 DBMS_XDB_REPOS Package Subprograms

Subprogram	Description
ACLCHECKPRIVILEGES Function	Checks access privileges granted to the current user by specified ACL document on a resource whose owner is specified by the 'owner' parameter.
APPENDRESOURCEMETADATA Procedure	Takes in user-defined metadata either as a REF to XMLTYPE or an XMLTYPE and adds it to the desired resource
CHANGEOWNER Procedure	Changes the owner of the resource/s to the specified owner.
CHANGEPRIVILEGES Function	Adds a specified ACE to a specified resource's ACL
CHECKPRIVILEGES Function	Checks access privileges granted to the current user on the specified resource
CREATEFOLDER Function	Creates a new folder resource in the hierarchy
CREATEOIDPATH Function	Creates a virtual path to the resource based on object ID

Table 223-2 (Cont.) DBMS_XDB_REPOS Package Subprograms

Subprogram	Description
CREATERESOURCE Functions	Creates a new resource
DELETERESOURCE Procedure	Deletes a resource from the hierarchy
DELETERESOURCEMETADATA Procedures	Deletes metadata from a resource (can be used for schema-based or nonschema-based metadata)
EXISTSRESOURCE Function	Determines if a resource is the hierarchy, based on its absolute path
GETACLDOCUMENT Function	Retrieves ACL document that protects resource given its path name
GETCONTENTBLOB Function	Retrieves the contents of a resource returned as a BLOB
GETCONTENTCLOB Function	Retrieves the contents of a resource returned as a CLOB
GETCONTENTVARCHAR2 Function	Retrieves the contents of a resource returned as a string
GETCONTENTXMLREF Function	Retrieves the contents of a resource returned as a a `REF` to an `XMLTYPE`
GETCONTENTXMLTYPE Function	Retrieves the contents of a resource returned as an `XMLTYPE`
GETLOCKTOKEN Procedure	Returns that resource's lock token for the current user given a path to a resource
GETPRIVILEGES Function	Gets all privileges granted to the current user on a specified resource
GETRESOID Function	Returns the object ID of the resource from its absolute path
GETXDB_TABLESPACE Function	Returns the current tablespace of the XDB (user)
HASBLOBCONTENT Function	Returns `TRUE` if the resource has BLOB content
HASCHARCONTENT Function	Returns `TRUE` if the resource has character content
HASXMLCONTENT Function	Returns `TRUE` if the resource has XML content
HASXMLREFERENCE Function	Returns `TRUE` if the resource has `REF` to XML content
ISFOLDER Function	Returns `TRUE` if the resource is a folder or container
LINK Procedures	Creates a link to an existing resource
LOCKRESOURCE Function	Gets a WebDAV-style lock on that resource given a path to that resource
PROCESSLINKS Procedure	Processes document links in the specified resource
PURGERESOURCEMETADATA Procedure	Deletes all user metadata from a resource
RENAMERESOURCE Procedure	Renames the XDB resource
SETACL Procedure	Sets the ACL on a specified resource
SPLITPATH Procedure	Splits the path into a parentpath and childpath
TOUCHRESOURCE Procedure	Changes the modification time of the resource to the current time
UNLOCKRESOURCE Function	Unlocks the resource given a lock token and resource path
UPDATERESOURCEMETADATA Procedures	Updates metadata for a resource

223.4.1 ACLCHECKPRIVILEGES Function

This function checks access privileges granted to the current user by specified ACL document by the `OWNER` of the resource. Returns positive integer if all privileges are granted.

Syntax

```
DBMS_XDB_REPOS.ACLCHECKPRIVILEGES(
   acl_path  IN  VARCHAR2,
   owner     IN  VARCHAR2,
   privs     IN  xmltype)
 RETURN PLS_INTEGER;
```

Parameters

Table 223-3 *ACLCHECKPRIVILEGES Function Parameters*

Parameter	Description
`acl_path`	Absolute path in the Hierarchy for ACL document
`owner`	Resource owner name; the pseudo user "DAV:owner" is replaced by this user during ACL privilege resolution
`privs`	An `XMLType` instance of the privilege element specifying the requested set of access privileges. See description for CHECKPRIVILEGES Function.

223.4.2 APPENDRESOURCEMETADATA Procedure

This procedure takes in user-defined metadata either as a `REF` to `XMLTYPE` or an `XMLTYPE` and adds it to the desired resource.

Syntax

```
DBMS_XDB_REPOS.APPENDRESOURCEMETADATA (
 abspath   IN  VARCHAR2,
 metadata  IN  XMLTYPE);

DBMS_XDB_REPOS.APPENDRESOURCEMETADATA (
 abspath   IN  VARCHAR2,
 metadata  IN  REF SYS.XMLTYPE);
```

Parameters

Table 223-4 APPENDRESOURCEMETADATA Procedure

Parameter	Description
`abspath`	Absolute path of the resource
`metadata`	Metadata can be schema based or nonschema-based. Schema-based metadata is stored in its own table.

Usage Notes

- In the case in which a `REF` is passed in, the procedure stores the `REF` in the resource, and the metadata is stored in a separate table. In this case you are responsible for populating the `RESID` column for the metadata table. Note that the`REF` passed in must be unique. In

other words, there must not be a`REF` with the same value in the resource metadata, as this would violate uniqueness of properties. An error is thrown if users attempt to add a `REF` that already exists.

- In the case where the XMLTYPE is passed in, the data is parsed to determine if it is schema-based or not and stored accordingly.

223.4.3 CHANGEOWNER Procedure

This procedure changes the owner of the resource/s to the specified owner.

Syntax

```
DBMS_XDB_REPOS.CHANGEOWNER(
     abspath    IN   VARCHAR2,
     owner      IN   VARCHAR2,
     recurse    IN   BOOLEAN := FALSE);
```

Parameters

Table 223-5 CHANGEOWNER Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource
`owner`	New owner for the resource
`recurse`	If `TRUE`, recursively change owner of all resources in the folder tree

223.4.4 CHANGEPRIVILEGES Function

This function adds a specified ACE to a specified resource's ACL.

Syntax

```
DBMS_XDB_REPOS.CHANGEPRIVILEGES(
   res_path   IN    VARCHAR2,
   ace        IN    xmltype)
 RETURN PLS_INTEGER;
```

Parameters

Table 223-6 CHANGEPRIVILEGES Function Parameters

Parameter	Description
`res_path`	Path name of the resource for which privileges need to be changed
`ace`	An `XMLType` instance of the `<ace>` element which specifies the `<principal>`, the operation `<grant>` and the list of privileges

Return Values

A positive integer if the ACL was successfully modified.

Usage Notes

If no ACE with the same principal and the same operation (`grant`/`deny`) already exists in the ACL, the new ACE is added at the end of the ACL.

223.4.5 CHECKPRIVILEGES Function

This function checks access privileges granted to the current user on the specified resource.

Syntax

```
DBMS_XDB_REPOS.CHECKPRIVILEGES(
   res_path   IN  VARCHAR2,
   privs      IN  xmltype)
 RETURN PLS_INTEGER;
```

Parameters

Table 223-7 CHECKPRIVILEGES Function Parameters

Parameter	Description
`res_path`	Absolute path in the Hierarchy for resource
`privs`	An `XMLType` instance of the privilege element specifying the requested set of access privileges

Return Values

A positive integer if all requested privileges granted.

223.4.6 CREATEFOLDER Function

This function creates a new folder resource in the hierarchy.

Syntax

```
DBMS_XDB_REPOS.CREATEFOLDER(
   path   IN  VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-8 CREATEFOLDER Function Parameters

Parameter	Description
`path`	Path name for the new folder

Return Values

`TRUE` if operation successful; `FALSE`, otherwise.

Usage Notes

The given path name's parent folder must already exist in the hierarchy: if `'/folder1/folder2'` is passed as the path parameter, then `'/folder1'` must already exist.

223.4.7 CREATEOIDPATH Function

This function creates a virtual path to the resource based on object ID.

Syntax

```
DBMS_XDB_REPOS.CREATEOIDPATH(
   oid    IN   RAW)
 RETURN VARCHAR2;
```

Parameters

Table 223-9 CREATEOIDPATH Function Parameters

Parameter	Description
`oid`	Object ID of the resource

223.4.8 CREATERESOURCE Functions

The functions create a new resource. The description of the overload options precede each version of the syntax

Syntax

Creates a new resource with a specified string as its contents:

```
DBMS_XDB_REPOS.CREATERESOURCE(
     abspath        IN  VARCHAR2,
     data           IN  VARCHAR2)
 RETURN BOOLEAN;
```

Creates a new resource with a specified `XMLType` data as its contents:

```
DBMS_XDB_REPOS.CREATERESOURCE(
     abspath        IN  VARCHAR2,
     data           IN  SYS.XMLTYPE)
 RETURN BOOLEAN;
```

Given a `REF` to an existing `XMLType` row, creates a resource whose contents point to that row. That row should not already exist inside another resource:

```
DBMS_XDB_REPOS.CREATERESOURCE(
     abspath        IN  VARCHAR2,
     datarow        IN  REF SYS.XMLTYPE)
 RETURN BOOLEAN;
```

Creates a resource with a specified `BLOB` as its contents, and specifies character set of the source `BLOB`:

```
DBMS_XDB_REPOS.CREATERESOURCE(
     abspath        IN  VARCHAR2,
     data           IN  BLOB,
     csid           IN  NUMBER :=0)
 RETURN BOOLEAN;
```

Creates a resource with a specified `BFILE` as its contents, and specifies character set of the source `BFILE`:

```
DBMS_XDB_REPOS.CREATERESOURCE (
     abspath        IN  VARCHAR2,
     data           IN  BFILE,
     csid           IN  NUMBER :=0)
 RETURN BOOLEAN;
```

Creates a resource with a specified `CLOB` as its contents:

```
DBMS_XDB_REPOS.CREATERESOURCE (
     abspath        IN  VARCHAR2,
     data           IN  CLOB)
 RETURN BOOLEAN;
```

Given a string, inserts a new resource into the hierarchy with the string as the contents:

```
DBMS_XDB_REPOS.CREATERESOURCE (
   abspath        IN  VARCHAR2,
   data           IN  VARCHAR2,
   schemaurl      IN  VARCHAR2 := NULL,
   elem           IN  VARCHAR2 := NULL)
 RETURN BOOLEAN;
```

Given an `XMLTYPE` and a schema URL, inserts a new resource into the hierarchy with the `XMLTYPE` as the contents:

```
DBMS_XDB_REPOS.CREATERESOURCE (
   abspath        IN  VARCHAR2,
   data           IN  SYS.XMLTYPE,
   schemaurl      IN  VARCHAR2 := NULL,
   elem           IN  VARCHAR2 := NULL)
 RETURN BOOLEAN;
```

Parameters

Table 223-10 CREATERESOURCE Function Parameters

Parameter	Description
`abspath`	Absolute path of the resource to create. The path name's parent folder must already exist in the hierarchy. In other words, if `/foo/bar.txt` is passed in, then folder `/foo` must already exist.
`data`	String buffer containing new resource's contents. The data is parsed to check if it contains a schema-based XML document, and the contents are stored as schema-based in the schema's default table. Otherwise, it is saved as binary data.
`datarow`	`REF` to an `XMLType` row to be used as the contents
`csid`	Character set id of the document. Must be a valid Oracle ID; otherwise returns an error. If CSID is not specified, or if a zero CSID is specified, then the character set id of the document is determined as follows: • From the abspath extension, determine the resource's MIME type. • If the MIME type is */xml, then the encoding is detected based on Appendix F of the W3C XML 1.0 Reference at http://www.w3.org/TR/2000/REC-xml-20001006; • Otherwise, it is defaulted to the database character set.
`schemaurl`	For XML data, schema URL data conforms to (default `NULL`)
`elem`	Element name (default `NULL`)

Return Values

`TRUE` if operation successful; `FALSE`, otherwise.

223.4.9 DELETERESOURCE Procedure

This procedure deletes a resource from the hierarchy.

Syntax

```
DBMS_XDB_REPOS.DELETERESOURCE(
   path          IN      VARCHAR2,
   delete_option IN      PLS_INTEGER);
```

Parameters

Table 223-11 DELETERESOURCE Procedure Parameters

Parameter	Description
`path`	Path name of the resource to delete
`delete_option`	The option that controls how a a resource is deleted: • `DELETE_RESOURCE` • `DELETE_RECURSIVE` • `DELETE_FORCE` • `DELETE_RECURSIVE_FORCE`

223.4.10 DELETERESOURCEMETADATA Procedures

This procedure takes in a resource by absolute path and removes either the schema-based metadata identified by the REF, or the metadata identified by the namespace and name combination, which can be either schema-based or non-schema based. It also takes an additional (optional) parameter that specifies how to delete it. This parameter is only relevant for schema-based resource metadata that needs to be deleted. For non-schema based metadata, this parameter is ignored.

Syntax

Can be used only for schema-based metadata:

```
DBMS_XDB_REPOS.DELETERESOURCEMETADATA (
   abspath        IN  VARCHAR2,
   metadata       IN  REF SYS.XMLTYPE,
   delete_option  IN  pls_integer := DBMS_XDB_REPOS.DELETE_RESOURCE_METADATA_CASCADE);
```

Can be used for schema-based or nonschema-based metadata:

```
DBMS_XDB_REPOS.DELETERESOURCEMETADATA (
   abspath        IN VARCHAR2,
   metadatans     IN VARCHAR2,
   metadataname   IN VARCHAR2,
   delete_option  IN pls_integer := DBMS_XDB_REPOS.DELETE_RESOURCE_METADATA_CASCADE);
```

Parameters

Table 223-12 DELETERESOURCEMETADATA Procedure Parameters

Parameter	Description
abspath	Absolute path of the resource
metadata	REF to the piece of metadata (schema based) to be deleted
mettadatans	Namespace of the metadata fragment to be removed
mettadataname	Local name of the metadata fragment to be removed
delete_option	Only applicable for schema-based metadata, this can be one of the following: • DELETE_RES_METADATA_CASCADE - deletes the corresponding row in the metadata table • DELETE_RES_METADATA_NOCASCADE - does not delete the row in the metadata table

223.4.11 EXISTSRESOURCE Function

This function indicates if a resource is in the hierarchy. It matches the resource by a string that represents its absolute path.

Syntax

```
DBMS_XDB_REPOS.EXISTSRESOURCE(
   abspath    IN    VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-13 EXISTSRESOURCE Function Parameters

Parameter	Description
abspath	Path name of the resource whose ACL document is required

Return Values

TRUE if the resource is found.

223.4.12 GETACLDOCUMENT Function

This function retrieves ACL document that protects resource given its path name.

Syntax

```
DBMS_XDB_REPOS.GETACLDOCUMENT(
   abspath  IN  VARCHAR2)
  RETURN sys.xmltype;
```

Parameters

Table 223-14 GETACLDOCUMENT Function Parameters

Parameter	Description
abspath	Path name of the resource whose ACL document is required

Return Values

The XMLType for ACL document.

223.4.13 GETCONTENTBLOB Function

This function retrieves the contents of a resource returned as a BLOB.

Syntax

```
DBMS_XDB_REPOS.GETCONTENTBLOB(
     abspath    IN     VARCHAR2,
     csid       OUT    PLS_INTEGER,
     locksrc    IN     BOOLEAN := FALSE)
  RETURN BLOB;
```

Parameters

Table 223-15 GETCONTENTBLOB Function Parameters

Parameter	Description
abspath	Absolute path of the resource
csid	If TRUE, lock and return the source LOB. If FALSE, return a temp LOB copy.
locksrc	Contents of the resource as a BLOB

Return Values

The contents of the resource as a BLOB.

223.4.14 GETCONTENTCLOB Function

This function gets the contents of a resource returned as a CLOB.

Syntax

```
DBMS_XDB_REPOS.GETCONTENTCLOB(
     abspath    IN     VARCHAR2,
  RETURN CLOB;
```

Parameters

Table 223-16 GETCONTENTCLOB Function Parameters

Parameter	Description
abspath	Absolute path of the resource

Return Values

The contents of the resource as a CLOB.

223.4.15 GETCONTENTVARCHAR2 Function

This function gets the contents of a resource returned as a string.

Syntax

```
DBMS_XDB_REPOS.GETCONTENTVARCHAR2(
     abspath    IN     VARCHAR2,
  RETURN BLOB;
```

Parameters

Table 223-17 GETCONTENTVARCHAR2 Function Parameters

Parameter	Description
abspath	Absolute path of the resource

Return Values

The contents of the resource as a string.

223.4.16 GETCONTENTXMLREF Function

This function retrieves the contents of a resource returned as a a `REF` to an `XMLTYPE`.

Syntax

```
DBMS_XDB_REPOS.GETCONTENTXMLREF(
     abspath    IN     VARCHAR2,
  RETURN SYS.XMLTYPE;
```

Parameters

Table 223-18 GETCONTENTXMLREF Function Parameters

Parameter	Description
abspath	Absolute path of the resource

Return Values

The contents of the resource as a `REF` to an `XMLTYPE`.

223.4.17 GETCONTENTXMLTYPE Function

This function retrieves the contents of a resource returned as an `XMLTYPE`.

Syntax

```
DBMS_XDB_REPOS.GETCONTENTXMLTYPE(
     abspath     IN      VARCHAR2,
  RETURN SYS.XMLTYPE;
```

Parameters

Table 223-19 GETCONTENTXMLTYPE Function Parameters

Parameter	Description
`abspath`	Absolute path of the resource

Return Values

The contents of the resource as an XMLTYPE.

223.4.18 GETLOCKTOKEN Procedure

Given a path to a resource, this procedure returns that resource's lock token for the current user.

Syntax

```
DBMS_XDB_REPOS.GETLOCKTOKEN(
   path          IN      VARCHAR2,
   locktoken     OUT     VARCHAR2);
```

Parameters

Table 223-20 GETLOCKTOKEN Procedure Parameters

Parameter	Description
`path`	Path name to the resource
`locktoken`	Logged-in user's lock token for the resource

Usage Notes

The user must have `READPROPERTIES` privilege on the resource.

223.4.19 GETPRIVILEGES Function

This function gets all privileges granted to the current user on a specified resource.

Syntax

```
DBMS_XDB_REPOS.GETPRIVILEGES(
   res_path     IN      VARCHAR2)
 RETURN sys.xmltype;
```

Parameters

Table 223-21 GETPRIVILEGES Function Parameters

Parameter	Description
res_path	Absolute path in the hierarchy of the resource

Return Values

An XMLType instance of <privilege> element, which contains the list of all leaf privileges granted on this resource to the current user.

223.4.20 GETRESOID Function

The GETRESOID function returns the object ID of the resource from its absolute path.

Syntax

```
DBMS_XDB_REPOS.GETRESOID(
   abspath IN VARCHAR2)
RETURN RAW;
```

Parameters

Table 223-22 GETRESOID Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

NULL if the resource is not present.

223.4.21 GETXDB_TABLESPACE Function

This function returns the current tablespace of the XDB (user).

Syntax

```
DBMS_XDB_REPOS.GETXDB_TABLESPACE
 RETURN VARCHAR2;
```

223.4.22 HASBLOBCONTENT Function

This function returns TRUE if the resource has BLOB content.

Syntax

```
DBMS_XDB_REPOS.HASBLOBCONTENT
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-23 HASBLOBCONTENT Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

TRUE if the resource has BOB content.

223.4.23 HASCHARCONTENT Function

This function returns TRUE if the resource has character content.

Syntax

```
DBMS_XDB_REPOS.HASCHARCONTENT
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-24 HASCHARCONTENT Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

TRUE if the resource has character content.

223.4.24 HASXMLCONTENT Function

This function returns TRUE if the resource has XML content.

Syntax

```
DBMS_XDB_REPOS.HASXMLCONTENT
     abspath    IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-25 HASXMLCONTENT Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

TRUE if the resource has XML content.

223.4.25 HASXMLREFERENCE Function

This function returns `TRUE` if the resource has a `REF` to XML content.

Syntax

```
DBMS_XDB_REPOS.HASXMLREFERENCE
     abspath     IN      VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-26 HASXMLREFERENCE Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

`TRUE` resource has a `REF` to XML content.

223.4.26 ISFOLDER Function

This function returns `TRUE` if the resource is a folder or container.

Syntax

```
DBMS_XDB_REPOS.ISFOLDER
     abspath     IN      VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-27 ISFOLDER Function Parameters

Parameter	Description
abspath_path	Absolute path of the resource

Return Values

`TRUE` if the resource is a folder or container.

223.4.27 LINK Procedures

This procedure creates from a specified folder to a specified resource.

Syntax

```
DBMS_XDB_REPOS.LINK(
   srcpath      IN   VARCHAR2,
   linkfolder   IN   VARCHAR2,
   linkname     IN   VARCHAR2);
```

```
DBMS_XDB_REPOS.LINK(
   srcpath      IN   VARCHAR2,
   linkfolder   IN   VARCHAR2,
   linkname     IN   VARCHAR2,
   linktype     IN   PLS_INTEGER := DBMS_XDB_REPOS.LINK_TYPE_HARD);
```

Parameters

Table 223-28 LINK Procedure Parameters

Parameter	Description
`srcpath`	Path name of the resource to which a link is created
`linkfolder`	Folder in which the new link is placed
`linkname`	Name of the new link
`linktype`	Type of link to be created: • `DBMS_XDB.LINK_TYPE_HARD` (default) • `DBMS_XDB.LINK_TYPE_WEAK` • `DBMS_XDB.LINK_TYPE_SYMBOLIC`

223.4.28 LOCKRESOURCE Function

Given a path to a resource, this function gets a WebDAV-style lock on that resource.

Syntax

```
DBMS_XDB_REPOS.LOCKRESOURCE(
   path      IN  VARCHAR2,
   depthzero IN  BOOLEAN,
   shared    IN  boolean)
RETURN BOOLEAN;
```

Parameters

Table 223-29 LOCKRESOURCE Function Parameters

Parameter	Description
`path`	Path name of the resource to lock.
`depthzero`	Currently not supported
`shared`	Passing `TRUE` obtains a shared write lock

Return Values

`TRUE` if successful.

Usage Notes

The user must have `UPDATE` privileges on the resource.

223.4.29 PROCESSLINKS Procedure

This procedure processes document links in the specified resource.

Syntax

```
DBMS_XDB_REPOS.PURGERESOURCEMETADATA(
 abspath  IN  VARCHAR2,
 recurse  IN  BOOLEAN := FALSE);
```

Parameters

Table 223-30 PROCESSLINKS Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource. If the path is a folder, use the `recurse` flag.
`recurse`	Used only if `abspath` specifies a folder. If `TRUE`, process links of all resources in the folder hierarchy rooted at the specified resource. If `FALSE`, process links of all documents in this folder only.

223.4.30 PURGERESOURCEMETADATA Procedure

This procedure deletes all user metadata from a resource. Schema-based metadata is removed in cascade mode, rows being deleted from the corresponding metadata tables.

Syntax

```
DBMS_XDB_REPOS.PURGERESOURCEMETADATA(
 abspath  IN  VARCHAR2);
```

Parameters

Table 223-31 PURGERESOURCEMETADATA Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource

223.4.31 RENAMERESOURCE Procedure

This procedure renames the XDB resource.

Syntax

```
DBMS_XDB_REPOS.RENAMERESOURCE(
   srcpath    IN  VARCHAR2,
   destfolder IN  CARCHAR2,
   newname    IN  VARCHAR2);
```

Parameters

Table 223-32 RENAMERESOURCE Procedure Parameters

Parameter	Description
srcpath	Absolute path in the Hierarchy for the source resource destination folder
destfolder	Absolute path in the Hierarchy for the destination folder
newname	Name of the child in the destination folder

223.4.32 SETACL Procedure

This procedure sets the ACL on a specified resource to be the ACL specified by path.

Syntax

```
DBMS_XDB_REPOS.SETACL(
   res_path   IN  VARCHAR2,
   acl_path   IN  VARCHAR2);
```

Parameters

Table 223-33 SETACL Procedure Parameters

Parameter	Description
res_path	Absolute path in the Hierarchy for resource
acl_path	Absolute path in the Hierarchy for ACL

Usage Notes

The user must have <write-acl> privileges on the resource.

223.4.33 SPLITPATH Procedure

This procedure splits the path into a parentpath and childpath.

Syntax

```
DBMS_XDB_REPOS.SPLITPATH(
      abspath     IN  VARCHAR2,
      parentpath  OUT VARCHAR2,
     childpath    OUT VARCHAR2);
```

Parameters

Table 223-34 SPLITPATH Procedure Parameters

Parameter	Description
abspath	Absolute path to be split
parentpath	Parentpath
childpath	Childpath

223.4.34 TOUCHRESOURCE Procedure

This procedure changes the modification time of the resource to the current time.

Syntax

```
DBMS_XDB_REPOS.TOUCHRESOURCE
     abspath     IN      VARCHAR2);
```

Parameters

Table 223-35 TOUCHRESOURCE Procedure Parameters

Parameter	Description
abspath_path	Absolute path of the resource

223.4.35 UNLOCKRESOURCE Function

This function unlocks the resource given a lock token and a path to the resource.

Syntax

```
DBMS_XDB_REPOS.UNLOCKRESOURCE(
   path     IN  VARCHAR2,
   deltoken IN  VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 223-36 UNLOCKRESOURCE Function Parameters

Parameter	Description
path	Path name to the resource
deltoken	Lock token to be removed

Return Values

TRUE if operation successful.

Usage Notes

The user must have UPDATE privileges on the resource.

223.4.36 UPDATERESOURCEMETADATA Procedures

This procedure updates metadata for a resource.

The procedure takes in a resource identified by absolute path and the metadata in it to replace identified by its REF. It replaces that piece of metadata with user-defined metadata which is either in the form of a REF to XMLTYPE or an XMLTYPE.

Syntax

Can be used to update schema-based metadata only. The new metadata must be schema-based:

```
DBMS_XDB_REPOS.UPDATERESOURCEMETADATA(
   abspath  IN VARCHAR2,
   oldmetadata    IN   REF SYS.XMLTYPE,
   newmetadata    IN   REF SYS.XMLTYPE)
```

Can be used to update schema-based metadata only. The new metadata must be schema-based or nonschema-based:

```
DBMS_XDB_REPOS.UPDATERESOURCEMETADATA(
   abspath        IN   VARCHAR2,
   oldmetadata    IN   REF SYS.XMLTYPE,
   newmetadata    IN   XMLTYPE);
```

Can be used for both schema-based and nonschema-based metadata:

```
DBMS_XDB_REPOS.UPDATERESOURCEMETADATA(
   abspath        IN   VARCHAR2,
   oldns          IN   VARCHAR2,
   oldname IN VARCHAR,
   newmetadata IN XMLTYPE);
```

Can be used for both schema-based or nonschema-based metadata. New metadata must be schema-based:

```
DBMS_XDB_REPOS.UPDATERESOURCEMETADATA(
   abspath       IN   VARCHAR2,
   oldns         IN   VARCHAR2,
   oldname       IN   VARCHAR,
   newmetadata   IN   REF SYS.XMLTYPE);
```

Parameters

Table 223-37 UPDATERESOURCEMETADATA Procedure Parameters

Parameter	Description
`abspath`	Absolute path of the resource
`oldmetadata`	`REF` to the old of metadata
`newmetadata`	`REF` to the new, replacement metadata (can be either schema-based or nonschema-based depending on the overload)
`oldns`	Namespace identifying old metadata
`oldname`	Local name identifying old metadata

Usage Notes

In the case of `REF`, it stores the `REF` in the resource and the metadata is stored in a separate table. Uniqueness of `REFs` is enforced. In the case where the `XMLTYPE` is passed in, data is parsed to determine if it is schema-based or not and is stored accordingly.

224
DBMS_XDB_VERSION

Oracle XML DB versioning interfaces are found in the `DBMS_XBD_VERSION` package. Functions and procedures of `DBMS_XDB_VERSION` help to create a VCR and manage the versions in the version history.

> **Note:**
>
> The Oracle XML DB Repository is deprecated with Oracle Database 23ai.
>
> Oracle recommends that you replace any functionality used in XML DB Repository with alternative technologies.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_XDB_VERSION Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

224.1 DBMS_XDB_VERSION Security Model

Owned by `XDB`, the `DBMS_XDB_VERSION` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`. Subprograms in this package are executed using the privileges of the current user.

224.2 Summary of DBMS_XDB_VERSION Subprograms

This table lists the `DBMS_XDB_VERSION` subprograms and briefly describes them.

Table 224-1 ***DBMS_XDB_VERSION Package Subprograms***

Method	Description
CHECKIN Function	Checks in a checked-out VCR and returns the resource id of the newly-created version
CHECKOUT Procedure	Checks out a VCR before updating or deleting it
GETCONTENTSBLOBBYRESID Function	Obtain contents as a `BLOB`
GETCONTENTSCLOBBYRESID Function	Obtain contents as a `CLOB`

Table 224-1 (Cont.) *DBMS_XDB_VERSION Package Subprograms*

Method	Description
GETCONTENTSXMLBYRESID Function	Obtain contents as an `XMLType`
GETPREDECESSORS Function	Retrieves the list of predecessors by path name
GETPREDSBYRESID Function	Retrieves the list of predecessors by resource id
GETRESOURCEBYRESID Function	Obtains the resource as an `XMLType`, given the resource object ID
GETSUCCESSORS Function	Retrieves the list of successors by path name
GETSUCCSBYRESID Function	Retrieves the list of successors by resource id
MAKEVERSIONED Function	Turns a regular resource whose path name is given into a version-controlled resource
UNCHECKOUT Function	Checks in a checked-out resource, returns the resource id of the version before the resource is checked out

224.2.1 CHECKIN Function

This function checks in a checked-out VCR and returns the resource id of the newly-created version.

Syntax

```
DBMS_XDB_VERSION.CHECKIN(
   pathname VARCHAR2)
 RETURN DBMS_XDB.resid_type;
```

Parameters

Table 224-2 CHECKIN Function Parameters

Parameter	Description
`pathname`	The path name of the checked-out resource.

Usage Notes

This is not an auto-commit SQL operation. CHECKIN Function doesn't have to take the same path name that was passed to CHECKOUT Procedure operation. However, the CHECKIN Function path name and the CHECKOUT Procedure path name must be of the same resource for the operations to function correctly. If the resource has been renamed, the new name must be used to CHECKIN Function because the old name is either invalid or is currently bound with a different resource. Exception is raised if the path name does not exist. If the path name has been changed, the new path name must be used to CHECKIN Function the resource.

224.2.2 CHECKOUT Procedure

This procedure checks out a VCR before updating or deleting it.

Syntax

```
DBMS_XDB_VERSION.Checkout(
   pathname    VARCHAR2);
```

Parameters

Table 224-3 CHECKOUT Procedure Parameters

Parameter	Description
pathname	The path name of the VCR to be checked out.

Usage Notes

This is not an auto-commit SQL operation. Two users of the same workspace cannot CHECKOUT Procedure the same VCR at the same time. If this happens, one user must rollback. As a result, it is good practice to commit the CHECKOUT Procedure operation before updating a resource and avoid loss of the update if the transaction is rolled back. An exception is raised if the given resource is not a VCR, if the VCR is already checked out, if the resource doesn't exist.

224.2.3 GETCONTENTSBLOBBYRESID Function

This function obtain contents as a `BLOB`.

Syntax

```
DBMS_XDB_VERSION.GETCONTENTSBLOBYRESID(
   resid      DBMS_XDB.resid_type)
 RETURN BLOB;
```

Parameters

Table 224-4 GETCONTENTSBLOBYRESID Function Parameters

Parameter	Description
resid	The resource id.

224.2.4 GETCONTENTSCLOBBYRESID Function

This function obtains contents as a `CLOB`.

Syntax

```
DBMS_XDB_VERSION.GETCONTENTSCLOBYRESID(
   resid     DBMS_XDB.resid_type)
 RETURN CLOB;
```

Parameters

Table 224-5 GETCONTENTSCLOBYRESID Function Parameters

Parameter	Description
resid	The resource id.

224.2.5 GETCONTENTSXMLBYRESID Function

This function obtains contents as an XMLType.

Syntax

```
DBMS_XDB_VERSION.GETCONTENTSXMLBYRESID(
   resid      DBMS_XDB.resid_type)
 RETURN XMLType;
```

Parameters

Table 224-6 GETCONTENTSXMLBYRESID Function Parameters

Parameter	Description
resid	The resource id.

Return Values

If the contents are not valid XML, returns NULL.

224.2.6 GETPREDECESSORS Function

This function retrieves the list of predecessors by the path name.

Syntax

```
DBMS_XDB_VERSION.GETPREDECESSORS(
   pathname       VARCHAR2)
 RETURN resid_list_type;
```

Parameters

Table 224-7 GETPREDECESSORS Function Parameters

Parameter	Description
pathname	The path name of the resource.

Return Values

An exception is raised if pathname is illegal.

224.2.7 GETPREDSBYRESID Function

This function retrieves the list of predecessors by resource id.

Syntax

```
DBMS_XDB_VERSION.GETPREDSBYRESID(
   resid      resid_type)
 RETURN resid_list_type;
```

Parameters

Table 224-8 GETPREDSBYRESID Function Parameters

Parameter	Description
`resid`	The resource id.

Usage Notes

Getting predecessors by `RESID` is more efficient than by `pathname`.

Exceptions

An exception is raised if the `RESID` is illegal.

224.2.8 GETRESOURCEBYRESID Function

This function obtains the resource as an `XMLType`, given the resource object ID. Because the system does not create a path name for versions, this function is useful for retrieving the resource using its resource id.

Syntax

```
DBMS_XDB_VERSION.GETRESOURCEBYRESID(
   resid      resid_type)
 RETURN XMLType;
```

Parameters

Table 224-9 GETRESOURCEBYRESID Function Parameters

Parameter	Description
`resid`	The resource id.

224.2.9 GETSUCCESSORS Function

Given a version resource or a VCR, this function retrieves the list of the successors of the resource by the path name.

Syntax

```
DBMS_XDB_VERSION.GETSUCCESSORS(
   pathname VARCHAR2)
 RETURN resid_list_type;
```

Parameters

Table 224-10 GETSUCCESSORS Function Parameters

Parameter	Description
pathname	The path name of the resource.

Usage Notes

Getting successors by RESID is more efficient than by pathname.

Exceptions

An exception is raised if the pathname is illegal.

224.2.10 GETSUCCSBYRESID Function

This function retrieves the list of the successors of the resource by resource id using version resource or VCR.

Syntax

```
DBMS_XDB_VERSION.GETSUCCSBYRESID(
   resid    resid_type)
  RETURN resid_list_type;
```

Parameters

Table 224-11 GETSUCCSBYRESID Function Parameters

Parameter	Description
resid	The resource id.

Usage Notes

Getting successors by RESID is more efficient than by pathname.

Exceptions

An exception is raised if the pathname is illegal.

224.2.11 MAKEVERSIONED Function

This function turns a regular resource whose path name is given into a version-controlled resource. This new resource is then put under version control. All other path names continue to refer to the original resource.

Syntax

```
DBMS_XDB_VERSION.MAKEVERSIONED(
   pathname   VARCHAR2)
 RETURN DBMS_XDB.resid_type;
```

Parameters

Table 224-12 MAKEVERSIONED Function Parameters

Parameter	Description
pathname	The path name of the resource to be put under version control.

Return Values

This function returns the resource ID of the first version, or root, of the VCR.

Usage Notes

If two or more path names are bound with the same resource, a copy of the resource is created, and the given path name is bound with the newly-created copy.

This is not an auto-commit SQL operation. An exception is raised if the resource doesn't exist.

- This call is legal for VCR, and neither exception nor warning is raised.
- This call is illegal for folder, version history, version resource, and ACL.
- No support for Schema-based resources is provided.

224.2.12 UNCHECKOUT Function

This function checks-in a checked-out resource and returns the resource id of the version before the resource is checked out.

Syntax

```
DBMS_XDB_VERSION.UNCHECKOUT(
   pathname    VARCHAR2)
 RETURN DBMS_XDB.resid_type;
```

Parameters

Table 224-13 UNCHECKOUT Function Parameters

Parameter	Description
pathname	The path name of the checked-out resource.

Usage Notes

This is not an auto-commit SQL operation. The UNCHECKOUT Function does not have to take the same path name that was passed to the operation by the CHECKOUT Procedure. However, the UNCHECKOUT Function path name and the CHECKOUT Procedure path name must be of the same resource for the operations to function correctly. If the resource has been renamed, the new name must be used to UNCHECKOUT Function, because the old name is either invalid or is currently bound with a different resource. If the path name has been changed, the new path name must be used to UNCHECKOUT Function the resource.

Exceptions

An exception is raised if the path name doesn't exist.

225
DBMS_XDBRESOURCE

The DBMS_XDBRESOURCE package provides the interface to operate on the resource's metadata and contents.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_XDBRESOURCE Subprograms

> **See Also:**
>
> Oracle XML DB Developer's Guide for examples of "Using DBMS_XDBRESOURCE"

225.1 DBMS_XDBRESOURCE Overview

The DBMS_XDBRESOURCE package provides routines to get and set the resource's metadata and contents. To take advantage of the DOM traversal facility, provided in DBMS_XMLDOM package, an XDBResource instance could be converted to a DOMDocument type by using DBMS_XDBRESOURCE.MAKEDOCUMENT routine.

225.2 DBMS_XDBRESOURCE Security Model

Owned by XDB, the DBMS_XDBRESOURCE package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

225.3 Summary of DBMS_XDBRESOURCE Subprograms

This table lists the DBMS_XDBRESOURCE subprograms and briefly describes them.

Table 225-1 ***DBMS_XDBRESOURCE Package Subprograms***

Subprogram	Description
FREERESOURCE Procedure	Frees any memory associated with an XDBResource
GETACL Function	Given an XDBResource, returns its ACL as string
GETACLDOCFROMRES Function	Returns the ACL Document for the given resource as XMLType
GETAUTHOR Function	Given an XDBResource, returns its author
GETCHARACTERSET Function	Given an XDBResource, returns its character set
GETCOMMENT Function	Given an XDBResource, returns its comment

Table 225-1 (Cont.) *DBMS_XDBRESOURCE Package Subprograms*

Subprogram	Description
GETCONTENTBLOB Function	Returns the contents of the resource as a `BLOB`
GETCONTENTCLOB Function	Returns the contents of the resource as a `CLOB`
GETCONTENTREF Function	Returns the contents of the resource as an `XMLTypeRef`
GETCONTENTTYPE Function	Given an XDBResource, returns its content-type
GETCONTENTXML Function	Returns the contents of the resource as XML
GETCONTENTVARCHAR2 Function	Returns the contents of the resource as a string
GETCREATIONDATE Function	Given an XDBResource, returns its creation date
GETCREATOR Function	Given an XDBResource, returns its creator
GETCUSTOMMETADATA Function	Returns the requested custom metadata given the xpath and namespace to the metadata
GETDISPLAYNAME Function	Given an XDBResource, returns its display name
GETLANGUAGE Function	Given an XDBResource, returns its language
GETLASTMODIFIER Function	Given an XDBResource, returns its last modifier
GETMODIFICATIONDATE Function	Given an XDBResource, returns its modification date
GETOWNER Function	Given an XDBResource, returns its owner.
GETREFCOUNT Function	Given an XDBResource, returns its reference count
GETVERSIONID Function	Given an XDBResource, returns its version ID.
HASACLCHANGED Function	Returns `TRUE` if the ACL of the given resource has changed, `FALSE` otherwise
HASAUTHORCHANGED Function	Returns `TRUE` if the ACL of the given resource has changed `FALSE` otherwise
HASCHANGED Function	Returns `TRUE` if the element or attribute represented by the given XPath has changed, `FALSE` otherwise
HASCHARACTERSETCHANGED Function	Returns `TRUE` if the character set of the given resource has changed, `FALSE` otherwise
HASCOMMENTCHANGED Function	Returns `TRUE` if the comment of the given resource has changed, `FALSE` otherwise
HASCONTENTCHANGED Function	Returns `TRUE` if the contents of the given resource has changed, `FALSE` otherwise
HASCONTENTTYPECHANGED Function	Returns `TRUE` if the content-type of the given resource has changed, `FALSE` otherwise
HASCREATIONDATECHANGED Function	Returns `TRUE` if the creation date of the given resource has changed, `FALSE` otherwise
HASCREATORCHANGED Function	Returns `TRUE` if the creator of the given resource has changed, `FALSE` otherwise
HASCUSTOMMETADATACHANGED Function	Returns `TRUE` if custom-metadata for this XPath has changed, `FALSE` otherwise
HASDISPLAYNAMECHANGED Function	Returns `TRUE` if the display name of the given resource has changed, `FALSE` otherwise
HASLANGUAGECHANGED Function	Returns `TRUE` if the language of the given resource has changed, `FALSE` otherwise

Table 225-1 (Cont.) *DBMS_XDBRESOURCE Package Subprograms*

Subprogram	Description
HASLASTMODIFIERCHANGED Function	Returns `TRUE` if the last modifier of the given resource has changed, `FALSE` otherwise
HASMODIFICATIONDATECHANGED Function	Returns `TRUE` if the modification date of the given resource has changed, `FALSE` otherwise
HASOWNERCHANGED Function	Returns `TRUE` if the owner of the given resource has changed, `FALSE` otherwise
HASREFCOUNTCHANGED Function	Returns `TRUE` if the reference count of the given resource has changed, `FALSE` otherwise
HASVERSIONIDCHANGED Function	Returns `TRUE` if the version ID of the given resource has changed, `FALSE` otherwise
ISFOLDER Function	Returns `TRUE` if the given resource is a folder, `FALSE` otherwise
ISNULL Function	Returns `TRUE` if input resource is `NULL`, `FALSE` otherwise
MAKEDOCUMENT Function	Converts the XDBResource to a DOMDocument which can be operated on using the XMLDOM interface
SAVE Procedure	Updates the resource with any modifications
SETACL Procedure	Sets the ACL of the given XDBResource to the path specified
SETAUTHOR Procedure	Sets the author of the given XDBResource to the specified string
SETCHARACTERSET Procedure	Sets the character set of the given XDBResource to a specified character set
SETCOMMENT Procedure	Sets a comment associated with the given XDBResource
SETCONTENT Procedures	Replaces the contents of the given resource with the given `CLOB`
SETCONTENTTYPE Procedure	Sets the content-type of the given XDBResource
SETCUSTOMMETADATA Procedure	Sets the custom metadata specified by the XPath and namespace to new data
SETDISPLAYNAME Procedure	Sets the display name of the given XDBResource
SETLANGUAGE Procedure	Sets the language of the given XDBResource
SETOWNER Procedure	Sets the owner of the given XDBResource

225.3.1 FREERESOURCE Procedure

This procedure frees any memory associated with an XDBResource.

Syntax

```
DBMS_XDBRESOURCE.FREERESOURCE (
   res    IN    XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-2 FREERESOURCE Procedure Parameters

Parameter	Description
res	XDBResource to free

225.3.2 GETACL Function

Given an XDBResource, this function returns its ACL as string.

Syntax

```
DBMS_XDBRESOURCE.GETACL (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-3 GETACL Function Parameters

Parameter	Description
res	XDBResource

225.3.3 GETACLDOCFROMRES Function

This function returns the ACL Document for the given resource as `XMLType`.

Syntax

```
DBMS_XDBRESOURCE.GETACLDOCFROMRES (
   res    IN     XDBResource)
 RETURN SYS.XMLTYPE;
```

Parameters

Table 225-4 GETACLDOCFROMRES Function Parameters

Parameter	Description
res	XDBResource

225.3.4 GETAUTHOR Function

Given an XDBResource, this function returns its author.

Syntax

```
DBMS_XDBRESOURCE.GETAUTHOR (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-5 GETAUTHOR Function Parameters

Parameter	Description
res	XDBResource

225.3.5 GETCHARACTERSET Function

Given an XDBResource, this function returns its characterset.

Syntax

```
DBMS_XDBRESOURCE.GETCHARACTERSET (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-6 GETCHARACTERSET Function Parameters

Parameter	Description
res	XDBResource

225.3.6 GETCOMMENT Function

Given an XDBResource, this function returns its comment.

Syntax

```
DBMS_XDBRESOURCE.GETCOMMENT (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-7 GETCOMMENT Function Parameters

Parameter	Description
res	XDBResource

225.3.7 GETCONTENTBLOB Function

This function returns the contents of the resource as a `BLOB`.

Syntax

```
DBMS_XDBRESOURCE.GETCONTENTBLOB (
   res     IN     XDBResource,
   csid    OUT    PLS_INTEGER)
 RETURN BLOB;
```

Parameters

Table 225-8 GETCONTENTBLOB Function Parameters

Parameter	Description
res	XDBResource
csid	Character set ID of the BLOB returned

225.3.8 GETCONTENTCLOB Function

This function returns the contents of the resource as a CLOB.

Syntax

```
DBMS_XDBRESOURCE.GETCONTENTCLOB (
   res    IN    XDBResource)
 RETURN CLOB;
```

Parameters

Table 225-9 GETCONTENTCLOB Function Parameters

Parameter	Description
res	XDBResource

225.3.9 GETCONTENTREF Function

This function returns the contents of the resource as an XMLTypeRef.

Syntax

```
DBMS_XDBRESOURCE.GETCONTENTREF (
   res    IN    XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-10 GETCONTENTREF Function Parameters

Parameter	Description
res	XDBResource

225.3.10 GETCONTENTTYPE Function

Given an XDBResource, this function returns its content-type.

Syntax

```
DBMS_XDBRESOURCE.GETCONTENTTYPE (
   res    IN    XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-11 GETCONTENTTYPE Function Parameters

Parameter	Description
res	XDBResource

225.3.11 GETCONTENTXML Function

This function returns the contents of the resource as an XMLTypeRef.

Syntax

```
DBMS_XDBRESOURCE.GETCONTENTXML (
   res    IN     XDBResource)
 RETURN XMLType;
```

Parameters

Table 225-12 GETCONTENTXML Function Parameters

Parameter	Description
res	XDBResource

225.3.12 GETCONTENTVARCHAR2 Function

This function returns the contents of the resource as a string.

Syntax

```
DBMS_XDBRESOURCE.GETCONTENTVARCHAR2 (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-13 GETCONTENTVARCHAR2 Function Parameters

Parameter	Description
res	XDBResource

225.3.13 GETCREATIONDATE Function

Given an XDBResource, this function returns its creation date.

Syntax

```
DBMS_XDBRESOURCE.GETCREATIONDATE (
   res    IN     XDBResource)
 RETURN TIMESTAMP;
```

Parameters

Table 225-14 GETCREATIONDATE Function Parameters

Parameter	Description
res	XDBResource

225.3.14 GETCREATOR Function

Given an XDBResource, this function returns its creator.

Syntax

```
DBMS_XDBRESOURCE.GETCREATOR (
   res    IN    XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-15 GETCREATOR Function Parameters

Parameter	Description
res	XDBResource

225.3.15 GETCUSTOMMETADATA Function

This function returns the requested custom metadata given the xpath and namespace to the metadata.

Syntax

```
DBMS_XDBRESOURCE.GETCUSTOMMETADATA (
   res         IN    XDBResource,
   xpath       IN    VARCHAR2,    namespace  IN    VARCHAR2)
 RETURN XMLType;
```

Parameters

Table 225-16 GETCUSTOMMETADATA Function Parameters

Parameter	Description
res	XDBResource
xpath	XPath for custom metadata
namespace	Namespace

Usage Notes

The first component of the XPath expression must be "Resource".

225.3.16 GETDISPLAYNAME Function

Given an XDBResource, this function returns its display name.

Syntax

```
DBMS_XDBRESOURCE.GETDISPLAYNAME (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-17 GETDISPLAYNAME Function Parameters

Parameter	Description
res	XDBResource

225.3.17 GETLANGUAGE Function

Given an XDBResource, this function returns its language.

Syntax

```
DBMS_XDBRESOURCE.GETLANGUAGE (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-18 GETLANGUAGE Function Parameters

Parameter	Description
res	XDBResource

225.3.18 GETLASTMODIFIER Function

Given an XDBResource, this function returns its last modifier.

Syntax

```
DBMS_XDBRESOURCE.GETLASTMODIFIER (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-19 GETLASTMODIFIER Function Parameters

Parameter	Description
res	XDBResource

225.3.19 GETMODIFICATIONDATE Function

Given an XDBResource, this function returns its modification date.

Syntax

```
DBMS_XDBRESOURCE.GETMODIFICATIONDATE (
   res    IN     XDBResource)
 RETURN TIMESTAMP;
```

Parameters

Table 225-20 GETMODIFICATIONDATE Function Parameters

Parameter	Description
res	XDBResource

225.3.20 GETOWNER Function

Given an XDBResource, this function returns its owner.

Syntax

```
DBMS_XDBRESOURCE.GETOWNER (
   res    IN     XDBResource)
 RETURN VARCHAR2;
```

Parameters

Table 225-21 GETOWNER Function Parameters

Parameter	Description
res	XDBResource

225.3.21 GETREFCOUNT Function

Given an XDBResource, this function returns its reference count.

Syntax

```
DBMS_XDBRESOURCE.GETREFCOUNT (
   res    IN     XDBResource)
 RETURN PLS_INTEGER;
```

Parameters

Table 225-22 GETREFCOUNT Function Parameters

Parameter	Description
res	XDBResource

225.3.22 GETVERSIONID Function

Given an XDBResource, this function returns its version ID.

Syntax

```
DBMS_XDBRESOURCE.GETVERSIONID (
   res    IN     XDBResource)
 RETURN PLS_INTEGER;
```

Parameters

Table 225-23 GETVERSIONID Function Parameters

Parameter	Description
res	XDBResource

225.3.23 HASACLCHANGED Function

This function returns `TRUE` if the ACL of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASACLCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-24 HASACLCHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.24 HASAUTHORCHANGED Function

This function returns `TRUE` if the author of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASAUTHORCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-25 HASAUTHORCHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.25 HASCHANGED Function

Given an XPath, this function determines whether the element or attribute represented by the XPath has changed.

Syntax

```
DBMS_XDBRESOURCE.HASCHANGED (
   res          IN    XDBResource,
   xpath        IN    VARCHAR2,
   namespace    IN    VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 225-26 HASCHANGED Function Parameters

Parameter	Description
`res`	XDBResource
`xpath`	XPath to check
`bnamespace`	Namespace to use

225.3.26 HASCHARACTERSETCHANGED Function

This function returns `TRUE` if the character set of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASCHARACTERSETCHANGED (
   res    IN    XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-27 HASCHARACTERSETCHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.27 HASCOMMENTCHANGED Function

This function returns `TRUE` if the comment of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASCOMMENTCHANGED (
   res    IN    XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-28 HASCOMMENTCHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.28 HASCONTENTCHANGED Function

This function returns TRUE if the contents of the given resource has changed, FALSE otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASCONTENTCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-29 HASCONTENTCHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.29 HASCONTENTTYPECHANGED Function

This function returns TRUE if the content-type of the given resource has changed, FALSE otherwise

Syntax

```
DBMS_XDBRESOURCE.HASCONTENTTYPECHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-30 HASCONTENTTYPECHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.30 HASCREATIONDATECHANGED Function

This function returns `TRUE` if the creation date of the given resource has changed, `FALSE` otherwise

Syntax

```
DBMS_XDBRESOURCE.HASCREATIONDATECHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-31 HASCREATIONDATECHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.31 HASCREATORCHANGED Function

This function returns `TRUE` if the creator of the given resource has changed, `FALSE` otherwise

Syntax

```
DBMS_XDBRESOURCE.HASCREATORCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-32 HASCREATORCHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.32 HASCUSTOMMETADATACHANGED Function

This function checks whether the custom-metadata for a given resource has changed.

Syntax

```
DBMS_XDBRESOURCE.HASCUSTOMMETADATACHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-33 HASCUSTOMMETADATACHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.33 HASDISPLAYNAMECHANGED Function

This function returns `TRUE` if the display name of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASDISPLAYNAMECHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-34 HASDISPLAYNAMECHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.34 HASLANGUAGECHANGED Function

This function returns `TRUE` if the language of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASLANGUAGECHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-35 HASLANGUAGECHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.35 HASLASTMODIFIERCHANGED Function

This function returns `TRUE` if the last modifier of the given resource has changed, `FALSE` otherwise

Syntax

```
DBMS_XDBRESOURCE.HASLASTMODIFIERCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-36 HASLASTMODIFIERCHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.36 HASMODIFICATIONDATECHANGED Function

This function returns TRUE if the modification date of the given resource has changed, FALSE otherwise

Syntax

```
DBMS_XDBRESOURCE.HASMODIFICATIONDATECHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-37 HASMODIFICATIONDATECHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.37 HASOWNERCHANGED Function

This function returns TRUE if the owner of the given resource has changed, FALSE otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASOWNERCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-38 HASOWNERCHANGED Function Parameters

Parameter	Description
res	XDBResource

225.3.38 HASREFCOUNTCHANGED Function

This function returns `TRUE` if the reference count of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASREFCOUNTCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-39 HASREFCOUNTCHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.39 HASVERSIONIDCHANGED Function

This function returns `TRUE` if the version ID of the given resource has changed, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.HASVERSIONIDCHANGED (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-40 HASVERSIONIDCHANGED Function Parameters

Parameter	Description
`res`	XDBResource

225.3.40 ISFOLDER Function

This function returns `TRUE` if the given resource is a folder, `FALSE` otherwise.

Syntax

```
DBMS_XDBRESOURCE.ISFOLDER (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-41 ISFOLDER Function Parameters

Parameter	Description
res	XDBResource

225.3.41 ISNULL Function

This function returns TRUE if input resource is NULL.

Syntax

```
DBMS_XDBRESOURCE.ISNULL (
   res    IN     XDBResource)
 RETURN BOOLEAN;
```

Parameters

Table 225-42 ISNULL Function Parameters

Parameter	Description
res	Input resource

225.3.42 MAKEDOCUMENT Function

This function converts the XDBResource to a DOMDocument which can be operated on using the XMLDOM interface.

> **See Also:**
>
> The DBMS_XMLDOM package

Syntax

```
DBMS_XDBRESOURCE.MAKEDOCUMENT (
   res    IN     XDBResource)
 RETURN DBMS_XMLDOM.DOMDocument;
```

Parameters

Table 225-43 MAKEDOCUMENT Function Parameters

Parameter	Description
res	XDBResource

225.3.43 SAVE Procedure

This procedures updates the resource with any modifications.

Syntax

```
DBMS_XDBRESOURCE.SAVE (
   res    IN    XDBResource);
```

Parameters

Table 225-44 SAVE Procedure Parameters

Parameter	Description
res	XDBResource

225.3.44 SETACL Procedure

This procedure sets the ACL of the given XDBResource to the path specified.

Syntax

```
DBMS_XDBRESOURCE.SETACL (
   res      IN OUT  XDBResource,
   ACLPath  IN      VARCHAR2);
```

Parameters

Table 225-45 SETACL Procedure Parameters

Parameter	Description
res	XDBResource
ACLPath	Absolute path of the new ACL

225.3.45 SETAUTHOR Procedure

This procedure sets the author of the given XDBResource to the specified string.

Syntax

```
DBMS_XDBRESOURCE.SETAUTHOR (
   res     IN OUT  XDBResource,
   author  IN      VARCHAR2);
```

Parameters

Table 225-46 SETAUTHOR Procedure Parameters

Parameter	Description
res	XDBResource
author	Author

225.3.46 SETCHARACTERSET Procedure

This procedure sets the character set of the given XDBResource to a specified character set.

Syntax

```
DBMS_XDBRESOURCE.SETCHARACTERSET (
   res      IN OUT  XDBResource,
   charSet  IN      VARCHAR2);
```

Parameters

Table 225-47 SETCHARACTERSET Procedure Parameters

Parameter	Description
res	XDBResource
charset	New character set

225.3.47 SETCOMMENT Procedure

This procedure sets a comment associated with the given XDBResource.

Syntax

```
DBMS_XDBRESOURCE.SETCOMMENT (
   res      IN OUT  XDBResource,
   comment  IN      VARCHAR2);
```

Parameters

Table 225-48 SETCOMMENT Procedure Parameters

Parameter	Description
res	XDBResource
comment	New comment

225.3.48 SETCONTENT Procedures

This procedure replaces the contents of the given resource with the given datatype.

Syntax

```
DBMS_XDBRESOURCE.SETCONTENT (
   res      IN OUT  XDBResource,
   data     IN      BFILE,
   csid     IN      NUMBER);

DBMS_XDBRESOURCE.SETCONTENT (
   res      IN OUT  XDBResource,
   data     IN      BLOB,
   csid     IN      PLS_INTEGER);

DBMS_XDBRESOURCE.SETCONTENT (
```

```
   res      IN OUT  XDBResource,
   data     IN      CLOB);

DBMS_XDBRESOURCE.SETCONTENT (
   res      IN OUT  XDBResource,
   data     IN      REF SYS.XMLType,
   sticky   IN      BOOLEAN := TRUE);

DBMS_XDBRESOURCE.SETCONTENT (
   res      IN OUT  XDBResource,
   data     IN      VARCHAR2);

DBMS_XDBRESOURCE.SETCONTENT (
   res      IN OUT  XDBResource,
   data     IN      SYS.XMLType);
```

Parameters

Table 225-49 SETCONTENT Procedure Parameters

Parameter	Description
res	XDBResource
data	Data input as BFILE, BLOB, CLOB, string, XMLType
csid	Character set ID of the BFILE, BLOB
sticky	If TRUE creates a sticky REF, otherwise non-sticky

225.3.49 SETCONTENTTYPE Procedure

This procedure sets the content-type of the given XDBResource.

Syntax

```
DBMS_XDBRESOURCE.SETCONTENTTYPE (
   res         IN OUT  XDBResource,
   conttype    IN      VARCHAR2);
```

Parameters

Table 225-50 SETCONTENTTYPE Procedure Parameters

Parameter	Description
res	XDBResource
conttype	New content-type

225.3.50 SETCUSTOMMETADATA Procedure

This procedure sets the custom metadata specified by the xpath and namespace to new data.

Syntax

```
DBMS_XDBRESOURCE.SETCUSTOMMETADATA (
   res          IN OUT  XDBResource,
   xpath        IN      VARCHAR2,
```

```
   namespace    IN      VARCHAR2,
   newMetadata  IN      XMLType);
```

Parameters

Table 225-51 SETCUSTOMMETADATA Procedure Parameters

Parameter	Description
res	XDBResource
xpath	XPath to change
namespace	Namespace to use
newMetadata	New data that should replace the metadata at the given XPath

Usage Notes

The first component of the XPath expression must be "Resource".

225.3.51 SETDISPLAYNAME Procedure

This procedure sets the display name of the given XDBResource.

Syntax

```
DBMS_XDBRESOURCE.SETDISPLAYNAME (
   res      IN OUT  XDBResource,
   name     IN      VARCHAR2);
```

Parameters

Table 225-52 SETDISPLAYNAME Procedure Parameters

Parameter	Description
res	XDBResource
name	New display name

225.3.52 SETLANGUAGE Procedure

This procedure sets the language of the given `XDBResource`.

Syntax

```
DBMS_XDBRESOURCE.SETLANGUAGE (
   res      IN OUT  XDBResource,
   ACLPath  IN      VARCHAR2);
```

Parameters

Table 225-53 SETLANGUAGE Procedure Parameters

Parameter	Description
res	XDBResource

Table 225-53 (Cont.) SETLANGUAGE Procedure Parameters

Parameter	Description
ACLPath	New path

225.3.53 SETOWNER Procedure

This procedure sets the owner of the given XDBResource.

Syntax

```
DBMS_XDBRESOURCE.SETOWNER (
   res      IN OUT  XDBResource,
   owner    IN      VARCHAR2);
```

Parameters

Table 225-54 SETOWNER Procedure Parameters

Parameter	Description
res	XDBResource
owner	New owner

Usage Notes

The user must have the XDBADMIN privilege to call this subprogram.

226
DBMS_XDBZ

The DBMS_XDBZ package controls the Oracle XML DB repository security, which is based on Access Control Lists (ACLs).

This chapter contains the following topics:

- Security Model
- Constants
- Summary of DBMS_XDBZ Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

226.1 DBMS_XDBZ Security Model

Owned by `XDB`, the `DBMS_XDBZ` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`. Subprograms in this package are executed using the privileges of the current user.

226.2 DBMS_XDBZ Constants

The `DBMS_XDBZ` package defines several constants that can be used for specifying parameter values.

The package uses the constants shown in following tables.

Table 226-1 *DBMS_XDBZ Constants* - Name Format

Constant	Type	Value	Description
`NAME_FORMAT_SHORT`	`PLS_INTEGER`	`1`	DB user name or LDAP nickname
`NAME_FORMAT_DISTINGUISHED`	`PLS_INTEGER`	`2`	LDAP distinguished name

Table 226-2 *DBMS_XDBZ Constants* - Enable Option

Constant	Type	Value	Description
`ENABLE_CONTENTS`	`PLS_INTEGER`	`1`	Enables hierarchy for contents and is used by users when calling the ENABLE_HIERARCHY Procedure

Table 226-2 (Cont.) *DBMS_XDBZ Constants* - Enable Option

Constant	Type	Value	Description
ENABLE_RESMETADATA	PLS_INTEGER	2	Enables hierarchy for resource metadata, that is, this table will store schema based custom metadata for resources

Table 226-3 *DBMS_XDBZ Constants* - Enable Option Exercised

Constant	Type	Value	Description
IS_ENABLED_CONTENTS	PLS_INTEGER	1	If hierarchy was enabled for contents, that is, the ENABLE_HIERARCHY Procedure was called with hierarchy_type as ENABLE_CONTENTS
IS_ENABLED_RESMETADATA	PLS_INTEGER	2	If hierarchy was enabled for resource metadata, that is, the ENABLE_HIERARCHY Procedure was called with hierarchy_type as ENABLE_RESMETADATA

226.3 Summary of DBMS_XDBZ Subprograms

This table lists the DBMS_XDBZ subprograms and briefly describes them.

Table 226-4 *DBMS_XDBZ Package Subprograms*

Method	Description
DISABLE_HIERARCHY Procedure	Disables repository support for the specified XMLTYPE table or view
ENABLE_HIERARCHY Procedure	Enables repository support for the specified XMLType table or view
GET_ACLOID Function	Retrieves the ACL Object ID for the specified resource
GET_USERID Function	Retrieves the user ID for the specified user
IS_HIERARCHY_ENABLED Function	Determines if repository support for the specified XMLType table or view is enabled
PURGELDAPCACHE Function	Purges the LDAP nickname cache

226.3.1 CREATENONCEKEY Procedure

This procedure generates a nonce value for use in digest authentication.

Syntax

```
DBMS_XDBZ.CREATENONCEKEY;
```

226.3.2 DISABLE_HIERARCHY Procedure

This procedure disables repository support for a particular `XMLType` table or view.

Syntax

```
DBMS_XDBZ.DISABLE_HIERARCHY(
   object_schema IN VARCHAR2,
   object_name   IN VARCHAR2);
```

Parameters

Table 226-5 DISABLE_HIERARCHY Procedure Parameters

Parameter	Description
`object_schema`	Schema name of the `XMLType` table or view
`object_name`	Name of the `XMLType` table or view

226.3.3 ENABLE_HIERARCHY Procedure

This procedure enables repository support for a particular `XMLType` table or view. This allows the use of a uniform ACL-based security model across all documents in the repository.

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for more information about

Syntax

```
DBMS_XDBZ.ENABLE_HIERARCHY(
   object_schema   IN   VARCHAR2,
   object_name     IN   VARCHAR2,
   hierarchy_type  IN   PLS_INTEGER := DBMS_XDBZ.ENABLE_CONTENTS);
```

Parameters

Table 226-6 ENABLE_HIERARCHY Procedure Parameters

Parameter	Description
`object_schema`	Schema name of the `XMLType` table or view
`object_name`	Name of the `XMLType` table or view

Table 226-6 (Cont.) ENABLE_HIERARCHY Procedure Parameters

Parameter	Description
hierarchy_type	How to enable the hierarchy. • ENABLE_CONTENTS - enable hierarchy for contents, that is, this table will store contents of resources in the repository • ENABLE_RESMETADATA - enable hierarchy for resource metadata, that is, this table will store schema based custom metadata for resources If this subprogram is called on a table, another call will have no effect. Note that you cannot enable hierarchy for both contents and resource metadata.

226.3.4 GET_ACLOID Function

This function retrieves the ACL Object ID for the specified resource, if the repository path is known.

Syntax

```
DBMS_XDBZ.GET_ACLOID(
   aclpath    IN   VARCHAR2,
   acloid     OUT  RAW)
 RETURN BOOLEAN;
```

Parameters

Table 226-7 GET_ACLOID Function Parameters

Parameter	Description
aclpath	ACL resource path for the repository
acloid	Returned Object ID

Return Values

Returns TRUE if successful.

226.3.5 GET_USERID Function

This function retrieves the user ID for the specified user name. The local database is searched first, and if found, the USERID is returned in 4-byte database format. Otherwise, the LDAP directory is searched, if available, and if found, the USERID is returned in 4-byte database format.

Syntax

```
DBMS_XDBZ.GET_USERID(
   username IN  VARCHAR2,
   userid   OUT RAW,
   format   IN  BINARY_INTEGER := NAME_FORMAT_SHORT)
 RETURN BOOLEAN;
```

Parameters

Table 226-8 GET_USERID Function Parameters

Parameter	Description
username	Name of the database or LDAP user.
userid	Return parameter for the matching user id.
format	Format of the specified user name; valid options are: • DBMS_XDBZ.NAME_FORMAT_SHORT (default) -- DB user name or LDAP nickname • DBMS_XDBZ.NAME_FORMAT_DISTINGUISHIED -- LDAP distinguished name.

Return Values

Returns TRUE if successful.

226.3.6 IS_HIERARCHY_ENABLED Function

This function determines if repository support for the specified XMLType table or view is enabled.

Syntax

```
DBMS_XDBZ.IS_HIERARCHY_ENABLED(
   object_schema    IN  VARCHAR2,
   object_name      IN  VARCHAR2,
   hierarchy_type   IN  PLS_INTEGER := IS_ENABLED_CONTENTS)
 RETURN BOOLEAN;
```

Parameters

Table 226-9 IS_HIERARCHY_ENABLED Function Parameters

Parameter	Description
object_schema	Schema name of the XMLType table or view
object_name	Name of the XMLType table or view
hierarchy_type	Type of hierarchy to check for: • IS_ENABLED_CONTENTS - if hierarchy was enabled for contents, that is, the ENABLE_HIERARCHY Procedure was called with hierarchy_type as ENABLE_CONTENTS • IS_ENABLED_RESMETADATA - if hierarchy was enabled for resource metadata, that is, the ENABLE_HIERARCHY Procedure was called with hierarchy_type as ENABLE_RESMETADATA

Return Values

Returns TRUE if the given XMLTYPE table or view has the XDB Hierarchy enabled with the specified type.

226.3.7 PURGELDAPCACHE Function

This function purges the LDAP nickname cache.

It returns `TRUE` if successful.

Syntax

```
DBMS_XDBZ.PURGELDAPCACHE
 RETURN BOOLEAN;
```

227
DBMS_XEVENT

The DBMS_XEVENTpackage provides event-related types and supporting subprograms.

This chapter contains the following topics:

- Security Model
- Constants
- Subprogram Groups
- Summary of DBMS_XEVENT Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for more information about "Oracle XML DB Repository Events"

227.1 DBMS_XEVENT Security Model

Owned by `XDB`, the `DBMS_XEVENT` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`. Subprograms in this package are executed using the privileges of the current user.

227.2 DBMS_XEVENT Constants

The `DBMS_XEVENT` package defines several constants that can be used for specifying parameter values.

The constants are defined in the following table.

Table 227-1 DBMS_XEVENT Constants

Name	Type	Value	Description
RENDER_EVENT	PLS_INTEGER	1	
PRE_CREATE_EVENT	PLS_INTEGER	2	
POST_CREATE_EVENT	PLS_INTEGER	3	
PRE_DELETE_EVENT	PLS_INTEGER	4	
POST_DELETE_EVENT	PLS_INTEGER	5	
PRE_UPDATE_EVENT	PLS_INTEGER	6	
POST_UPDATE_EVENT	PLS_INTEGER	7	
PRE_LOCK_EVENT	PLS_INTEGER	8	
POST_LOCK_EVENT	PLS_INTEGER	9	

Table 227-1 (Cont.) DBMS_XEVENT Constants

Name	Type	Value	Description
PRE_UNLOCK_EVENT	PLS_INTEGER	10	
POST_UNLOCK_EVENT	PLS_INTEGER	11	
PRE_LINKIN_EVENT	PLS_INTEGER	12	
POST_LINKIN_EVENT	PLS_INTEGER	13	
PRE_LINKTO_EVENT	PLS_INTEGER	14	
POST_LINKTO_EVENT	PLS_INTEGER	15	
PRE_UNLINKIN_EVENT	PLS_INTEGER	16	
POST_UNLINKIN_EVENT	PLS_INTEGER	17	
PRE_UNLINKFROM_EVENT	PLS_INTEGER	18	
POST_UNLINKFROM_EVENT	PLS_INTEGER	19	
PRE_CHECKIN_EVENT	PLS_INTEGER	20	
POST_CHECKIN_EVENT	PLS_INTEGER	21	
PRE_CHECKOUT_EVENT	PLS_INTEGER	22	
POST_CHECKOUT_EVENT	PLS_INTEGER	23	
PRE_UNCHECKOUT_EVENT	PLS_INTEGER	24	
POST_UNCHECKOUT_EVENT	PLS_INTEGER	25	
PRE_VERSIONCONTROL_EVENT	PLS_INTEGER	26	
POST_VERSIONCONTROL_EVENT	PLS_INTEGER	27	
PRE_OPEN_EVENT	PLS_INTEGER	28	
POST_OPEN_EVENT	PLS_INTEGER	29	
PRE_INCONSISTENT_UPDATE_EVENT	PLS_INTEGER	30	
POST_INCONSISTENT_UPDATE_EVENT	PLS_INTEGER	21	
POST_CHECKIN_EVENT	PLS_INTEGER	21	

227.3 Subprogram Groups

The subprograms in the DBMS_XEVENT package can be divided in groups.

- XDBEvent Type Subprograms
- XDBRepositoryEvent Type Subprograms
- XDBHandlerList Type Subprograms
- XDBHandler Type Subprograms
- XDBPath Type Subprograms
- XDBLink Type Subprograms

227.3.1 DBMS_XEVENT XDBEvent Type Subprograms

This subprogram group provides an interface for use with the XDBEvent type.

Table 227-2 XDBEvent Subprograms

Subprogram	Description
GETCURRENTUSER Function	Returns the name of the user executing the operation that triggers the event
GETEVENT Function	Returns a value identifying the triggering event
ISNULL Functions	Returns `TRUE` if input argument is `NULL`

The Summary of DBMS_XEVENT Subprograms contains a complete listing of all subprograms in the package.

227.3.2 DBMS_XEVENT XDBRepositoryEvent Type Subprograms

This subprogram group provides an interface for use in conjunction with the XDBRepositoryEvent type.

Table 227-3 XDBRepositoryEvent Subprograms

Subprogram	Description
GETAPPLICATIONDATA Function	Returns the `<applicationData>` element extracted from the resource configuration that defines the invoking handler
GETHANDLERLIST Function	Returns an `XDBHandlerList` object containing the list of handlers that will be executed after the currently executing handler
GETINTERFACE Function	Returns the top-level interface used to initiate the operation that triggered the event
GETLINK Function	Returns an `XDBLink` object for the target resource
GETLOCK Function	Returns the lock object corresponding to the current operation
GETOLDRESOURCE Function	Returns the original `XDBResource` object before the operation was executed
GETOPENACCESSMODE Function	Returns the access mode for the open operation
GETOPENDENYMODE Function	Returns the deny mode for the open operation
GETOUTPUTSTREAM Function	Returns the output BLOB in which the handler can write the rendered data
GETPARAMETER Function	Returns the value of a request or session-specific parameter
GETPARENT Function	Returns the resource object corresponding to a parent folder of the target resource
GETPATH Function	Returns the `XDBPath` object representing the path of the resource for which the event was fired
GETRESOURCE Function	Returns an `XDBResource` object that provides methods to access and modify the contents and metadata of the target resource
GETUPDATEBYTECOUNT Function	If the current operation is a byte-range write, returns the byte count
GETUPDATEBYTEOFFSET Function	If the current operation is a byte-range write, function returns the byte offset at which the range begins

Table 227-3 (Cont.) XDBRepositoryEvent Subprograms

Subprogram	Description
GETXDBEVENT Function	Converts an `XDBRepositoryEvent` object to an `XDBEvent` type
ISNULL Functions	Returns `TRUE` if input argument is `NULL`
SETRENDERPATH Procedure	Specifies the path of the resource that contains the rendered contents
SETRENDERSTREAM Procedure	sets the `BLOB` from which the rendered contents can be read

The Summary of DBMS_XEVENT Subprograms contains a complete listing of all subprograms in the package.

227.3.3 DBMS_XEVENT XDBHandlerList Type Subprograms

This subprogram group provides an interface for use in conjunction with the XDBHandlerList type.

Table 227-4 XDBHandlerList Subprograms

Subprogram	Description
CLEAR Procedure	Clears the handler list
GETFIRST Function	Returns the first handler in the list
GETNAME Function	Returns the next handler in the list
ISNULL Functions	Returns `TRUE` if input argument is `NULL`
REMOVE Procedure	Removes the specified handler from the handler list

The Summary of DBMS_XEVENT Subprograms contains a complete listing of all subprograms in the package.

227.3.4 DBMS_XEVENT XDBHandler Type Subprograms

This subprogram group provides an interface for use in conjunction with the XDBHandler type.

Table 227-5 XDBHandler Type Subprograms

Subprogram	Description
GETLANGUAGE Function	Returns the implementation language of the handler
GETSCHEMA Function	Returns the schema of the handler's source
GETSOURCE Function	Returns the name of the Java class, PL/SQL package or object type implementing the handler
ISNULL Functions	Returns `TRUE` if input argument is `NULL`

The Summary of DBMS_XEVENT Subprograms contains a complete listing of all subprograms in the package.

227.3.5 DBMS_XEVENT XDBPath Type Subprograms

This subprogram group provides an interface for use in conjunction with the XDBPath type.

Table 227-6 XDBPath Type Subprograms

Subprogram	Description
GETNAME Function	Returns the string representation of the path
GETPARENTPATH Function	Returns the parent's path
ISNULL Functions	Returns `TRUE` if input argument is `NULL`

The Summary of DBMS_XEVENT Subprograms contains a complete listing of all subprograms in the package.

227.3.6 DBMS_XEVENT XDBLink Type Subprograms

This subprogram group provides an interface for use in conjunction with the XDBLink type.

Table 227-7 XDBLink Type Subprograms

Subprogram	Description
GETCHILDOID Function	Returns the OID of the resource to which the link is pointing
GETPARENTNAME Function	Returns the link's parent folder's name
GETPARENTOID Function	Returns the link's parent folder's OID
ISNULL Functions	Returns `TRUE` if input argument is `NULL`

The Summary of DBMS_XEVENT Subprograms contains a complete listing of all subprograms in the package.

227.4 Summary of DBMS_XEVENT Subprograms

This table lists the DBMS_XEVENT subprograms and briefly describes them.

Table 227-8 DBMS_XEVENT Package Subprograms

Subprogram	Description	Group
CLEAR Procedure	Clears the handler list	XDBHandlerList Type Subprograms
GETAPPLICATIONDATA Function	Returns the `<applicationData>` element extracted from the resource configuration that defines the invoking handler	XDBRepositoryEvent Type Subprograms
GETCHILDOID Function	Returns the OID of the resource to which the link is pointing	XDBLink Type Subprograms
GETCURRENTUSER Function	Returns the name of the user executing the operation that triggers the event	XDBEvent Type Subprograms
GETEVENT Function	Returns a value identifying the triggering event	XDBEvent Type Subprograms

Table 227-8 (Cont.) DBMS_XEVENT Package Subprograms

Subprogram	Description	Group
GETFIRST Function	Returns the first handler in the list	XDBHandlerList Type Subprograms
GETHANDLERLIST Function	Returns an `XDBHandlerList` object containing the list of handlers that will be executed after the currently executing handler	XDBRepositoryEvent Type Subprograms
GETINTERFACE Function	Returns the top-level interface used to initiate the operation that triggered the event	XDBRepositoryEvent Type Subprograms
GETLANGUAGE Function	Returns the implementation language of the handler	XDBHandler Type Subprograms
GETLINK Function	Returns an `XDBLink` object for the target resource	XDBRepositoryEvent Type Subprograms
GETLOCK Function	Returns the lock object corresponding to the current operation	XDBRepositoryEvent Type Subprograms
GETNAME Function	Returns the string representation of the path	XDBPath Type Subprograms
GETNAME Function	Returns the next handler in the list	XDBHandlerList Type Subprograms
GETOLDRESOURCE Function	Returns the original `XDBResource` object before the operation was executed	XDBRepositoryEvent Type Subprograms
GETOPENACCESSMODE Function	Returns the access mode for the open operation	XDBRepositoryEvent Type Subprograms
GETOPENDENYMODE Function	Returns the deny mode for the open operation	XDBRepositoryEvent Type Subprograms
GETOUTPUTSTREAM Function	Returns the output BLOB in which the handler can write the rendered data	XDBRepositoryEvent Type Subprograms
GETPARAMETER Function	Returns the value of a request or session-specific parameter	XDBRepositoryEvent Type Subprograms
GETPARENT Function	Returns the resource object corresponding to a parent folder of the target resource	XDBRepositoryEvent Type Subprograms
GETPARENTNAME Function	Returns the link's parent folder's name	XDBLink Type Subprograms
GETPARENTOID Function	Returns the link's parent folder's OID	XDBLink Type Subprograms
GETPARENTNAME Function	Returns the parent's path	XDBPath Type Subprograms
GETPATH Function	Returns the `XDBPath` object representing the path of the resource for which the event was fired	XDBRepositoryEvent Type Subprograms
GETRESOURCE Function	Returns an `XDBResource` object that provides methods to access and modify the contents and metadata of the target resource	XDBRepositoryEvent Type Subprograms
GETSCHEMA Function	Returns the schema of the handler's source	XDBHandler Type Subprograms

Table 227-8 (Cont.) DBMS_XEVENT Package Subprograms

Subprogram	Description	Group
GETSOURCE Function	Returns the name of the Java class, PL/SQL package or object type implementing the handler	XDBHandler Type Subprograms
GETUPDATEBYTECOUNT Function	If the current operation is a byte-range write, returns the byte count	XDBRepositoryEvent Type Subprograms
GETUPDATEBYTEOFFSET Function	If the current operation is a byte-range write, function returns the byte offset at which the range begins	XDBRepositoryEvent Type Subprograms
GETXDBEVENT Function	Converts an `XDBRepositoryEvent` object to an `XDBEvent` type	XDBRepositoryEvent Type Subprograms
ISNULL Functions	Returns `TRUE` if input argument is `NULL`	XDBEvent Type Subprograms XDBRepositoryEvent Type Subprograms XDBHandlerList Type Subprograms XDBHandler Type Subprograms XDBPath Type Subprograms
REMOVE Procedure	Removes the specified handler from the handler list	XDBHandlerList Type Subprograms
SETRENDERPATH Procedure	Specifies the path of the resource that contains the rendered contents	XDBRepositoryEvent Type Subprograms
SETRENDERSTREAM Procedure	sets the `BLOB` from which the rendered contents can be read	XDBRepositoryEvent Type Subprograms

227.4.1 CLEAR Procedure

This procedure clears the handler list.

> **See Also:**
>
> XDBHandlerList Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.CLEAR (
    hl       IN OUT   XDBHandlerList);
```

Parameters

Table 227-9 CLEAR Procedure Parameters

Parameter	Description
`hl`	Handler list

227.4.2 GETAPPLICATIONDATA Function

This function returns the `<applicationData>` element extracted from the resource configuration that defines the invoking handler.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETAPPLICATIONDATA (
    ev   IN   XDBRepositoryEvent)
  RETURN XMLType;
```

Parameters

Table 227-10 GETAPPLICATIONDATA Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.3 GETCHILDOID Function

This function returns the OID of the resource to which the link is pointing.

> **See Also:**
>
> XDBLink Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETCHILDOID (
    link   IN   XDBLink)
  RETURN RAW;
```

Parameters

Table 227-11 GETCHILDOID Function Parameters

Parameter	Description
link	Link

227.4.4 GETCURRENTUSER Function

This function returns the name of the user executing the operation that triggers the event.

> **See Also:**
>
> XDBEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETCURRENTUSER (
    ev   IN   XDBEvent)
  RETURN VARCHAR2;
```

Parameters

Table 227-12 GETCURRENTUSER Function Parameters

Parameter	Description
ev	Event of XDBEvent type

227.4.5 GETEVENT Function

This function returns the name of the user executing the operation that triggers the event.

> **See Also:**
>
> XDBEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETEVENT (
    ev   IN   XDBEvent)
  RETURN XDBEventID;
```

Parameters

Table 227-13 GETEVENT Function Parameters

Parameter	Description
ev	Event of XDBEvent type

227.4.6 GETFIRST Function

This function returns the first handler in the list.

> **See Also:**
>
> XDBHandlerList Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETFIRST (
    hl   IN   XDBHandlerList)
  RETURN XDBHandler;
```

Parameters

Table 227-14 GETFIRST Function Parameters

Parameter	Description
hl	Handler list

227.4.7 GETHANDLERLIST Function

This function returns an `XDBHandlerList` object containing the list of handlers that will be executed after the currently executing handler.

The current handler can then filter out some of the subsequent handlers if necessary, subject to security checks. An insufficient privilege exception is raised if the executing user does not have the required access privilege to any of the resource configuration associating with a handler in the list.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETHANDLERLIST (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBHandlerList;
```

Parameters

Table 227-15 GETHANDLERLIST Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.8 GETINTERFACE Function

This function returns the top-level interface used to initiate the operation that triggered the event. This could be HTTP, FTP or SQL.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETINTERFACE (
    ev   IN   XDBRepositoryEvent)
  RETURN VARCHAR2;
```

Parameters

Table 227-16 GETINTERFACE Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.9 GETLANGUAGE Function

This function returns the implementation language of the handler.

> **See Also:**
>
> XDBHandler Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETLANGUAGE (
    handler   IN   XDBHandler)
  RETURN VARCHAR2;
```

Parameters

Table 227-17 GETLANGUAGE Function Parameters

Parameter	Description
handler	Handler

227.4.10 GETLINK Function

This function returns an `XDBLink` object for the target resource.

For a link* or unlink* event, this will be the link involved in the operation. For other events, an error is returned. Using this object the handler can access link properties, such as, `ParentName`, `ParentOID`, `ChildOID` and `LinkName`.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETLINK (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBLink;
```

Parameters

Table 227-18 GETLINK Function Parameters

Parameter	Description
`ev`	Event of XDBRepositoryEvent type

227.4.11 GETLINKNAME Function

This function returns the name of the link.

> **See Also:**
>
> XDBLink Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETLINKNAME (
    link   IN   XDBLink)
  RETURN VARCHAR2;
```

Parameters

Table 227-19 GETLINKNAME Function Parameters

Parameter	Description
`link`	Link

227.4.12 GETLOCK Function

This function returns the lock object corresponding to the current operation.I t is only valid for lock and unlock events.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETLOCK  (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBLock;
```

Parameters

Table 227-20 GETLOCK Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.13 GETLANGUAGE Function

This function returns the implementation language of the handler.

> **See Also:**
>
> XDBHandler Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETLANGUAGE (
    handler   IN   XDBHandler)
  RETURN VARCHAR2;
```

Parameters

Table 227-21 GETLANGUAGE Function Parameters

Parameter	Description
handler	Handler

227.4.14 GETNAME Function

This function returns the string representation of the path.

> **See Also:**
>
> XDBPath Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETNAME (
    path   IN   XDBPath)
  RETURN VARCHAR2;
```

Parameters

Table 227-22 GETNAME Function Parameters

Parameter	Description
path	Path

227.4.15 GETNEXT Function

This function returns the next handler in the list.

> **See Also:**
>
> XDBHandlerList Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETNEXT (
    hl   IN   XDBHandlerList)
  RETURN XDBHandler;
```

Parameters

Table 227-23 GETNEXT Function Parameters

Parameter	Description
hl	Handler list

227.4.16 GETOLDRESOURCE Function

This function returns the original `XDBResource` object before the operation was executed.

This method applies only to update event. For other events, an error is returned. This is a read-only object, and consequently none of the modifier methods will work on this object.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETOLDRESOURCE (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBResource;
```

Parameters

Table 227-24 GETOLDRESOURCE Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.17 GETOPENACCESSMODE Function

This function returns the access mode for the open operation.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETOPENACCESSMODE (
    ev   IN   XDBRepositoryEvent)
  RETURN PLS_INTEGER;
```

Parameters

Table 227-25 GETOPENACCESSMODE Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

Return Values

- `XDBRepositoryEvent.OPEN_ACCESS_READ` (value 1)

- XDBRepositoryEvent.OPEN_ACCESS_WRITE (value 2)
- XDBRepositoryEvent.OPEN_ACCESS_READ_WRITE (value 3)

227.4.18 GETOPENDENYMODE Function

This function returns the deny mode for the open operation. It is only valid for the open event.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETOPENDENYMODE (
    ev   IN   XDBRepositoryEvent)
  RETURN PLS_INTEGER;
```

Parameters

Table 227-26 GETOPENDENYMODE Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

Return Values

- XDBRepositoryEvent.OPEN_DENY_NONE (value 0)
- XDBRepositoryEvent.OPEN_DENY_READ (value 1)
- XDBRepositoryEvent.OPEN_DENY_READ_WRITE (value 2)

227.4.19 GETOUTPUTSTREAM Function

This function returns the output BLOB in which the handler can write the rendered data. It is only valid for the render event.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETOUTPUTSTREAM (
    ev   IN   XDBRepositoryEvent)
  RETURN BLOB;
```

Parameters

Table 227-27 GETOUTPUTSTREAM Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.20 GETPARAMETER Function

This function returns the value of a request or session-specific parameter. The definition of the key parameter can be found in RFC 2616 (HTTP/1.1). They will be mapped to equivalent SQL session parameters (if any).

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETPARAMETER (
    ev   IN   XDBRepositoryEvent,
    key  IN   VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 227-28 GETPARAMETER Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type
key	Supported parameters: • ACCEPT • ACCEPT-LANGUAGE • ACCEPT-CHARSET • ACCEPT_ENCODING

227.4.21 GETPARENT Function

This function returns the resource object corresponding to a parent folder of the target resource. Note that this could be any folder that contains a link to the target resource. This is a read-only object, and consequently none of the modifier methods will work on this object. For a link* or unlink* event, this method returns the link's parent folder.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETPARENT (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBResource;
```

Parameters

Table 227-29 GETPARENT Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.22 GETPARENTNAME Function

This function returns the link's parent folder's name.

> **See Also:**
>
> XDBLink Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETPARENTNAME (
    link   IN   XDBLink)
  RETURN VARCHAR2;
```

Parameters

Table 227-30 GETPARENTNAME Function Parameters

Parameter	Description
link	Link

227.4.23 GETPARENTOID Function

This function returns the link's parent folder's OID.

> **See Also:**
>
> XDBLink Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETPARENTOID (
    link   IN   XDBLink)
  RETURN RAW;
```

Parameters

Table 227-31 GETPARENTOID Function Parameters

Parameter	Description
link	Link

227.4.24 GETPARENTPATH Function

This function returns the parent's path. The level indicates the number of levels up the hierarchy. This value must be greater than zero. Level 1 means the immediate parent. If level exceeds the height of the tree then a NULL is returned.

> **See Also:**
>
> XDBPath Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETPARENTPATH (
    path   IN   XDBPath,
   level   IN   INTEGER)
  RETURN XDBPath;
```

Parameters

Table 227-32 GETPARENTPATH Function Parameters

Parameter	Description
path	Path
level	Number of levels up the hierarchy

227.4.25 GETPATH Function

This function returns the XDBPath object representing the path of the resource for which the event was fired. From this object, functions are provided to get the different path segments.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETPATH (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBPath;
```

Parameters

Table 227-33 GETPATH Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.26 GETRESOURCE Function

This function returns an `XDBResource` object that provides methods to access and modify the contents and metadata of the target resource. This object reflects any changes made by previous handlers to the resource.

The modifier methods will work only in the pre-create and pre-update event handlers. For a link* or unlink* event, this method returns the resource that the link is pointing to. For a create event, this method returns the resource that is being created.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETRESOURCE (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBResource;
```

Parameters

Table 227-34 GETRESOURCE Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.27 GETSCHEMA Function

This function returns the schema of the handler's source.

> **See Also:**
>
> XDBHandler Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETSCHEMA (
    handler   IN   XDBHandler)
  RETURN VARCHAR2;
```

Parameters

Table 227-35 GETSCHEMA Function Parameters

Parameter	Description
handler	Handler

227.4.28 GETSOURCE Function

This function returns the name of the Java class, PL/SQL package or object type implementing the handler.

> **See Also:**
>
> XDBHandler Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETSOURCE (
    handler   IN   XDBHandler)
  RETURN VARCHAR2;
```

Parameters

Table 227-36 GETSOURCE Function Parameters

Parameter	Description
handler	Handler

227.4.29 GETUPDATEBYTECOUNT Function

If the current operation is a byte-range write, the `GETUPDATEBYTECOUNT` function returns the byte count. It is only valid for the inconsistent-update event.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETUPDATEBYTECOUNT (
    ev   IN   XDBRepositoryEvent)
  RETURN NUMBER;
```

Parameters

Table 227-37 GETUPDATEBYTECOUNT Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.30 GETUPDATEBYTEOFFSET Function

If the current operation is a byte-range write, the `GETUPDATEBYTEOFFSET` function returns the byte offset at which the range begins. It is only valid for the inconsistent-update event.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETUPDATEBYTEOFFSET (
    ev   IN   XDBRepositoryEvent)
  RETURN NUMBER;
```

Parameters

Table 227-38 GETUPDATEBYTEOFFSET Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.31 GETXDBEVENT Function

This function converts an `XDBRepositoryEvent` object to an `XDBEvent` type.

> **See Also:**
>
> XDBRepositoryEvent Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.GETXDBETEVENT (
    ev   IN   XDBRepositoryEvent)
  RETURN XDBEvent;
```

Parameters

Table 227-39 GETXDBEVENT Function Parameters

Parameter	Description
ev	Event of XDBRepositoryEvent type

227.4.32 ISNULL Functions

This function returns `TRUE` if input argument is `NULL`.

> **See Also:**
>
> - XDBEvent Type Subprograms for other subprograms in this group
> - XDBRepositoryEvent Type Subprograms for other subprograms in this group
> - XDBHandlerList Type Subprograms for other subprograms in this group
> - XDBHandler Type Subprograms for other subprograms in this group
> - XDBPath Type Subprograms for other subprograms in this group
> - XDBLink Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.ISNULL (
    ev          IN   XDBEvent)
  RETURN BOOLEAN;

DBMS_XEVENT.ISNULL (
    ev          IN   XDBRepositoryEvent)
  RETURN BOOLEAN;

DBMS_XEVENT.ISNULL (
    hl          IN   XDBHandlerList)
  RETURN BOOLEAN;

DBMS_XEVENT.ISNULL (
    handler     IN   XDBHandler)
  RETURN BOOLEAN;
  RETURN BOOLEAN;

DBMS_XEVENT.ISNULL (
    path        IN   XDBPath)
  RETURN BOOLEAN;

DBMS_XEVENT.ISNULL (
    link        IN   XDBLink)
  RETURN BOOLEAN;
```

Parameters

Table 227-40 ISNULL Function Parameters

Parameter	Description
ev	Event of specified type
hl	Handler list
handler	Handler
path	Path

227.4.33 REMOVE Procedure

This procedure removes the specified handler from the handler list.

> **See Also:**
>
> XDBHandlerList Type Subprograms for other subprograms in this group

Syntax

```
DBMS_XEVENT.REMOVE (
    hl       IN OUT   XDBHandlerList,
    handler  IN       XDBHandler);
```

Parameters

Table 227-41 REMOVE Procedure Parameters

Parameter	Description
hl	Handler list
handler	Handler

227.4.34 SETRENDERPATH Procedure

This procedure specifies the path of the resource that contains the rendered contents.

This should not be called after the stream returned by GETOUTPUTSTREAM Function is written to or after the SETRENDERSTREAM Procedure is called; doing so will result in an error. This is only valid for the render event.

Syntax

```
DBMS_XEVENT.SETRENDERPATH (
    ev     IN   XDBRepositoryEvent,
    path   IN   VARCHAR2);
```

Parameters

Table 227-42 SETRENDERPATH Procedure Parameters

Parameter	Description
ev	XDB Repository Event object
path	Path of the resource containing the rendered contents

227.4.35 SETRENDERSTREAM Procedure

This procedure sets the BLOB from which the rendered contents can be read.

This should not be called after the stream returned by GETOUTPUTSTREAM is written to or after SETRENDERPATH is called; doing so will result in an error. This is only valid for the render event.

Syntax

```
DBMS_XEVENT.SETRENDERSTREAM (
    ev    IN   XDBRepositoryEvent,
    istr  IN   BLOB);
```

Parameters

Table 227-43 SETRENDERSTREAM Procedure Parameters

Parameter	Description
ev	XDBRepositoryEvent object
istr	Input stream from which to get the rendered contents

228
DBMS_XMLDOM

The `DBMS_XMLDOM` package is used to access `XMLType` objects, and implements the Document Object Model (DOM), an application programming interface for HTML and XML documents.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Types
- Exceptions
- Subprogram Groups
- Subprogram Groups
- Summary of DBMS_XMLDOM Subprograms

> **See Also:**
>
> *Oracle XML Developer's Kit Programmer's Guide*

228.1 DBMS_XMLDOM Overview

The Document Object Model (DOM) is an application programming interface (API) for HTML and XML documents. It defines the logical structure of documents, and the manner in which they are accessed and manipulated

In the DOM specification, the term "document" is used in the broad sense. XML is being increasingly used to represent many different kinds of information that may be stored in diverse systems. This information has been traditionally be seen as "data"; nevertheless, XML presents this data as documents, and the `DBMS_XMLDOM` package allows you access to both schema-based and non schema-based documents.

> **Note:**
>
> Read-from and write-to files must be on the server file system.

With DOM, anything found in an HTML or XML document can be accessed, changed, deleted, or added using the Document Object Model, with a few exceptions. In particular, the DOM interfaces for the XML internal and external subsets have not yet been specified.

One important objective of the W3C DOM specification is to provide a standard programming interface that can be used in a wide variety of environments, programming languages, and

applications. Because the DOM standard is object-oriented while PL/SQL is essentially a procedural language, some changes had to be made:

- Various DOM interfaces such as `Node`, `Element`, and others have equivalent PL/SQL types `DOMNode`, `DOMElement`, respectively.
- Various `DOMException` codes such as `WRONG_DOCUMENT_ERR`, `HIERARCHY_REQUEST_ERR`, and others, have similarly named PL/SQL exceptions.
- Various DOM Node type codes such as `ELEMENT_NODE`, `ATTRIBUTE_NODE`, and others, have similarly named PL/SQL constants.
- Subprograms defined on a DOM type become functions or procedures that accept it as a parameter. For example, to perform APPENDCHILD Function on a `DOMNode n`, the APPENDCHILD FunctionPL/SQL function is provided.
- To perform `setAttribute` on a `DOMElement elem`SETATTRIBUTE Procedures, use PL/SQL procedure .

DOM defines an inheritance hierarchy. For example, `Document`, `Element`, and `Attr` are defined to be subtypes of `Node` (see Figure 228-1). Thus, a method defined in the `Node` interface should be available in these as well. Since such inheritance is not supported in PL/SQL, it is implemented through direct invocation of the `MAKENODE` function. Calling `MAKENODE` on various DOM types converts these types into a `DOMNode`. The appropriate functions or procedures that accept `DOMNode`s can then be called to operate on these types. If, subsequently, type specific functionality is desired, the `DOMNode` can be converted back into the original type by the `make`*`XXX`* functions, where `DOM`*`XXX`* is the desired DOM type.

Figure 228-1 Inheritance Diagram for DOM Types

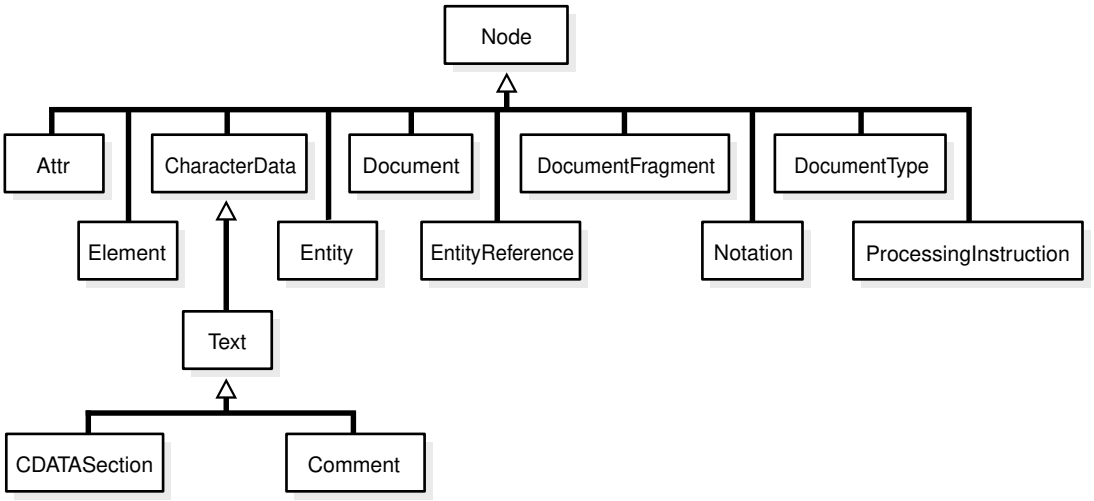

The implementation of this interface follows the REC-DOM-Level-1-19981001.

228.2 DBMS_XMLDOM Security Model

Owned by `XDB`, the `DBMS_XMLDOM` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`.

Subprograms in this package are executed using the privileges of the current user.

228.3 DBMS_XMLDOM Constants

The DBMS_XMLDOM package defines several constants that can be used for specifying parameter values.

These constants are listed in the following table.

Table 228-1 Defined Constants for DBMS_XMLDOM

Constant	Type	Value	Description
ELEMENT_NODE	PLS_INTEGER	1	The Node is an Element.
ATTRIBUTE_NODE	PLS_INTEGER	2	The Node is an Attribute.
TEXT_NODE	PLS_INTEGER	3	The Node is a Text node.
CDATA_SECTION_NODE	PLS_INTEGER	4	The Node is a CDataSection.
ENTITY_REFERENCE_NODE	PLS_INTEGER	5	The Node is an Entity Reference.
ENTITY_NODE	PLS_INTEGER	6	The Node is an Entity.
PROCESSING_INSTRUCTION_NODE	PLS_INTEGER	7	The Node is a Processing Instruction.
COMMENT_NODE	PLS_INTEGER	8	The Node is a Comment.
DOCUMENT_NODE	PLS_INTEGER	9	The Node is a Document.
DOCUMENT_TYPE_NODE	PLS_INTEGER	10	The Node is a Document Type Definition.
DOCUMENT_FRAGMENT_NODE	PLS_INTEGER	11	The Node is a Document fragment.
NOTATION_NODE	PLS_INTEGER	12	The Node is a Notation.

228.4 DBMS_XMLDOM Types

This table lists and briefly describes the types for the DBMS_XMLDOM.DOMTYPE package.

Table 228-2 XDB_XMLDOM Types

Type	Description
DOMATTR	Implements the DOM Attribute interface.
DOMCDATASECTION	Implements the DOM CDataSection interface.
DOMCHARACTERDATA	Implements the DOM Character Data interface.
DOMCOMMENT	Implements the DOM Comment interface.
DOMDOCUMENT	Implements the DOM Document interface.
DOMDOCUMENTFRAGMENT	Implements the DOM DocumentFragment interface.
DOMDOCUMENTTYPE	Implements the DOM Document Type interface.
DOMELEMENT	Implements the DOM Element interface.
DOMENTITY	Implements the DOM Entity interface.
DOMENTITYREFERENCE	Implements the DOM EntityReference interface.

Table 228-2 (Cont.) XDB_XMLDOM Types

Type	Description
DOMIMPLEMENTATION	Implements the DOM Implementation interface.
DOMNAMEDNODEMAP	Implements the DOM Named Node Map interface.
DOMNODE	Implements the DOM Node interface.
DOMNODELIST	Implements the DOM NodeList interface.
DOMNOTATION	Implements the DOM Notation interface.
DOMPROCESSINGINSTRUCTION	Implements the DOM Processing instruction interface.
DOMTEXT	Implements the DOM Text interface.

228.5 DBMS_XMLDOM Exceptions

DBMS_XMLDOM generates an exception when it encounters an issue.

This table lists the exceptions defined for DBMS_XMLDOM:

Table 228-3 Exceptions for DBMS_XMLDOM

Exception	Description
DOMSTRING_SIZE_ERR	If the specified range of text does not fit into a DOMString.
HIERARCHY_REQUEST_ERR	If any node is inserted somewhere it doesn't belong.
INDEX_SIZE_ERR	If index or size is negative, or greater than the allowed value.
INUSE_ATTRIBUTE_ERR	If an attempt is made to add an attribute that is already in use elsewhere.
INVALID_CHARACTER_ERR	If an invalid or illegal character is specified, such as in a name. See production 2 in the XML specification for the definition of a legal character, and production 5 for the definition of a legal name character.
NO_DATA_ALLOWED_ERROR	If data is specified for a node that does not support data.
NOT_FOUND_ERR	If an attempt is made to reference a node in a context where it does not exist.
NO_MODIFICATION_ALLOWED_ERR	If an attempt is made to modify an object where modifications are not allowed.
NOT_SUPPORTED_ERR	If the implementation does not support the requested type of object or operation.
WRONG_DOCUMENT_ERR	If a node is used in a different document than the one that created it (that doesn't support it).

228.6 DBMS_XMLDOM Subprogram Groups

DBMS_XMLDOM subprograms are divided into groups according to W3C Interfaces.

- DOMNode Subprograms
- DOMAttr Subprograms
- DOMCDataSection Subprograms

- DOMCharacterData Subprograms
- DOMComment Subprograms
- DOMDocument Subprograms
- DOMDocumentFragment Subprograms
- DOMDocumentType Subprograms
- DOMElement Subprograms
- DOMEntity Subprograms
- DOMEntityReference Subprograms
- DOMImplementation Subprograms
- DOMNamedNodeMap Subprograms
- DOMNodeList Subprograms
- DOMNotation Subprograms
- DOMProcessingInstruction Subprograms
- DOMText Subprograms

228.6.1 DBMS_XMLDOM DOMNode Subprograms

This table lists and briefly describes the `DOMNode` subprograms of `DBMS_XMLDOM`.

Table 228-4 Summary of DOMNode Subprograms; DBMS_XMLDOM

Subprogram	Description
ADOPTNODE Function	Adopts a node from another document
APPENDCHILD Function	Appends a new child to the node
CLONENODE Function	Clones the node
FREENODE Procedure	Frees all resources associated with the node
GETATTRIBUTES Function	Retrieves the attributes of the node
GETCHILDNODES Function	Retrieves the children of the node
GETEXPANDEDNAME Procedure and Functions	Retrieves the expanded name of the node
GETFIRSTCHILD Function	Retrieves the first child of the node
GETLASTCHILD Function	Retrieves the last child of the node
GETLOCALNAME Procedure and Functions	Retrieves the local part of the qualified name
GETNAMESPACE Procedure and Functions	Retrieves the node's namespace URI
GETNEXTSIBLING Function	Retrieves the next sibling of the node
GETNODENAME Function	Retrieves the Name of the Node
GETNODETYPE Function	Retrieves the Type of the node
GETNODEVALUE Function	Retrieves the Value of the Node
GETNODEVALUEASBINARYSTREAM Function & Procedure	Retrieves Node Value as binary stream
GETNODEVALUEASCHARACTERSTREAM Function & Procedure	Retrieves Node Value as character stream
GETOWNERDOCUMENT Function	Retrieves the owner document of the node

Table 228-4 (Cont.) Summary of DOMNode Subprograms; DBMS_XMLDOM

Subprogram	Description
GETPARENTNODE Function	Retrieves the parent of this node
GETPREFIX Function	Retrieves the namespace prefix
GETPREVIOUSSIBLING Function	Retrieves the previous sibling of the node
GETSCHEMANODE Function	Retrieves the associated schema URI
HASATTRIBUTES Function	Tests if the node has attributes
HASCHILDNODES Function	Tests if the node has child nodes
IMPORTNODE Function	Imports a node from another document
INSERTBEFORE Function	Inserts a child before the reference child
ISNULL Functions	Tests if the node is `NULL`
MAKEATTR Function	Casts the node to an Attribute
MAKECDATASECTION Function	Casts the node to a CData Section
MAKECHARACTERDATA Function	Casts the node to Character Data
MAKECOMMENT Function	Casts the node to a Comment
MAKEDOCUMENT Function	Casts the node to a DOM Document
MAKEDOCUMENTFRAGMENT Function	Casts the node to a DOM Document Fragment
MAKEDOCUMENTTYPE Function	Casts the node to a DOM Document Type
MAKEELEMENT Function	Casts the node to a DOM Element
MAKEENTITY Function	Casts the node to a DOM Entity
MAKEENTITYREFERENCE Function	Casts the node to a DOM Entity Reference
MAKENOTATION Function	Casts the node to a DOM Notation
MAKEPROCESSINGINSTRUCTION Function	Casts the node to a DOM Processing Instruction
MAKETEXT Function	Casts the node to a DOM Text
REMOVECHILD Function	Removes a specified child from a node
REPLACECHILD Function	Replaces the old child with a new child
SETNODEVALUE Procedure	Sets the Value of the node
SETNODEVALUEASBINARYSTREAM Function & Procedure	Sets the Node Value as binary stream
SETNODEVALUEASCHARACTERSTREAM Function & Procedure	Sets the Node Value as a character stream
SETPREFIX Procedure	Sets the namespace prefix
USEBINARYSTREAM Function	Establishes that the stream is valid
WRITETOBUFFER Procedures	Writes the contents of the node to a buffer
WRITETOCLOB Procedures	Writes the contents of the node to a `CLOB`
WRITETOFILE Procedures	Writes the contents of the node to a file

228.6.2 DBMS_XMLDOM DOMAttr Subprograms

This table lists the DOMAttr subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-5 Summary of DOMAttr Subprograms; DBMS_XMLDOM

Method	Description
GETEXPANDEDNAME Procedure and Functions	Retrieves the expanded name of the attribute
GETLOCALNAME Procedure and Functions	Retrieves the local name of the attribute
GETNAME Functions	Retrieves the name of the attribute
GETNAMESPACE Procedure and Functions	Retrieves the NS URI of the attribute
GETOWNERELEMENT Function	Retrieves the Element node, parent of the attribute
GETQUALIFIEDNAME Functions	Retrieves the Qualified Name of the attribute
GETSPECIFIED Function	Tests if attribute was specified in the element
GETVALUE Function	Retrieves the value of the attribute
ISNULL Functions	Tests if the Attribute node is `NULL`
MAKENODE Functions	Casts the Attribute to a node
SETVALUE Procedure	Sets the value of the attribute

228.6.3 DBMS_XMLDOM DOMCDataSection Subprograms

This table lists the DOMCdata subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-6 Summary of DOMCdata Subprograms; DBMS_XMLDOM

Method	Description
ISNULL Functions	Tests if the `CDataSection` is `NULL`
MAKENODE Functions	Casts the `CDatasection` to a node

228.6.4 DBMS_XMLDOM DOMCharacterData Subprograms

This table lists the DOMCharacterData subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-7 Summary of DOMCharacterData Subprograms; DBMS_XMLDOM

Method	Description
APPENDDATA Procedure	Appends the specified data to the node data
DELETEDATA Procedure	Deletes the data from the specified offSets
GETDATA Functions	Retrieves the data of the node
GETLENGTH Functions	Retrieves the length of the data
INSERTDATA Procedure	Inserts the data in the node at the specified offSets
ISNULL Functions	Tests if the CharacterData is `NULL`
MAKENODE Functions	Casts the CharacterData to a node
REPLACEDATA Procedure	Changes a range of characters in the node

Table 228-7 (Cont.) Summary of DOMCharacterData Subprograms; DBMS_XMLDOM

Method	Description
SETDATA Procedures	Sets the data to the node
SUBSTRINGDATA Function	Retrieves the substring of the data

228.6.5 DBMS_XMLDOM DOMComment Subprograms

The table lists the DOMComment subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-8 Summary of DOMComment Subprograms; DBMS_XMLDOM

Method	Description
ISNULL Functions	Tests if the comment is `NULL`
MAKENODE Functions	Casts the Comment to a node

228.6.6 DBMS_XMLDOM DOMDocument Subprograms

This table lists the DOMDocument subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-9 Summary of DOMDocument Subprograms; DBMS_XMLDOM

Method	Description
CREATEATTRIBUTE Functions	Creates an Attribute
CREATECDATASECTION Function	Creates a `CDataSection` node
CREATECOMMENT Function	Creates a Comment node
CREATEDOCUMENT Function	Creates a new Document
CREATEDOCUMENTFRAGMENT Function	Creates a new Document Fragment
CREATEELEMENT Functions	Creates a new Element
CREATEENTITYREFERENCE Function	Creates an Entity reference
CREATEPROCESSINGINSTRUCTION Function	Creates a Processing Instruction
CREATETEXTNODE Function	Creates a Text node
FREEDOCFRAG Procedure	Frees the document fragment
FREEDOCUMENT Procedure	Frees the document
GETCHARSET Function	Retrieves the characterset of the `DOM` document
GETDOCTYPE Function	Retrieves the DTD of the document
GETDOCUMENTELEMENT Function	Retrieves the root element of the document
GETELEMENTSBYTAGNAME Functions	Retrieves • the elements in the DOMNODELIST by tag name • elements in the subtree of a DOMNODELIST by tagname
GETIMPLEMENTATION Function	Retrieves the DOM implementation

Table 228-9 (Cont.) Summary of DOMDocument Subprograms; DBMS_XMLDOM

Method	Description
GETSTANDALONE Function	Retrieves the standalone property of the document
GETVERSION Function	Retrieves the version of the document
GETXMLTYPE Function	Retrieves the `XMLType` associated with the DOM Document
ISNULL Functions	Tests if the document is `NULL`
MAKENODE Functions	Casts the document to a node
NEWDOMDOCUMENT Functions	Creates a new document
SETCHARSET Procedure	Sets the characterset of the `DOM` document
SETDOCTYPE Procedure	Sets the DTD of the document
SETSTANDALONE Procedure	Sets the standalone property of the document
SETVERSION Procedure	Sets the version of the document
WRITETOBUFFER Procedures	Writes the document to a buffer
WRITETOCLOB Procedures	Writes the document to a `CLOB`
WRITETOFILE Procedures	Writes the document to a file

228.6.7 DBMS_XMLDOM DOMDocumentFragment Subprograms

This table lists the DOMDocumentFragment subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-10 Summary of DOMDocumentFragment Subprograms; DBMS_XMLDOM

Method	Description
FREEDOCFRAG Procedure	Frees the specified document fragment
ISNULL Functions	Tests if the DocumentFragment is `NULL`
MAKENODE Functions	Casts the Document Fragment to a node
WRITETOBUFFER Procedures	Writes the contents of a document fragment into a buffer

228.6.8 DBMS_XMLDOM DOMDocumentType Subprograms

This table lists the DOMDocumentType subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-11 Summary of DOMDocumentType Subprograms; DBMS_XMLDOM

Method	Description
FINDENTITY Function	Finds the specified entity in the document type
FINDNOTATION Function	Finds the specified notation in the document type
GETENTITIES Function	Retrieves the nodemap of entities in the Document type
GETNAME Functions	Retrieves the name of the Document type

Table 228-11 (Cont.) Summary of DOMDocumentType Subprograms; DBMS_XMLDOM

Method	Description
GETNOTATIONS Function	Retrieves the nodemap of the notations in the Document type
GETPUBLICID Functions	Retrieves the public ID of the document type
GETSYSTEMID Functions	Retrieves the system ID of the document type
ISNULL Functions	Tests if the Document Type is `NULL`
MAKENODE Functions	Casts the document type to a node

228.6.9 DBMS_XMLDOM DOMElement Subprograms

This table lists the DOMElement subprograms of DBMS_XMLDOM in alphabetical order and briefly describes them.

Table 228-12 Summary of DOMElement Subprograms; DBMS_XMLDOM

Method	Description
FREEELEMENT Procedure	Frees memory allocated to a DOMElement handle
GETATTRIBUTE Functions	Retrieves the attribute node by name
GETATTRIBUTENODE Functions	Retrieves the attribute node by name
GETCHILDRENBYTAGNAME Functions	Retrieves children of the element by tag name
GETELEMENTSBYTAGNAME Functions	Retrieves • the elements in the DOMNODELIST by tag name • elements in the subtree of a DOMNODELIST by tagname
GETEXPANDEDNAME Procedure and Functions	Retrieves the expanded name of the element
GETLOCALNAME Procedure and Functions	Retrieves the local name of the element
GETNAMESPACE Procedure and Functions	Retrieves the NS URI of the element
GETQUALIFIEDNAME Functions	Retrieves the qualified name of the element
GETTAGNAME Function	Retrieves the Tag name of the element
HASATTRIBUTE Functions	Tests if an attribute exists
ISNULL Functions	Tests if the Element is `NULL`
MAKENODE Functions	Casts the Element to a node
NORMALIZE Procedure	Normalizes the text children of the element
REMOVEATTRIBUTE Procedures	Removes the attribute specified by the name
REMOVEATTRIBUTENODE Function	Removes the attribute node in the element
RESOLVENAMESPACEPREFIX Function	Resolve the prefix to a namespace URI
SETATTRIBUTE Procedures	Sets the attribute specified by the name
SETATTRIBUTENODE Functions	Sets the attribute node in the element

228.6.10 DBMS_XMLDOM DOMEntity Subprograms

This table lists and briefly describes the `DOMEntity` subprograms of `DBMS_XMLDOM`.

Table 228-13 Summary of DOMEntity Subprograms; DBMS_XMLDOM

Method	Description
GETNOTATIONNAME Function	Retrieves the notation name of the entity
GETPUBLICID Functions	Retrieves the public Id of the entity
GETSYSTEMID Functions	Retrieves the system Id of the entity
ISNULL Functions	Tests if the Entity is `NULL`
MAKENODE Functions	Casts the Entity to a node

228.6.11 DBMS_XMLDOM DOMEntityReference Subprograms

This table lists and briefly describes the `DOMEntityReference` subprograms of `DBMS_XMLDOM`.

Table 228-14 Summary of DOMEntityReference Subprograms; DBMS_XMLDOM

Method	Description
ISNULL Functions	Tests if the `DOMEntityReference` is `NULL`
MAKENODE Functions	Casts the `DOMEntityReference` to `NULL`

228.6.12 DBMS_XMLDOM DOMImplementation Subprograms

This table lists and briefly describes the `DOMImplementation` subprograms of `DBMS_XMLDOM`.

Table 228-15 Summary of DOMImplementation Subprograms; DBMS_XMLDOM

Method	Description
ISNULL Functions	Tests if the `DOMImplementation` node is `NULL`
HASFEATURE Function	Tests if the `DOMImplementation` implements a feature

228.6.13 DBMS_XMLDOM DOMNamedNodeMap Subprograms

This table lists and briefly describes the `DOMNamedNodeMap` subprograms of `DBMS_XMLDOM`.

Table 228-16 Summary of DOMNamedNodeMap Subprograms; DBMS_XMLDOM

Method	Description
GETLENGTH Functions	Retrieves the number of items in the map
GETNAMEDITEM Function	Retrieves the item specified by the name
ISNULL Functions	Tests if the `NamedNodeMap` is `NULL`
ITEM Functions	Retrieves the item given the index in the map
REMOVENAMEDITEM Function	Removes the item specified by name
SETNAMEDITEM Function	Sets the item in the map specified by the name

228.6.14 DBMS_XMLDOM DOMNodeList Subprograms

This table lists and briefly describes the `DOMNodeList` subprograms of `DBMS_XMLDOM`.

Table 228-17 Summary of DOMNodeList Subprograms; DBMS_XMLDOM

Method	Description
FREENODELIST Procedure	Frees all resources associated with a nodelist
GETLENGTH Functions	Retrieves the number of items in the list
ISNULL Functions	Tests if the `NodeList` is `NULL`
ITEM Functions	Retrieves the item given the index in the list

228.6.15 DBMS_XMLDOM DOMNotation Subprograms

This table lists and briefly describes the `DOMNotation` subprograms of `DBMS_XMLDOM`.

Table 228-18 Summary of DOMNotation Subprograms; DBMS_XMLDOM

Method	Description
GETPUBLICID Functions	Retrieves the public Id of the notation
GETSYSTEMID Functions	Retrieves the system Id of the notation
ISNULL Functions	Tests if the Notation is `NULL`
MAKENODE Functions	Casts the notation to a node

228.6.16 DBMS_XMLDOM DOMProcessingInstruction Subprograms

This table lists and briefly describes the `DOMProcessingInstruction` subprograms of `DBMS_XMLDOM`.

Table 228-19 Summary of DOMProcessingInstruction Subprograms; DBMS_XMLDOM

Method	Description
GETDATA Functions	Retrieves the data of the processing instruction
GETTARGET Function	Retrieves the target of the processing instruction
ISNULL Functions	Tests if the Processing Instruction is `NULL`
MAKENODE Functions	Casts the Processing Instruction to a node
SETDATA Procedures	Sets the data of the processing instruction

228.6.17 DBMS_XMLDOM DOMText Subprograms

This table lists and briefly describes the `DOMText` subprograms of `DBMS_XMLDOM`.

Table 228-20 Summary of DOMText Subprograms; DBMS_XMLDOM

Method	Description
ISNULL Functions	Tests if the text is `NULL`
MAKENODE Functions	Casts the text to a node
SPLITTEXT Function	Splits the contents of the text node into 2 text nodes

228.7 Summary of DBMS_XMLDOM Subprograms

This table lists the DBMS_XMLDOM subprograms and briefly describes them.

Table 228-21 Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
ADOPTNODE Function	Adopts a node from another document	DOMNode Subprograms
APPENDCHILD Function	Appends a new child to the node	DOMNode Subprograms
APPENDDATA Procedure	Appends the specified data to the node data	DOMCharacterData Subprograms
CLONENODE Function	Clones the node	DOMNode Subprograms
CREATEATTRIBUTE Functions	Creates an Attribute	DOMDocument Subprograms
CREATECDATASECTION Function	Creates a `CDataSection` node	DOMDocument Subprograms
CREATECOMMENT Function	Creates a Comment node	DOMDocument Subprograms
CREATEDOCUMENT Function	Creates a new Document	DOMDocument Subprograms
CREATEDOCUMENTFRAGMENT Function	Creates a new Document Fragment	DOMDocument Subprograms
CREATEELEMENT Functions	Creates a new Element	DOMDocument Subprograms
CREATEENTITYREFERENCE Function	Creates an Entity reference	DOMDocument Subprograms
CREATEPROCESSINGINSTRUCTION Function	Creates a Processing Instruction	DOMDocument Subprograms
CREATETEXTNODE Function	Creates a Text node	DOMDocument Subprograms
DELETEDATA Procedure	Deletes the data from the specified offSets	DOMCharacterData Subprograms
FINDENTITY Function	Finds the specified entity in the document type	DOMDocumentType Subprograms
FINDNOTATION Function	Finds the specified notation in the document type	DOMDocumentType Subprograms
FREEDOCFRAG Procedure	Frees the document fragment	DOMDocument Subprograms and DOMDocumentFragment Subprograms
FREEDOCUMENT Procedure	Frees the document	DOMDocument Subprograms
FREEELEMENT Procedure	Frees memory allocated to a DOMElement handle	DOMElement Subprograms
FREENODE Procedure	Frees all resources associated with the node	DOMNode Subprograms

Table 228-21 (Cont.) Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
FREENODELIST Procedure	Frees all resources associated with a nodelist	DOMNodeList Subprograms
GETATTRIBUTE Functions	Retrieves the attribute node by name	DOMElement Subprograms
GETATTRIBUTENODE Functions	Retrieves the attribute node by name	DOMElement Subprograms
GETATTRIBUTES Function	Retrieves the attributes of the node	DOMNode Subprograms
GETCHARSET Function	Retrieves the characterset of the `DOM` document	DOMDocument Subprograms
GETCHILDNODES Function	Retrieves the children of the node	DOMNode Subprograms
GETCHILDRENBYTAGNAME Functions	Retrieves children of the element by tag name	DOMCharacterData Subprograms
GETDATA Functions	Retrieves • the data of the node • the data of the processing instruction	• DOMCharacterData Subprograms • DOMProcessingInstruction Subprograms
GETDOCTYPE Function	Retrieves the DTD of the document	DOMDocument Subprograms
GETDOCUMENTELEMENT Function	Retrieves the root element of the document	DOMDocument Subprograms
GETELEMENTSBYTAGNAME Functions	Retrieves • the elements in the DOMNODELIST by tag name • elements in the subtree of a DOMNODELIST by tagname	• DOMDocument Subprograms • DOMElement Subprograms
GETENTITIES Function	Retrieves the nodemap of entities in the Document type	DOMDocumentType Subprograms
GETEXPANDEDNAME Procedure and Functions	Retrieves • the expanded name of the node • the expanded name of the attribute • the expanded name of the element	• DOMNode Subprograms • DOMAttr Subprograms • DOMElement Subprograms
GETFIRSTCHILD Function	Retrieves the first child of the node	DOMNode Subprograms
GETIMPLEMENTATION Function	Retrieves the DOM implementation	DOMDocument Subprograms
GETLASTCHILD Function	Retrieves the last child of the node	DOMNode Subprograms
GETLENGTH Functions	Retrieves • the length of the data • the number of items in the map • the number of items in the list	• DOMCharacterData Subprograms • DOMNamedNodeMap Subprograms • DOMNodeList Subprograms
GETLOCALNAME Procedure and Functions	Retrieves • the local part of the qualified name • the local name of the attribute • the local name of the element	• DOMNode Subprograms • DOMAttr Subprograms • DOMElement Subprograms

Table 228-21 (Cont.) Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
GETNAME Functions	Retrieves • the name of the attribute • the name of the Document type	• DOMAttr Subprograms • DOMDocumentType Subprograms
GETNAMEDITEM Function	Retrieves • an item specified by name • and namespace URI)	• DOMNamedNodeMap Subprograms • DOMNamedNodeMap Subprograms
GETNAMESPACE Procedure and Functions	Retrieves • the node's namespace URI • the NS URI of the attribute • the NS URI of the element	• DOMNode Subprograms • DOMAttr Subprograms • DOMElement Subprograms
GETNEXTSIBLING Function	Retrieves the next sibling of the node	DOMNode Subprograms
GETNODENAME Function	Retrieves the Name of the Node	DOMNode Subprograms
GETNODETYPE Function	Retrieves the Type of the node	DOMNode Subprograms
GETNODEVALUE Function	Retrieves the Value of the Node	DOMNode Subprograms
GETNODEVALUEASBINARYSTREAM Function & Procedure	Retrieves the Node Value as binary stream	DOMNode Subprograms
GETNODEVALUEASCHARACTERSTREAM Function & Procedure	Retrieves the Node Value as character stream	DOMNode Subprograms
GETNOTATIONNAME Function	Retrieves the notation name of the entity	DOMEntity Subprograms
GETNOTATIONS Function	Retrieves the nodemap of the notations in the Document type	DOMDocumentType Subprograms
GETTARGET Function	Retrieves the target of the processing instruction	DOMProcessingInstruction Subprograms
GETOWNERDOCUMENT Function	Retrieves the owner document of the node	DOMNode Subprograms
GETOWNERELEMENT Function	Retrieves the Element node, parent of the attribute	DOMAttr Subprograms
GETPARENTNODE Function	Retrieves the parent of this node	DOMNode Subprograms
GETPREFIX Function	Retrieves the namespace prefix)	DOMNode Subprograms
GETPREVIOUSSIBLING Function	Retrieves the previous sibling of the node	DOMNode Subprograms
GETPUBLICID Functions	Retrieves • the public ID of the document type • the public Id of the entity • the public Id of the notation	• DOMDocumentType Subprograms • DOMEntity Subprograms • DOMNotation Subprograms
GETQUALIFIEDNAME Functions	Retrieves • the Qualified Name of the attribute • the qualified name of the element	• DOMAttr Subprograms • DOMElement Subprograms
GETSCHEMANODE Function	Retrieves the associated schema URI	DOMNode Subprograms

Table 228-21 (Cont.) Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
GETSPECIFIED Function	Tests if attribute was specified in the element.	DOMAttr Subprograms
GETSTANDALONE Function	Retrieves the standalone property of the document	DOMDocument Subprograms
GETSYSTEMID Functions	Retrieves • the system ID of the document type • the system Id of the entity • the system Id of the notation	• DOMDocumentType Subprograms • DOMEntity Subprograms • DOMNotation Subprograms
GETTAGNAME Function	Retrieves the Tag name of the element	DOMElement Subprograms
GETVALUE Function	Retrieves the value of the attribute	DOMAttr Subprograms
GETVERSION Function	Retrieves the version of the document	DOMDocument Subprograms)
GETXMLTYPE Function	Retrieves the `XMLType` associated with the DOM Document	DOMDocument Subprograms
HASATTRIBUTES Function	Tests if the node has attributes	DOMNode Subprograms
HASATTRIBUTE Functions	Tests if an attribute exists	DOMElement Subprograms
HASCHILDNODES Function	Tests if the node has child nodes	DOMNode Subprograms
HASFEATURE Function	Tests if the `DOMImplementation` implements a feature	DOMImplementation Subprograms
IMPORTNODE Function	Imports a node from another document	DOMNode Subprograms
INSERTBEFORE Function	Inserts a child before the reference child	DOMNode Subprograms
INSERTDATA Procedure	Inserts the data in the node at the specified offSets	DOMCharacterData Subprograms

Table 228-21 (Cont.) Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
ISNULL Functions	Tests • if the node is NULL • if the Attribute node is NULL • if the CDataSection is NULL • if the CharacterData is NULL • if the comment is NULL • if the document is NULL • if the DocumentFragment is NULL • if the Document Type is NULL • if the Element is NULL • if the Entity is NULL • if the DOMEntityReference is NULL • if the DOMImplementation node is NULL • if the NamedNodeMap is NULL • if the NodeList is NULL • if the Notation is NULL • if the Processing Instruction is NULL • if the text is NULL	• DOMNode Subprograms • DOMAttr Subprograms • DOMCDataSection Subprograms • DOMCharacterData Subprograms • DOMComment Subprograms • DOMDocument Subprograms • DOMDocumentFragment Subprograms • DOMDocumentType Subprograms • DOMElement Subprograms • DOMEntity Subprograms • DOMEntityReference Subprograms • DOMImplementation Subprograms • DOMNamedNodeMap Subprograms • DOMNodeList Subprograms • DOMNotation Subprograms • DOMProcessingInstruction Subprograms • DOMText Subprograms
ITEM Functions	Retrieves • the item given the index in the map • the item given the index in the NodeList	• DOMNamedNodeMap Subprograms • DOMNodeList Subprograms
MAKEATTR Function	Casts the node to an Attribute	DOMNode Subprograms
MAKECDATASECTION Function	Casts the node to a CData Section	DOMNode Subprograms
MAKECHARACTERDATA Function	Casts the node to Character Data	DOMNode Subprograms
MAKECOMMENT Function	Casts the node to a Comment	DOMNode Subprograms
MAKEDOCUMENT Function	Casts the node to a DOM Document	DOMNode Subprograms
MAKEDOCUMENTFRAGMENT Function	Casts the node to a DOM Document Fragment	DOMNode Subprograms)
MAKEDOCUMENTTYPE Function	Casts the node to a DOM Document Type	DOMNode Subprograms
MAKEELEMENT Function	Casts the node to a DOM ElemenT	DOMNode Subprograms
MAKEENTITY Function	Casts the node to a DOM Entity	DOMNode Subprograms
MAKEENTITYREFERENCE Function	Casts the node to a DOM Entity Reference	DOMNode Subprograms

Table 228-21 (Cont.) Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
MAKENODE Functions	Casts • the Attribute to a node • the `CDatasection` to a node • the CharacterData to a node • the Comment to a node • the document to a node • the Document Fragment to a node • the document type to a node • the Element to a node • the Entity to a node • • the `DOMEntityReference` to `NULL` • the notation to a node • the Processing Instruction to a node • the text to a node	• DOMAttr Subprograms • DOMCDataSection Subprograms • DOMCharacterData Subprograms • DOMComment Subprograms • DOMDocument Subprograms • DOMDocumentFragment Subprograms • DOMDocumentType Subprograms • DOMElement Subprograms • DOMEntity Subprograms • DOMEntityReference Subprograms • DOMNotation Subprograms • DOMProcessingInstruction Subprograms • DOMText Subprograms
MAKENOTATION Function	Casts the node to a DOM Notation	DOMNode Subprograms
MAKEPROCESSINGINSTRUCTION Function	Casts the node to a DOM Processing Instruction	DOMNode Subprograms
MAKETEXT Function	Casts the node to a DOM Text	DOMNode Subprograms
NEWDOMDOCUMENT Functions	Creates a new document	DOMDocument Subprograms
NORMALIZE Procedure	Normalizes the text children of the element	DOMElement Subprograms
REMOVEATTRIBUTE Procedures	Removes the attribute specified by the name	DOMElement Subprograms
REMOVEATTRIBUTENODE Function	Removes the attribute node in the element	DOMElement Subprograms
REMOVECHILD Function	Removes a specified child from a node	DOMNode Subprograms
REMOVENAMEDITEM Function	Removes the item specified by name	DOMNamedNodeMap Subprograms
REPLACECHILD Function	Replaces the old child with a new child	DOMNode Subprograms
REPLACEDATA Procedure	Changes a range of characters in the node	DOMCharacterData Subprograms
RESOLVENAMESPACEPREFIX Function	Resolve the prefix to a namespace URI	DOMElement Subprograms
SETATTRIBUTE Procedures	Sets the attribute specified by the name	DOMElement Subprograms
SETATTRIBUTENODE Functions	Sets the attribute node in the element	DOMElement Subprograms

Table 228-21 (Cont.) Summary of DBMS_XMLDOM Package Subprogram

Subprogram	Description	Group
SETCHARSET Procedure	Sets the characterset of the DOM document	DOMDocument Subprograms
SETDATA Procedures	Sets • the data to the node • the data of the processing instruction	• DOMCharacterData Subprograms • DOMProcessingInstruction Subprograms
SETDOCTYPE Procedure	Sets the DTD of the document.	DOMDocument Subprograms
SETNAMEDITEM Function	Sets the item in the map specified by the name	DOMNamedNodeMap Subprograms
SETNODEVALUE Procedure	Sets the Value of the node	DOMNode Subprograms
SETNODEVALUEASBINARYSTREAM Function & Procedure	Sets the Node Value as a binary stream	DOMNode Subprograms
SETNODEVALUEASCHARACTERSTREAM Function & Procedure	Sets the Node Value as a character stream	DOMNode Subprograms
SETPREFIX Procedure	Sets the namespace prefix	DOMNode Subprograms
SETSTANDALONE Procedure	Sets the standalone property of the document	DOMDocument Subprograms
SETVALUE Procedure	Sets the value of the attribute	DOMAttr Subprograms
SETVERSION Procedure	Sets the version of the document	DOMDocument Subprograms
SPLITTEXT Function	Splits the contents of the text node into 2 text nodes	DOMText Subprograms
SUBSTRINGDATA Function	Retrieves the substring of the data	DOMCharacterData Subprograms
USEBINARYSTREAM Function	Strabismus that the stream is valid for use	DOMNode Subprograms
WRITETOBUFFER Procedures	Writes • the contents of the node to a buffer • the document to a buffer • the contents of a document fragment into a buffer	• DOMNode Subprograms • DOMDocument Subprograms • DOMDocumentFragment Subprograms
WRITETOCLOB Procedures	Writes • the contents of the node to a CLOB • the document to a CLOB	• DOMNode Subprograms • DOMDocument Subprograms
WRITETOFILE Procedures	Writes • the contents of the node to a file • the document to a file	• DOMNode Subprograms • DOMDocument Subprograms

228.7.1 ADOPTNODE Function

This function adopts a node from another document, and returns this new node.

> **See Also:**
>
> DOMNode Subprograms for other subprograms in this group

Syntax

```
DBMS_XMLDOM.ADOPTNODE(
   doc             IN   DOMDocument,
   importedNode    IN   DOMNode)
 RETURN DOMNODE;
```

Parameters

Table 228-22 ADOPTNODE Function Parameters

Parameter	Description
`doc`	Document that is adopting the node
`importedNode`	Node to adopt

Usage Notes

Note that the ADOPTNODE Function removes the node from the source document while the IMPORTNODE Function clones the node in the source document.

228.7.2 APPENDCHILD Function

This function adds the node `newchild` to the end of the list of children of this node, and returns the newly added node. If the `newchild` is already in the tree, it is first removed.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.APPENDCHILD(
   n           IN    DOMNode,
   newchild    IN    DOMNode)
 RETURN DOMNODE;
```

Parameters

Table 228-23 APPENDCHILD Function Parameters

Parameter	Description
n	DOMNode
newchild	The child to be appended to the list of children of node n

228.7.3 APPENDDATA Procedure

This procedure appends the string to the end of the character data of the node. Upon success, data provides access to the concatenation of data and the specified string argument.

> **See Also:**
>
> DOMCharacterData Subprograms

Syntax

```
DBMS_XMLDOM.APPENDDATA(
   cd      IN    DOMCHARACTERDATA,
   arg     IN    VARCHAR2);
```

Parameters

Table 228-24 APPENDDATA Procedure Parameters

Parameter	Description
cd	DOMCHARACTERDATA
arg	The data to append to the existing data

228.7.4 CLONENODE Function

This function returns a duplicate of this node, and serves as a generic copy constructor for nodes. The duplicate node has no parent, its parent node is NULL.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.CLONENODE(
   n       IN    DOMNODE,
   deep    IN    BOOLEAN)
 RETURN DOMNODE;
```

Parameters

Table 228-25 CLONENODE Function Parameters

Parameter	Description
n	DOMNODE
deep	Determines if children are to be cloned

Usage Notes

- Cloning an Element copies all attributes and their values, including those generated by the XML processor to represent defaulted attributes, but this method does not copy any text it contains unless it is a deep clone, since the text is contained in a child Text node.
- Cloning an Attribute directly, as opposed to be cloned as part of an Element cloning operation, returns a specified attribute (specified is TRUE).
- Cloning any other type of node simply returns a copy of this node.

228.7.5 CREATEATTRIBUTE Functions

This function creates a DOMATTR node.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

Creates a DOMATTR with the specified name:

```
DBMS_XMLDOM.CREATEATTRIBUTE(
   doc      IN     DOMDOCUMENT,
   name     IN     VARCHAR2)
 RETURN DOMATTR;
```

Creates a DOMATTR with the specified name and namespace URI:

```
DBMS_XMLDOM.CREATEATTRIBUTE(
   doc      IN     DOMDOCUMENT,
   qname     IN     VARCHAR2,
   ns       IN     VARCHAR2)
RETURN DOMATTR;
```

Parameters

Table 228-26 CREATEATTRIBUTE Function Parameters

Parameter	Description
doc	DOMDOCUMENT
qname	New attribute qualified name

Table 228-26 (Cont.) CREATEATTRIBUTE Function Parameters

Parameter	Description
ns	Namespace

228.7.6 CREATECDATASECTION Function

This function creates a DOMCDATASECTION node.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATECDATASECTION(
   doc      IN       DOMDOCUMENT,
   data     IN       VARCHAR2)
 RETURN DOMCDATASECTION;
```

Parameters

Table 228-27 CREATECDATASECTION Function Parameters

Parameter	Description
doc	DOMDOCUMENT
data	Content of the DOMCDATASECTION node

228.7.7 CREATECOMMENT Function

This function creates a DOMCOMMENT node.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATECOMMENT(
   doc       IN       DOMDOCUMENT,
   data      IN       VARCHAR2)
 RETURN DOMCOMMENT;
```

Parameters

Table 228-28 CREATECOMMENT Function Parameters

Parameter	Description
doc	DOMDOCUMENT
data	Content of the DOMComment node

228.7.8 CREATEDOCUMENT Function

This function creates a DOMDOCUMENT with specified namespace URI, root element name, DTD.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATEDOCUMENT(
   namespaceURI      IN     VARCHAR2,
   qualifiedName     IN     VARCHAR2,
   doctype           IN     DOMTYPE := NULL)
 RETURN DOMDOCUMENT;
```

Parameters

Table 228-29 CREATEDOCUMENT Function Parameters

Parameter	Description
namespaceURI	Namespace URI
qualifiedName	Root element name
doctype	Document type

228.7.9 CREATEDOCUMENTFRAGMENT Function

This function creates a DOMDOCUMENTFRAGMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATEDOCUMENTFRAGMENT(
   doc       IN      DOMDOCUMENT)
 RETURN DOMDOCUMENTFRAGMENT;
```

Parameters

Table 228-30 CREATEDOCUMENTFRAGMENT Function Parameters

Parameter	Description
doc	DOMDocument

228.7.10 CREATEELEMENT Functions

This function creates a DOMELEMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

Creates a DOMElement with specified name:

```
DBMS_XMLDOM.CREATEELEMENT(
   doc         IN      DOMDOCUMENT,
   tagName     IN      VARCHAR2)
 RETURN DOMELEMENT;
```

Creates a DOMElement with specified name and namespace URI:

```
DBMS_XMLDOM.CREATEELEMENT(
   doc         IN     DOMDOCUMENT,
   tagName     IN     VARCHAR2,
   ns          IN     VARCHAR2)
 RETURN DOMELEMENT;
```

Parameters

Table 228-31 CREATEELEMENT Function Parameters

Parameter	Description
doc	DOMDOCUMENT
tagName	Tagname for new DOMELEMENT
ns	Namespace

228.7.11 CREATEENTITYREFERENCE Function

This function creates a `DOMENTITYREFERENCE` node.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATEENTITYREFERENCE(
   doc         IN     DOMDOCUMENT,
   name        IN     VARCHAR2)
 RETURN DOMENTITYREFERENCE;
```

Parameters

Table 228-32 CREATEENTITYREFERENCE Function Parameters

Parameter	Description
`doc`	`DOMDOCUMENT`
`name`	New entity reference name

228.7.12 CREATEPROCESSINGINSTRUCTION Function

This function creates a `DOMPROCESSINGINSTRUCTION` node.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATEPROCESSINGINSTRUCTION(
   doc       IN      DOMDocument,
   target    IN      VARCHAR2,
   data      IN      VARCHAR2)
 RETURN DOMPROCESSINGINSTRUCTION;
```

Parameters

Table 228-33 CREATEPROCESSINGINSTRUCTION Function Parameters

Parameter	Description
`doc`	`DOMDOCUMENT`
`target`	Target of the new processing instruction

Table 228-33 (Cont.) CREATEPROCESSINGINSTRUCTION Function Parameters

Parameter	Description
data	Content data of the new processing instruction

228.7.13 CREATETEXTNODE Function

This function creates a DOMTEXT node.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.CREATETEXTNODE(
   doc      IN     DOMDocument,
   data     IN     VARCHAR2)
 RETURN DOMTEXT;
```

Parameters

Table 228-34 CREATETEXTNODE Function Parameters

Parameter	Description
doc	DOMDOCUMENT
data	Content of the DOMText node

228.7.14 DELETEDATA Procedure

This procedure removes a range of characters from the node. Upon success, data and length reflect the change.

> **See Also:**
>
> DOMCharacterData Subprograms

Syntax

```
DBMS_XMLDOM.DELETEDATA(
   cd         IN     DOMCHARACTERDATA,
   offset     IN     NUMBER,
   cnt        IN     NUMBER);
```

Parameters

Table 228-35 DELETEDATA PROCEDURE Parameters

Parameter	Description
cd	DOMCHARACTERDATA
offset	The offset from which to delete the data
cnt	The number of characters (starting from offset) to delete

228.7.15 FINDENTITY Function

This function finds an entity in the specified DTD, and returns that entity if found.

> **See Also:**
>
> DOMDocumentType Subprograms

Syntax

```
DBMS_XMLDOM.FINDENTITY(
   dt      IN      DOMDOCUMENTTYPE,
   name    IN      VARCHAR2,
   par     IN      BOOLEAN)
 RETURN  DOMENTITY;
```

Parameters

Table 228-36 FINDENTITY Function Parameters

Parameter	Description
dt	The DTD
name	Entity to find
par	Flag to indicate type of entity; TRUE for parameter entity and FALSE for normal entity

228.7.16 FINDNOTATION Function

This function finds the notation in the specified DTD, and returns it, if found.

> **See Also:**
>
> DOMDocumentType Subprograms

Syntax

```
DBMS_XMLDOM.FINDNOTATION(
   dt         IN     DOMDocumentType,
   name       IN     VARCHAR2)
 RETURN DOMNOTATION;
```

Parameters

Table 228-37 FINDNOTATION Function Parameters

Parameter	Description
dt	The DTD
name	The notation to find

228.7.17 FREEDOCFRAG Procedure

This procedure frees the specified document fragment.

> **See Also:**
>
> DOMDocument Subprograms and DOMDocumentFragment Subprograms

Syntax

```
DBMS_XMLDOM.FREEDOCFRAG(
   df     IN     DOMDOCUMENTFRAGMENT);
```

Parameters

Table 228-38 FREEDOCFRAG Procedure Parameters

Parameter	Description
df	DOM document fragment

228.7.18 FREEDOCUMENT Procedure

This procedure frees `DOMDOCUMENT` object.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.FREEDOCUMENT(
   doc      IN      DOMDOCUMENT);
```

Parameters

Table 228-39 FREEDOCUMENT Procedure Parameters

Parameter	Description
doc	DOMDOCUMENT

228.7.19 FREEELEMENT Procedure

This procedure frees memory allocated to a DOMElement handle.

> **See Also:**
>
> DBMS_XMLDOM DOMElement Subprograms

Syntax

```
DBMS_XMLDOM.FREEELEMENT(
    elem IN DOMELEMENT);
```

Parameters

Table 228-40 FREEELEMENT Procedure Parameters

Parameter	Description
elem	Of type DOMELEMENT

228.7.20 FREENODE Procedure

This procedure frees all resources associated with a DOMNODE.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.FREENODE(
   n       IN      DOMNODE);
```

Parameters

Table 228-41 FREENODE Procedure Parameters

Parameter	Description
n	DOMNODE

228.7.21 FREENODELIST Procedure

This procedure frees all resources associated with a nodelist.

> **See Also:**
>
> DBMS_XMLDOM DOMNodeList Subprograms

Syntax

```
DBMS_XMLDOM.FREENODELIST(
   nl IN DOMNodeList);
```

Parameters

Table 228-42 FREENODELIST Procedure Parameters

Parameter	Description
nl	Of type DOMNODELIST

228.7.22 GETATTRIBUTE Functions

This function returns the value of an attribute of an DOMELEMENT by name.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Returns the value of a DOMELEMENT's attribute by name:

```
DBMS_XMLDOM.GETATTRIBUTE(
   elem       IN      DOMELEMENT,
   name       IN      VARCHAR2)
 RETURN VARCHAR2;
```

Returns the value of a DOMELEMENT's attribute by name and namespace URI:

```
DBMS_XMLDOM.GETATTRIBUTE(
   elem      IN     DOMELEMENT,
   name      IN     VARCHAR2,
   ns        IN     VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 228-43 GETATTRIBUTE Function Parameters

Parameter	Description
elem	The DOMELEMENT
name	Attribute name
ns	Namespace

228.7.23 GETATTRIBUTENODE Functions

This function returns an attribute node from the DOMELEMENT by name. The function is overloaded. The specific forms of functionality are described along with the syntax declarations.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Returns an attribute node from the DOMELEMENT by name:

```
DBMS_XMLDOM.GETATTRIBUTENODE(
   elem      IN     DOMElement,
   name      IN     VARCHAR2)
 RETURN DOMATTR;
```

Returns an attribute node from the DOMELEMENT by name and namespace URI:

```
DBMS_XMLDOM.GETATTRIBUTENODE(
   elem      IN     DOMElement,
   name      IN     VARCHAR2,
   ns        IN     VARCHAR2)
RETURN DOMATTR;
```

Parameters

Table 228-44 GETATTRIBUTENODE Function Parameters

Parameter	Description
elem	The DOMELEMENT
name	Attribute name; * matches any attribute
ns	Namespace

228.7.24 GETATTRIBUTES Function

This function retrieves a `NAMEDNODEMAP` containing the attributes of this node (if it is an Element) or `NULL` otherwise.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETATTRIBUTES(
   n       IN       DOMNode)
 RETURN DOMNAMEDNODEMAP;
```

Parameters

Table 228-45 GETATTRIBUTES Function Parameters

Parameter	Description
n	DOMNODE

228.7.25 GETCHARSET Function

This function retrieves the characterset of the DOM document.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETCHARSET(
   doc IN     DOMDocument)
 RETURN VARCHAR2;
```

Parameters

Table 228-46 GETCHARSET Function Parameters

Parameter	Description
doc	DOM document

Usage Notes

For a newly parsed document, we return the database characterset. Once the `SETCHARSET` Procedure is called with a non-`NULL` value for `charset`, that `charset` is returned.

228.7.26 GETCHILDNODES Function

This function retrieves a `DOMNODELIST` that contains all children of this node. If there are no children, this is a `DOMNODELIST` containing no nodes.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETCHILDNODES(
   n       IN     DOMNode)
 RETURN DOMNodeList;
```

Parameters

Table 228-47 GETCHILDNODES Function Parameters

Parameter	Description
n	DOMNODE

228.7.27 GETCHILDRENBYTAGNAME Functions

This function returns the children of the `DOMELEMENT`.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Returns children of the `DOMELEMENT` given the tag name:

```
DBMS_XMLDOM.GETCHILDRENBYTAGNAME(
   elem       IN       DOMElement,
   name       IN       VARCHAR2)
 RETURN DOMNODELIST;
```

Returns children of the `DOMELEMENT` given the tag name and namespace:

```
DBMS_XMLDOM.GETCHILDRENBYTAGNAME(
   elem       IN       DOMElement,
   name       IN       VARCHAR2,
   ns         IN       VARCHAR2)
RETURN DOMNODELIST;
```

Parameters

Table 228-48 GETCHILDRENBYTAGNAME Function Parameters

Parameter	Description
elem	DOMELEMENT
name	Tag name
ns	Namespace

228.7.28 GETDATA Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Gets the character data of the node that implements this interface (See Also: DOMCharacterData Subprograms):

```
DBMS_XMLDOM.GETDATA(
   cd      IN    DOMCHARACTERDATA)
 RETURN VARCHAR2;
```

Returns the content data of the DOMProcessingInstruction (See Also: DOMProcessingInstruction Subprograms):

```
DBMS_XMLDOM.GETDATA(
   pi      IN    DOMPROCESSINGINSTRUCTION)
 RETURN VARCHAR2;
```

Parameters

Table 228-49 GETDATA Function Parameters

Parameter	Description
cd	DOMCHARACTERDATA
pi	The DOMPROCESSINGINSTRUCTION

228.7.29 GETDOCTYPE Function

This function returns the DTD associated to the DOMDOCUMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETDOCTYPE(
   doc      IN     DOMDOCUMENT)
RETURN DOMDOCUMENTTYPE;
```

Parameters

Table 228-50 GETDOCTYPE Function Parameters

Parameter	Description
doc	DOMDOCUMENT

228.7.30 GETDOCUMENTELEMENT Function

This function returns the root element of the DOMDOCUMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETDOCUMENTELEMENT(
   doc      IN      DOMDOCUMENT)
 RETURN DOMELEMENT;
```

Parameters

Table 228-51 GETDOCUMENTELEMENT Function Parameters

Parameter	Description
doc	DOMDOCUMENT

228.7.31 GETELEMENTSBYTAGNAME Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Returns a DOMNODELIST of all the elements with a specified tagname (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.GETELEMENTSBYTAGNAME(
   doc         IN      DOMDOCUMENT,
   tagname     IN      VARCHAR2)
 RETURN DOMNODELIST;
```

Returns the element children of the DOMELEMENT given the tag name (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.GETELEMENTSBYTAGNAME(
   elem      IN     DOMELEMENT,
   name      IN     VARCHAR2)
 RETURN DOMNODELIST;
```

Returns the element children of the DOMELEMENT given the tag name and namespace (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.GETELEMENTSBYTAGNAME(
   elem      IN     DOMELEMENT,
   name      IN     VARCHAR2,
   ns        IN     VARCHAR2)
 RETURN DOMNODELIST;
```

Parameters

Table 228-52 GETELEMENTSBYTAGNAME Function Parameters

Parameter	Description
doc	DOMDOCUMENT
tagname	Name of the tag to match on
elem	The DOMELEMENT
name	Tag name; using a wildcard(*) would match any tag
ns	Namespace

228.7.32 GETENTITIES Function

This function retrieves a DOMNAMEDNODEMAP containing the general entities, both external and internal, declared in the DTD.

> **See Also:**
>
> DOMDocumentType Subprograms

Syntax

```
DBMS_XMLDOM.GETENTITIES(
   dt       IN     DOMDocumentType)
 RETURN DOMNAMEDNODEMAP;
```

Parameters

Table 228-53 GETENTITIES Function Parameters

Parameter	Description
dt	DOMDOCUMENTTYPE

228.7.33 GETEXPANDEDNAME Procedure and Functions

This subprogram is overloaded as a procedure and two functions. The specific forms of functionality are described along with the syntax declarations.

Syntax

Retrieves the expanded name of the `Node` if is in an `Element` or `Attribute` type; otherwise, returns `NULL` (See Also: DOMNode Subprograms)

```
DBMS_XMLDOM.GETEXPANDEDNAME(
   n        IN       DOMNODE
   data     OUT      VARCHAR);
```

Returns the expanded name of the `DOMAttr` (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.GETEXPANDEDNAME(
   a        IN      DOMAttr)
 RETURN VARCHAR2;
```

Returns the expanded name of the `DOMElement` (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.GETEXPANDEDNAME(
   elem       IN     DOMELEMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-54 GETEXPANDEDNAME Procedure and Function Parameters

Parameter	Description
n	DOMNODE
data	Returned expanded name of the `Node`
a	DOMATTR
elem	DOMELEMENT

228.7.34 GETFIRSTCHILD Function

This function retrieves the first child of this node. If there is no such node, this returns `NULL`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETFIRSTCHILD(
   n       IN       DOMNODE)
 RETURN DOMNODE;
```

Parameters

Table 228-55 GETFIRSTCHILD Function Parameters

Parameter	Description
n	DOMNODE

228.7.35 GETIMPLEMENTATION Function

This function returns the `DOMIMPLEMENTATION` object that handles this `DOMDOCUMENT`.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETIMPLEMENTATION(
   doc       IN      DOMDOCUMENT)
 RETURN DOMIMPLEMENTATION;
```

Parameters

Table 228-56 GETIMPLEMENTATION Function Parameters

Parameter	Description
doc	DOMDOCUMENT

228.7.36 GETLASTCHILD Function

This function retrieves the last child of this node. If there is no such node, this returns `NULL`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETLASTCHILD(
   n      IN    DOMNODE)
 RETURN DOMNODE;
```

Parameters

Table 228-57 GETLASTCHILD Function Parameters

Parameter	Description
n	DOMNODE

228.7.37 GETLENGTH Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Gets the number of characters in the data. This may have the value zero, because CharacterData nodes may be empty (See Also: DOMCharacterData Subprograms):

```
DBMS_XMLDOM.GETLENGTH(
   cd      IN      DOMCHARACTERDATA)
 RETURN NUMBER;
```

Gets the number of nodes in this map. The range of valid child node indexes is `0` to `length-1`, inclusive (See Also: DOMNamedNodeMap Subprograms):

```
DBMS_XMLDOM.GETLENGTH(
   nnm       IN      DOMNAMEDNODEMAP)
 RETURN NUMBER;
```

Gets the number of nodes in the list. The range of valid child node indexes is `0` to `length-1`, inclusive (See Also: DOMNodeList Subprograms):

```
DBMS_XMLDOM.GETLENGTH(
   nl      IN     DOMNODELIST)
 RETURN NUMBER;
```

Parameters

Table 228-58 GETLENGTH Function Parameters

Parameter	Description
cd	DOMCHARACTERDATA
nnm	DOMNAMEDNODEMAP
nl	DOMNODELIST

228.7.38 GETLOCALNAME Procedure and Functions

This function is overloaded as a procedure and two functions. The specific forms of functionality are described alongside the syntax declarations.

Syntax

Retrieves the local part of the node's qualified name (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.GETLOCALNAME(
   n        IN     DOMNODE,
   data     OUT    VARCHAR2);
```

Returns the local name of the `DOMAttr` (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.GETLOCALNAME(
   a        IN     DOMATTR)
 RETURN VARCHAR2;
```

Returns the local name of the `DOMElement` (See Also: DOMElement Subprograms)

```
DBMS_XMLDOM.GETLOCALNAME(
   elem        IN     DOMELEMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-59 GETLOCALNAME Procedure and Function Parameters

Parameter	Description
n	DOMNode
data	Returned local name.
a	DOMAttr.
elem	DOMElement.

228.7.39 GETNAME Functions

This function is overloaded. The specific forms of functionality are described with the syntax declarations.

Syntax

Returns the name of this attribute (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.GETNAME(
   a        IN     DOMATTR)
 RETURN VARCHAR2;
```

Retrieves the name of DTD, or the name immediately following the `DOCTYPE` keyword (See Also: DOMDocumentType Subprograms):

```
DBMS_XMLDOM.GETNAME(
   dt        IN     DOMDOCUMENTTYPE)
 RETURN VARCHAR2;
```

Parameters

Table 228-60 GETNAME Function Parameters

Parameter	Description
a	DOMATTR
dt	DOMDOCUMENTTYPE

228.7.40 GETNAMEDITEM Function

`GETNAMEDITEM` retrieves a node specified by name.

> **See Also:**
>
> DOMNamedNodeMap Subprograms

Syntax

Retrieves a node specified by name:

```
DBMS_XMLDOM.GETNAMEDITEM(
   nnm    IN  DOMNAMEDNODEMAP,
   name   IN  VARCHAR2)
 RETURN DOMNODE;
```

Retrieves a node specified by name and namespace URI:

```
DBMS_XMLDOM.GETNAMEDITEM(
   nnm    IN  DOMNAMEDNODEMAP,
   name   IN  VARCHAR2,
   ns     IN  VARCHAR2)
 RETURN DOMNODE;
```

Parameters

Table 228-61 GETNAMEDITEM Function Parameters

Parameter	Description
nnm	DOMNAMEDNODEMAP
name	Name of the item to be retrieved
ns	Namespace

228.7.41 GETNAMESPACE Procedure and Functions

This subprogram is overloaded as a procedure and two functions. The specific forms of functionality are described alongside the syntax declarations.

Syntax

Retrieves the namespace URI associated with the node (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.GETNAMESPACE(
   n        IN     DOMNODE,
   data     OUT    VARCHAR2);
```

Retrieves the namespace of the `DOMATTR` (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.GETNAMESPACE(
   a        IN     DOMATTR)
 RETURN VARCHAR2;
```

Retrieves the namespace of the `DOMELEMENT` (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.GETNAMESPACE(
   elem       IN     DOMELEMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-62 GETNAMESPACE Procedure and Function Parameters

Parameter	Description
n	DOMNODE
data	Returned namespace URI
a	DOMATTR
elem	DOMELEMENT

228.7.42 GETNEXTSIBLING Function

This function retrieves the node immediately following this node. If there is no such node, this returns `NULL`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETNEXTSIBLING(
   n        IN     DOMNODE)
 RETURN DOMNode;
```

Parameters

Table 228-63 GETNEXTSIBLING Function Parameters

Parameter	Description
n	DOMNODE

228.7.43 GETNODETYPE Function

This function retrieves a code representing the type of the underlying object.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETNODETYPE(
   n        IN      DOMNODE)
 RETURN NUMBER;
```

Parameters

Table 228-64 GETNODETYPE Function Parameters

Parameter	Description
n	DOMNODE

228.7.44 GETNODENAME Function

This function gets the name of the node depending on its type.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETNODENAME(
   n        IN      DOMNODE)
 RETURN VARCHAR2;
```

Parameters

Table 228-65 GETNODENAME Function Parameters

Parameter	Description
n	DOMNODE

228.7.45 GETNODEVALUE Function

This function gets the value of this node, depending on its type.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETNODEVALUE(
   n        IN      DOMNODE)
 RETURN VARCHAR2;*
```

Parameters

Table 228-66 GETNODEVALUE Function Parameters

Parameter	Description
n	DOMNODE

228.7.46 GETNODEVALUEASBINARYSTREAM Function & Procedure

The operation of these subprograms is described with each syntax implementation.

> **See Also:**
>
> DOMNode Subprograms

Syntax

This function returns an instance of the PL/SQL `XMLBinaryInputStream`. The node datatype must be `RAW` or `BLOB` – if not an exception is raised.

```
DBMS_XMLDOM.GETNODEVALUEASBINARYSTREAM (
   n      IN     DOMNODE)
 RETURN SYS.UTL_BINARYINPUTSTREAM;
```

Using this procedure, the application passes an implementation of `SYS.UTL_BINARYOUTPUTSTREAM` into which XDB writes the contents of the node. The datatype of the node must be `RAW` or `CLOB` – if not an exception is raised.

```
DBMS_XMLDOM.GETNODEVALUEASBINARYSTREAM (
   n        in   DOMNODE,
   value    in   SYS.UTL_BINARYOUTPUTSTREAM);
```

Parameters

Table 228-67 GETNODEVALUEASBINARYSTREAM Function & Procedure Parameters

Parameter	Description
n	DOMNODE
value	BINARYOUTPUTSTREAM

228.7.47 GETNODEVALUEASCHARACTERSTREAM Function & Procedure

The operation of these subprograms is described with each syntax implementation.

> **See Also:**
>
> DOMNode Subprograms

Syntax

This function returns an instance of the PL/SQL `XMLCharacterInputStream`. If the node data is character it is converted to the current session character set. If the node data is not character data, it is first converted to character data.

```
DBMS_XMLDOM.GETNODEVALUEASCHARACTERSTREAM  (
   n         IN     DOMNODE)
 RETURN SYS.UTL_CHARACTERINPUTSTREAM;
```

Using this procedure, the node data is converted, as necessary, to the session character set and then "pushed" into the `SYS.UTL_CHARACTEROUTPUTSTREAM`.

```
DBMS_XMLDOM.GETNODEVALUEASCHARACTERSTREAM  (
   n         IN    DOMNODE,
   value     IN    SYS.UTL_CHARACTEROUTPUTSTREAM);
```

Parameters

Table 228-68 GETNODEVALUEASCHARACTERSTREAM Function & Procedure Parameters

Parameter	Description
n	DOMNODE
value	CHARACTEROUTPUTSTREAM

228.7.48 GETNOTATIONNAME Function

This function returns the notation name of the `DOMENTITY`.

> **See Also:**
>
> DOMEntity Subprograms

Syntax

```
DBMS_XMLDOM.GETNOTATIONNAME(
   ent        IN     DOMENTITY)
 RETURN VARCHAR2;
```

Parameters

Table 228-69 GETNOTATIONNAME Function Parameters

Parameter	Description
ent	DOMENTITY

228.7.49 GETNOTATIONS Function

This function retrieves a DOMNAMEDNODEMAP containing the notations declared in the DTD.

> **See Also:**
>
> DOMDocumentType Subprograms

Syntax

```
DBMS_XMLDOM.GETNOTATIONS(
   dt        IN      DOMDOCUMENTTYPE)
 RETURN DOMNAMEDNODEMAP;
```

Parameters

Table 228-70 GETNOTATIONS Function Parameters

Parameter	Description
dt	DOMDOCUMENTTYPE

228.7.50 GETTARGET Function

This function returns the target of the DOMPROCESSINGINSTRUCTION.

> **See Also:**
>
> DOMProcessingInstruction Subprograms

Syntax

```
DBMS_XMLDOM.GETTARGET(
   pi        IN      DOMPROCESSINGINSTRUCTION)
 RETURN VARCHAR2;
```

Parameters

Table 228-71 GETTARGET Function Parameters

Parameter	Description
pi	DOMPROCESSINGINSTRUCTION

228.7.51 GETOWNERDOCUMENT Function

This function retrieves the Document object associated with this node. This is also the Document object used to create new nodes. When this node is a Document or a Document Type that is not used with any Document yet, this is `NULL`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETOWNERDOCUMENT(
   n        IN      DOMNODE)
 RETURN DOMDOCUMENT;
```

Parameters

Table 228-72 GETOWNERDOCUMENT Function Parameters

Parameter	Description
n	DOMNODE

228.7.52 GETOWNERELEMENT Function

This function retrieves the Element node to which the specified Attribute is attached.

> **See Also:**
>
> DOMAttr Subprograms

Syntax

```
DBMS_XMLDOM.GETOWNERELEMENT(
   a        IN      DOMATTR)
 RETURN DOMElement;
```

Parameters

Table 228-73 GETOWNERELEMENT Function Parameters

Parameter	Description
a	Attribute

228.7.53 GETPARENTNODE Function

This function retrieves the parent of this node. All nodes, except `Attr`, `Document`, `DocumentFragment`, `Entity`, and `Notation` may have a parent. However, if a node has just been created and not yet added to the tree, or if it has been removed from the tree, this is `NULL`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETPARENTNODE(
   n        IN      DOMNODE)
 RETURN DOMNODE;
```

Parameters

Table 228-74 GETPARENTNODE Function Parameters

Parameter	Description
n	DOMNODE

228.7.54 GETPREFIX Function

This function retrieves the namespace prefix of the node.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETPREFIX(
   n        IN      DOMNODE)
 RETURN VARCHAR2;
```

Parameters

Table 228-75 GETPREFIX Function Parameters

Parameter	Description
n	DOMNODE

228.7.55 GETPREVIOUSSIBLING Function

This function retrieves the node immediately preceding this node. If there is no such node, this returns NULL.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETPREVIOUSSIBLING(
   n        IN     DOMNODE)
 RETURN DOMNODE;
```

Parameters

Table 228-76 GETPREVIOUSSIBLING Function Parameters

Parameter	Description
n	DOMNODE

228.7.56 GETPUBLICID Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Returns the public identifier of the specified DTD (See Also: DOMDocumentType Subprograms):

```
DBMS_XMLDOM.GETPUBLICID(
   dt        IN     DOMDOCUMENTTYPE)
 RETURN VARCHAR2;
```

Returns the public identifier of the DOMENTITY (See Also: DOMEntity Subprograms):

```
DBMS_XMLDOM.GETPUBLICID(
   ent       IN     DOMENTITY)
 RETURN VARCHAR2;
```

Returns the public identifier of the DOMNOTATION (See Also: DOMNotation Subprograms):

```
DBMS_XMLDOM.GETPUBLICID(
   n         IN      DOMNOTATION)
 RETURN VARCHAR2;
```

Parameters

Table 228-77 GETPUBLICID Function Parameters

Parameter	Description
dt	The DTD
ent	DOMENTITY
n	DOMNOTATION

228.7.57 GETQUALIFIEDNAME Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Returns the qualified name of the `DOMATTR` (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.GETQUALIFIEDNAME(
   a         IN      DOMATTR)
 RETURN VARCHAR2;
```

Returns the qualified name of the `DOMElement` (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.GETQUALIFIEDNAME(
   elem      IN      DOMELEMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-78 GETQUALIFIEDNAME Functions Parameters

Parameter	Description
a	DOMATTR
elem	DOMELEMENT

228.7.58 GETSCHEMANODE Function

This function retrieves the schema URI associated with the node.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.GETSCHEMANODE(
   n        IN      DOMNODE)
 RETURN DOMNODE;
```

Parameters

Table 228-79 GETSCHEMANODE Function Parameters

Parameter	Description
n	DOMNODE

228.7.59 GETSPECIFIED Function

If this attribute was explicitly specified, a value in the original document, this is true; otherwise, it is false.

> **See Also:**
>
> DOMAttr Subprograms

Syntax

```
DBMS_XMLDOM.GETSPECIFIED(
   a        IN      DOMATTR)
 RETURN BOOLEAN;
```

Parameters

Table 228-80 GETSPECIFIED Function Parameters

Parameter	Description
a	DOMATTR

228.7.60 GETSTANDALONE Function

This function returns the standalone property associated with the `DOMDOCUMENT`.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETSTANDALONE(
   doc        IN      DOMDOCUMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-81 GETSTANDALONE Function Parameters

Parameter	Description
doc	DOMDOCUMENT.

228.7.61 GETSYSTEMID Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Returns the system id of the specified DTD (See Also: DOMDocumentType Subprograms):

```
DBMS_XMLDOM.GETSYSTEMID(
   dt        IN     DOMDOCUMENTTYPE)
 RETURN VARCHAR2;
```

Returns the system identifier of the DOMENTITY (See Also: DOMEntity Subprograms):

```
DBMS_XMLDOM.GETSYSTEMID(
   ent       IN     DOMENTITY)
 RETURN VARCHAR2;
```

Returns the system identifier of the DOMNOTATION (See Also: DOMNotation Subprograms):

```
DBMS_XMLDOM.GETSYSTEMID(
   n         IN     DOMNOTATION)
 RETURN VARCHAR2;
```

Parameters

Table 228-82 GETSYSTEMID Function Parameters

Parameter	Description
dt	The DTD.
ent	DOMEntity.
n	DOMNotation.

228.7.62 GETTAGNAME Function

This function returns the name of the DOMELEMENT.

> **See Also:**
>
> DOMElement Subprograms

Syntax

```
DBMS_XMLDOM.GETTAGNAME(
   elem       IN     DOMELEMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-83 GETTAGNAME Function Parameters

Parameter	Description
elem	The DOMELEMENT

228.7.63 GETVALUE Function

This function retrieves the value of the attribute.

> **See Also:**
>
> DOMAttr Subprograms

Syntax

```
DBMS_XMLDOM.GETVALUE(
   a       IN     DOMATTR)
 RETURN VARCHAR2;
```

Parameters

Table 228-84 GETVALUE Function Parameters

Parameter	Description
a	DOMATTR

228.7.64 GETVERSION Function

This function returns the version of the DOMDOCUMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETVERSION(
   doc       IN     DOMDOCUMENT)
 RETURN VARCHAR2;
```

Parameters

Table 228-85 GETVERSION Function Parameters

Parameter	Description
doc	DOMDOCUMENT

228.7.65 GETXMLTYPE Function

This function returns the `XMLType` associated with the `DOMDOCUMENT`.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.GETXMLTYPE(
   doc        IN     DOMDOCUMENT)
 RETURN SYS.XMLTYPE;
```

Parameters

Table 228-86 GETXMLTYPE Function Parameters

Parameter	Description
doc	DOMDOCUMENT

228.7.66 HASATTRIBUTE Functions

Verifies whether an attribute has been defined for `DOMELEMENT`, or has a default value.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Verifies whether an attribute with the specified name has been defined for `DOMElement`:

```
DBMS_XMLDOM.HASATTRIBUTE(
   elem      IN  DOMELEMENT,
   name      IN  VARCHAR2)
 RETURN VARCHAR2;
```

Verifies whether an attribute with specified name and namespace URI has been defined for `DOMELEMENT`; namespace enabled:

```
DBMS_XMLDOM.HASATTRIBUTE(
   elem      IN  DOMELEMENT,
   name      IN  VARCHAR2,
   ns        IN  VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 228-87 HASATTRIBUTE Function Parameters

Parameter	Description
`elem`	The `DOMELEMENT`
`name`	Attribute name; * matches any attribute
`ns`	Namespace

228.7.67 HASATTRIBUTES Function

This function returns whether this node has any attributes.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.HASATTRIBUTES(
   n        IN      DOMNODE)
 RETURN BOOLEAN;
```

Parameters

Table 228-88 HASATTRIBUTES Function Parameters

Parameter	Description
`n`	`DOMNODE`

228.7.68 HASCHILDNODES Function

This function determines whether this node has any children.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.HASCHILDNODES(
   n        IN     DOMNODE)
 RETURN BOOLEAN;
```

Parameters

Table 228-89 HASCHILDNODES Function Parameters

Parameter	Description
n	DOMNODE

228.7.69 HASFEATURE Function

This function tests if the DOMIMPLEMENTATION implements a specific feature.

> **See Also:**
>
> DOMImplementation Subprograms

Syntax

```
DBMS_XMLDOM.HASFEATURE(
   di       IN     DOMIMPLEMENTATION,
   feature  IN     VARCHAR2,
   version  IN     VARCHAR2)
 RETURN BOOLEAN;
```

Parameters

Table 228-90 HASFEATURE Function Parameters

Parameter	Description
di	DOMIMPLEMENTATION
feature	The feature to check for
version	The version of the DOM to check in

228.7.70 IMPORTNODE Function

This function imports a node from an external document and returns this new node.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.IMPORTNODE(
   doc            IN  DOMDOCUMENT,
   importedNode   IN  DOMNODE,
   deep           IN  BOOLEAN)
  RETURN DOMNODE;
```

Parameters

Table 228-91 IMPORTNODE Function Parameters

Parameter	Description
doc	Document from which the node is imported
importedNode	Node to import
deep	Setting for recursive import. • If this value is TRUE, the entire subtree of the node will be imported with the node. • If this value is FALSE, only the node itself will be imported.

Usage Notes

Note that the ADOPTNODE Function removes the node from the source document while the IMPORTNODE Function clones the node in the source document.

228.7.71 INSERTBEFORE Function

This function inserts the node newchild before the existing child node refchild. If refchild is NULL, insert newchild at the end of the list of children.

If newchild is a DOCUMENTFRAGMENT object, all of its children are inserted, in the same order, before refchild. If the newchild is already in the tree, it is first removed.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.INSERTBEFORE(
   n          IN     DOMNODE,
   newchild   IN     DOMNODE,
   refchild   IN     DOMNODE)
  RETURN DOMNode;
```

Parameters

Table 228-92 INSERTBEFORE Function Parameters

Parameter	Description
n	DOMNODE

Table 228-92 (Cont.) INSERTBEFORE Function Parameters

Parameter	Description
newChild	The child to be inserted in the DOMNODE
refChild	The reference node before which the newchild is to be inserted

228.7.72 INSERTDATA Procedure

This procedure inserts a string at the specified character offset.

> **See Also:**
>
> DOMCharacterData Subprograms

Syntax

```
DBMS_XMLDOM.INSERTDATA(
   cd       IN     DOMCHARACTERDATA,
   offset   IN     NUMBER,
   arg      IN     VARCHAR2);
```

Parameters

Table 228-93 INSERTDATA Procedure Parameters

Parameter	Description
cd	DOMCHARACTERDATA
offset	The offset at which to insert the data
arg	The value to be inserted

228.7.73 ISNULL Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Checks if the specified DOMNODE is NULL. Returns TRUE if it is NULL, FALSE otherwise (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.ISNULL(
  n         IN     DOMNODE)
 RETURN BOOLEAN;
```

Checks that the specified DOMATTR is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.ISNULL(
   a        IN     DOMATTR)
 RETURN BOOLEAN;
```

Checks that the specified DOMCDATASECTION is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMCDataSection Subprograms):

```
DBMS_XMLDOM.ISNULL(
   cds       IN      DOMCDATASECTION)
 RETURN BOOLEAN;
```

Checks that the specified DOMCHARACTERDATA is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMCharacterData Subprograms):

```
DBMS_XMLDOM.ISNULL(
   cd        IN      DOMCHARACTERDATA)
 RETURN BOOLEAN;
```

Checks that the specified DOMCOMMENT is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMComment Subprograms):

```
DBMS_XMLDOM.ISNULL(
   com        IN      DOMCOMMENT)
 RETURN BOOLEAN;
```

Checks that the specified DOMDOCUMENT is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.ISNULL(
   doc        IN      DOMDOCUMENT)
 RETURN BOOLEAN;
```

Checks that the specified DOMDOCUMENTFRAGMENT is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMDocumentFragment Subprograms):

```
DBMS_XMLDOM.ISNULL(
   df        IN      DOMDOCUMENTFRAGMENT)
 RETURN BOOLEAN;
```

Checks that the specified DOMDOCUMENTTYPE is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMDocumentType Subprograms):

```
DBMS_XMLDOM.ISNULL(
   dt        IN      DOMDOCUMENTTYPE)
 RETURN BOOLEAN;
```

Checks that the specified DOMELEMENT is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.ISNULL(
   elem      IN      DOMELEMENT)
 RETURN BOOLEAN;
```

Checks that the specified DOMENTITY is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMEntity Subprograms):

```
DBMS_XMLDOM.ISNULL(
   ent        IN      DOMENTITY)
 RETURN BOOLEAN;
```

Checks that the specified DOMENTITYREFERENCE is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMEntityReference Subprograms):

```
DBMS_XMLDOM.ISNULL(
   EREF       IN      DOMENTITYREFERENCE)
 RETURN BOOLEAN;
```

Checks that the specified DOMIMPLEMENTATION is NULL; returns TRUE if it is NULL (See Also: DOMImplementation Subprograms):

```
DBMS_XMLDOM.ISNULL(
   di        IN      DOMIMPLEMENTATION)
 RETURN BOOLEAN;
```

Checks that the specified DOMNAMEDNODEMAP is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMNamedNodeMap Subprograms):

```
DBMS_XMLDOM.ISNULL(
   nnm        IN      DOMNAMEDNODEMAP)
 RETURN BOOLEAN;
```

Checks that the specified DOMNODELIST is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMNodeList Subprograms):

```
DBMS_XMLDOM.ISNULL(
   nl        IN      DOMNODELIST)
 RETURN BOOLEAN;
```

Checks that the specified DOMNOTATION is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMNotation Subprograms):

```
DBMS_XMLDOM.ISNULL(
   n        IN      DOMNOTATION)
 RETURN BOOLEAN;
```

Checks that the specified DOMPROCESSINGINSTRUCTION is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMProcessingInstruction Subprograms):

```
DBMS_XMLDOM.ISNULL(
   pi        IN      DOMPROCESSINGINSTRUCTION)
 RETURN BOOLEAN;
```

Checks that the specified DOMTEXT is NULL; returns TRUE if it is NULL, FALSE otherwise (See Also: DOMText Subprograms):

```
DBMS_XMLDOM.ISNULL(
   t        IN      DOMTEXT)
 RETURN BOOLEAN;
```

Parameters

Table 228-94 ISNULL Function Parameters

Parameter	Description
n	DOMNODE to check
a	DOMATTR to check
cds	DOMCDATASECTION to check
cd	DOMCHARACTERDATA to check
com	DOMCOMMENT to check
doc	DOMDOCUMENT to check
dF	DOMDOCUMENTFRAGMENT to check
dt	DOMDOCUMENTTYPE to check
elem	DOMELEMENT to check

Table 228-94 (Cont.) ISNULL Function Parameters

Parameter	Description
ent	DOMENTITY to check
eref	DOMENTITYREFERENCE to check
di	DOMIMPLEMENTATION to check
nnm	DOMNAMENODEMAP to check
nl	DOMNODELIST to check
n	DOMNOTATION to check
pi	DOMPROCESSINGINSTRUCTION to check
t	DOMTEXT to check

228.7.74 ITEM Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Returns the item in the map which corresponds to the INDEX parameter. If INDEX is greater than or equal to the number of nodes in this map, this returns NULL (See Also: DOMNamedNodeMap Subprograms):

```
DBMS_XMLDOM.ITEM(
   nnm       IN     DOMNAMEDNODEMAP,
   index     IN     NUMBER)
 RETURN DOMNODE;
```

Returns the item in the collection which corresponds to the INDEX parameter. If index is greater than or equal to the number of nodes in the list, this returns NULL (See Also: DOMNodeList Subprograms):

```
DBMS_XMLDOM.ITEM(
   nl        IN     DOMNODELIST,
   index     IN     NUMBER)
 RETURN DOMNODE;
```

Parameters

Table 228-95 ITEM Function Parameters

Parameter	Description
nnm	DOMNAMEDNODEMAP
index	The index in the node map at which the item is to be retrieved
nl	DOMNODELIST
index	The index in the NodeList used to retrieve the item

228.7.75 MAKEATTR Function

This function casts a specified `DOMNODE` to a `DOMATTR`, and returns the `DOMATTR`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEATTR(
   n        IN      DOMNODE)
 RETURN DOMATTR;
```

Parameters

Table 228-96 MAKEATTR Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.76 MAKECDATASECTION Function

This function casts a specified `DOMNODE` to a `DOMCDATASECTION`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKECDATASECTION(
   n        IN      DOMNODE)
 RETURN DOMCDATASECTION;
```

Parameters

Table 228-97 MAKECDATASECTION Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.77 MAKECHARACTERDATA Function

This function casts a specified `DOMNODE` to a `DOMCHARACTERDATA`, and returns the `DOMCHARACTERDATA`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKECHARACTERDATA(
   n        IN      DOMNode)
 RETURN DOMCharacterData;
```

Parameters

Table 228-98 MAKECHARACTERDATA Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.78 MAKECOMMENT Function

This function casts a specified `DOMNODE` to a `DOMCOMMENT`, and returns the `DOMCOMMENT`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKECOMMENT(
   n        IN      DOMNODE)
 RETURN DOMCOMMENT;
```

Parameters

Table 228-99 MAKECOMMENT Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.79 MAKEDOCUMENT Function

This function casts a specified `DOMNODE` to a `DOMDOCUMENT`, and returns the `DOMDOCUMENT`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEDOCUMENT(
   n        IN      DOMNODE)
 RETURN DOMDocument;
```

Parameters

Table 228-100 MAKEDOCUMENT Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.80 MAKEDOCUMENTFRAGMENT Function

This function casts a specified `DOMNODE` to a `DOMDOCUMENTFRAGMENT`, and returns the `DOMDOCUMENTFRAGMENT`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEDOCUMENTFRAGMENT(
   n        IN      DOMNODE)
 RETURN DOMDOCUMENTFRAGMENT;
```

Parameters

Table 228-101 MAKEDOCUMENTFRAGMENT Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.81 MAKEDOCUMENTTYPE Function

This function casts a specified `DOMNODE` to a `DOMDOCUMENTTYPE` and returns the `DOMDOCUMENTTYPE`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEDOCUMENTTYPE(
   n        IN      DOMNODE)
 RETURN DOMDOCUMENTTYPE;
```

Parameters

Table 228-102 MAKEDOCUMENTTYPE Function Parameters

Parameter	Description
n	`DOMNODE` to cast.

228.7.82 MAKEELEMENT Function

This function casts a specified `DOMNODE` to a `DOMELEMENT`, and returns the `DOMELEMENT`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEELEMENT(
   n        IN      DOMNODE)
 RETURN DOMELEMENT;
```

Parameters

Table 228-103 MAKEELEMENT Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.83 MAKEENTITY Function

This function casts a specified `DOMNODE` to a `DOMENTITY`, and returns the `DOMENTITY`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEENTITY(
   n        IN      DOMNODE)
 RETURN DOMENTITY;
```

Parameters

Table 228-104 MAKEENTITY Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.84 MAKEENTITYREFERENCE Function

This function casts a specified `DOMNODE` to a `DOMENTITYREFERENCE`, and returns the `DOMENTITYREFERENCE`.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEENTITYREFERENCE(
   n        IN      DOMNODE)
 RETURN DOMENTITYREFERENCE;
```

Parameters

Table 228-105 MAKEENTITYREFERENCE Function Parameters

Parameter	Description
n	`DOMNODE` to cast

228.7.85 MAKENODE Functions

This function is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Casts specified DOMATTR to a DOMNODE, and returns the DOMNODE (See Also: DOMAttr Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   a         IN     DOMATTR)
 RETURN DOMNODE;
```

Casts the DOMCDATASECTION to a DOMNODE, and returns that DOMNODE (See Also: DOMCDataSection Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   cds       IN     DOMCDATASECTION)
 RETURN DOMNODE;
```

Casts the specified DOMCHARACTERDATA as a DOMNODE, and returns that DOMNODE (See Also: DOMCharacterData Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   cd        IN     DOMCHARACTERDATA)
 RETURN DOMNODE;
```

Casts the specified DOMCOMMENT to a DOMNODE, and returns that DOMNODE (See Also: DOMComment Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   com       IN     DOMCOMMENT)
 RETURN DOMNODE;
```

Casts the DOMDOCUMENT to a DOMNODE, and returns that DOMNODE (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   doc       IN     DOMDOCUMENT)
 RETURN DOMNODE;
```

Casts the specified DOMDOCUMENTFRAGMENT to a DOMNODE, and returns that DOMNODE (See Also: DOMDocumentFragment Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   df        IN     DOMDOCUMENTFRAGMENT)
 RETURN DOMNode;
```

Casts the specified DOMDOCUMENTTYPE to a DOMNODE, and returns that DOMNODE (See Also: DOMDocumentType Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   dt        IN     DOMDOCUMENTTYPE)
 RETURN DOMNODE;
```

Casts the specified DOMELEMENT to a DOMNODE, and returns that DOMNODE (See Also: DOMElement Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   elem       IN     DOMELEMENT)
 RETURN DOMNODE;
```

Casts specified DOMENTITY to a DOMNODE, and returns that DOMNODE (See Also: DOMEntity Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   ent        IN     DOMENTITY)
 RETURN DOMNODE;
```

Casts the DOMENTITYREFERENCE to a DOMNODE, and returns that DOMNODE (See Also: DOMEntityReference Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   eref       IN     DOMENTITYREFERENCE)
 RETURN DOMNODE;
```

Casts the DOMNOTATION to a DOMNODE, and returns that DOMNODE (See Also: DOMNotation Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   n        IN     DOMNOTATION)
 RETURN DOMNODE;
```

Casts the DOMPROCESSINGINSTRUCTION to a DOMNODE, and returns the DOMNODE (See Also: DOMProcessingInstruction Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   pi       IN     DOMPROCESSINGINSTRUCTION)
 RETURN DOMNODE;
```

Casts the DOMTEXT to a DOMNODE, and returns that DOMNODE (See Also: DOMText Subprograms):

```
DBMS_XMLDOM.MAKENODE(
   t        IN     DOMTEXT)
 RETURN DOMNODE;
```

Parameters

Table 228-106 MAKENODE Function Parameters

Parameter	Description
a	DOMATTR to cast
cds	DOMCDATASECTION to cast
cd	DOMCHARACTERDATA to cast
com	DOMCOMMENT to cast
doc	DOMDOCUMENT to cast
df	DOMDOCUMENTFRAGMENT to cast
dt	DOMDOCUMENTTYPE to cast
elem	DOMELEMENT to cast
ent	DOMENTITY to cast
eref	DOMENTITYREFERENCE to cast
n	DOMNOTATION to cast

Table 228-106 (Cont.) MAKENODE Function Parameters

Parameter	Description
pi	DOMPROCESSINGINSTRUCTION to cast
t	DOMTEXT to cast

228.7.86 MAKENOTATION Function

This function casts a specified DOMNODE to a DOMNOTATION, and returns the DOMNOTATION.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKENOTATION(
   n        IN      DOMNODE)
 RETURN DOMNOTATION;
```

Parameters

Table 228-107 MAKENOTATION Function Parameters

Parameter	Description
n	DOMNODE to cast

228.7.87 MAKEPROCESSINGINSTRUCTION Function

This function casts a specified DOMNODE to a DOMPROCESSINGINSTRUCTION, and returns the Domprocessinginstruction.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKEPROCESSINGINSTRUCTION(
   n        IN      DOMNODE)
 RETURN DOMPROCESSINGINSTRUCTION;
```

Parameters

Table 228-108 MAKEPROCESSINGINSTRUCTION Function Parameters

Parameter	Description
n	DOMNODE to cast

228.7.88 MAKETEXT Function

This function casts a specified DOMNODE to a DOMTEXT, and returns the DOMTEXT.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.MAKETEXT(
   n        IN      DOMNODE)
 RETURN DOMTEXT;
```

Parameters

Table 228-109 MAKETEXT Function Parameters

Parameter	Description
n	DOMNODE to cast

228.7.89 NEWDOMDOCUMENT Functions

This function returns a new DOMDOCUMENT instance.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

Returns a new DOMDOCUMENT instance:

```
DBMS_XMLDOM.NEWDOMDOCUMENT
 RETURN DOMDOCUMENT;
```

Returns a new DOMDOCUMENT instance created from the specified XMLType object:

```
DBMS_XMLDOM.NEWDOMDOCUMENT(
   xmldoc     IN SYS.XMLTYPE)
 RETURN DOMDOCUMENT;
```

Returns a new DOMDOCUMENT instance created from the specified CLOB:

```
DBMS_XMLDOM.NEWDOMDOCUMENT(
   cl        IN     CLOB)
 RETURN DOMDOCUMENT;
```

Parameters

Table 228-110 NEWDOMDOCUMENT Function Parameters

Parameter	Description
xmldoc	XMLType source for the DOMDOCUMENT
cl	CLOB source for the DOMDOCUMENT

228.7.90 NORMALIZE Procedure

This procedure normalizes the text children of the DOMELEMENT.

> **See Also:**
>
> DOMElement Subprograms

Syntax

```
DBMS_XMLDOM.NORMALIZE(
   elem       IN     DOMELEMENT);
```

Parameters

Table 228-111 NORMALIZE Procedure Parameters

Parameter	Description
elem	The DOMELEMENT

228.7.91 REMOVEATTRIBUTE Procedures

This procedure removes an attribute from the DOMELEMENT by name.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Removes the value of a DOMELEMENT's attribute by name:

```
DBMS_XMLDOM.REMOVEATTRIBUTE(
   elem      IN     DOMELEMENT,
   name      IN     VARCHAR2);
```

Removes the value of a DOMELEMENT's attribute by name and namespace URI.

```
DBMS_XMLDOM.REMOVEATTRIBUTE(
   elem     IN    DOMELEMENT,
   name     IN    VARCHAR2,
   ns       IN    VARCHAR2);
```

Parameters

Table 228-112 REMOVEATTRIBUTE Procedure Parameters

Parameter	Description
elem	The DOMELEMENT
name	Attribute name
ns	Namespace

228.7.92 REMOVEATTRIBUTENODE Function

This function removes the specified attribute node from the DOMELEMENT. The method returns the removed node.

> **See Also:**
>
> DOMElement Subprograms

Syntax

```
DBMS_XMLDOM.REMOVEATTRIBUTENODE(
   elem       IN     DOMELEMENT,
   oldAttr    IN     DOMATTR)
  RETURN DOMAttr;
```

Parameters

Table 228-113 REMOVEATTRIBUTENODE Function Parameters

Parameter	Description
elem	The DOMELEMENT.
oldAttr	The old DOMATTR.

228.7.93 REMOVECHILD Function

This function removes the child node indicated by `oldchild` from the list of children, and returns it.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.REMOVECHILD(
   n          IN     DOMNode,
   oldchild   IN     DOMNode)
 RETURN DOMNODE;
```

Parameters

Table 228-114 REMOVECHILD Function Parameters

Parameter	Description
`n`	`DOMNODE`
`oldCHild`	The child of the node `n` to be removed

228.7.94 REMOVENAMEDITEM Function

This function removes a node, specified by name, from the map and returns this node.

When this map contains the attributes attached to an element, if the removed attribute is known to have a default value, an attribute immediately appears containing the default value as well as the corresponding namespace URI, local name, and prefix when applicable.

> **See Also:**
>
> DOMNamedNodeMap Subprograms

Syntax

Removes a node specified by name:

```
DBMS_XMLDOM.REMOVENAMEDITEM(
   nnm      IN     DOMNamedNodeMap,
   name     IN     VARCHAR2)
 RETURN DOMNode;
```

Removes a node specified by name and namespace URI:

```
DBMS_XMLDOM.REMOVENAMEDITEM(
   nnm      IN     DOMNamedNodeMap,
   name     IN     VARCHAR2,
```

```
   ns       IN     VARCHAR2)
 RETURN DOMNode;
```

Parameters

Table 228-115 REMOVENAMEDITEM Function Parameters

Parameter	Description
nnm	DOMNamedNodeMap
name	The name of the item to be removed from the map
ns	Namespace

228.7.95 REPLACECHILD Function

This function replaces the child node `oldchild` with `newchild` in the list of children, and returns the `oldchild` node.

If `newchild` is a `DocumentFragment` object, `oldchild` is replaced by all of the `DocumentFragment` children, which are inserted in the same order. If the `newchild` is already in the tree, it is first removed.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.REPLACECHILD(
   n           IN     DOMNode,
   newchild    IN     DOMNode,
   oldchild    IN     DOMNode)
 RETURN DOMNode;
```

Parameters

Table 228-116 REPLACECHILD Function Parameters

Parameter	Description
n	DOMNode
newchild	The new child which is to replace the old child
oldchild	The child of the node n which is to be replaced

228.7.96 REPLACEDATA Procedure

This procedure changes a range of characters in the node. Upon success, data and length reflect the change.

> **See Also:**
>
> DOMCharacterData Subprograms

Syntax

```
DBMS_XMLDOM.REPLACEDATA(
   cd        IN     DOMCHARACTERDATA,
   offset    IN     NUMBER,
   cnt       IN     NUMBER,
   arg       IN     VARCHAR2);
```

Parameters

Table 228-117 REPLACEDATA Procedure Parameters

Parameter	Description
cd	DOMCHARACTERDATA
offset	The offset at which to replace
cnt	The number of characters to replace
arg	The value to replace with

228.7.97 RESOLVENAMESPACEPREFIX Function

This function resolves the specified namespace prefix, and returns the resolved namespace.

> **See Also:**
>
> DOMElement Subprograms

Syntax

```
DBMS_XMLDOM.RESOLVENAMESPACEPREFIX(
   elem       IN     DOMELEMENT,
   prefix     IN     VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 228-118 RESOLVENAMESPACEPREFIX Function Parameters

Parameter	Description
elem	The DOMELEMENT
prefix	Namespace prefix

228.7.98 SETATTRIBUTE Procedures

This procedure sets the value of a DOMELEMENT's attribute by name.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Sets the value of a DOMELEMENT's attribute by name:

```
DBMS_XMLDOM.SETATTRIBUTE(
   elem       IN  DOMELEMENT,
   name       IN  VARCHAR2,
   newvalue   IN  VARCHAR2);
```

Sets the value of a DOMElement's attribute by name and namespace URI:

```
DBMS_XMLDOM.SETATTRIBUTE(
   elem       IN  DOMELEMENT,
   name       IN  VARCHAR2,
   newvalue   IN  VARCHAR2,
   ns         IN  VARCHAR2);
```

Parameters

Table 228-119 SETATTRIBUTE Procedure Parameters

Parameter	Description
elem	The DOMELEMENT
name	Attribute name
newvalue	Attribute value
ns	Namespace

228.7.99 SETATTRIBUTENODE Functions

This function adds a new attribute node to the `DOMELEMENT`.

> **See Also:**
>
> DOMElement Subprograms

Syntax

Adds a new attribute node to the `DOMELEMENT`:

```
DBMS_XMLDOM.SETATTRIBUTENODE(
   elem      IN  DOMELEMENT,
   newAttr   IN  DOMATTR)
 RETURN DOMATTR;
```

Adds a new attribute node to the `DOMElement`; namespace enabled:

```
DBMS_XMLDOM.SETATTRIBUTENODE(
   elem      IN  DOMELEMENT,
   newAttr   IN  DOMATTR,
   ns        IN  VARCHAR2)
 RETURN DOMATTR;
```

Parameters

Table 228-120 SETATTRIBUTENODE Function Parameters

Parameter	Description
`elem`	The `DOMELEMENT`
`newAttr`	The new `DOMATTR`
`ns`	The namespace

228.7.100 SETCHARSET Procedure

This function sets the characterset of the DOM document.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.SETCHARSET(
   doc      IN    DOMDocument,
   charset  IN    VARCHAR2);
```

Parameters

Table 228-121 SETCHARSET Procedure Parameters

Parameter	Description
doc	DOM document
charset	Characterset

Usage Notes

This is used for WRITETOFILE Procedures if not explicitly specified at that time.

228.7.101 SETDATA Procedures

This overloaded procedure sets character data or DOMPROCESSINGINSTRUCTION content data. The specific functionality is described in the syntax declarations.

Syntax

Sets the character data of the node that implements this interface (See Also: DOMCharacterData Subprograms):

```
DBMS_XMLDOM.SETDATA(
   cd       IN     DOMCHARACTERDATA,
   data     IN     VARCHAR2);
```

Sets the content data of the DOMPROCESSINGINSTRUCTION (See Also: DOMProcessingInstruction Subprograms):

```
DBMS_XMLDOM.SETDATA(
   pi       IN     DOMPROCESSINGINSTRUCTION,
   data     IN     VARCHAR2);
```

Parameters

Table 228-122 SETDATA Procedure Parameters

Parameter	Description
cd	DOMCHARACTERDATA
data	The data to which the node is set
pi	DOMPROCESSINGINSTRUCTION
data	New processing instruction content data

228.7.102 SETDOCTYPE Procedure

Given a DOM document, this procedure creates a new DTD with the specified name, system id and public id and sets it in the document.

This DTD can later be retrieved using the GETDOCTYPE Function.

Syntax

```
DBMS_XMLDOM.SETDOCTYPE(
  doc      IN   DOMDocument,
  name     IN   VARCHAR2,
  sysid    IN   VARCHAR2,
  pubid    IN   VARCHAR2);
```

Parameters

Table 228-123 SETDOCTYPE Procedure Parameters

Parameter	Description
doc	The document whose DTD has to be set
name	The name that the doctype needs to be initialized with
sysid	The system ID that the doctype needs to be initialized with
pubid	The public ID that the doctype needs to be initialized with

228.7.103 SETNAMEDITEM Function

This function adds a node using its NodeName attribute.

If a node with that name is already present in this map, it is replaced by the new one. The old node is returned on replacement; if no replacement is made, NULL is returned.

As the NodeName attribute is used to derive the name under which the node must be stored, multiple nodes of certain types, those that have a "special" string value, cannot be stored because the names would clash. This is seen as preferable to allowing nodes to be aliased.

> **See Also:**
>
> DOMNamedNodeMap Subprograms

Syntax

Adds a node using its NodeName attribute:

```
DBMS_XMLDOM.SETNAMEDITEM(
   nnm      IN      DOMNAMEDNODEMAP,
   arg      IN      DOMNODE)
 RETURN DOMNode;
```

Adds a node using its NodeName attribute and namespace URI:

```
DBMS_XMLDOM.SETNAMEDITEM(
   nnm      IN      DOMNAMEDNODEMAP,
   arg       IN     DOMNODE,
   ns       IN      VARCHAR2)
 RETURN DOMNode;
```

Parameters

Table 228-124 SETNAMEDITEM Function Parameters

Parameter	Description
nnm	DOMNAMEDNODEMAP
arg	The Node to be added using its NodeName attribute
ns	Namespace

228.7.104 SETNODEVALUE Procedure

This procedure sets the value of this node, depending on its type. When it is defined to be NULL, setting it has no effect.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.SETNODEVALUE(
   n         IN     DOMNODE,
   nodeValue IN     VARCHAR2);
```

Parameters

Table 228-125 SETNODEVALUE Procedure Parameters

Parameter	Description
n	DOMNode
nodeValue	The value to which node is set

228.7.105 SETNODEVALUEASBINARYSTREAM Function & Procedure

The operation of these subprograms is described in the syntax section.

> **See Also:**
>
> DOMNode Subprograms

Syntax

This function returns an instance of the PL/SQL XMLBINARYOUTPUTSTREAM into which the caller can write the node value. The datatype of the node must be RAW or BLOB – if not, an exception is raised.

```
DBMS_XMLDOM.SETNODEVALUEASBINARYSTREAM (
   n       IN      DOMNODE)
 RETURN SYS.UTL_BINARYOUTPUTSTREAM;
```

Using this procedure, the application passes in an implementation of `sys.utl_BinaryInputStream` from which XDB reads data to populate the node. The datatype of the node must be `RAW` or BLOB – if not an exception is raised.

```
DBMS_XMLDOM.SETNODEVALUEASBINARYSTREAM (
   n         in    DOMNODE,
   value     in    SYS.UTL_BINARYINPUTSTREAM);
```

Parameters

Table 228-126 SETNODEVALUEASBINARYSTREAM Function & Procedure Parameters

Parameter	Description
n	DOMNODE
value	BINARYINPUTSTREAM

228.7.106 SETNODEVALUEASCHARACTERSTREAM Function & Procedure

The operation of these subprograms is described in the syntax section.

> **See Also:**
>
> DOMNode Subprograms

Syntax

This function returns an instance of the PL/SQL `XMLCHARACTEROUTPUTSTREAM` type into which the caller can write the node value. The datatype of the node can be any valid XDB datatype. If the type is not character or `CLOB`, the character data written to the stream is converted to the node datatype. If the datatype of the node is character or `CLOB`, then the character data written to the stream is converted from PL/SQL session character set to the character set of the node.

```
DBMS_XMLDOM.SETNODEVALUEASCHARACTERSTREAM  (
   n         IN      DOMNODE)
 RETURN SYS.UTL_CHARACTEROUTPUTSTREAM;
```

Using this procedure, the application passes in an implementation of `SYS.UTL_CHARACTERINPUTSTREAM` from which XDB reads to populate the node. The datatype of the node may be any valid type supported by XDB. If a non-character datatype, the character data read from the stream is converted to the datatype of the node. If the datatype of the node is either character or `CLOB`, then no conversion occurs and the character set of the node becomes the character set of the PL/SQL session.

```
DBMS_XMLDOM.SETNODEVALUEASCHARACTERSTREAM  (
   n         IN    DOMNODE,
   value     IN    SYS.UTL_CHARACTERINPUTSTREAM);
```

Parameters

Table 228-127 SETNODEVALUEASCHARACTERSTREAM Function & Procedure Parameters

Parameter	Description
n	DOMNODE
value	CHARACTERINPUTSTREAM

228.7.107 SETPREFIX Procedure

This procedure sets the namespace prefix for this node to the specified value.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.SETPREFIX(
   n       IN     DOMNODE,
   prefix  IN     VARCHAR2);
```

Parameters

Table 228-128 SETPREFIX Procedure Parameters

Parameter	Description
n	DOMNODE
prefix	The value for the namespace prefix of the node

228.7.108 SETSTANDALONE Procedure

This procedure sets the standalone property of the DOMDOCUMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.SETSTANDALONE(
   doc         IN     DOMDOCUMENT,
   newvalue    IN     VARCHAR2);
```

Parameters

Table 228-129 SETSTANDALONE Procedure Parameters

Parameter	Description
doc	DOMDOCUMENT
newvalue	Value of the standalone property of the document

228.7.109 SETVALUE Procedure

This procedure sets the value of the attribute.

> **See Also:**
>
> DOMAttr Subprograms

Syntax

```
DBMS_XMLDOM.SETVALUE(
   a        IN     DOMATTR,
   value    IN     VARCHAR2);
```

Parameters

Table 228-130 SETVALUE Procedure Parameters

Parameter	Description
a	DOMATTR
value	The value to which to set the attribute

228.7.110 SETVERSION Procedure

This procedure sets the version of the DOMDOCUMENT.

> **See Also:**
>
> DOMDocument Subprograms

Syntax

```
DBMS_XMLDOM.SETVERSION(
   doc         IN     DOMDOCUMENT,
   version     IN     VARCHAR2);
```

Parameters

Table 228-131 SETVERSION Procedure Parameters

Parameter	Description
doc	DOMDOCUMENT
version	The version of the document

228.7.111 SPLITTEXT Function

This function breaks this DOMTEXT node into two DOMTEXT nodes at the specified offset.

> **See Also:**
>
> DBMS_XMLDOM DOMText Subprograms

Syntax

```
DBMS_XMLDOM.SPLITTEXT(
   t         IN     DOMTEXT,
   offset    IN     NUMBER)
 RETURN DOMText;
```

Parameters

Table 228-132 SPLITTEXT Function Parameters

Parameter	Description
t	DOMTEXT
offset	Offset at which to split

228.7.112 SUBSTRINGDATA Function

This function extracts a range of data from the node.

> **See Also:**
>
> DOMCharacterData Subprograms

Syntax

```
DBMS_XMLDOM.SUBSTRINGDATA(
   cd        IN     DOMCHARACTERDATA,
   offset    IN     NUMBER,
```

```
   cnt       IN     NUMBER)
 RETURN VARCHAR2;
```

Parameters

Table 228-133 SUBSTRINGDATA Function Parameters

Parameter	Description
cd	DOMCHARACTERDATA
offset	The starting offset of the data from which to get the data
cnt	The number of characters (from the offset) of the data to get

228.7.113 USEBINARYSTREAM Function

This function returns TRUE if the datatype of the node is RAW or BLOB, so that the node value may be read or written using an UTL_BINARYINPUTSTREAM or UTL_BINARYOUTPUTSTREAM.

If a value of FALSE is returned, the node value may only be accessed through an UTL_CHARACTERINPUTSTREAM or UTL_CHARACTEROUTPUTSTREAM.

> **See Also:**
>
> DOMNode Subprograms

Syntax

```
DBMS_XMLDOM.USEBINARYSTREAM   (
   n         IN     DOMNODE)
 RETURN BOOLEAN;
```

Parameters

Table 228-134 USEBINARYSTREAM Function Parameters

Parameter	Description
n	DOMNODE

228.7.114 WRITETOBUFFER Procedures

WRITETOBUFFER is an overloaded procedure that writes an XML node, XML document, or a document fragment to a specified buffer.

This procedure is overloaded. The specific forms of functionality are described along with the syntax declarations.

Syntax

Writes XML node to specified buffer using the database character set (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.WRITETOBUFFER(
   n         IN       DOMNODE,
   buffer    IN OUT   VARCHAR2);
```

Writes XML document to a specified buffer using database character set (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.WRITETOBUFFER(
   doc       IN       DOMDOCUMENT,
   buffer    IN OUT   VARCHAR2);
```

Writes the contents of the specified document fragment into a buffer using the database character set (See Also: DOMDocumentFragment Subprograms):

```
DBMS_XMLDOM.WRITETOBUFFER(
   df        IN       DOMDOCUMENTFRAGMENT,
   buffer    IN OUT   VARCHAR2);
```

Parameters

Table 228-135 WRITETOBUFFER Procedure Parameters

Parameter	Description
n	DOMNODE
buffer	Buffer to which to write
doc	DOMDOCUMENT
df	DOM document fragment

228.7.115 WRITETOCLOB Procedures

`WRITETOCLOB` is an overloaded procedure that writes an XML node or document to a specified `CLOB`.

The specific forms of functionality are described along with the syntax declarations.

Syntax

Writes XML node to specified `CLOB` using the database character set (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.WRITETOCLOB(
   n         IN       DOMNODE,
   cl        IN OUT   CLOB);
```

Writes XML document to a specified `CLOB` using database character set (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.WRITETOCLOB(
   doc       IN       DOMDOCUMENT,
   cl        IN OUT   CLOB);
```

Parameters

Table 228-136 WRITETOCLOB Procedure Parameters

Parameter	Description
n	DOMNODE
cl	CLOB to which to write
doc	DOMDOCUMENT

228.7.116 WRITETOFILE Procedures

This overloaded procedure writes an XML node or XML document to a specified node.

The specific forms of functionality are described along with the syntax declarations.

Syntax

Writes XML node to specified file using the database character set (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.WRITETOFILE(
   n          IN      DOMNODE,
   fileName   IN      VARCHAR2);
```

Writes XML node to specified file using the specified character set, which is passed in as a separate parameter (See Also: DOMNode Subprograms):

```
DBMS_XMLDOM.WRITETOFILE(
   n          IN      DOMNODE,
   fileName   IN      VARCHAR2,
   charset    IN      VARCHAR2);
```

Writes an XML document to a specified file using database character set (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.WRITETOFILE(
   doc        IN   DOMDOCUMENT,
   filename   IN   VARCHAR2);
```

Writes an XML document to a specified file using specified character set (See Also: DOMDocument Subprograms):

```
DBMS_XMLDOM.WRITETOFILE(
   doc       IN   DOMDOCUMENT,
   fileName  IN   VARCHAR2,
   charset   IN   VARCHAR2);
```

Parameters

Table 228-137 WRITETOFILE Procedure Parameters

Parameter	Description
n	DOMNODE

Table 228-137 (Cont.) WRITETOFILE Procedure Parameters

Parameter	Description
`fileName`	File to which to write. The filename should be in the format of `database_directory_object_name/filename`, for example `mydir/filename` (on windows, use `\` instead of `/`).
`charset`	specified character set
`doc`	`DOMDOCUMENT`
`charset`	Character set

229
DBMS_XMLGEN

The DBMS_XMLGEN package converts the results of a SQL query to a canonical XML format.

> **Note:**
>
> The PL/SQL package DBMS_XMLGEN is deprecated in Oracle Database 23ai.
>
> DBMS_XMLGEN is a non-standard Oracle-proprietary package that is provided to generate and convert XML documents from SQL queries or with PL/SQL. This package is deprecated, and can be desupported in a future release. Oracle recommends that you use SQL/XML operators to generate XML from relational columns instead. Using ANSI SQL/XML operators for any generation and modification of XML documents provides a standardized and future-proof way to work with XML documents.

The package takes an arbitrary SQL query as input, converts it to XML format, and returns the result as a CLOB. This package is similar to the DBMS_XMLQUERY package, except that it is written in C and compiled into the kernel. This package can only be run on the database.

This chapter contains the following topic:

- Security Model
- Summary of DBMS_XMLGEN Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide,* or more information on XML support and on examples of using DBMS_XMLGEN

229.1 DBMS_XMLGEN Security Model

Owned by XDB, the DBMS_XMLGEN package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

229.2 Summary of DBMS_XMLGEN Subprograms

This table lists the DBMS_XMLGEN subprograms and briefly describes them.

Table 229-1 Summary of DBMS_XMLGEN Package Subprograms

Subprogram	Description
CLOSECONTEXT Procedure	Closes the context and releases all resources
CONVERT Functions	Converts the XML into the escaped or unescaped XML equivalent
GETNUMROWSPROCESSED Function	Gets the number of SQL rows that were processed in the last call to GETXML Functions
GETXML Functions	Gets the XML document
GETXMLTYPE Functions	Gets the XML document and returns it as `XMLType`
NEWCONTEXT Functions	Creates a new context handle
NEWCONTEXTFROMHIERARCHY Function	Obtains a handle to use in the GETXML Functions and other functions to get a hierarchical XML with recursive elements from the result
RESTARTQUERY Procedure	Restarts the query to start fetching from the beginning
SETCONVERTSPECIALCHARS Procedure	Sets whether special characters such as `$,` which are non-XML characters, should be converted or not to their escaped representation
SETMAXROWS Procedure	Sets the maximum number of rows to be fetched each time
SETNULLHANDLING Procedure	Sets `NULL` handling options
SETROWSETTAG Procedure	Sets the name of the element enclosing the entire result
SETROWTAG Procedure	Sets the name of the element enclosing each row of the result
SETSKIPROWS Procedure	Sets the number of rows to skip every time before generating the XML.
USEITEMTAGSFORCOLL Procedure	Forces the use of the collection column name appended with the tag `_ITEM` for collection elements
USENULLATTRIBUTEINDICATOR Procedure	Specifies whether to use an XML attribute to indicate `NULL`ness, or to do it by omitting the inclusion of the particular entity in the XML document.

229.2.1 CLOSECONTEXT Procedure

This procedure closes a given context and releases all resources associated with it, including the SQL cursor and bind and define buffers. After this call, the handle cannot be used for a subsequent function call.

Syntax

```
DBMS_XMLGEN.CLOSECONTEXT (
   ctx  IN ctxHandle);
```

Parameters

Table 229-2 CLOSECONTEXT Procedure Parameters

Parameter	Description
`ctx`	The context handle to close.

229.2.2 CONVERT Functions

This function converts the XML data into the escaped or unescapes XML equivalent, and returns XML `CLOB` data in encoded or decoded format. There are several version of the function.

Syntax

Uses `XMLDATA` in string form (`VARCHAR2`):

```
DBMS_XMLGEN.CONVERT (
   xmlData IN VARCHAR2,
   flag    IN NUMBER := ENTITY_ENCODE)
RETURN VARCHAR2;
```

Uses `XMLDATA` in `CLOB` form:

```
DBMS_XMLGEN.CONVERT (
   xmlData IN CLOB,
   flag    IN NUMBER := ENTITY_ENCODE)
 RETURN CLOB;
```

Parameters

Table 229-3 CONVERT Function Parameters

Parameter	Description
`xmlData`	The XML `CLOB` data to be encoded or decoded.
`flag`	The flag setting; `ENTITY_ENCODE` (default) for encode, and `ENTITY_DECODE` for decode.

Usage Notes

This function escapes the XML data if the `ENTITY_ENCODE` is specified. For example, the escaped form of the character < is `<`. Unescaping is the reverse transformation.

229.2.3 GETNUMROWSPROCESSED Function

This function retrieves the number of SQL rows processed when generating the XML using the GETXML Functions call. This count does not include the number of rows skipped before generating the XML.

Note that GETXML Functions always generates an XML document, even if there are no rows present.

Syntax

```
DBMS_XMLGEN.GETNUMROWSPROCESSED (
   ctx     IN    ctxHandle)
RETURN NUMBER;
```

Parameters

Table 229-4 GETNUMROWSPROCESSED Function Parameters

Parameter	Description
ctx	The context handle obtained from the NEWCONTEXT Functions call.

Usage Notes

This function is used to determine the terminating condition if calling GETXML Functions in a loop.

Related Topics

- GETXML Functions
  This function gets the XML document. The function is overloaded.

229.2.4 GETXML Functions

This function gets the XML document. The function is overloaded.

Syntax

Gets the XML document by fetching the maximum number of rows specified. It appends the XML document to the CLOB passed in. Use this version of GETXML Functions to avoid any extra CLOB copies and to reuse the same CLOB for subsequent calls. Because of the CLOB reuse, this GETXML Functionscall is potentially more efficient:

```
DBMS_XMLGEN.GETXML (
   ctx          IN ctxHandle,
   tmpclob      IN OUT NCOPY CLOB,
   dtdOrSchema  IN number := NONE)
 RETURN BOOLEAN;
```

Generates the XML document and returns it as a temporary CLOB. The temporary CLOB obtained from this function must be freed using the DBMS_LOB.FREETEMPORARY call:

```
DBMS_XMLGEN.GETXML (
   ctx          IN ctxHandle,
   dtdOrSchema  IN number := NONE)
 RETURN CLOB;
```

Converts the results from the SQL query string to XML format, and returns the XML as a temporary CLOB, which must be subsequently freed using the DBMS_LOB.FREETEMPORARY call:

```
DBMS_XMLGEN.GETXML (
   sqlQuery     IN VARCHAR2,
   dtdOrSchema  IN number := NONE)
 RETURN CLOB;
```

Parameters

Table 229-5 GETXML Function Parameters

Parameter	Description
ctx	The context handle obtained from the newContext call.
tmpclob	The CLOB to which the XML document is appended.
sqlQuery	The SQL query string.
dtdOrSchema	Generate a DTD or a schema? Only NONE is supported.

Usage Notes

When the rows indicated by the SETSKIPROWS Procedure call are skipped, the maximum number of rows as specified by the SETMAXROWS Procedure call (or the entire result if not specified) is fetched and converted to XML. Use the GETNUMROWSPROCESSED Function to check if any rows were retrieved.

229.2.5 GETXMLTYPE Functions

This function gets the XML document and returns it as an XMLTYPE. XMLTYPE operations can be performed on the results. This function is overloaded.

Syntax

Generates the XML document and returns it as a sys.XMLType:

```
DBMS_XMLGEN.GETXMLTYPE (
   ctx           IN ctxhandle,
   dtdOrSchema   IN number := NONE)
 RETURN sys.XMLType;
```

Converts the results from the SQL query string to XML format, and returns the XML as a sys.XMLType:

```
DBMS_XMLGEN.GETXMLTYPE (
   sqlQuery     IN VARCHAR2,
   dtdOrSchema  IN number := NONE)
 RETURN sys.XMLType
```

Parameters

Table 229-6 GETXMLTYPE Function Parameters

Parameter	Description
ctx	The context handle obtained from the newContext call.
sqlQuery	The SQL query string.
dtdOrSchema	Generate a DTD or a schema? Only NONE is supported.

229.2.6 NEWCONTEXT Functions

This function generates and returns a new context handle.

This context handle is used in GETXML Functions and other functions to get XML back from the result. There are several version of the function.

Syntax

Generates a new context handle from a query:

```
DBMS_XMLGEN.NEWCONTEXT (
      query     IN VARCHAR2)
 RETURN ctxHandle;
```

Generates a new context handle from a query string in the form of a PL/SQL ref cursor:

```
DBMS_XMLGEN.NEWCONTEXT (
   queryString  IN SYS_REFCURSOR)
 RETURN ctxHandle;
```

Parameters

Table 229-7 NEWCONTEXT Function Parameters

Parameter	Description
`query`	The query, in the form of a `VARCHAR`, the result of which must be converted to XML.
`queryString`	The query string in the form of a PL/SQL ref cursor, the result of which must be converted to XML.

229.2.7 NEWCONTEXTFROMHIERARCHY Function

This function obtains a handle to use in the GETXML Functions and other functions to get a hierarchical XML with recursive elements from the result.

Syntax

```
DBMS_XMLGEN.NEWCONTEXTFROMHIERARCHY (
   queryString IN VARCHAR2)
 RETURN ctxHandle;
```

Parameters

Table 229-8 NEWCONTEXTFROMHIERARCHY Function Parameters

Parameter	Description
`queryString`	The query string, the result of which must be converted to XML. The query is a hierarchical query typically formed using a `CONNECT BY` clause, and the result must have the same property as the result set generated by a `CONNECT BY` query. The result set must have only two columns, the level number and an XML value. The level number is used to determine the hierarchical position of the XML value within the result XML document.

Related Topics

- GETXML Functions
  This function gets the XML document. The function is overloaded.

229.2.8 RESTARTQUERY Procedure

This procedure restarts the query and generates the XML from the first row.

It can be used to start executing the query again, without having to create a new context.

Syntax

```
DBMS_XMLGEN.RESTARTQUERY (
ctx  IN ctxHandle);
```

Parameters

Table 229-9 RESTARTQUERY Procedure Parameters

Parameter	Description
ctx	The context handle corresponding to the current query.

229.2.9 SETCONVERTSPECIALCHARS Procedure

This procedure sets whether or not special characters in the XML data must be converted into their escaped XML equivalent. For example, the < sign is converted to <.

The default is to perform conversions.

This function improves performance of XML processing when the input data cannot contain any special characters such as <, >, ", ', which must be escaped. It is expensive to scan the character data to replace the special characters, particularly if it involves a lot of data.

Syntax

```
DBMS_XMLGEN.SETCONVERTSPECIALCHARS (
ctx   IN ctxHandle,
conv  IN BOOLEAN);
```

Parameters

Table 229-10 SETCONVERTSPECIALCHARS Procedure Parameters

Parameter	Description
ctx	The context handle obtained from one of the NEWCONTEXT Functions call.
conv	TRUE indicates that conversion is needed.

229.2.10 SETMAXROWS Procedure

This procedure sets the maximum number of rows to fetch from the SQL query result for every invocation of the GETXML Functions call.

It is used when generating paginated results. For example, when generating a page of XML or HTML data, restrict the number of rows converted to XML or HTML by setting the `maxrows` parameter.

Syntax

```
DBMS_XMLGEN.SETMAXROWS (
ctx      IN ctxHandle,
maxRows  IN NUMBER);
```

Parameters

Table 229-11 SETMAXROWS Procedure Parameters

Parameter	Description
`ctx`	The context handle corresponding to the query executed.
`maxRows`	The maximum number of rows to get for each call to GETXML Functions

Related Topics

- GETXML Functions
  This function gets the XML document. The function is overloaded.

229.2.11 SETNULLHANDLING Procedure

This procedure sets `NULL` handling options, handled through the `flag` parameter setting.

Syntax

```
DBMS_XMLGEN.SETNULLHANDLING(
ctx  IN ctx,
flag IN NUMBER);
```

Parameters

Table 229-12 SETNULLHANDLING Procedure Parameters

Parameter	Description
`ctx`	The context handle corresponding to the query executed.
`flag`	The `NULL` handling option set. • `DROP_NULLS CONSTANT NUMBER:= 0;` (Default) Leaves out the tag for `NULL` elements. • `NULL_ATTR CONSTANT NUMBER:= 1;` Sets `xsi:nil="true"`. • `EMPTY_TAG CONSTANT NUMBER:= 2;` Sets, for example, `<foo/>`.

229.2.12 SETROWSETTAG Procedure

This procedure sets the name of the root element of the document. The default name is `ROWSET`.

Syntax

```
DBMS_XMLGEN.SETROWSETTAG (
ctx            IN ctxHandle,
rowSetTagName  IN VARCHAR2);
```

Parameters

Table 229-13 SETROWSETTAG Procedure Parameters

Parameter	Description
ctx	The context handle obtained from the NEWCONTEXT Functions call.
rowSetTagName	The name of the document element. Passing `NULL` indicates that you do not want the `ROWSET` element present.

Usage Notes

The user can set the `rowSetTag` to `NULL` to suppress the printing of this element. However, an error is produced if both the row and the rowset are `NULL` and there is more than one column or row in the output. This is because the generated XML would not have a top-level enclosing tag, and so would be invalid.

229.2.13 SETROWTAG Procedure

This procedure sets the name of the element separating all the rows. The default name is `ROW`.

Syntax

```
DBMS_XMLGEN.SETROWTAG (
ctx         IN ctxHandle,
rowTagName  IN VARCHAR2);
```

Parameters

Table 229-14 SETROWTAG Procedure Parameters

Parameter	Description
`ctx`	The context handle obtained from the NEWCONTEXT Functions call.
`rowTagName`	The name of the `ROW` element. Passing `NULL` indicates that you do not want the `ROW` element present.

Usage Notes

The user can set the name of the element to `NULL` to suppress the `ROW` element itself. However, an error is produced if both the row and the rowset are `NULL` and there is more than one column or row in the output. This is because the generated XML would not have a top-level enclosing tag, and so would be invalid.

229.2.14 SETSKIPROWS Procedure

This procedure skips a given number of rows before generating the XML output for every call to the GETXML Functions. It is used when generating paginated results for stateless Web pages using this utility.

For example, when generating the first page of XML or HTML data, set `skiprows` to zero. For the next set, set the `skiprows` to the number of rows obtained in the first case. See GETNUMROWSPROCESSED Function.

Syntax

```
DBMS_XMLGEN.SETSKIPROWS (
ctx       IN ctxHandle,
skipRows  IN NUMBER);
```

Parameters

Table 229-15 SETSKIPROWS Procedure Parameters

Parameter	Description
`ctx`	The context handle corresponding to the query executed.
`skipRows`	The number of rows to skip for each call to `getXML`.

Related Topics

- GETXML Functions
  This function gets the XML document. The function is overloaded.

229.2.15 USEITEMTAGSFORCOLL Procedure

This procedure overrides the default name of the collection elements. The default name for collection elements is the type name itself.

Syntax

```
DBMS_XMLGEN.USEITEMTAGSFORCOLL (
   ctx  IN ctxHandle);
```

Parameters

Table 229-16 USEITEMTAGSFORCOLL Procedure Parameters

Parameter	Description
`ctx`	The context handle.

Usage Notes

Using this procedure, you can override the default to use the name of the column with the `_ITEM` tag appended to it. If there is a collection of `NUMBER`, the default tag name for the collection elements is `NUMBER`.

229.2.16 USENULLATTRIBUTEINDICATOR Procedure

This procedure specifies whether to use an XML attribute to indicate `NULL`, or to do it by omitting the inclusion of the particular entity in the XML document.

It is used as a shortcut for the SETNULLHANDLING Procedure.

Syntax

```
DBMS_XMLGEN.USENULLATTRIBUTEINDICATOR(
ctx       IN   ctxType,
attrind   IN   BOOLEAN := TRUE);
```

Parameters

Table 229-17 USENULLATTRIBUTEINDICATOR Procedure Parameters

Parameter	Description
`ctx`	Context handle.
`attrind`	Use attribute to indicate `NULL`?

230
DBMS_XMLINDEX

The DBMS_XMLINDEX package provides an interface to implement asynchronous indexing.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_XMLINDEX Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide* for more information about "XMLIndex"

230.1 DBMS_XMLINDEX Overview

DBMS_XMLINDEX provides a mechanism for asynchronous index maintenance.

Asynchronous Index Maintenance

The basic XMLIndex is maintained on every DML operation. However, given the computing costs, in many cases the availability of stale result is adequate. In such situations, it is desirable to defer index updates to a more convenient time, for example when the load on the database is low. DBMS_XMLINDEX provides this mechanism.

230.2 DBMS_XMLINDEX Security Model

Owned by XDB, the DBMS_XMLINDEX package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC.

Subprograms in this package are executed using the privileges of the current user.

230.3 Summary of DBMS_XMLINDEX Subprograms

This table lists the DBMS_XMLINDEX subprograms and briefly describes them.

Table 230-1 DBMS_XMLINDEX Package Subprograms

Subprogram	Description
CREATEDATEINDEX Procedure	Creates a secondary index for date values in the VALUE column of a PATH TABLE which is the storage table of an XMLIndex
CREATENUMBERINDEX Procedure	Creates a secondary index for number values in the VALUE column of a PATH TABLE which is the storage table of an XMLIndex

Table 230-1 (Cont.) DBMS_XMLINDEX Package Subprograms

Subprogram	Description
DROPPARAMETER Procedure	Drops the XMLIndex parameter string that is associated with a given parameter identifier.
MODIFYPARAMETER Procedure	Modifies the XMLIndex parameter string that is associated with a given parameter name
PROCESS_PENDING Procedure	Processes pending rows for a NONBLOCKING ALTER INDEX OPERATION on an XMLIndex
REGISTERPARAMETER Procedure	Registers a parameter string and XMLIndex parameter string pair in XDB
SYNCINDEX Procedure	Synchronizes the index manually

230.3.1 CREATEDATEINDEX Procedure

This procedure creates a secondary index for date values in the VALUE column of a PATH TABLE which is the storage table of an XMLIndex. The second form of the procedure allows for the date_index_clause to be set to an empty string.

Syntax

```
DBMS_XMLINDEX.CREATEDATEINDEX  (
   xml_index_schema   IN   VARCHAR2,
   xml_index_name     IN   VARCHAR2,
   date_index_name    IN   VARCHAR2,
   xmltypename        IN   VARCHAR2,
   date_index_clause  IN   VARCHAR2);

DBMS_XMLINDEX.CREATEDATEINDEX  (
   xml_index_schema   IN   VARCHAR2
   xml_index_name     IN   VARCHAR2,
   date_index_name    IN   VARCHAR2,
   xmltypename        IN   VARCHAR2);
```

Parameters

Table 230-2 CREATEDATEINDEX Procedure Parameters

Parameter	Description
xml_index_schema	Name of the owner of the XMLIndex
xml_index_name	Name of the XMLIndex
date_index_name	Name of the secondary index to be created for date values in the VALUE column of the PATH TABLE of XMLIndex named xml_index_name and owned by xml_index_schema
xmltypename	The type to which values in the VALUE column of the path table are to be cast. Acceptable values are the following strings: DATETIME, TIME, DATE, GDAY, GMONTH, GYEAR, GYEARMONTH, GMONTHDAY.
date_index_clause	Storage clause to be applied to the date index during its creation. This is a string argument appended to the CREATE INDEX statement for creating the date index

230.3.2 CREATENUMBERINDEX Procedure

This procedure creates a secondary index for number values in the `VALUE` column of a `PATH TABLE` which is the storage table of an `XMLIndex`.

Syntax

```
DBMS_XMLINDEX.CREATENUMBERINDEX (
   xml_index_schema    IN   VARCHAR2,
   xml_index_name      IN   VARCHAR2,
   num_index_name      IN   VARCHAR2,
   num_index_clause    IN   VARCHAR2,
   xmltypename         IN   VARCHAR2);
```

Parameters

Table 230-3 CREATENUMBERINDEX Procedure Parameters

Parameter	Description
`xml_index_schema`	Name of the owner of the `XMLIndex`
`xml_index_name`	Name of the `XMLIndex`
`num_index_name`	Name of the secondary index to be created for number values in the `VALUE` column of the `PATH TABLE` of `XMLIndex` named `xml_index_name` and owned by `xml_index_schema`
`num_index_clause`	Storage clause to be applied to the number index during its creation. This is a string argument appended to the `CREATE INDEX` statement for creating the number index.
`xmltypename`	The type to which values in the `VALUE` column of the path table are to be cast. Acceptable values are the following strings: `FLOAT`, `DOUBLE`, `DECIMAL`, `INTEGER`, `NONPOSITIVEINTEGER`, `NEGATIVEINTEGER`, `LONG`, `INT`, `SHORT`, `BYTE`, `NONNEGATIVEINTEGER`, `UNSIGNEDLONG`, `UNSIGNEDINT`, `UNSIGNEDSHORT`, `UNSIGNEDBYTE`, `POSITIVEINTEGER`.

230.3.3 DROPPARAMETER Procedure

This procedure drops the `XMLIndex` parameter string that is associated with a given parameter identifier.

Syntax

```
DBMS_XMLINDEX.DROPPARAMETER (
   name        IN      VARCHAR2);
```

Parameters

Table 230-4 DROPPARAMETER Procedure Parameters

Parameter	Description
`name`	Identifier for parameter string

Examples

```
DBMS_XMLINDEX.DROPPARAMETER (
   'myIndexParam');
```

230.3.4 MODIFYPARAMETER Procedure

This procedure modifies the `XMLIndex` parameter string that is associated with a given parameter identifier.

Syntax

```
DBMS_XMLINDEX.MODIFYPARAMETER (
   name        IN      VARCHAR2,
   parameter   IN      CLOB);
```

Parameters

Table 230-5 MODIFYPARAMETER Procedure Parameters

Parameter	Description
`name`	Identifier for parameter string
`parameter`	`XMLIndex` parameter clause that can appear in a `CREATE INDEX` or an `ALTER INDEX` statement

Examples

```
DBMS_XMLINDEX.MODIFYPARAMETER (
   'myIndexParam',
   'PATH TABLE po_ptab
    PATH ID INDEX po_pidx
    ORDER KEY INDEX po_oidx
    VALUE INDEX po_vidx');
```

230.3.5 PROCESS_PENDING Procedure

This procedure processes executes DMLs required to complete a `NONBLOCKING ALTER INDEX ADD_GROUP/ADD_COLUMN` operation on an `XMLIndex`.

Syntax

```
DBMS_XMLINDEX.PROCESS_PENDING (
   xml_index_schema    IN     VARCHAR2,
   xml_index_name      IN     VARCHAR2,
   pending_row_count   OUT    BINARY_INTEGER,
   error_row_count     OUT    BINARY_INTEGER);
```

Parameters

Table 230-6 PROCESS_PENDING Procedure Parameters

Parameter	Description
`xml_index_schema`	Name of the owner of the `XMLIndex`

Table 230-6 (Cont.) PROCESS_PENDING Procedure Parameters

Parameter	Description
xml_index_name	Name of the XMLIndex to be altered using NONBLOCKING ALTER INDEX OPERATION
pending_row_count	Number of pending rows to be processed
error_row_count	Number of rows for which indexing may have failed because of an error

Usage Notes

- This procedure will iteratively attempt to index all necessary rows in small batches while skipping rows that are locked and rows for which index maintenance fails with an error. Therefore, it may have to be executed multiple times for an XMLIndex until all pending rows are processed. Once all pending rows are processed, user can complete the NONBLOCKING ALTER INDEX OPERATION.
- If it is not possible process all the pending rows after multiple trials, the user will have to manually triage the locking or error issues by examining unprocessed rows in SYS_AIXSXI_######_PENDINGTAB and errors in SYS_AIXSXI_#####_ERRORTAB.. Keeping track of rows and the errors is useful in triaging issues.

Examples

```
EXEC DBMS_XMLINDEX.PROCESS_PENDING(
   'SCOTT', 'PO_XMLINDEX_IX', out_param1, out_param2);
```

230.3.6 REGISTERPARAMETER Procedure

This procedure registers a parameter identifier and XMLIndex parameter string pair in XDB.

Syntax

```
DBMS_XMLINDEX.REGISTERPARAMETER (
   name        IN      VARCHAR2,
   parameter   IN      CLOB);
```

Parameters

Table 230-7 REGISTERPARAMETER Procedure Parameters

Parameter	Description
name	Identifier for parameter string
parameter	XMLIndex parameter clause that can appear in a CREATE INDEX or an ALTER INDEX statement

Examples

```
DBMS_XMLINDEX.REGISTERPARAMETER (
   'myIndexParam',
   'PATH TABLE po_ptab
    PATH ID INDEX po_pidx
    ORDER KEY INDEX po_oidx
    VALUE INDEX po_vidx
```

```
PATHS(NAMESPACE MAPPING(xmlns:p="http://www.example.com/IPO"))
GROUP MASTERGROUP XMLTABLE PO_TAB
('/p:PurchaseOrder'
    COLUMNS
       REFERENCE VARCHAR2(30) PATH 'p:Reference',
       REQUESTOR VARCHAR2(30) PATH 'p:Requestor' )
GROUP ITEMGROUP XMLTABLE ITEMGROUP_TAB
('/p:PurchaseOrder/p:LineItems/p:LineItem'
    COLUMNS
       LINENUMBER NUMBER(38) PATH '@p:ItemNumber',
       QUANTITY NUMBER(38) PATH '@p:Quantity',
       DESCRIPTION VARCHAR2(256) PATH 'p:Description'));
```

230.3.7 SYNCINDEX Procedure

This function synchronizes an asynchronously maintained XMLIndex.

It applies to the XMLIndex changes that are logged in the pending table, and brings the path table up-to-date with the base `XMLTYPE` column.

Syntax

```
DBMS_XMLINDEX.SYNCINDEX (
   xml_index_schema      IN VARCHAR2,
   xml_index_name        IN VARCHAR2,
   partition_name        IN VARCHAR2 DEFAULT NULL,
   reindex               IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 230-8 SYNCINDEX Procedure Parameters

Parameter	Description
`xml_index_schema`	Name of the owner of the XMLIndex
`xml_schema_name`	Name of the XMLIndex
`partition_name`	[Currently not supported]
`reindex`	Default is `FALSE`. If set to `TRUE`, this drops the secondary indexes and recreates them later so that they can be bulk-loaded.

Examples

```
EXEC DBMS_XMLINDEX.SYNCINDEX('USER1', 'SS_TAB_XMLI', REINDEX=>TRUE);
```

231
DBMS_XMLPARSER

Using `DBMS_XMLPARSER`, you can access the contents and structure of XML documents. XML describes a class of data XML document objects. It partially describes the behavior of computer programs which process them. By construction, XML documents are conforming SGML documents.

XML documents are made up of storage units called entities, which contain either parsed or unparsed data. Parsed data is made up of characters, some of which form character data, and some of which form markup. Markup encodes a description of the document's storage layout and logical structure. XML provides a mechanism to impose constraints on the storage layout and logical structure.

A software module called an XML processor is used to read XML documents and provide access to their content and structure. It is assumed that an XML processor is doing its work on behalf of another module, called the application. This PL/SQL implementation of the XML processor (or parser) follows the W3C XML specification REC-xml-19980210 and includes the required behavior of an XML processor in terms of how it must read XML data and the information it must provide to the application.

The default behavior for this PL/SQL XML parser is to build a parse tree that can be accessed by DOM APIs, validate it if a DTD is found (otherwise, it is non-validating), and record errors if an error log is specified. If parsing fails, an application error is raised.

This chapter contains the following topics:

- Security Model
- Summary of DBMS_XMLPARSER Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

231.1 DBMS_XMLPARSER Security Model

Owned by `XDB`, the `DBMS_XMLPARSER` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`.

Subprograms in this package are executed using the privileges of the current user.

231.2 Summary of DBMS_XMLPARSER Subprograms

This table lists the `DBMS_XMLPARSER` subprograms and briefly describes them.

Table 231-1 DBMS_XMLPARSER Package Subprograms

Method	Description
FREEPARSER	Frees a parser object.
GETDOCTYPE	Gets parsed DTD.
GETDOCUMENT	Gets DOM document.
GETRELEASEVERSION	Returns the release version of Oracle XML Parser for PL/SQL.
GETVALIDATIONMODE	Returns validation mode.
NEWPARSER	Returns a new parser instance
PARSE	Parses XML stored in the given url/file.
PARSEBUFFER	Parses XML stored in the given buffer
PARSECLOB	Parses XML stored in the given clob
PARSEDTD	Parses DTD stored in the given url/file
PARSEDTDBUFFER	Parses DTD stored in the given buffer
PARSEDTDCLOB	Parses DTD stored in the given clob
SETBASEDIR	Sets base directory used to resolve relative URLs.
SETDOCTYPE	Sets DTD.
SETERRORLOG	Sets errors to be sent to the specified file
SETPRESERVEWHITESPACE	Sets white space preserve mode
SETVALIDATIONMODE	Sets validation mode.
SHOWWARNINGS	Turns warnings on or off.

231.2.1 FREEPARSER

This procedures frees a parser object.

Syntax

```
PROCEDURE freeParser(
    p Parser);
```

Parameters

Table 231-2 FREEPARSER Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.

231.2.2 GETDOCTYPE

The `GETDOCTYPE` function returns the parsed DTD. This function must be called only after a DTD is parsed.

Syntax

```
FUNCTION getDoctype(
  p Parser)
RETURN DOMDocumentType;
```

Parameters

Table 231-3 GETDOCTYPE Function Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.

231.2.3 GETDOCUMENT

`GETDOCUMENT` returns the document node of a DOM tree document built by the parser. This function must be called only after a document is parsed.

Syntax

```
FUNCTION GETDOCUMENT(
  p Parser)
RETURN DOMDocument;
```

Parameters

Table 231-4 GETDOCUMENT Function Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.

231.2.4 GETRELEASEVERSION

`GETRELEASEVERSION` returns the release version of the Oracle XML parser for PL/SQL.

Syntax

```
FUNCTION getReleaseVersion
RETURN VARCHAR2;
```

231.2.5 GETVALIDATIONMODE

The `GETVALIDATIONMODE` function retrieves the validation mode: `TRUE` for validating, `FALSE` otherwise.

Syntax

```
FUNCTION GETVALIDATIONMODE(
  p Parser)
RETURN BOOLEAN;
```

Parameters

Table 231-5 GETVALIDATIONMODE Function Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.

231.2.6 NEWPARSER

This function returns a new parser instance.

This function must be called before the default behavior of Parser can be changed and if other parse methods need to be used.

Syntax

```
FUNCTION newParser
RETURN Parser;
```

231.2.7 PARSE

`PARSE` parses XML stored in the given URL or file. An application error is raised if parsing fails.

There are several versions of this method.

Syntax

Function. Use this when the default parser behavior is acceptable, and only a URL or file needs to be parsed. Returns the built DOM document.

```
FUNCTION parse(url VARCHAR2)
RETURN DOMDocument;
```

Procedure. Any changes to the default parser behavior should be effected before calling this procedure.

```
PROCEDURE parse(
  p Parser,
  url VARCHAR2);
```

Parameters

Table 231-6 PARSE Subprogram Parameters

Parameter	IN / OUT	Description
url	(IN)	Complete path of the url/file to be parsed.
p	(IN)	Parser instance.

231.2.8 PARSEBUFFER

PARSEBUFFER parses XML stored in the given buffer.

Any changes to the default parser behavior should be effected before calling this procedure. An application error is raised if parsing fails.

Syntax

```
PROCEDURE PARSEBUFFER(
  p   Parser,
doc VARCHAR2);
```

Parameters

Table 231-7 PARSEBUFFER Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
doc	(IN)	XML document buffer to parse.

231.2.9 PARSECLOB

PARSECLOB parses XML stored in the given clob.

Any changes to the default parser behavior should be effected before calling this procedure. An application error is raised if parsing fails.

Syntax

```
PROCEDURE PARSECLOB(
  p   Parser,
doc CLOB);
```

Parameters

Table 231-8 PARSECLOB Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.

Table 231-8 (Cont.) PARSECLOB Procedure Parameters

Parameter	IN / OUT	Description
doc	(IN)	XML document buffer to parse.

231.2.10 PARSEDTD

PARSEDTD parses the DTD stored in the given URL or file.

Any changes to the default parser behavior should be effected before calling this procedure. An application error is raised if parsing fails.

Syntax

```
PROCEDURE PARSEDTD(
  p     Parser,
  url   VARCHAR2,
  root  VARCHAR2);
```

Parameters

Table 231-9 PARSEDTD Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
url	(IN)	Complete path of the URL or file to be parsed.
root	(IN)	Name of the root element.

231.2.11 PARSEDTDBUFFER

PARSEDTDBUFFER parses the DTD stored in the given buffer.

Any changes to the default parser behavior should be effected before calling this procedure. An application error is raised if parsing fails.

Syntax

```
PROCEDURE PARSEDTDBUFFER(
  p    Parser,
  dtd  VARCHAR2,
  root VARCHAR2);
```

Parameters

Table 231-10 PARSEDTDBUFFER Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.

Table 231-10 (Cont.) PARSEDTDBUFFER Procedure Parameters

Parameter	IN / OUT	Description
dtd	(IN)	DTD buffer to parse.
root	(IN)	Name of the root element.

231.2.12 PARSEDTDCLOB

PARSEDTDCLOB parses the DTD stored in the given clob.

Any changes to the default parser behavior should be effected before calling this procedure. An application error is raised if parsing fails.

Syntax

```
PROCEDURE PARSEDTDCLOB(
  p    Parser,
  dtd  CLOB,
  root VARCHAR2);
```

Parameters

Table 231-11 PARSEDTDCLOB Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
dtd	(IN)	DTD Clob to parse.
root	(IN)	Name of the root element.

231.2.13 SETBASEDIR

This procedure sets the base directory used to resolve relative URLs. An application error is raised if parsing fails.

Syntax

```
PROCEDURE setBaseDir(
  p   Parser,
  dir VARCHAR2);
```

Parameters

Table 231-12 SETBASEDIR Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
dir	(IN)	Directory used as a base directory.

231.2.14 SETDOCTYPE Procedure

This procedure sets a DTD to be used by the parser for validation. This call should be made before the document is parsed.

Syntax

```
PROCEDURE setDoctype(
  p   Parser,
  dtd DOMDocumentType);
```

Parameters

Table 231-13 SETDOCTYPE Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
dtd	(IN)	DTD to set.

231.2.15 SETERRORLOG Procedure

This procedure sets errors to be sent to the specified file.

Syntax

```
PROCEDURE setErrorLog(
  p        Parser,
  fileName VARCHAR2);
```

Parameters

Table 231-14 SETERRORLOG Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
fileName	(IN)	Complete path of the file to use as the error log.

231.2.16 SETPRESERVEWHITESPACE

This procedure sets whitespace preserving mode.

Syntax

```
PROCEDURE setPreserveWhitespace(
  p   Parser,
  yes BOOLEAN);
```

Parameters

Table 231-15 SETPRESERVEWHITESPACE Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
yes	(IN)	Mode to set: TRUE - preserve, FALSE - don't preserve.

231.2.17 SETVALIDATIONMODE

This procedure sets the validation mode.

Syntax

```
PROCEDURE setValidationMode(
  p   Parser,
  yes BOOLEAN);
```

Parameters

Table 231-16 SETVALIDATIONMODE Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
yes	(IN)	Mode to set: TRUE - validate, FALSE - don't validate.

231.2.18 SHOWWARNINGS

This procedure turns warnings on or off.

Syntax

```
PROCEDURE showWarnings(
  p   Parser,
  yes BOOLEAN);
```

Paramters

Table 231-17 SHOWWARNINGS Procedure Parameters

Parameter	IN / OUT	Description
p	(IN)	Parser instance.
yes	(IN)	Mode to set: TRUE - show warnings, FALSE - don't show warnings.

232
DBMS_XMLSCHEMA

DBMS_XMLSCHEMA package provides procedures to manage XML schemas.

It is created by script dbmsxsch.sql during Oracle database installation.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Views
- Operational Notes
- Summary of DBMS_XMLSCHEMA Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

232.1 DBMS_XMLSCHEMA Overview

The DBMS_XMLSCHEMA package uses subprograms to manage XML schemas.

These subprograms provide the following XML schema management::

- Register an XML schema
- Delete a previously registered XML schema
- Re-compile a previously registered XML schema
- Generate an XML schema
- Evolves an XML schema

232.2 DBMS_XMLSCHEMA Security Model

Owned by XDB, the DBMS_XMLSCHEMA package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

232.3 DBMS_XMLSCHEMA Constants

The DBMS_XMLSCHEMA package defines several constants to use when specifying parameter values.

These constants are shown in following tables.

- Table 232-1
- Table 232-2
- Table 232-3

Table 232-1 *DBMS_XMLSCHEMA Constants* - Delete Option

Constant	Type	Value	Description
DELETE_RESTRICT	NUMBER	1	Deletion of an XML schema fails if there are any tables or XML schemas that depend on it
DELETE_INVALIDATE	NUMBER	2	Deletion of an XML schema does not fail if there are tables or XML schemas that depend on it. All dependent tables and schemas are invalidated.
DELETE_CASCADE	NUMBER	3	Deletion of an XML schema also drops all SQL types and default tables associated with it. SQL types are dropped only if gentypes argument was set to TRUE during registration of the XML schema. However, deletion of the XML schema fails if there are any instance documents conforming to the schema or any dependent XML schemas.
DELETE_CASCADE_FORCE	NUMBER	4	This option is similar to DELETE_CASCADE except that it does not check for any stored instance documents conforming to the schema or any dependent XML schemas. Also, it ignores any errors.

Table 232-2 *DBMS_XMLSCHEMA Constants* - Enable Hierarchy

Constant	Type	Value	Description
ENABLE_HIERARCHY_NONE	PLS_INTEGER	1	The ENABLE_HIERARCHY procedure of the DBMS_XDBZ package will not be called on any tables created while registering that schema
ENABLE_HIERARCHY_CONTENTS	PLS_INTEGER	2	The ENABLE_HIERARCHY procedure of the DBMS_XDBZ package will be called for all tables created during schema registration with hierarchy_type as DBMS_XDBZ.ENABLE_CONTENTS

Table 232-2 (Cont.) *DBMS_XMLSCHEMA Constants* - Enable Hierarchy

Constant	Type	Value	Description
ENABLE_HIERARCHY_RESMETADATA	PLS_INTEGER	3	The ENABLE_HIERARCHY procedure of the DBMS_XDBZ package will be called on all tables created during schema registration with hierarchy_type as DBMS_XDBZ.ENABLE_RESMETADATA. Users should pass in DBMS_XMLSCHEMA.ENABLE_RESMETADATA for schemas they intend to use as resource metadata tables.

Table 232-3 *DBMS_XMLSCHEMA Constants* - Register CSID

Constant	Type	Value	Description
REGISTER_NODOCID	NUMBER	1	If a schema is registered for metadata use (using the value ENABLE_HIER_RESMETADATA for parameter enablehierarchy during registration), a column named DOCID is added to all tables created during schema registration. This constant can be used in the options argument of REGISTERSCHEMA to prevent the creation of this column if the user wishes to optimize on storage
REGISTER_CSID_NULL	NUMBER	-1	If user wishes to not specify the character set of the input schema document when invoking REGISTERSCHEMA, this value can be used for the csid parameter

232.4 Views

This table lists the views used by the DBMS_XMLSCHEMA package

The columns of these views are described in detail in the *Oracle Database Reference*

Table 232-4 Summary of Views used by DBMS_XMLSCHEMA

Schema	Description
USER_XML_SCHEMAS	All registered XML Schemas owned by the user
ALL_XML_SCHEMAS	All registered XML Schemas usable by the current user
DBA_XML_SCHEMAS	All registered XML Schemas in the database
DBA_XML_TABLES	All XMLType tables in the system

Table 232-4 (Cont.) Summary of Views used by DBMS_XMLSCHEMA

Schema	Description
USER_XML_TABLES	All XMLType tables owned by the current user
ALL_XML_TABLES	All XMLType tables usable by the current user
DBA_XML_TAB_COLS	All XMLType table columns in the system
USER_XML_TAB_COLS	All XMLType table columns in tables owned by the current user
ALL_XML_TAB_COLS	All XMLType table columns in tables usable by the current user
DBA_XML_VIEWS	All XMLType views in the system
USER_XML_VIEWS	All XMlType views owned by the current user
ALL_XML_VIEWS	All XMLType views usable by the current user
DBA_XML_VIEW_COLS	All XMLType view columns in the system
USER_XML_VIEW_COLS	All XMLType view columns in views owned by the current user
ALL_XML_VIEW_COLS	All XMLType view columns in views usable by the current user

232.5 DBMS_XMLSCHEMA Operational Notes

There are guidelines for using in-place XML schema evolution.

Before you perform an in-place XML-schema evolution, you should follow these preparatory steps:

1. Back up all existing data (instance documents) for the XML schema that will be evolved.
2. Perform a dry run using trace only, that is, without actually evolving the XML schema or updating any instance documents, to produce a trace of the update operations that would be performed during evolution. To do this, set the flag parameter value to only INPLACE_TRACE. Do not also use INPLACE_EVOLVE. After performing the dry run, examine the trace file, verifying that the listed DDL operations are in fact those that you intend.

232.6 Summary of DBMS_XMLSCHEMA Subprograms

This table lists the DBMS_XMLSCHEMA subprograms and briefly describes them.

Table 232-5 DBMS_XMLSCHEMA Package Subprograms

Method	Description
COMPILESCHEMA Procedure	Used to re-compile an already registered XML schema. This is useful for bringing a schema in an invalid state to a valid state.
COPYEVOLVE Procedure	Evolves registered schemas so that existing XML instances remain valid
DELETESCHEMA Procedure	Removes the schema from the database
INPLACEEVOLVE Procedure	Evolves registered schemas by propagating schema changes to object types and tables
PURGESCHEMA Procedure	Removes the XML schema
REGISTERSCHEMA Procedures	Registers the specified schema for use by Oracle. This schema can then be used to store documents conforming to this.

Table 232-5 (Cont.) DBMS_XMLSCHEMA Package Subprograms

Method	Description
REGISTERURI Procedure	Registers an XML schema specified by a URI name

232.6.1 COMPILESCHEMA Procedure

This procedure can be used to re-compile an already registered XML schema. This is useful for bringing a schema in an invalid state to a valid state. Can result in a `ORA-31001` exception: invalid resource handle or path name.

Syntax

```
DBMS_XMLSCHEMA.COMPILESCHEMA(
   schemaurl IN VARCHAR2);
```

Parameters

Table 232-6 COMPILESCHEMA Procedure Parameters

Parameter	Description
`schemaurl`	URL identifying the schema

232.6.2 COPYEVOLVE Procedure

This procedure evolves registered schemas so that existing XML instances remain valid.

This procedure is accomplished in according to the following basic scenario (alternative actions are controlled by the procedure's parameters):

- copies data in schema based `XMLType` tables to temporary table storage
- drops old tables
- deletes old schemas
- registers new schemas
- creates new `XMLType` tables
- Populates new tables with data in temporary storage; auxiliary structures (constraints, triggers, indexes, and others) are not preserved
- drops temporary tables

> **See Also:**
>
> – "Schema Evolution" chapter of the *Oracle XML DB Developer's Guide* for examples on how to evolve existing schemas
>
> – *Oracle Database Error Messages* for information on exceptions specific to schema evolution, ORA-30142 through ORA-30946.

Syntax

```
DBMS_XMLSCHEMA.COPYEVOLVE(
   schemaurls       IN  XDB$STRUBG_LIST_T,
   newschemas       IN  XMLSequenceType,
   transforms       IN  XMLSequenceType :=NULL,
   preserveolddocs  IN  BOOLEAN :=FALSE,
   maptablename     IN  VARCHAR2 :=NULL,
   generatetables   IN  BOOLEAN :=TRUE,
   force            IN  BOOLEAN :=FALSE,
   schemaowners     IN  XDB$STRING_LIST_T :=NULL
   parallelDegree   IN  PLS_INTEGER := 0,
   options          IN  PLS_INTEGER := 0);
```

Parameters

Table 232-7 COPYEVOLVE Procedure Parameters

Parameter	Description
schemaurls	VARRAY of URLs of all schemas to be evolved. Should include the dependent schemas. Unless the FORCE parameter is TRUE, URLs should be in the order of dependency.
newschemas	VARRAY of new schema documents. Should be specified in same order as the corresponding URLs.
transforms	VARRAY of transforming XSL documents to be applied to schema-based documents. Should be specified in same order as the corresponding URLs. Optional if no transformations are required.
preserveolddocs	Default is FALSE, and temporary tables with old data are dropped. If TRUE, these table are still available after schema evolution is complete.
maptabname	Specifies the name of the table mapping permanent to temporary tables during the evolution process. Valid columns are: • SCHEMA_URL - VARCHAR2(700) - URL of schema to which this table conforms • SCHEMA_OWNER -VARCHAR2(30) - Owner of the schema • ELEMENT_NAME - VARCHAR2(256)- Element to which this table conforms • TAB_NAME - VARCHAR2(65) - Qualified table name: *<owner_name>.<table_name>* • COL_NAME - VARCHAR2(4000) - Name of the column (NULL for XMLType tables) • TEMP_TABNAME - VARCHAR2(30) - Name of temporary tables which holds data for this table.
generatetables	Default is TRUE, and new tables will be generated. If FALSE: • new tables will not be generated after registration of new schemas • preserveolddocs must be TRUE • maptablename must be non-NULL
force	Default is FALSE. If TRUE, ignores errors generated during schema evolution. Used when there are circular dependencies among schemas to ensure that all schemas are stored despite possible errors in registration.

Table 232-7 (Cont.) COPYEVOLVE Procedure Parameters

Parameter	Description
schemaowners	VARRAY of names of schema owners. Should be specified in same order as the corresponding URLs. Default is NULL, assuming that all schemas are owned by the current user.
paralleldegree	Specifies the degree of parallelism to be used in a PARALLEL hint during the data copy stage of the evolution. If this is 0 (default), the PARALLEL hint will not be given in the data copy statements.
options	Currently, the only supported option is COPYEVOLVE_BINARY_XML which lets you register the new schemas for binary XML and create the new tables/columns with binary XML as the storage type.

Usage Notes

You should back up all schemas and documents prior to invocation because COPYEVOLVE Procedure deletes all conforming documents prior to implementing the schema evolution.

232.6.3 DELETESCHEMA Procedure

This procedure deletes the XML Schema specified by the URL.

Syntax

```
DBMS_XMLSCHEMA.DELETESCHEMA(
   schemaurl      IN  VARCHAR2,
   delete_option  IN  PLS_INTEGER := DELETE_RESTRICT);
```

> **See Also:**
>
> "XMLSCHEMA Storage and Query: Basic" chapter of the *Oracle XML DB Developer's Guide*

Parameters

Table 232-8 DELETESCHEMA Procedure Parameters

Parameter	Description
schemaurl	URL identifying the schema to be deleted

Table 232-8 (Cont.) DELETESCHEMA Procedure Parameters

Parameter	Description
delete_option	Delete options: • DELETE_RESTRICT - Schema deletion fails if there are any tables or schemas that depend on this schema • DELETE_INVALIDATE - Schema deletion does not fail if there are any dependencies. Instead, it simply invalidates all dependent objects. • DELETE_CASCADE - Schema deletion will also drop all default SQL types and default tables. However the deletion fails if there are any stored instances conforming to this schema. • DELETE_CASCADE_FORCE - Similar to DELETE_CASCADE except that it does not check for any stored instances conforming to this schema. Also, it ignores any errors.

Exceptions

Table 232-9 DELETESCHEMA Procedure Exceptions

Exception	Description
ORA-31001	Invalid resource handle or path name

232.6.4 INPLACEEVOLVE Procedure

This procedure evolves registered schemas by propagating schema changes to object types and tables.

Syntax

```
DBMS_XMLSCHEMA.INPLACEEVOLVE(
   schemaURL     IN   VARCHAR2,
   diffXML       IN   XMLType,
   flags         IN   NUMBER);
```

Parameters

Table 232-10 INPLACEEVOLVE Procedure Parameters

Parameter	Description
schemaurl	URL of the schema to evolve
diffXML	Changes to be applied to the schema. This is an XML document conforming to the XDIFF schema and specifies what changes need to be applied and the locations in the schema document where the changes are to be applied.

Table 232-10 (Cont.) INPLACEEVOLVE Procedure Parameters

Parameter	Description
`flags`	The following bits may be set in this parameter to control the behavior of this procedure: • `INPLACE_EVOLVE` (value 1, meaning that bit 1 is on) – Perform in-place XML schema evolution: construct a new XML schema and validate it (against the XML schema for XML schemas); construct the DDL statements needed to evolve the instance-document disk structures, execute the DDL statements, and replace the old XML schema with the new. • `INPLACE_TRACE` (value 2, meaning that bit 2 is on) – Perform all steps necessary for in-place evolution, except executing the DDL statements and overwriting the old XML schema with the new, then write both the DDL statements and the new XML schema to a trace file. That is, each of the bits constructs the new XML schema, validates it, and determines the steps needed to evolve the disk structures underlying the instance documents. In addition: • Bit `INPLACE_EVOLVE` carries out those evolution steps and replaces the old XML schema with the new. • Bit `INPLACE_TRACE` saves the evolution steps and the new XML schema in a trace file (it does not carry out the evolution steps)

Exceptions

The procedure raises exceptions in the following cases:

- An error will be raised for invalid `XPATH` expressions and for `XDIFF` documents that do not conform to the xdiff schema.
- Path expressions that are syntactically correct but result in an invalid node in the schema document will result in an error.
- If the schema change makes the schema an ill-formed XML document or an invalid XML schema, this will raise an error.
- Any errors resulting from `CREATE TYPE`, `ALTER TYPE` and like commands will generate error messages.

Usage Notes

- Users are required to backup all their data before attempting in-place evolution, as there is no rollback with this operation.
- A user must register their new XML schema with the database using the REGISTERSCHEMA Procedures and the REGISTERURI Procedure at a schema URL that is different from that of the one to be evolved. If the new schema registers successfully and is usable, only then should the user attempt to evolve the existing schema to the new schema by means of this subprogram. If the registration of the new schema is successful, then the user must delete this schema (and all its dependent objects) before attempting to evolve the schema at the old schema URL.

232.6.5 PURGESCHEMA Procedure

This procedure removes the XML schema.

> **See Also:**
>
> "XMLSCHEMA Storage and Query: Advanced" chapter of the *Oracle XML DB Developer's Guide*

Syntax

```
DBMS_XMLSCHEMA.PURGESCHEMA(
   schemaid   IN  RAW);
```

Parameters

Table 232-11 PURGESCHEMA Procedure Parameters

Parameter	Description
`schemaid`	ID of the schema to be purged

Usage Notes

- The schema should have been originally registered for binary encoding and should have been deleted in the `HIDE` mode.
- Once a schema has been deleted in `HIDE` mode, it continues to exist in the XML DB dictionary and is used for decoding already encoded documents. The user invokes this interface when there are no stored instances encoded with this schema.
- Once the schema is purged, any space used by that schema will be reclaimed and documents encoded using the schema will raise an error if an attempt is made to decode them.
- The Schema ID can be obtained from the catalog views.

232.6.6 REGISTERSCHEMA Procedures

This procedure registers the specified schema for use by the database.

The procedure is overloaded. The different functionality of each form of syntax is presented along with the definition.

> **Note:**
>
> As of Oracle Database 11*g* Release 2 (11.2) the `genbean` parameter is deprecated. Oracle recommends that you do not use this parameter in new applications. Support for this feature is for backward compatibility only.

> **See Also:**
>
> "XMLSCHEMA Storage and Query: Basic" chapter of the *Oracle XML DB Developer's Guide*

Syntax

Registers a schema specified as a `VARCHAR2`:

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
    schemaurl        IN  VARCHAR2,
    schemadoc        IN  VARCHAR2,
    local            IN  BOOLEAN := TRUE,
    gentypes         IN  BOOLEAN := TRUE,
    genbean          IN  BOOLEAN := FALSE,
    gentables        IN  BOOLEAN := TRUE,
    force            IN  BOOLEAN := FALSE,
    owner            IN  VARCHAR2 := NULL,
    enablehierarchy  IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
    options          IN  PLS_INTEGER := 0);
```

Registers the schema specified as a `BFILE`. The contents of the schema document must be in the database character set:

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl        IN  VARCHAR2,
   schemadoc        IN  BFILE,
   local            IN  BOOLEAN := TRUE,
   gentypes         IN  BOOLEAN := TRUE,
   genbean          IN  BOOLEAN := FALSE,
   force            IN  BOOLEAN := FALSE,
   owner            IN  VARCHAR2 := NULL,
   enablehierarchy  IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
   options          IN  PLS_INTEGER := 0);
```

Registers the schema specified as a `BFILE` and identifies the character set id of the schema document:

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl        IN  VARCHAR2,
   schemadoc        IN  BFILE,
   local            IN  BOOLEAN := TRUE,
   gentypes         IN  BOOLEAN := TRUE,
   genbean          IN  BOOLEAN := TRUE,
   gentables        IN  BOOLEAN := TRUE,
   force            IN  BOOLEAN := TRUE,
   owner            IN  VARCHAR2 := '',
   csid             IN  NUMBER,
   enablehierarchy  IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
   options          IN  PLS_INTEGER := 0);
```

Registers the schema specified as a `BLOB`. The contents of the schema document must be in the database character set:

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl        IN  VARCHAR2,
```

```
   schemadoc         IN  BLOB,
   local             IN  BOOLEAN := TRUE,
   genTypes          IN  BOOLEAN := TRUE,
   genBean           IN  BOOLEAN := FASLE,
   force             IN  BOOLEAN := FALSE,
   owner             IN  VARCHAR2 := NULL,
   enablehierarchy   IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
   options           IN  PLS_INTEGER := 0);
```

Registers the schema specified as a `BLOB` and identifies the character set id of the schema document:

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl         IN  VARCHAR2,
   schemadoc         IN  BLOB,
   local             IN  BOOLEAN := TRUE,
   gentypes          IN  BOOLEAN := TRUE,
   genbean           IN  BOOLEAN := TRUE,
   gentables         IN  BOOLEAN := TRUE,
   force             IN  BOOLEAN := TRUE,
   owner             IN  VARCHAR2 := '',
   csid              IN  NUMBER,
   enablehierarchy   IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
   options           IN  PLS_INTEGER := 0);
```

Registers the schema specified as a `CLOB`

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl         IN  VARCHAR2,
   schemadoc         IN  CLOB,
   local             IN  BOOLEAN := TRUE,
   gentypes          IN  BOOLEAN := TRUE,
   genbean           IN  BOOLEAN := FALSE,
   force             IN  BOOLEAN := FALSE,
   owner             IN  VARCHAR2 := NULL,
   options           IN  PLS_INTEGER := 0);
```

Registers the schema specified as an `XMLTYPE`.

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl         IN  VARCHAR2,
   schemadoc         IN  SYS.XMLTYPE,
   local             IN  BOOLEAN := TRUE,
   gentypes          IN  BOOLEAN := TRUE,
   genbean           IN  BOOLEAN := FALSE,
   force             IN  BOOLEAN := FALSE,
   owner             IN  VARCHAR2 := NULL,
   enablehierarchy   IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
   options           IN  PLS_INTEGER := 0);
```

Registers the schema specified as a `BLOB`. The contents of the schema document must be in the database character set:

```
DBMS_XMLSCHEMA.REGISTERSCHEMA(
   schemaurl         IN  VARCHAR2,
   schemadoc         IN  SYS.URIType,
   local             IN  BOOLEAN := TRUE,
   gentypes          IN  BOOLEAN := TRUE,
```

```
  genbean          IN  BOOLEAN := FALSE,
  force            IN  BOOLEAN := FALSE,
  owner            IN  VARCHAR2 := NULL,
  enablehierarchy  IN  PLS_INTEGER := DBMS_XMLSCHEMA.ENABLE_CONTENTS,
  options          IN  PLS_INTEGER := 0);
```

Parameters

Table 232-12 REGSITERSCHEMA Procedure Parameters

Parameter	Description
schemaurl	URL that uniquely identifies the schema document. This value is used to derive the path name of the schema document within the database hierarchy. Can be used inside schemalocation attribute of XML Schema import element.
schemadoc	A valid XML schema document
local	Is this a local or global schema? • By default, all schemas are registered as local schemas, under /sys/schemas/<username>/... • If a schema is registered as global, it is added under /sys/schemas/PUBLIC/... You need write privileges on the directory to be able to register a schema as global.
gentypes	Determines whether the schema compiler generates object types. By default, TRUE. If you use binary XML, you must be set gentypes to FALSE.
genbean	Determines whether the schema compiler generates Java beans. By default, FALSE. Oracle recommends that this parameter always be set to FALSE.
gentables	Determines whether the schema compiler generates default tables. By default, TRUE
force	If this parameter is set to TRUE, the schema registration will not raise errors. Instead, it creates an invalid XML schema object in case of any errors. By default, the value of this parameter is FALSE.
owner	This parameter specifies the name of the database user owning the XML schema object. By default, the user registering the schema owns the XML schema object. This parameter can be used to register a XML schema to be owned by a different database user.
csid	Identifies the character set of the input schema document. If this value is 0, the schema document's encoding is determined by the current rule for "text/xml" MIME type.
enablehierarchy	• ENABLE_HIERARCHY_NONE - enable hierarchy will not be called on any tables created while registering that schema • ENABLE_HIERARCHY_CONTENTS - enable hierarchy will be called for all tables created during schema registration with hierarchy_type as DBMS_XDBZ.ENABLE_CONTENTS. This is the default. • ENABLE_HIERARCHY_RESMETADATA - enable hierarchy will be called on all tables created during schema registration with hierarchy_type as DBMS_XDBZ.ENABLE_RESMETADATA. Users should pass in DBMS_XMLSCHEMA.ENABLE_RESMETADATA for schemas they intend to use as resource metadata tables.

Table 232-12 (Cont.) REGSITERSCHEMA Procedure Parameters

Parameter	Description
options	Additional options to specify how the schema should be registered. The various options are represented as bits of an integer and the options parameter should be constructed by doing a BITOR of the desired bits. Possible bits: • REGISTER_NODOCID - this will suppress the creation of the DOCID column for out of line tables. This is a storage optimization which might be desirable when we do not need to join back to the document table (for example if we do not care about rewriting certain queries that could be rewritten by making use of the DOCID column) • REGISTER_BINARYXML - Register the schema for Binary XML • REGISTER_NT_AS_IOT - Store nested tables created during schema registration as index organized tables. The default is to store nested tables as heap tables

232.6.7 REGISTERURI Procedure

This procedure registers an XML Schema specified by a URI name.

> **Note:**
>
> As of Oracle Database 11g Release 2 (11.2) the genbean parameter is deprecated. Oracle recommends that you do not use this parameter in new applications. Support for this feature is for backward compatibility only.

Syntax

```
DBMS_XMLSCHEMA.REGISTERURI(
   schemaurl      IN  VARCHAR2,
   schemadocuri   IN  VARCHAR2,
   local          IN  BOOLEAN := TRUE,
   gentypes       IN  BOOLEAN := TRUE,
   genbean        IN  BOOLEAN := FALSE,
   gentables      IN  BOOLEAN := TRUE,
   force          IN  BOOLEAN := FALSE,
   owner          IN  VARCHAR2 := NULL,
   options          IN  PLS_INTEGER := 0);
```

Parameters

Table 232-13 REGISTERURI Procedure Parameters

Parameter	Description
schemaurl	Uniquely identifies the schema document. Can be used inside schemaLocation attribute of XML Schema import element.

Table 232-13 (Cont.) REGISTERURI Procedure Parameters

Parameter	Description
schemadocuri	Pathname (URI) corresponding to the physical location of the schema document. The URI path could be based on HTTP, FTP, DB or Oracle XML DB protocols. This function constructs a URIType instance using the urifactory, and invokes the REGISTERSCHEMA Procedures.
local	Determines whether this is a local or global schema. By default, all schemas are registered as local schemas, under /sys/schemas/<username>/... If a schema is registered as global, it is added under /sys/schemas/PUBLIC/... The user needs write privileges on the directory to register a global schema.
gentypes	Determines whether the compiler generate object types. By default, TRUE.
genbean	Determines whether the compiler generate Java beans. By default, FALSE.
gentables	Determines whether the compiler generate default tables. TRUE by default.
force	TRUE: schema registration will not raise errors. Instead, it creates an invalid XML schema object in case of any errors. By default, the value of this parameter is FALSE.
owner	This parameter specifies the name of the database user owning the XML schema object. By default, the user registering the schema owns the XML schema object. This parameter can be used to register a XML schema to be owned by a different database user.
options	Additional options to specify how the schema should be registered. The various options are represented as bits of an integer and the options parameter should be constructed by doing a BITOR of the desired bits. Possible bits: • REGISTER_NODOCID - this will suppress the creation of the DOCID column for out of line tables. This is a storage optimization which might be desirable when we do not need to join back to the document table (for example if we do not care about rewriting certain queries that could be rewritten by making use of the DOCID column)

233
DBMS_XMLSCHEMA_ANNOTATE

The `DBMS_XMLSCHEMA_ANNOTATE` package provides an interface to manage and configure the structured storage model, mainly through the use of pre-registration schema annotations.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_XMLSCHEMA_ANNOTATE Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

233.1 DBMS_XMLSCHEMA_ANNOTATE Overview

The `DBMS_XMLSCHEMA_ANNOTATE` package contains procedures to manage and configure the structured storage model, mainly through the use of pre-registration schema annotations.

Schema annotations influence the way the XML data is stored. For example, the default table annotation assigns a user-provided name to an XML element instead of allowing the database to generate a system name. Consequently, query plans are more readable and it is easier to create constraints on that table.

233.2 DBMS_XMLSCHEMA_ANNOTATE Security Model

Owned by `XDB`, the `DBMS_XMLSCHEMA_ANNOTATE` package must be created by `SYS` or `XDB`. The `EXECUTE` privilege is granted to `PUBLIC`. Subprograms in this package are executed using the privileges of the current user.

233.3 Summary of DBMS_XMLSCHEMA_ANNOTATE Subprograms

This table lists and describes the `DBMS_XMLSCHEMA_ANNOTATE` package subprograms.

Table 233-1 DBMS_XMLSCHEMA_ANNOTATE Package Subprograms

Subprogram	Description
ADDXDBNAMESPACE Procedure	Adds the XDB namespace required for XDB annotation
DISABLEDEFAULTTABLECREATION Procedure	Prevents the creation of a table for the top-level element by adding a default table attribute with an empty value to the element

Table 233-1 (Cont.) DBMS_XMLSCHEMA_ANNOTATE Package Subprograms

Subprogram	Description
DISABLEMAINTAINDOM Procedure	Sets the DOM fidelity attribute to `FALSE`
ENABLEDEFAULTTABLECREATION Procedure	Enables the creation of `ALL` top level tables by removing the empty default table name annotation
ENABLEMAINTAINDOM Procedure	Sets the DOM fidelity attribute to `TRUE`
GETSCHEMAANNOTATIONS Function	Creates a document containing the differences between the annotated XML schema and the original XML schema
GETSIDXDEFFROMVIEW Function	Takes a XMLTABLE view definition on a xmltype column or table and it returns a `CLOB` which can be used as parameter to create a structured xmlindex that backs up the XMLTABLE view as relational table
PRINTWARNINGS Procedure	Lets a user raise or suppress a warning if an annotation maps to zero nodes in the XML schema
REMOVEANYSTORAGE Procedure	Removes the setting of the SQL type from the `ANY` child of the complex type with the given name
REMOVEDEFAULTTABLE Procedure	Removes any default table attribute given for the element. After calling this procedure, the system generates table names
REMOVEMAINTAINDOM Procedure	Removes all annotations used to maintain DOM from the given schema
REMOVEOUTOFLINE Procedure	Removes any existing `SQLInline` attributes to prevent out-of-line storage
REMOVESQLCOLLTYPE Procedure	Removes a SQL collection type.
REMOVESQLNAME Procedure	Removes a `SQLNAME` from a global element
REMOVESQLTYPE Procedure	Removes a SQL type
REMOVESQLTYPEMAPPING Procedure	Removes the SQL type mapping for the given schema type.
REMOVETABLEPROPS Procedure	Removes the table storage properties from the `CREATE TABLE` statement
REMOVETIMESTAMPWITHTIMEZONE Procedure	Removes he setting of the `TimeStampWithTimeZone` datatype from all `dateTime` typed elements in the XML schema
SETANYSTORAGE Procedure	Assigns a SQL datatype to the `ANY` child of the complex type with the given name
SETDEFAULTTABLE Procedure	Sets the name of the table for the specified global element
SETOUTOFLINE Procedure	Sets the `SQLInline` attribute to `FALSE`
SETSCHEMAANNOTATATIONS Procedure	Takes the annotated differences resulting from a call to `DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS` and patches them into the provided XML schema
SETSQLCOLLTYPE Procedure	Assigns a SQL type name for a collection
SETSQLNAME Procedure	Assigns a name to the SQL attribute that corresponds to an element defined in the XML schema
SETSQLTYPE Procedure	Assigns a SQL type to a global object
SETSQLTYPEMAPPING Procedure	Defines a mapping of schema type and SQL type
SETTABLEPROPS Procedure	Specifies properties in the `TABLE` storage clause that is appended to the default `CREATE TABLE` statement

Table 233-1 (Cont.) DBMS_XMLSCHEMA_ANNOTATE Package Subprograms

Subprogram	Description
SETTIMESTAMPWITHTIMEZONE Procedure	Sets the `TIMESTAMPWITHTIMEZONE` datatype to all `dateTime` typed elements in the XML schema

233.3.1 ADDXDBNAMESPACE Procedure

This procedure adds the XDB namespace required for XDB annotation.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.ADDXDBNAMESPACE (
   xmlschema      IN OUT XMLTYPE);
```

Parameters

Table 233-2 ADDXDBNAMESPACE Procedure Parameters

Parameter	Description
`xmschema`	Gets an XML Schema as `XMLTYPE`, performs the annotation and returns it

Usage Notes

This procedure is called implicitly by any other procedure that adds a schema annotation. Since there is no reason to add an XDB namespace without other annotations, this procedure is most likely called by other annotations procedures and not by the user directly.

233.3.2 DISABLEDEFAULTTABLECREATION Procedure

This procedure prevents the creation of a table for the top-level element by adding a default table attribute with an empty value to the element. The first overload applies to a specified top-level element and the second applies to all top-level elements. The procedure always overwrites. This is equivalent to using the schema annotation `xdb:defaultTable=""` for the top-level element or elements.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.DISABLEDEFAULTTABLECREATION (
   xmlschema         IN OUT XMLType,
   globalElementName IN VARCHAR2);


DBMS_XMLSCHEMA_ANNOTATE.DISABLEDEFAULTTABLECREATION (
   xmlschema         IN OUT XMLType);
```

Parameters

Table 233-3 DISABLEDEFAULTTABLECREATION Procedure Parameters

Parameter	Description
xmlschema	XML schema to be annotated
globalElementName	Name of the global element in the schema

Example

The purchaseOrder element will have an annotation similar to xdb:defaultTable="".

```
DECLARE
   xml_schema   XMLTYPE;
BEGIN
   SELECT out INTO xml_schema FROM annotation_tab;
   DBMS_XMLSCHEMA_ANNOTATE.DISABLEDEFAULTTABLECREATION(xml_schema,
                                                       'purchaseOrder');
   UPDATE annotation_tab SET out = xml_schema;
END;
/
```

233.3.3 DISABLEMAINTAINDOM Procedure

This procedure sets the DOM fidelity attribute to FALSE.

There are two overloads. The first sets DOM fidelity attribute to FALSE for all complex types, and the second sets it to FALSE for the named complex type. This is equivalent to adding xdb:maintainDOM="false" on all or specified complex types respectively.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.DISABLEMAINTAINDOM (
   xmlschema       IN OUT XMLType,
   overwrite       IN BOOLEAN default TRUE);


DBMS_XMLSCHEMA_ANNOTATE.DISABLEMAINTAINDOM (
   xmlschema       IN OUT XMLType,
   complexTypeName IN VARCHAR2,
   overwrite       IN BOOLEAN default TRUE);
```

Parameters

Table 233-4 DISABLEMAINTAINDOM Procedure Parameters

Parameter	Description
xmlschema	The XML schema to be annotated
complexTypeName	The name of the complex type
overwrite	A boolean that indicates whether or not the procedure overwrites element attributes. The default is TRUE

233.3.4 ENABLEDEFAULTTABLECREATION Procedure

This procedure enables the creation of ALL top level tables by removing the empty default table name annotation.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.ENABLEDEFAULTTABLECREATION (
   xmlschema            IN OUT  XMLTYPE);

DBMS_XMLSCHEMA_ANNOTATE.ENABLEDEFAULTTABLECREATION (
   xmlschema            IN OUT  XMLTYPE,
   globalElementName    IN      VARCHAR2);;);
```

Parameters

Table 233-5 ENABLEDEFAULTTABLECREATION Procedure Parameters

Parameter	Description
xmlschema	The XML schema to be annotated
gloablElementName	Name of the global element in the schema

Usage Notes

This procedure does not affect elements that have a default table name.

233.3.5 ENABLEMAINTAINDOM Procedure

This overloaded procedure sets the DOM fidelity attribute to `TRUE`.

There are two overloads. The first sets DOM fidelity attribute to `TRUE` for all complex types, and the second sets it to TRUE for the named complex type.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.ENABLEMAINTAINDOM (
   xmlschema       IN OUT XMLType,
   overwrite       IN BOOLEAN default TRUE);


DBMS_XMLSCHEMA_ANNOTATE.ENABLEMAINTAINDOM (
   xmlschema        IN OUT XMLType,
   complexTypeName  IN VARCHAR2,
   overwrite        IN BOOLEAN default TRUE);
```

Parameters

Table 233-6 ENABLEMAINTAINDOM Procedure Parameters

Parameter	Description
xmlschema	The XML schema to be annotated
complexTypeName	The name of the complex type

Table 233-6 (Cont.) ENABLEMAINTAINDOM Procedure Parameters

Parameter	Description
`overwrite`	A boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`

233.3.6 GETSCHEMAANNOTATIONS Function

This function creates a document containing the differences between the annotated XML schema and the original XML schema.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS (
   xmlschema IN xmlType)
  RETURN XMLType;
```

Parameters

Table 233-7 GETSCHEMAANNOTATIONS Function Parameters

Parameter	Description
`xmlschema`	The original XML schema

Return Values

This function returns the document `annotations.xml` as an `XMLType`.

Usage Notes

This function saves all annotations in one document, named `annotations`, and returns it. With this document, you can apply all annotations to a non-annotated schema, using `DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS`.

`DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS` is not available on Oracle Database release 10.2 (only Oracle Database release 11.*x*).

> **See Also:**
>
> SETSCHEMAANNOTATATIONS Procedure

Example

For an example of `DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS`, see the example in SETSCHEMAANNOTATATIONS Procedure.

233.3.7 GETSIDXDEFFROMVIEW Function

This function takes a XMLTABLE view definition on a xmltype column or table and it returns a CLOB which can be used as parameter to create a structured xmlindex that backs up the XMLTABLE view as relational table.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.GETSIDXDEFFROMVIEW (
   viewName   IN xmlType)
  RETURN CLOB;
```

Parameters

Table 233-8 GETSIDXDEFFROMVIEW Function Parameters

Parameter	Description
viewName	The original XML schema

Return Values

This function returns a CLOB which can be used as parameter to create a structured xmlindex that backs up the XMLTABLE view as relational table.

233.3.8 PRINTWARNINGS Procedure

This procedure lets a user raise or suppress a warning if an annotation maps to zero nodes in the XML schema.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.PRINTWARNINGS (
   value       IN   BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-9 PRINTWARNINGS Procedure Parameters

Parameter	Description
val	For the NO MATCHING ELEMENTS FOUND error message to be raised val must be set to TRUE. In cases in which user wishes to suppress this warning, set to FALSE.

Usage Notes

If an annotation maps to more than one node in the XML schema, this raise the error ANNOTATION MAPS TO MULTIPLE ELEMENTS. In this case no annotation is performed, and the user must correct the parameters to the procedure call to refer to a unique node in the XML schema.

233.3.9 REMOVEANYSTORAGE Procedure

This procedure removes the setting of the SQL type from the `ANY` child of the complex type with the given name.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVEANYSTORAGE (
   xmlschema       IN OUT XMLType,
   complexTypeName IN VARCHAR2);
```

Parameters

Table 233-10 REMOVEANYSTORAGE Procedure Parameters

Parameter	Description
`xmlschema`	The XML schema to be annotated.
`complexTypeName`	The name of the complex type.

Usage Notes

This procedure reverses the SETANYSTORAGE Procedure.

233.3.10 REMOVEDEFAULTTABLE Procedure

This procedure removes any default table attribute given for the element.

After calling this procedure, the system generates table names. This procedure always overwrites.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVEDEFAULTTABLE (
   xmlschema          IN OUT XMLTYPE,
   globalElementName  IN     VARCHAR2);
```

Parameters

Table 233-11 REMOVEDEFAULTTABLE Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`globalElementName`	Name of the global element in the schema

Example

Annotations can be verified anytime using "`select out from annotation_tab`".

```
--The purchaseOrder element will have no annotation for defaultTable.
DECLARE
   xml_schema XMLTYPE;
BEGIN
```

```
   SELECT out INTO xml_schema FROM annotation_tab;
   DBMS_XMLSCHEMA_ANNOTATE.REMOVEDEFAULTTABLE(xml_schema,
                                              'purchaseOrder');
   UPDATE annotation_tab SET out = xml_schema;
END;
/
```

233.3.11 REMOVEMAINTAINDOM Procedure

This procedure removes all annotations used to maintain DOM from the given schema.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVEMAINTAINDOM (
   xmlschema    IN OUT XMLType);
```

Parameters

Table 233-12 REMOVEMAINTAINDOM Procedure Parameters

Parameter	Description
xmlschema	The XML schema to be annotated

233.3.12 REMOVEOUTOFLINE Procedure

This procedure removes any existing `SQLInline` attributes to prevent out-of-line storage.

There are three overloads.

Syntax

Removes the `SQLInline` attribute for the named element.

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVEOUTOFLINE (
   xmlschema           IN OUT  XMLType,
   elementName         IN      VARCHAR2,
   elementType         IN      VARCHAR2,
   overwrite           IN      BOOLEAN default TRUE);
```

Removes the `SQLInline` attribute for the object specified by its global object and local element names.

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVEOUTOFLINE (
   xmlschema          IN OUT  XMLType,
   globalObject       IN      VARCHAR2,
   globalObjectName   IN      VARCHAR2,
   localElementName   IN      VARCHAR2);
```

Removes the `SQLInline` attribute for the referenced global element.

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVEOUTOFLINE (
   xmlschema         IN OUT  XMLType,
   reference         IN      VARCHAR2);
```

Parameters

Table 233-13 REMOVEOUTOFLINE Procedure Parameters

Parameter	Description
xmlschema	The XML schema to be annotated
elementName	The element name
elementType	The element type
globalObject	The global object (global complex type or global element)
globalObjectName	The name of the global object
localElementName	The name of a local element that descends from the global element
reference	A reference to a global element
overwrite	A boolean that indicates whether or not the procedure overwrites element attributes. The default is TRUE.

Usage Notes

This procedure reverses SETOUTOFLINE Procedure.

233.3.13 REMOVESQLCOLLTYPE Procedure

This procedure removes a SQL collection type.

The first overload removes the SQL collection type corresponding to the named element and the second overload removes the type from the XML element inside the complex type.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVESQLCOLLTYPE (
   xmlschema   IN OUT XMLType,
   elementName IN VARCHAR2);

DBMS_XMLSCHEMA_ANNOTATE.REMOVESQLCOLLTYPE (
   xmlschema        IN OUT XMLType,
   globalObject     IN VARCHAR2,
   globalName       IN VARCHAR2,
   localElementName IN VARCHAR2);
```

Parameters

Table 233-14 REMOVESQLCOLLTYPE Procedure Parameters

Parameter	Description
xmlschema	The XML schema to be annotated
elementName	The element name
globalObject	The global object (global complex type or global element)
globalName	The name of the global object
localElementName	The name of a local element that descends from the global element

Usage Notes

This procedure reverses the SETSQLCOLLTYPE Procedure.

233.3.14 REMOVESQLNAME Procedure

This procedure removes a `SQLNAME` from a global element.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVESQLNAME (
   xmlschema          IN OUT  XMLType,
   globalObject       IN      VARCHAR2,
   globalObjectName   IN      VARCHAR2,
   localObject        IN      VARCHAR2,
   localObjectName    IN      VARCHAR2,
   sqlName            IN      VARCHAR2,
   overwrite          IN      BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-15 REMOVESQLNAME Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`globalObject`	Global object (global complex type or global element)
`globalObjectName`	Name of the global object
`localObject`	Object descended from the global object
`localObjectName`	Name of the local object
`sqlName`	Name of the SQL attribute that corresponds to the element defined in the XML schema
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

Example

The `shipTo` element will have an annotation similar to `xdb:SQLName="SHIPTO_SQLNAME"`.

```
DECLARE
   xml_schema   XMLTYPE;
BEGIN
   SELECT out INTO xml_schema FROM annotation_tab;
   DBMS_XMLSCHEMA_ANNOTATE.SETSQLNAME (xml_schema,
                                       'element', 'purchaseOrder',
                                       'element', 'shipTo',
                                       'SHIPTO_SQLNAME');
   UPDATE annotation_tab SET out = xml_schema;
END;
/
```

233.3.15 REMOVESQLTYPE Procedure

This procedure removes a SQL type.

The first overload removes a SQL type from a global element and the second overload removes the type from a global element inside the complex type.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVESQLTYPE (
   xmlschema in out XMLType,
   globalElementName   IN      VARCHAR2);

DBMS_XMLSCHEMA_ANNOTATE.REMOVESQLTYPE (
   xmlschema           IN OUT  XMLTYPE,
   globalObject        IN      VARCHAR2,
   globalObjectName    IN      VARCHAR2,
   localObject         IN      VARCHAR2,
   localObjectName     IN      VARCHAR2);
```

Parameters

Table 233-16 REMOVESQLTYPE Procedure Parameters

Parameter	Description
xmlschema	XML schema to be annotated.
globalObject	Global object (global complex type or global element)
globalElementName	Name of the global element.
globalObjectName	Name of the global object
localObject	Object descended from the global object
localObjectName	Name of the local object

Usage Notes

This procedure reverses the SETSQLTYPE Procedure.

233.3.16 REMOVESQLTYPEMAPPING Procedure

This procedure removes the SQL type mapping for the given schema type.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVESQLTYPEMAPPING (
   xmlschema        IN OUT  XMLTYPE,
   schemaTypeName   IN      VARCHAR2);
```

Parameters

Table 233-17 REMOVESQLTYPEMAPPING Procedure Parameters

Parameter	Description
xmlschema	XML schema to be annotated

Table 233-17 (Cont.) REMOVESQLTYPEMAPPING Procedure Parameters

Parameter	Description
schemaTypeName	Name of the schema type

Usage Notes

This procedure reverses the SETSQLTYPEMAPPING Procedure.

233.3.17 REMOVETABLEPROPS Procedure

This procedure removes the table storage properties from the CREATE TABLE statement.

This procedure is overloaded. Each overload has different parameter requirements as indicated.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVETABLEPROPS (
   xmlschema            IN OUT  XMLTYPE,
   globalElementName    IN      VARCHAR2);

DBMS_XMLSCHEMA_ANNOTATE.REMOVETABLEPROPS (
   xmlschema            IN OUT  XMLTYPE,
   globalObject         IN      VARCHAR2,
   globalObjectName     IN      VARCHAR2,
   localElementName     IN      VARCHAR2);
```

Parameters

Table 233-18 REMOVETABLEPROPS Procedure Parameters

Parameter	Description
xmlschema	XML schema to be annotated
globalElementName	Name of the global element in the schema
globalObject	Global object (global complex type or global element)
globalObjectName	Name of the global object
localElementName	Name of a local element that descends from the global element

Usage Notes

This procedure reverses the SETTABLEPROPS Procedure.

233.3.18 REMOVETIMESTAMPWITHTIMEZONE Procedure

This procedure removes the setting of the TimeStampWithTimeZone datatype from all dateTime typed elements in the XML schema.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVETIMESTAMPWITHTIMEZONE (
   xmlschema          IN OUT  XMLTYPE);
```

```
DBMS_XMLSCHEMA_ANNOTATE.REMOVETIMESTAMPWITHTIMEZONE (
   xmlschema         IN OUT  XMLTYPE,
   schemaTypeName    IN      VARCHAR2);
```

Parameters

Table 233-19 REMOVETIMESTAMPWITHTIMEZONE Procedure Parameters

Parameter	Description
xmlschema	XML schema to be annotated
schemaTypeName	Name of the schema type

Usage Notes

This procedure reverses the SETTIMESTAMPWITHTIMEZONE Procedure.

233.3.19 SETANYSTORAGE Procedure

This procedure assigns a SQL datatype to the ANY child of the complex type with the given name.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETANYSTORAGE )
   xmlschema       IN OUT  XMLType,
   complexTypeName IN      VARCHAR2,
   sqlTypeName     IN      VARCHAR2,
   overwrite       IN      BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-20 SETANYSTORAGE Procedure Parameters

Parameter	Description
xmlschema	XML schema to be annotated
complexTypeName	Name of the complex type
sqlTypeName	Name of the SQL type
overwrite	Boolean that indicates whether or not the procedure overwrites element attributes. The default is TRUE.

Example

The xsd:any child of complex type Items is assigned an annotation similar to xdb:SQLType="VARCHAR".

```
DECLARE   xml_schema  XMLTYPE;BEGIN   SELECT out INTO xml_schema FROM annotation_tab;
DBMS_XMLSCHEMA_ANNOTATE.setAnyStorage
(xml_schema,
'Items',                                        'VARCHAR');   UPDATE annotation_tab SET
out = xml_schema;END;
/
```

233.3.20 SETDEFAULTTABLE Procedure

This procedure sets the name of the table for the specified global element. This is equivalent to using the schema annotation `xdb:defaultTable="<default_table_name>"` for the top-level element.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETDEFAULTTABLE (
   xmlschema           IN  OUT  XMLTYPE,
   globalElementName   IN       VARCHAR2,
   tableName           IN       VARCHAR2,
   overwrite           IN       BOOLEAN   DEFAULT TRUE);
```

Parameters

Table 233-21 SETDEFAULTTABLE Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`globalElementName`	Name of the global element in the schema
`tableName`	Name being assigned to the table
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

233.3.21 SETOUTOFLINE Procedure

This procedure sets the `SQLInline` attribute to `FALSE`, that is, it sets `xdb:SQLInLine=FALSE`.

This forces XDB to store the corresponding elements in the XML document out-of-line as rows in a separate `XMLType` table. XDB stores references to each row of the `XMLType` table in a link table that is maintained by the main table

This procedure can improve performance in some situations if the out-of-line table acts as the driver for the query. Storing elements in an out-of-line table also reduces the numbers of columns in the base table, thus avoiding `'4096 column limit'` errors during XML schema registration, when some elements have complex types with many elements.

> **Also See:**
>
> *Oracle XML DB Developer's Guide*

There are three overloads.

Syntax

Sets the `SQLInline` attribute to `FALSE`, forcing out-of-line storage for the named element.

```
DBMS_XMLSCHEMA_ANNOTATE.SETOUTOFLINE (
   xmlschema           IN OUT  XMLType,
   elementName         IN      VARCHAR2,
```

```
   elementType        IN      VARCHAR2,
   defaultTableName   IN      VARCHAR2,
   overwrite          IN      BOOLEAN DEFAULT TRUE);
```

Sets the `SQLInline` attribute to `FALSE`, forcing out-of-line storage for the element specified by its local and global name.

```
DBMS_XMLSCHEMA_ANNOTATE.SETOUTOFLINE (
   xmlschema          IN OUT  XMLType,
   globalObject       IN      VARCHAR2,
   globalObjectName   IN      VARCHAR2,
   localElementName   IN      VARCHAR2,
   defaultTableName   IN      VARCHAR2,
   overwrite          IN      BOOLEAN DEFAULT TRUE);
```

Sets the `SQLInline` attribute to `FALSE` to force out-of-line storage and sets the default table name for all references to a particular global element.

```
DBMS_XMLSCHEMA_ANNOTATE.SETOUTOFLINE (
   xmlschema         IN OUT  XMLType,
   reference         IN      VARCHAR2,
   defaultTableName  IN      VARCHAR2,
   overwrite         IN      BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-22 SETOUTOFLINE Procedure Parameters

Parameter	Description
`xmlschema`	The XML schema to be annotated.
`elementName`	The element name
`elementType`	The element type
`defaultTableName`	The name of the default table.
`globalObject`	The global object (global complex type or global element)
`globalObjectName`	The name of the global object
`localElementName`	The name of a local element that descends from the global element.
`reference`	A reference to a global element
`overwrite`	A boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

Usage Notes

After XML schema registration and before loading XML instance data, use `DBMS_XMLSTORAGE_MANAGE.SCOPEXMLREFERENCES()` to make these references scope to the out-of-line table only. This ensures better query performance later on.

Example

The following example illustrates the third overloaded method. The element comment will have an annotation similar to `xdb:defaultTable="CMMNT_DEFAULT_TABLE"`

```
DECLARE
   xml_schema xmltype;
BEGIN
SELECT OUT INTO xml_schema FROM annotation_tab;
```

```
DBMS_XMLSCHEMA_ANNOTATE.SETOUTOFLINE (xml_schema,
                                      'ipo:comment',
                                      'CMMNT_DEFAULT_TABLE');
UPDATE annotation_tab SET OUT = xml_schema;
END;
/
```

233.3.22 SETSCHEMAANNOTATATIONS Procedure

This procedure takes the annotated differences resulting from a call to `DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS` and patches them into the provided XML schema.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETSCHEMAANNOTATIONS (
   xmlschema     IN OUT xmlType,
   annotations   IN VARCHAR2);
```

Parameters

Table 233-23 SETSCHEMAANNOTATIONS Procedure Parameters

Parameter	Description
`xmlschema`	An XML schema to be patched.
`annotations`	The differences document produced by calling `DBMS_XMLSCHEMA_ANNOTATE.GETSCHEMAANNOTATIONS` on the original XML schema and an annotated XML schema.

Usage Notes

`DBMS_XMLSCHEMA_ANNOTATE.SETSCHEMAANNOTATATIONS` is not available on Oracle Database release 10.2 (only Oracle Database release 11.*x*).

> **See Also:**
>
> GETSCHEMAANNOTATIONS Function

Example

The following example illustrates `DBMS_XMLSCHEMA_ANNOTATE.SETSCHEMAANNOTATIONS` shown here and GETSCHEMAANNOTATIONS Function.

```
-- test getannotations and apply them
declare
 xml_schema xmltype;
 xml_schema2 xmltype;
 annotations xmltype;
begin
   select out into xml_schema from annotation_tab;

   -- get the annotations from the schema
   annotations := DBMS_XMLSCHEMA_ANNOTATE.getSchemaAnnotations (xml_schema);
```

```
   -- apply the annotations to the schema
   select inp into xml_schema2 from annotation_tab;

   DBMS_XMLSCHEMA_ANNOTATE.setSchemaAnnotations(xml_schema2, annotations);

   update annotation_tab t set t.out = xml_schema2;
end;
/
```

233.3.23 SETSQLCOLLTYPE Procedure

This procedure assigns a SQL type name for a collection. A collection is a global or local element with `maxOccurs>1`.

Using this procedure, XDB creates `SQLType`s with the user-defined names provided.

There are two overloads. The first sets the name of the SQL collection type corresponding to an XML element and the second to an XML element inside the specified complex type.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETSQLCOLLTYPE (
   xmlschema          IN OUT   XMLTYPE,
   elementName        IN       VARCHAR2,
   sqlCollType        IN       VARCHAR2,
   overwrite          IN       BOOLEAN DEFAULT TRUE);

DBMS_XMLSCHEMA_ANNOTATE.SETSQLCOLLTYPE (
   xmlschema          IN OUT XMLType,
   globalObject       IN VARCHAR2,
   globalObjectName   IN VARCHAR2,
   localElementName   IN VARCHAR2,
   sqlCollType        IN VARCHAR2,
   overwrite          IN BOOLEAN default TRUE );
```

Parameters

Table 233-24 SETSQLCOLLTYPE Procedure Parameters

Parameter	Description
`xmlschema`	The XML schema to be annotated
`elementName`	The element name
`sqlCollType`	The SQL collection type
`globalObject`	The global object (global complex type or global element)
`globalObjectName`	The name of the global object
`localElementName`	The name of a local element that descends from the global element
`overwrite`	A boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

Example

The `item` element will have an annotation similar to `xdb:SQLCollType="ITEM_SQL_COL_TYPE"`.

```
declare
   xml_schema xmltype;
begin
```

```
   SELECT out INTO xml_schema FROM annotation_tab;
   DBMS_XMLSCHEMA_ANNOTATE.setSQLCollType (xml_schema,
                                           'item',
                                           'ITEM_SQL_COL_TYPE',TRUE);
   UPDATE annotation_tab SET out = xml_schema;
end;
```

233.3.24 SETSQLNAME Procedure

This procedure assigns a name to the SQL attribute that corresponds to an element defined in the XML schema.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETSQLNAME (
   xmlschema          IN OUT  XMLType,
   globalObject       IN      VARCHAR2,
   globalObjectName   IN      VARCHAR2,
   localObject        IN      VARCHAR2,
   localObjectName    IN      VARCHAR2,
   sqlName            IN      VARCHAR2,
   overwrite          IN      BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-25 SETSQLNAME Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`globalObject`	Global object (global complex type or global element)
`globalObjectName`	Name of the global object
`localObject`	Object descended from the global object
`localObjectName`	Name of the local object
`sqlName`	Name of the SQL attribute that corresponds to the element defined in the XML schema
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

Example

The `shipTo` element will have an annotation similar to `xdb:SQLName="SHIPTO_SQLNAME"`.

```
DECLARE
   xml_schema   XMLTYPE;
BEGIN
   SELECT out INTO xml_schema FROM annotation_tab;
   DBMS_XMLSCHEMA_ANNOTATE.SETSQLNAME (xml_schema,
                                       'element', 'purchaseOrder',
                                       'element', 'shipTo',
                                       'SHIPTO_SQLNAME');
   UPDATE annotation_tab SET out = xml_schema;
END;
/
```

233.3.25 SETSQLTYPE Procedure

This procedure assigns a SQL type to a global object.

There are two overloads. The first overload assigns a SQL Type to a global object, such as a global element or global complex type and the second to a local object.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETSQLTYPE (
   xmlschema            IN OUT   XMLTYPE,
   globalElementName    IN       VARCHAR2,
   sqlType              IN       VARCHAR2,
   overwrite            IN       BOOLEAN DEFAULT TRUE);


DBMS_XMLSCHEMA_ANNOTATE.SETSQLTYPE (
   xmlschema           IN OUT   XMLTYPE,
   globalObject        IN       VARCHAR2,
   globalObjectName    IN       VARCHAR2,
   localObject         IN       VARCHAR2,
   localObjectName     IN       VARCHAR2,
   sqlType             IN       VARCHAR2,
   overwrite           IN       BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-26 SETSQLTYPE Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`globalObject`	Global object (global complex type or global element)
`globalObjectName`	Name of the global object
`globalElementName`	Name of the global element
`localObject`	Object descended from the global object
`localObjectName`	Name of the local object
`sqlType`	SQL type assigned to the named global element
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

Example

The `purchaseOrder` element will have an annotation similar to `xdb:SQLType="PO_SQLTYPE"` and the `shipTo` element has one similar to `xdb:SQLType="VARCHAR"`.

```
DECLARE
   xml_schema  xmltype;
BEGIN
  SELECT out INTO xml_schema FROM annotation_tab;
  DBMS_XMLSCHEMA_ANNOTATE.setSQLType (xml_schema,
                                      'purchaseOrder',
                                      'PO_SQLTYPE');
  UPDATE annotation_tab SET out = xml_schema;
END;
/
```

```
DECLARE   xml_schema   xmltype;BEGIN  SELECT out INTO xml_schema FROM annotation_tab;
DBMS_XMLSCHEMA_ANNOTATE.setSQLType (xml_schema,
'element','purchaseOrder',
'element' ,'shipTo',                                          'VARCHAR');  UPDATE
annotation_tab SET out = xml_schema;END;
/
```

233.3.26 SETSQLTYPEMAPPING Procedure

This procedure defines a mapping of schema type and SQL type.

If you use this procedure, you do not need to call the `SETSQLTYPE` procedure on all instances of the schema type; instead the procedure traverses the schema and assigns the SQL type automatically.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETSQLTYPEMAPPING (
   xmlschema         IN OUT   XMLType,
   schemaTypeName    IN       VARCHAR2,
   sqlTypeName       IN       VARCHAR2,
   overwrite         IN       BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-27 SETSQLTYPEMAPPING Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`schemaTypeName`	Schema type
`sqlTypeName`	Name of the SQL type
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`

Example

The attribute `orderDate` will have an annotation similar to `xdb:SQLType="DATE"`.

```
declare   xml_schema xmltype;beginSELECT out INTO xml_schema FROM
annotation_tab;DBMS_XMLSCHEMA_ANNOTATE.setSQLTypeMapping
(xml_schema,
'date',                                          'DATE');UPDATE annotation_tab SET out =
xml_schema;end;
/
```

233.3.27 SETTABLEPROPS Procedure

This procedure specifies properties in the `TABLE` storage clause that is appended to the default `CREATE TABLE` statement.

There are two overloads with different parameter requirements, as indicated:

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETTABLEPROPS (
   xmlschema            IN OUT  XMLType,
```

```
   globalElementName    IN      VARCHAR2,
   tableProps           IN      VARCHAR2,
   overwrite            IN      BOOLEAN DEFAULT TRUE);

DBMS_XMLSCHEMA_ANNOTATE.SETTABLEPROPS (
   xmlschema            IN OUT  XMLTYPE,
   globalObject         IN      VARCHAR2,
   globalObjectName     IN      VARCHAR2,
   localElementName     IN      VARCHAR2,
   tableProps           IN      VARCHAR2,
   overwrite            IN      BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-28 SETTABLEPROPS Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`globalElementName`	Name of the global element in the schema
`tableProps`	Table properties
`globalObject`	Global object (global complex type or global element)
`globalObjectName`	Name of the global object
`localElementName`	Name of a local element that descends from the global element
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

Example

The `purchaseOrder` element will have an annotation similar to `xdb:tableProps="CACHE"`.

```
DECLARE   xml_schema XMLTYPE;BEGIN   SELECT out INTO xml_schema FROM annotation_tab;
DBMS_XMLSCHEMA_ANNOTATE.SETTABLEPROPS(xml_schema,
 'purchaseOrder' ,'CACHE');   UPDATE annotation_tab SET out = xml_schema;END;
/
```

233.3.28 SETTIMESTAMPWITHTIMEZONE Procedure

This procedure sets the `TIMESTAMPWITHTIMEZONE` datatype to all `dateTime` typed elements in the XML schema.

This is equivalent to adding `xdb:SQLType="TIMESTAMP WITH TIME ZONE"` to all `dateTime` objects.

Syntax

```
DBMS_XMLSCHEMA_ANNOTATE.SETTIMESTAMPWITHTIMEZONE (
   xmlschema      IN OUT   XMLTYPE,
   overwrite      IN       BOOLEAN DEFAULT TRUE);
```

Parameters

Table 233-29 SETTIMESTAMPWITHTIMEZONE Procedure Parameters

Parameter	Description
`xmlschema`	XML schema to be annotated
`overwrite`	Boolean that indicates whether or not the procedure overwrites element attributes. The default is `TRUE`.

234
DBMS_XMLSCHEMA_UTIL

The `DBMS_XMLSCHEMA_UTIL` package provides an interface for XML schema validation.

This chapter contains the following topics:

- DBMS_XMLSCHEMA_UTIL Security Model
- Summary of DBMS_XMLSCHEMA_UTIL Subprograms

> **See Also:**
>
> - *Oracle XML DB Developer's Guide*

234.1 DBMS_XMLSCHEMA_UTIL Security Model

Applications must have the `EXECUTE` privilege on the `DBMS_XMLSCHEMA_UTIL` package

234.2 Summary of DBMS_XMLSCHEMA_UTIL Subprograms

This topic lists the `DBMS_XMLSCHEMA_UTIL` subprograms in alphabetical order and briefly describes them.

Table 234-1 *DBMS_XMLSCHEMA_UTIL Package Subprograms*

Subprogram	Description
CONFORMING Function	Validates an XML document against a schema
VALIDATE Procedure	Validates an XML document against a schema

234.2.1 CONFORMING Function

This function is used to validate an XML document against a schema.

Syntax

```
DBMS_XMLSCHEMA_UTIL.CONFORMING (
   doc             IN NUMBER,
   sch             IN NUMBER)
RETURN NUMBER;
```

Parameters

Table 234-2 CONFORMING Function Parameters

Parameter	Description
doc	The XML instance
sch	The XML schema

Function Exceptions

The exceptions of `DBMS_XMLSCHEMA_UTIL.CONFORMING` function are as follows:

- The function returns zero if the schema is legal and the document conforms to the schema; otherwise it returns an LSX error code from the schema validation.

> **Note:**
>
> The function takes an xml data instance and xml schema instance. The user does not have to register the schema.

Example: The document conforms to the schema

```
SELECT DBMS_XMLSCHEMA_UTIL.CONFORMING(
     XMLType('<A/>'),
     XMLType(
    '<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
      <xs:element name="A" type="xs:string"/>
     </xs:schema>')) "LSX CODE"
    FROM DUAL;
```

234.2.2 VALIDATE Procedure

This procedure is used to validate an XML document against a schema.

Syntax

```
DBMS_XMLSCHEMA_UTIL.VALIDATE (
   doc             IN XMLTYPE,
   sch             IN XMLTYPE);
```

Parameters

Table 234-3 VALIDATE Procedure Parameters

Parameter	Description
doc	The XML instance
sch	The XML schema

Procedure Exceptions

The exceptions of `DBMS_XMLSCHEMA_UTIL.VALIDATE` procedure are as follows:

- The procedure raises ORA-31154 if either the schema is not legal, or the document does not conform to the schema.

> **Note:**
>
> The procedure takes an xml data instance and xml schema instance. The user does not have to register the schema.

Example: The document conforms to the schema

```
BEGIN
      DBMS_XMLSCHEMA_UTIL.VALIDATE(
         XMLType('<A/>'),
         XMLType(
                 '<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
                  <xs:element name="A" type="xs:string"/>
                  </xs:schema>'));
END;
/
```

235
DBMS_XMLSTORAGE_MANAGE

The DBMS_XMLSTORAGE_MANAGE package provides an interface to manage and modify XML storage after schema registration has been completed.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of DBMS_XMLSTORAGE_MANAGE Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

235.1 DBMS_XMLSTORAGE_MANAGE Overview

DBMS_XMLSTORAGE_MANAGE contains procedures to manage and modify XML storage after schema registration has been completed.

Use subprograms from this package to improve the performance of bulk load operations. You can disable indexes and constraints before doing a bulk load process and to enable them afterwards.

235.2 DBMS_XMLSTORAGE_MANAGE Security Model

Owned by XDB, the DBMS_XMLSTORAGE_MANAGE package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC.

Subprograms in this package are executed using the privileges of the current user.

235.3 Summary of DBMS_XMLSTORAGE_MANAGE Subprograms

This table lists and describes the DBMS_XML_STORAGE package subprograms.

Table 235-1 DBMS_XMLSTORAGE_MANAGE Package Subprograms

Subprogram	Description
DISABLEINDEXESANDCONSTRAINTS Procedure	Disables the indexes and constraints for XMLType tables and XMLType columns
ENABLEINDEXESANDCONSTRAINTS Procedure	Rebuilds all indexes and enables the constraints on an XMLType table including its child tables and out-of-line tables

Table 235-1 (Cont.) DBMS_XMLSTORAGE_MANAGE Package Subprograms

Subprogram	Description
EXCHANGEPOSTPROC Procedure	Enable constraints after exchange partition
EXCHANGEPREPROC Procedure	Disable constraints before exchange partition
INDEXXMLREFERENCES Procedure	Creates unique indexes on the `REF` columns of the given XML type table or the XML type column of a given table
REFPARTITIONEXCHANGEIN Procedure	This is an auxiliary procedure to load data through exchange partition operation into a partitioned table and its reference-partitioned child table provided the child table has an `xmltype` column with a local `xmlindex`.
REFPARTITIONEXCHANGEOUT Procedure	Auxiliary procedure to load data through exchange partition operation out of a partitioned table and its reference-partitioned child table provided that the child table has an `xmltype` column with a local `xmlindex`.
RENAMECOLLECTIONTABLE Procedure	Renames a collection table to the given table name
SCOPEXMLREFERENCES Procedure	Scopes all XML references. Scoped `REF` types require less storage space and allow more efficient access than unscoped `REF` types
XPATH2TABCOLMAPPING Function	Maps a path expression (in XPath notation or DOT notations) to the corresponding table name and column name

235.3.1 DISABLEINDEXESANDCONSTRAINTS Procedure

This procedure disables the indexes and constraints for `XMLType` tables and `XMLType` columns.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.DISABLEINDEXESANDCONSTRAINTS (
   owner_name    IN  VARCHAR2 DEFAULT USER,
   table_name    IN  VARCHAR2,
   column_name   IN  VARCHAR2 DEFAULT NULL,
   clear         IN  BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 235-2 DISABLEINDEXESANDCONSTRAINTS Procedure Parameters

Parameter	Description
`owner_name`	Owner's name
`table_name`	Name of the `XMLType` table that the procedure is being performed on
`column_name`	`XMLType` column name
`clear`	Boolean that when set to `TRUE` clears all stored index and constraint data for the table before the procedure executes. The default is `FALSE`, which does not clear them.

Usage Notes

Passing XMLTYPE tables

For XMLType tables, you must pass the XMLType table name on which the bulk load operation is to be performed. For XMLType columns, you must pass the relational table name and the corresponding XMLType column name.

Using clear to Enable and Disable Indexes and Constraints

> **Note:**
>
> If the DISABLEINDEXESANDCONTRAINTS procedure is called with clear set to TRUE, it removes any index or constraint information about the XMLTYPE table or column memorized during earlier executions of the procedure.

Therefore, you must ensure that all disabled indexes and constraints are re-enabled on the table or column before you call the DISABLEINDEXESANDCONTRAINTS procedure with clear set to TRUE.

Ideally, it is recommended that you set clear set to TRUE for the first execution. For any subsequent executions (due to errors while disabling or enabling indexes) clear should be set to FALSE, the default value. Once you have successfully re-enabled all the indexes and constraints following the bulk load operation, you can call this procedure again with clear set to TRUE for the next bulk load operation.

Example

The following example illustrates the use of clear in the DISABLEINDEXESANDCONSTRAINTS procedure and the ENABLEINDEXESANDCONSTRAINTS Procedure.

First, add a not-NULL constraint on comment element of the PURCHASEORDER_TAB table:

```
ALTER TABLE PURCHASEORDER_TAB ADD CONSTRAINT c1 check            ("XMLDATA"."comment" IS
NOT NULL);
```

Then, disable all the indexes and constraints by passing the clear as TRUE, by calling the DISABLEINDEXESANDCONSTRAINTS procedure:

```
BEGIN
     XDB.DBMS_XMLSTORAGE_MANAGE.DISABLEINDEXESANDCONSTRAINTS
                                ( USER,'PURCHASEORDER_TAB',NULL,TRUE );
END;
/
```

Next, perform a bulk load operation (such as datapump import) which violates constraint c1 in the ALTER table statement. This does not raise an error because the constraint is disabled:

```
host impdp orexample/orexample directory=dir dumpfile=dmp.txt
   tables=OREXAMPLE.PURCHASEORDER_TAB content = DATA_ONLY;
```

NOTE: To view the disabled constraints and indexes use:

```
SELECT constraint_name,table_name,status FROM all_constraints
  WHERE owner = user;
```

Finally, try to enable the constraint using the ENABLEINDEXESANDCONSTRAINTS procedure. It raises an error because c1, the not null constraint, is violated by the bulk load operation:

```
BEGIN
     XDB.DBMS_XMLSTORAGE_MANAGE.ENABLEINDEXESANDCONSTRAINTS
                                          ( USER,'PURCHASEORDER_TAB');
END;
/
```

To disable all the indexes and constraints, again use `DISABLEINDEXESANDCONSTRAINTS`, but set `clear= FALSE` (because the `ENABLEINDEXESANDCONSTRAINTS` failed to complete successfully). Note: `clear = FALSE` by default, so we do not need to pass it explicitly in the next call.

```
BEGIN
     xdb.DBMS_XMLSTORAGE_MANAGE.DISABLEINDEXESANDCONSTRAINTS
                                ( USER,'PURCHASEORDER_TAB');
END;
/
```

Then, delete the incorrect rows entered into the table

```
DELETE FROM purchaseorder_tab p
     WHERE p.xmldata."comment" IS NULL;
```

Re-enable the indexes and constraints using `ENABLEINDEXESANDCONSTRAINTS`, which completes successfully.

```
BEGIN
     xdb.DBMS_XMLSTORAGE_MANAGE.ENABLEINDEXESANDCONSTRAINTS
                                          ( USER,'PURCHASEORDER_TAB');
END;
/
```

235.3.2 ENABLEINDEXESANDCONSTRAINTS Procedure

This procedure rebuilds all indexes and enables the constraints on an `XMLType` table including its child tables and out-of-line tables.

When `column_name` is passed, it does the same for this `XMLType` column.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.ENABLEINDEXESANDCONSTRAINTS (
   owner_name   IN VARCHAR2 DEFAULT USER,
   table_name   IN VARCHAR2,
   column_name  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 235-3 ENABLEINDEXESANDCONSTRAINTS Procedure Parameters

Parameter	Description
`owner_user`	Owner's name
`table_name`	Name of the table that the indexes and constraints are being removed from
`column_name`	Column name

Usage Notes

This procedure reverses DISABLEINDEXESANDCONSTRAINTS Procedure.

Example

See DISABLEINDEXESANDCONSTRAINTS Procedure

235.3.3 EXCHANGEPOSTPROC Procedure

This procedure enable constraints after exchange partition.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.EXCHANGEPOSTPROC (
   owner_name   IN VARCHAR2 DEFAULT USER,
   table_name   IN VARCHAR2);
```

Parameters

Table 235-4 EXCHANGEPOSTPROC Procedure Parameters

Parameter	Description
owner_user	Owner's name
table_name	Name of the table that the indexes and constraints are being removed from

235.3.4 EXCHANGEPREPROC Procedure

This procedure disable constraints before exchange partition.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.EXCHANGEPREPROC (
   owner_name   IN VARCHAR2 DEFAULT USER,
   table_name   IN VARCHAR2);
```

Parameters

Table 235-5 EXCHANGEPREPROC Procedure Parameters

Parameter	Description
owner_user	Owner's name
table_name	Name of the table that the indexes and constraints are being removed from

235.3.5 INDEXXMLREFERENCES Procedure

This procedure creates unique indexes on the `REF` columns of the given XML type table or the XML type column of a given table.

If the procedure creates multiple `REF` columns, it appends `_1`, `_2`, and so on to their names.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.INDEXXMLREFERENCES (
   owner_name     IN VARCHAR2 DEFAULT USER,
   table_name     IN VARCHAR2,
   column_name    IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 235-6 INDEXXMLREFERENCES Procedure Parameters

Parameter	Description
owner_name	The owner's name
table_name	The table being indexed
column_name	A column name. Not needed for XML type tables.
index_name	The name of the newly created index

Usage Notes

This procedure is only used if the REFs are scoped. See SCOPEXMLREFERENCES Procedure.

Indexed REFs lead to better performance when joins between the base table and a child table occur in the query plan.

- If the base table has a higher selectivity than the child table, there is no need to index the REFs.
- If the selectivity of the child table is higher than that of the base table and if no indexes are present, then the join of one row in the child table with the base table leads to a full table scan of the base table.

INDEXXMLREFERENCES does not index REFs recursively in child tables of a table it is called on. To do this, Oracle recommends calling the procedure from within a loop over the XML_OUT_OF_LINE_TABLES or XML_NESTED_TABLES view. This creates the index names from the current value of a column in the view.

> **Note:**
>
> This procedure is limited to the structured storage model.

235.3.6 REFPARTITIONEXCHANGEIN Procedure

This is an auxiliary procedure to load data through exchange partition operation into a partitioned table and its reference-partitioned child table provided that the child table has an xmltype column with a local xmlindex.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.REFPARTITIONEXCHANGEIN (
   owner_name                IN VARCHAR2,
   parent_table_name         IN VARCHAR2,
```

```
   child_table_name           IN VARCHAR2,
   parent_exchange_table_name IN VARCHAR2,
   child_exchange_table_name  IN VARCHAR2,
   parent_exchange_stmt       IN CLOB,
   child_exchange_stmt        IN CLOB);
```

Parameters

Table 235-7 REFPARTITIONEXCHANGEIN Parameters

Parameter	Description
owner_name	owner's name
parent_table_name	the partitioned base table
child_table_name	a partitioned table with reference partitioning based on the table named parent_table_name
parent_exchange_table_name	an exchange table for the partitioned base table
child_exchange_table_name	an exchange table for the table named child_table_name
parent_exchange_stmt	SQL statement to execute exchange partition operation between the table named parent_table_name and the table named parent_exchange_table_name
child_exchange_stmt	SQL statement to execute exchange partition operation between the table named child_table_name and the table named child_exchange_table_name

235.3.7 REFPARTITIONEXCHANGEOUT Procedure

This is an auxiliary procedure to load data through exchange partition operation out of a partitioned table and its reference-partitioned child table provided that the child table has an xmltype column with a local xmlindex.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.REFPARTITIONEXCHANGEOUT (
   owner_name                 IN VARCHAR2,
   parent_table_name          IN VARCHAR2,
   child_table_name           IN VARCHAR2,
   parent_exchange_table_name IN VARCHAR2,
   child_exchange_table_name  IN VARCHAR2,
   parent_exchange_stmt       IN CLOB,
   child_exchange_stmt        IN CLOB);
```

Parameters

Table 235-8 REFPARTITIONEXCHANGEOUT Parameters

Parameter	Description
owner_name	owner's name

Table 235-8 (Cont.) REFPARTITIONEXCHANGEOUT Parameters

Parameter	Description
parent_table_name	the partitioned base table
child_table_name	a partitioned table with reference partitioning based on the table named parent_table_name
parent_exchange_table_name	an exchange table for the partitioned base table
child_exchange_table_name	an exchange table for the table named child_table_name
parent_exchange_stmt	SQL statement to execute exchange partition operation between the table named parent_table_name and the table named parent_exchange_table_name
child_exchange_stmt	SQL statement to execute exchange partition operation between the table named child_table_name and the table named child_exchange_table_name

235.3.8 RENAMECOLLECTIONTABLE Procedure

This procedure renames a collection table to the given table name.

An XPath expression specifies the collection table, starting from the XMLtype base table or an XMLType column of the base table.

This procedure provides the only way to derive a collection table name from the corresponding collection type name because there is no direct schema annotation for the purpose.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.RENAMECOLLECTIONTABLE (
   owner_name             IN VARCHAR2 DEFAULT USER,
   table_name             IN VARCHAR2,
   column_name            IN VARCHAR2 DEFAULT NULL,
   xpath                  IN VARCHAR2,
   collection_table_name  IN VARCHAR2
   namespaces             IN VARCHAR2 default NULL);  // For release 11.2 only
```

Parameters

Table 235-9 RENAMECOLLECTIONTABLE Procedure Parameters

Parameter	Description
owner_name	The name of the owner
table_name	The name of a base table that can be used as the starting point for specifying the collection table
column_name	An XMLType column that can be the starting point for specifying the collection table
xpath	The XPath expression that specifies the collection table
collection_table_name	The name of the collection table

Table 235-9 (Cont.) RENAMECOLLECTIONTABLE Procedure Parameters

Parameter	Description
`namespaces`	For Oracle Database 11*g* Release 2 (11.2) and higher. The namespaces used in XPath.

Usage Notes

Call this procedure after registering the XML schema.

The table name serves as a prefix to the index names.

Oracle recommends using this function because it makes query execution plans more readable.

Report errors that occur while this procedure runs to the user that called the procedure.

> **Note:**
>
> This procedure is limited to the structured storage model.

For Oracle Database 11*g* Release 2 (11.2) and higher, only, this function accepts XPath notation as well as DOT notation. If XPath notation is used, a `namespaces` parameter may also be required.

Example

The collection table name will be `EMP_TAB_NAMELIST`. You can verify this using `SELECT * FROM user_nested_tables`.

Using DOT Notation:

```
call  XDB.DBMS_XMLSTORAGE_MANAGE.RENAMECOLLECTIONTABLE (
            USER,
            'EMP_TAB',
            NULL,
            '"XMLDATA"."EMPLOYEE"."NAME"',
           ''EMP_TAB_NAMELIST);
```

Using XPath Notation:

XPath notation is available with Oracle Database 11g Release 2 (11.2) and higher.

```
call  XDB.DBMS_XMLSTORAGE_MANAGE.RENAMECOLLECTIONTABLE (
               USER,
               'EMP_TAB',
               NULL,
               '/e:Employee/Name',
               'EMP_TAB_NAMELIST',
               '''http://www.oracle.com/emp.xsd'' as "e"');
```

235.3.9 SCOPEXMLREFERENCES Procedure

This procedure scopes all XML references. Scoped `REF` types require less storage space and allow more efficient access than unscoped `REF` types.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.SCOPEXMLREFERENCES;
```

Usage Notes

- If you have used SETOUTOFLINE Procedure in the DBMS_XMLSTORAGE_MANAGE package to avoid raising `'4096 column limit'` errors during XML schema registration, you should also use SCOPEXMLREFERENCES Procedure.
- Using `SCOPEXMLREFERENCES` after XML schema registration and before loading XML instance data, makes these reference scoped to the out-of-line table only.

> **Note:**
>
> This procedure is limited to the structured storage model.

235.3.10 XPATH2TABCOLMAPPING Function

This function maps a path expression (in XPath notation or DOT notations) to the corresponding table name and column name. This is necessary in cases in which the user wants to create an index on this table, or to add a constraint, or to rename a table to make query execution plans more readable.

Syntax

```
DBMS_XMLSTORAGE_MANAGE.XPATH2TABCOLMAPPING (
   owner_name    IN  VARCHAR2 DEFAULT USER,
   table_name    IN  VARCHAR2,
   column_name   IN  VARCHAR2 DEFAULT NULL,
   xpath         IN  VARCHAR2,
   namespaces    IN  VARCHAR2 DEFAULT NULL)
 RETURN XMLTYPE;
```

Parameters

Table 235-10 XPATH2TABCOLMAPPING Procedure Parameters

Parameter	Description
`owner_user`	Owner's name
`table_name`	Name of the base table
`column_name`	Optional name of the XML type column if `table_name` is not an XMLtype table. If table_name refers to XMLType table then `column_name` should be NULL.
`xpath`	Path expression in DOT notation or XPath notation (see examples below)

Table 235-10 (Cont.) XPATH2TABCOLMAPPING Procedure Parameters

Parameter	Description
namespaces	Optional namespace definitions for path expression

Examples

XPath2TablColMapping evaluated on XMLType table with Xpath Notation, namespaces provided

```
SELECT XDB.DBMS_XMLSTORAGE_MANAGE.XPATH2TABCOLMAPPING (
   USER, 'XML_TAB', '', '//n1:item/n1:location',''''xdbXmark'' as "n1"'')
FROM DUAL;
```

This produces a result, for example:

```
<Result>
<Mapping TableName="SYS_NT12345" ColumnName="location"/>
</Result>
```

This allows us to define an index or constraint on table `SYS_NT12345` and column `location`.

XPath2TablColMapping evaluated on table not of XMLType but with XMLType column by means of DOT notation

```
SELECT XDB.DBMS_XMLSTORAGE_MANAGE.XPATH2TABCOLMAPPING (
   USER,'PurchaseOrderTab','XMLCOL','xmldata.LineItems.LineItem', '')
FROM DUAL;
```

236
DBMS_XMLSTORE

DBMS_XMLSTORE provides the ability to store XML data in relational tables.

> **Note:**
>
> The PL/SQL package DBMS_XMLSTORE is deprecated in Oracle Database 23ai.
>
> DBMS_XMLSTORE is a non-standard Oracle-proprietary package that enables you to store and manipulate XML data in Oracle Database. This package is deprecated, and can be desupported in a future release. Oracle recommends that you use regular SQL DML and with standard XQuery and SQL/XML to store and manage XML data. Using standard functionality provides future-proof way to store and manipulate XML data.

This chapter contains the following sections:

- Security Model
- Types
- Summary of DBMS_XMLSTORE Subprograms

> **See Also:**
>
> *Oracle XML DB Developer's Guide*

236.1 DBMS_XMLSTORE Security Model

Owned by XDB, the DBMS_XMLSTORE package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

236.2 Types

The DBMS_XMLSTORE subprograms use the ctxType Type.

Table 236-1 ***Types of DBMS_XMLSTORE***

Type	Description
ctxType	The type of the query context handle. This is the return type of NEWCONTEXT.

236.3 Summary of DBMS_XMLSTORE Subprograms

This table lists the `DBMS_XMLSTORE` subprograms and briefly describes them.

Table 236-2 DBMS_XMLSTORE Package Subprograms

Method	Description
CLEARKEYCOLUMNLIST	Clears the key column list.
CLEARUPDATECOLUMNLIST	Clears the update column list.
CLOSECONTEXT	It closes/deallocates a particular save context.
DELETEXML	Deletes records specified by data from the XML document, from the table specified at the context creation time.
INSERTXML	Inserts the XML document into the table specified at the context creation time.
NEWCONTEXT	Creates a save context, and returns the context handle.
SETKEYCOLUMN	This method adds a column to the key column list.
SETROWTAG	Names the tag used in the XML document., to enclose the XML elements corresponding to the database.
SETUPDATECOLUMN	Adds a column to the "update column list".
UPDATEXML	Updates the table given the XML document.

236.3.1 CLEARKEYCOLUMNLIST

This procedure clears the key column list.

Syntax

```
PROCEDURE clearKeyColumnList(
   ctxHdl IN ctxType);
```

Table 236-3 CLEARKEYCOLUMNLIST Procedure Parameters

Parameter	IN / OUT	Description
ctxHdl	(IN)	Context handle.

236.3.2 CLEARUPDATECOLUMNLIST

This procedure clears the update column list.

Syntax

```
PROCEDURE clearUpdateColumnList(
   ctxHdl IN ctxType);
```

Table 236-4 CLEARUPDATECOLUMNLIST Procedure Parameters

Parameter	IN / OUT	Description
ctxHdl	(IN)	Context handle.

236.3.3 CLOSECONTEXT

This procedure closes/deallocates a particular save context.

Syntax

```
PROCEDURE closeContext(ctxHdl IN ctxType);
```

Table 236-5 CLOSECONTEXT Procedure Parameters

Parameter	IN / OUT	Description
ctxHdl	(IN)	Context handle.

236.3.4 DELETEXML

DELETEXML deletes records specified by data from the XML document from the table specified at the context creation time, and returns the number of rows deleted.

Syntax

The following syntax uses a VARCHAR2 type for the xDoc parameter.

```
FUNCTION deleteXML(
  ctxHdl IN ctxPType,
  xDoc IN VARCHAR2)
RETURN NUMBER;
```

The following syntax uses a CLOB type for the xDoc parameter.

```
FUNCTION deleteXML(
  ctxHdl IN ctxType,
  xDoc IN CLOB)
RETURN NUMBER;
```

The following syntax uses an XMLType type for the xDoc parameter.

```
FUNCTION deleteXML(
  ctxHdl IN ctxType,
  xDoc IN XMLType)
RETURN NUMBER;
```

Parameters

Table 236-6 DELETEXML Function Parameters

Parameter	IN / OUT	Description
ctxHdl	(IN)	Context handle.
xDoc	(IN)	String containing the XML document.

236.3.5 INSERTXML

Inserts the XML document into the table specified at the context creation time, and returns the number of rows inserted.

Note that if a user passes an XML file for `insertXML` to `DBMS_XMLSTORE` that contains extra elements (elements that do not match any columns in the table), Oracle tries to insert into those columns unless SETUPDATECOLUMN is used. The use of `setUpdateColumn` is optional only if the elements in the XML file match up to the columns in the table.

Syntax

```
FUNCTION insertXML(
  ctxHdl IN ctxType,
  xDoc IN VARCHAR2)
RETURN NUMBER;


FUNCTION insertXML(
  ctxHdl IN ctxType,
  xDoc IN CLOB)
RETURN NUMBER;


FUNCTION insertXML(
  ctxHdl IN ctxType,
  xDoc IN XMLType)
RETURN NUMBER;
```

Parameters

Table 236-7 INSERTXML Function Parameters

Parameter	IN / OUT	Description
`ctxHdl`	`(IN)`	Context handle.
`xDoc`	`(IN)`	String containing the XML document.

236.3.6 NEWCONTEXT

`NEWCONTEXT` creates a save context and returns the context handle.

Syntax

```
FUNCTION newContext(
   targetTable IN VARCHAR2)
RETURN ctxType;
```

Table 236-8 NEWCONTEXT Function Parameters

Parameter	IN / OUT	Description
`targetTable`	`(IN)`	The target table into which to load the XML document.

236.3.7 SETKEYCOLUMN

This method adds a column to the "key column list".

The value for the column cannot be `NULL`. In case of update or delete, the columns in the key column list make up the `WHERE` clause of the statement. The key columns list must be specified before updates can complete; this is optional for delete operations

Syntax

```
PROCEDURE setKeyColumn(
   ctxHdl IN ctxType,
   colName IN VARCHAR2);
```

Table 236-9 SETKEYCOLUMN Procedure Parameters

Parameter	IN / OUT	Description
`ctxHdl`	`(IN)`	Context handle.
`colName`	`(IN)`	Column to be added to the key column list; cannot be `NULL`.

236.3.8 SETROWTAG

This procedure names the tag used in the XML document, to enclose the XML elements corresponding to database records.

Syntax

```
PROCEDURE setRowTag(
   ctxHdl IN ctxType,
   tag IN VARCHAR2);
```

Table 236-10 SETROWTAG Procedure Parameters

Parameter	IN / OUT	Description
`ctxHdl`	`(IN)`	Context handle.
`tag`	`(IN)`	Tag name.

236.3.9 SETUPDATECOLUMN

`SETUPDATECOLUMN` adds a column to the update column list.

In case of insert, the default is to insert values to all the columns in the table. In case of updates, the default is to only update the columns corresponding to the tags present in the `ROW` element of the XML document. When the update column list is specified, the columns making up this list alone will get updated or inserted into.

Note that if a user passes an XML file for INSERTXML to `DBMS_XMLSTORE` which contains extra elements (ones that do not match up to any columns in the table), Oracle will try to insert into those columns unless setUpdateColumn is used. The use of setUpdateColumn is optional only if the elements in the XML file match up to the columns in the table.

Syntax

```
PROCEDURE setUpdateColumn(
   ctxHdl IN ctxType,
   colName IN VARCHAR2);
```

Table 236-11 SETUPDATECOLUMN Procedure Parameters

Parameter	IN / OUT	Description
ctxHdl	(IN)	Context handle.
colName	(IN)	Column to be added to the update column list.

236.3.10 UPDATEXML

Updates the table specified at the context creation time with data from the XML document, and returns the number of rows updated.

The options are described in the following table.

Syntax

The following syntax passes the `xDoc` parameter as a `VARCHAR2`.

```
FUNCTION updateXML(
  ctxHdl IN ctxType,
  xDoc IN VARCHAR2)
RETURN NUMBER;
```

The following syntax passes the `xDoc` parameter as a `CLOB`.

```
FUNCTION updateXML(
  ctxHdl IN ctxType,
  xDoc IN CLOB)
RETURN NUMBER;
```

The following syntax passes the `xDoc` parameter as a `XMLType`.

```
FUNCTION updateXML(
  ctxHdl IN ctxType,
  xDoc IN XMLType)
RETURN NUMBER;
```

Parameters

Table 236-12 UPDATEXML Function Parameters

Parameter	IN / OUT	Description
ctxHdl	(IN)	Context handle.
xDoc	(IN)	String containing the XML document.

237
DBMS_XPLAN

The `DBMS_XPLAN` package provides an easy way to display the output of the `EXPLAIN PLAN` command in several, predefined formats.

You can also use the `DBMS_XPLAN` package to display the plan of a statement stored in the Automatic Workload Repository (AWR) or stored in a SQL tuning set. It further provides a way to display the SQL execution plan and SQL execution runtime statistics for cached SQL cursors based on the information stored in the `V$SQL_PLAN` and `V$SQL_PLAN_STATISTICS_ALL` fixed views. Finally, it displays plans from a SQL plan baseline.

> **See Also:**
>
> - For more information on the `EXPLAIN PLAN` command, the AWR, and SQL tuning set, see *Oracle Database SQL Tuning Guide*.
> - For more information on the `V$SQL_PLAN` fixed view, see *Oracle Database Reference*
> - For more information on the `V$SQL_PLAN_STATISTICS` fixed view, see *Oracle Database Reference*

This chapter contains the following topics:

- Overview
- Security Model
- Examples
- Summary of DBMS_XPLAN Subprograms

237.1 DBMS_XPLAN Overview

The `DBMS_XPLAN` package supplies five table functions.

These functions are listed below:

- `DISPLAY` - to format and display the contents of a plan table.
- `DISPLAY_AWR` - to format and display the contents of the execution plan of a stored SQL statement in the AWR.
- `DISPLAY_CURSOR` - to format and display the contents of the execution plan of any loaded cursor.
- `DISPLAY_SQL_PLAN_BASELINE` - to display one or more execution plans for the SQL statement identified by SQL handle
- `DISPLAY_SQLSET` - to format and display the contents of the execution plan of statements stored in a SQL tuning set.

237.2 DBMS_XPLAN Security Model

This package runs with the privileges of the calling user, not the package owner (SYS). The table function DISPLAY_CURSOR requires SELECT or READ privileges on the following fixed views: VSQL_PLAN, VSESSION and V$SQL_PLAN_STATISTICS_ALL. This function also requires SELECT/READ permissions on V$SQL.

DISPLAY_AWR Function requires the user to have SELECT or READ privileges on DBA_HIST_SQL_PLAN, AWR_ROOT_SQL_PLAN, AWR_PDB_SQL_PLAN, AWR_ROOT_SQLTEXT, DBA_HIST_SQLTEXT, AWR_PDB_SQLTEXT, and V$DATABASE.

DISPLAY_SQLSET Function requires the user to have the SELECT or READ privilege on ALL_SQLSET_STATEMENTS and ALL_SQLSET_PLANS.

DISPLAY_SQL_PLAN_BASELINE Function requires the user to have the SELECT or READ privilege on DBA_SQL_PLAN_BASELINES as well as the privileges to execute the SQL statement for which the user is trying to get the plan.

The preceding privileges are granted automatically as part of SELECT_CATALOG_ROLE.

237.3 DBMS_XPLAN Data Structures

The DBMS_XPLAN package defines a TABLE type.

Table Types

- DBMS_XPLAN PLAN_OBJECT_LIST Table Type

237.3.1 DBMS_XPLAN PLAN_OBJECT_LIST Table Type

This type allows for a list of generic objects as input to the COMPARE_PLANS function.

Syntax

```
TYPE plan_object_list IS TABLE OF generic_plan_object;
```

The generic object abstracts the common attributes of plans from all plan sources. Every plan source is a subclass of the generic_plan_object superclass. The following table summarizes the different plan sources. Note that when an optional parameter is null, it can correspond to multiple objects. For example, if you do not specify a child number for cursor_cache_object, then it matches all cursor cache statements with the specified SQL ID.

Table 237-1 Plan Sources for PLAN_OBJECT_LIST

Plan Source	Specification	Description
Plan table	`plan_table_object(owner, plan_table_name, statement_id, plan_id)`	The parameters are as follows: • `owner`—The owner of the plan table • `plan_table_name`—The name of the plan table • `statement_id`—The ID of the statement (optional) • `plan_id`—The ID of the plan (optional)
Cursor cache	`cursor_cache_object(sql_id, child_number)`	The parameters are as follows: • `sql_id`—The SQL ID of the plan • `child_number`—The child number of the plan in the cursor cache (optional)
AWR	`awr_object(sql_id, dbid, con_dbid, plan_hash_value)`	The parameters are as follows: • `sql_id`—The SQL ID of the plan • `dbid`—The database ID (optional) • `con_dbid`—The CDB ID (optional) • `plan_hash_value`—The hash value of the plan (optional)
SQL tuning set	`sqlset_object (sqlset_owner, sqlset_name, sql_id, plan_hash_value)`	The parameters are as follows: • `sqlset_owner`—The owner of the SQL tuning set • `sqlset_name`—The name of the SQL tuning set • `sql_id`—The SQL ID of the plan • `plan_hash_value`—The hash value of the plan (optional)
SQL plan management	`spm_object (sql_handle, plan_name)`	The parameters are as follows: • `sql_handle`—The SQL handle of plans protected by SQL plan management • `plan_name`—The name of the SQL plan baseline (optional)
SQL profile	`sql_profile_object (profile_name)`	The `profile_name` parameter specifies the name of the SQL profile.
Advisor	`advisor_object (task_name, execution_name, sql_id, plan_id)`	The parameters are as follows: • `task_name`—The name of the advisor task • `execution_name`—The name of the task execution • `sql_id`—The SQL ID of the plan • `plan_id`—The advisor plan ID (optional)

237.4 Examples

These examples show sample uses of DBMS_XPLAN.

Displaying a Plan Table Using DBMS_XPLAN.DISPLAY

Execute an explain plan command on a `SELECT` statement:

```
EXPLAIN PLAN FOR
SELECT * FROM emp e, dept d
   WHERE e.deptno = d.deptno
   AND e.ename='benoit';
```

Display the plan using the `DBMS_XPLAN.DISPLAY` table function

```
SET LINESIZE 130
SET PAGESIZE 0
SELECT * FROM TABLE(DBMS_XPLAN.DISPLAY);
```

This query produces the following output:

```
Plan hash value: 3693697075
-------------------------------------------------------------------------------
| Id  | Operation          | Name | Rows  | Bytes | Cost (%CPU)| Time     |
-------------------------------------------------------------------------------
|   0 | SELECT STATEMENT   |      |     1 |    57 |     6  (34)| 00:00:01 |
|*  1 |  HASH JOIN         |      |     1 |    57 |     6  (34)| 00:00:01 |
|*  2 |   TABLE ACCESS FULL| EMP  |     1 |    37 |     3  (34)| 00:00:01 |
|   3 |   TABLE ACCESS FULL| DEPT |     4 |    80 |     3  (34)| 00:00:01 |
-------------------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------------
   1 - access("E"."DEPTNO"="D"."DEPTNO")
   2 - filter("E"."ENAME"='benoit')

15 rows selected.
```

Displaying a Cursor Execution Plan Using DBMS_XPLAN.DISPLAY_CURSOR

By default, the table function `DISPLAY_CURSOR` formats the execution plan for the last SQL statement executed by the session. For example:

```
SELECT ename  FROM  emp e, dept d
   WHERE   e.deptno = d.deptno
   AND   e.empno=7369;

ENAME
----------
SMITH
```

To display the execution plan of the last executed statement for that session:

```
SET PAGESIZE 0
SELECT * FROM DBMS_XPLAN.DISPLAY_CURSOR();
```

This query produces the following output:

```
Plan hash value: 3693697075, SQL hash value: 2096952573, child number: 0
-------------------------------------------------------------------
SELECT ename FROM emp e, dept d WHERE e.deptno = d.deptno
```

```
AND e.empno=7369

-----------------------------------------------------------------------------
| Id  | Operation          | Name | Rows  | Bytes | Cost (%CPU)| Time     |
-----------------------------------------------------------------------------
|   0 | SELECT STATEMENT   |      |       |       |            |          |
|*  1 |  HASH JOIN         |      |     1 |    16 |     6  (34)| 00:00:01 |
|*  2 |   TABLE ACCESS FULL| EMP  |     1 |    13 |     3  (34)| 00:00:01 |
|   3 |   TABLE ACCESS FULL| DEPT |     4 |    12 |     3  (34)| 00:00:01 |
-----------------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------------
   1 - access("E"."DEPTNO"="D"."DEPTNO")
   2 - filter("E"."EMPNO"=7369)

21 rows selected.
```

You can also use the table function `DISPLAY_CURSOR` to display the execution plan for any loaded cursor stored in the cursor cache. In that case, you must supply a reference to the child cursor to the table function. This includes the SQL ID of the statement and optionally the child number.

Run a query with a distinctive comment:

```
SELECT /* TOTO */ ename, dname
FROM dept d join emp e USING (deptno);
```

Get `sql_id` and `child_number` for the preceding statement:

```
SELECT sql_id, child_number
FROM v$sql
WHERE sql_text LIKE '%TOTO%';

SQL_ID         CHILD_NUMBER
----------     ------------------------------
gwp663cqh5qbf   0
```

Display the execution plan for the cursor:

```
SELECT * FROM DBMS_XPLAN.DISPLAY_CURSOR('gwp663cqh5qbf',0);

Plan hash value: 3693697075, SQL ID: gwp663cqh5qbf, child number: 0
-------------------------------------------------------------
SELECT /* TOTO */ ename, dname
FROM dept d JOIN emp e USING (deptno);

-----------------------------------------------------------------------------
| Id  | Operation           | Name | Rows  | Bytes | Cost (%CPU)| Time     |
-----------------------------------------------------------------------------
|   0 | SELECT STATEMENT    |      |       |       |     7 (100)|          |
|   1 |  SORT GROUP BY      |      |     4 |    64 |     7  (43)| 00:00:01 |
|*  2 |   HASH JOIN         |      |    14 |   224 |     6  (34)| 00:00:01 |
|   3 |    TABLE ACCESS FULL| DEPT |     4 |    44 |     3  (34)| 00:00:01 |
|   4 |    TABLE ACCESS FULL| EMP  |    14 |    70 |     3  (34)| 00:00:01 |
-----------------------------------------------------------------------------

Predicate Information (identified by operation id):
---------------------------------------------------
   2 - access("E"."DEPTNO"="D"."DEPTNO")
```

Instead of issuing two queries, one to the get the sql_id and child_number pair and one to display the plan, you can combine these in a single query:

Display the execution plan of all cursors matching the string 'TOTO':

```
SELECT t.*
FROM v$sql s, DBMS_XPLAN.DISPLAY_CURSOR(s.sql_id, s.child_number) t WHERE sql_text LIKE
'%TOTO%';
```

Displaying a Plan Table with Parallel Information

By default, only relevant information is reported by the display and `display_cursor` table functions. In Displaying a Plan Table Using DBMS_XPLAN.DISPLAY, the query does not execute in parallel. Hence, information related to the parallelization of the plan is not reported. As shown in the following example, parallel information is reported only if the query executes in parallel.

```
ALTER TABLE emp PARALLEL;
EXPLAIN PLAN for
SELECT * FROM emp e, dept d
   WHERE e.deptno = d.deptno
   AND e.ename    ='hermann'
   ORDER BY e.empno;
```

Display the plan using the `DBMS_XPLAN.DISPLAY` table function

```
SET LINESIZE 130
SET PAGESIZE 0
SELECT * FROM DBMS_XPLAN.DISPLAY();
Plan hash value: 3693697345
```

```
-----------------------------------------------------------------------------------------------------------------
| Id | Operation          | Name    | Rows | Bytes | Cost (%CPU)| Time     | TQ    |INOUT |PQ Distrib   |
-----------------------------------------------------------------------------------------------------------------
| 0  | SELECT STATEMENT   |         | 1    | 117   | 6 (50)     | 00:00:01 |       |      |             |
| 1  | PX COORDINATOR     |         |      |       |            |          |       |      |             |
| 2  | PX SEND QC (ORDER) |:TQ10003 | 1    | 117   | 6 (50)     | 00:00:01 | Q1,03 | P->S | QC (ORDER)  |
| 3  | SORT ORDER BY      |         | 1    | 117   | 6 (50)     | 00:00:01 | Q1,03 | PCWP |             |
| 4  | PX RECEIVE         |         | 1    | 117   | 5 (40)     | 00:00:01 | Q1,03 | PCWP |             |
| 5  | PX SEND RANGE      |:TQ10002 | 1    | 117   | 5 (40)     | 00:00:01 | Q1,02 | P->P | RANGE       |
|* 6 | HASH JOIN          |         | 1    | 117   | 5 (40)     | 00:00:01 | Q1,02 | PCWP |             |
| 7  | PX RECEIVE         |         | 1    | 87    | 2 (50)     | 00:00:01 | Q1,02 | PCWP |             |
| 8  | PX SEND HASH       |:TQ10001 | 1    | 87    | 2 (50)     | 00:00:01 | Q1,01 | P->P | HASH        |
| 9  | PX BLOCK ITERATOR  |         | 1    | 87    | 2 (50)     | 00:00:01 | Q1,01 | PCWC |             |
|* 10| TABLE ACCESS FULL  | EMP     | 1    | 87    | 2 (50)     | 00:00:01 | Q1,01 | PCWP |             |
| 11 | BUFFER SORT        |         |      |       |            |          | Q1,02 | PCWC |             |
| 12 | PX RECEIVE         |         | 4    | 120   | 3 (34)     | 00:00:01 | Q1,02 | PCWP |             |
| 13 | PX SEND HASH       |:TQ10000 | 4    | 120   | 3 (34)     | 00:00:01 |       | S->P | HASH        |
| 14 | TABLE ACCESS FULL  | DEPT    | 4    | 120   | 3 (34)     | 00:00:01 |       |      |             |
-----------------------------------------------------------------------------------------------------------------
Predicate Information (identified by operation id):
---------------------------------------------------
6 - access("E"."DEPTNO"="D"."DEPTNO")
10 - filter("E"."ENAME"='hermann')
---------------------------------------------------
```

When the query is parallel, information related to parallelism is reported: table queue number (`TQ` column), table queue type (`INOUT`) and table queue distribution method (`PQ Distrib`).

By default, if several plans in the plan table match the `statement_id` parameter passed to the display table function (default value is `NULL`), only the plan corresponding to the last `EXPLAIN PLAN` command is displayed. Hence, there is no need to purge the plan table after each `EXPLAIN PLAN`. However, you should purge the plan table regularly to ensure good performance in the execution of the `DISPLAY` table function. If no plan table is created, Oracle uses a global

temporary table to store any plan information for individual users and preserves its content throughout the lifespan of a session. Note that you cannot truncate the content of a global temporary table.

For ease of use, you can define a view on top of the display table function and then use that view to display the output of the `EXPLAIN PLAN` command:

Using a View to Display Last Explain Plan

```
# define plan view
CREATE VIEW PLAN AS SELECT * FROM TABLE(DBMS_XPLAN.DISPLAY);

# display the output of the last explain plan command
SELECT * FROM PLAN;
```

237.5 Summary of DBMS_XPLAN Subprograms

This table lists the DBMS_XPLAN subprograms and briefly describes them.

Table 237-2 ***DBMS_XPLAN Package Subprograms***

Subprogram	Description
COMPARE_PLANS Function	Compares each plan in a list with a reference plan and returns the report
DIFF_PLAN Function	Compares plans
DISPLAY Function	Displays the contents of the plan table
DISPLAY_AWR Function	Displays the contents of an execution plan stored in the AWR
DISPLAY_CURSOR Function	Displays the execution plan of any cursor in the cursor cache
DISPLAY_PLAN Function	Displays the contents of the plan table in a variety of formats with `CLOB` output type
DISPLAY_SQL_PLAN_BASELINE Function	Displays one or more execution plans for the specified SQL handle of a SQL plan baseline
DISPLAY_SQLSET Function	Displays the execution plan of a given statement stored in a SQL tuning set

237.5.1 COMPARE_PLANS Function

This function compares each plan in a list with a reference plan and returns the report.

Syntax

```
DBMS_XPLAN.COMPARE_PLANS(
   reference_plan    IN generic_plan_object,
   compare_plan_list IN plan_object_list,
   type              IN VARCHAR2 := 'TEXT',
   level             IN VARCHAR2 := 'TYPICAL',
   section           IN VARCHAR2 := 'ALL')
 RETURN CLOB;
```

Parameters

Table 237-3 COMPARE_PLANS Function Parameters

Parameter	Description
reference_plan	The reference plan. This plan should always evaluate to a single plan.
compare_plan_list	List of plans to compare with reference plan. The compare_plan_list is a list of generic_object and each generic_object could correspond to one or more plans.
type	Type of the report. Possible values are: • TEXT • HTML • XML
level	Format of the report. Possible values are: • BASIC • TYPICAL • ALL
section	A particular section in the report. Possible values are: • SUMMARY • FINDINGS • PLANS • INFORMATION • ERRORS

Example 237-1 Examples

The following examples illustrate the usage of COMPARE_PLANS Function.

```
var report clob;
exec :report := dbms_xplan.compare_plans(cursor_cache_object('8mkxm7ur07za0', 2),
                - plan_object_list (cursor_cache_object('8mkxm7ur07za0', 4)));

print report
```

The above example compares the plan of child cursor number 2 for the SQL ID '8mkxm7ur07za0' with that of the child cursor number 4 for the same SQL ID. Returns the report in the text format (default).

```
var report clob;
exec :report := dbms_xplan.compare_plans(cursor_cache_object('8mkxm7ur07za0', 2),
                plan_object_list( spm_object('SQL_024d0f7d21351f5d',
'SQL_PLAN_sdfjkd')));

print report
```

The above example compares the plan of child cursor number 2 for the SQL ID '8mkxm7ur07za0' with that of the plan baseline captured by SPM for query whose SQL handle is 'SQL_024d0f7d21351f5d' and plan name is 'SQL_PLAN_sdfjkd'. Returns the report in the text format (default).

```
var report clob;
exec :report = dbms_xplan.compare_plans(cursor_cache_object('8mkxm7ur07za0', 2),
```

```
              plan_object_list( cursor_cache_object('8mkxm7ur07za0'),
sqlset_object('SH', 'SQLT_WORKLOAD', '               6vfqvav0rgyad'),
awr_object('6vfqvav0rgyad', 5), spm_object('SQL_024d0f7d21351f5d',
'SQL_PLAN_               sdfjkd'), plan_table_object('SH', 'plan_table', NULL, 38),
sql_profile_object('pe3r3ejsfd'), a              dvisor_object('TASK_1228',
'EXEC_1928', '8mkxm7ur07za0')), type => 'XML',
              level => 'ALL', section => 'SUMMARY');
```

The above example compares the plan of child cursor number 2 for the SQL ID '8mkxm7ur07za0' with each of the plans in the following list:

- cursor_cache_object('8mkxm7ur07za0'): All the plans in the cursor cache that are generated for the SQL ID '8mkxm7ur07za0'.
- sqlset_object('SH', 'SQLT_WORKLOAD', '6vfqvav0rgyad'): All the plans generated in the SQL tuning set SH. SQLT_WORKLOAD for the SQL ID '6vfqvav0rgyad'.
- awr_object('6vfqvav0rgyad', 5): All the plans in AWR that are captured for database ID 5 and SQL ID '6vfqvav0rgyad'.
- spm_object ('SQL_024d0f7d21351f5d, 'SQL_PLAN_sdfjkd'): The plan baseline for the query with SQL handle 'SQL_024d0f7d21351f5d' with name 'SQL_PLAN_sdfjkd'.
- plan_table_object('SH', 'plan_table', NULL, 38): Plan stored in SH.plan_table identified by plan_id=38.
- sql_profile_object('pe3r3ejsfd'): Plan identified by the SQL profile name 'pe3r3ejsfd'.
- advisor_object('TASK_1228', 'EXEC_1928', '8mkxm7ur07za0')): All the plans stored in SQL advisor identified by task name 'TASK_1228', execution name 'EXEC_1928' and SQL ID '8mkxm7ur07za0'.

237.5.2 DIFF_PLAN Function

This function compares two sql plans, the reference plan and the target plan. This function returns a `task_id` that can be used to retrieve the report of findings.

Syntax

```
DBMS_XPLAN.DIFF_PLAN(
   sql_text     IN   CLOB,
   outline      IN   CLOB,
   user_name    IN   VARCHAR2 := 'NULL')
 RETURN VARCHAR2;
```

Parameters

Table 237-4 DIFF_PLAN Function Parameters

Parameter	Description
`sql_text`	The text of the SQL statement.
`outline`	Used to generate the target plan.
`user_name`	The parsing schema name default to current user.

237.5.3 DISPLAY Function

This table function displays the contents of the plan table.

In addition, you can use this table function to display any plan (with or without statistics) stored in a table as long as the columns of this table are named the same as columns of the plan table (or `V$SQL_PLAN_STATISTICS_ALL` if statistics are included). You can apply a predicate on the specified table to select rows of the plan to display.

Syntax

```
DBMS_XPLAN.DISPLAY(
   table_name     IN  VARCHAR2  DEFAULT 'PLAN_TABLE',
   statement_id   IN  VARCHAR2  DEFAULT  NULL,
   format         IN  VARCHAR2  DEFAULT  'TYPICAL',
   filter_preds   IN  VARCHAR2 DEFAULT NULL);
```

Parameters

Table 237-5 DISPLAY Function Parameters

Parameter	Description
`table_name`	Specifies the table name where the plan is stored. This parameter defaults to `PLAN_TABLE`, which is the default plan table for the `EXPLAIN PLAN` command. If `NULL` is specified it also defaults to `PLAN_TABLE`.
`statement_id`	Specifies the `statement_id` of the plan to be displayed. This parameter defaults to `NULL`, which is the default when the `EXPLAIN PLAN` command is executed without a `set statement_id` clause. If no `statement_id` is specified, the function shows you the plan of the most recent explained statement.

Table 237-5 (Cont.) DISPLAY Function Parameters

Parameter	Description
format	Controls the level of details for the plan. It accepts the following values: • BASIC: Displays the minimum information in the plan—the operation ID, the operation name and its option. • TYPICAL: This is the default. Displays the most relevant information in the plan (operation id, name and option, #rows, #bytes and optimizer cost). Pruning, parallel and predicate information are only displayed when applicable. Excludes only PROJECTION, ALIAS, and REMOTE SQL information (see below). • SERIAL: Like TYPICAL except that the parallel information is not displayed, even if the plan executes in parallel. • ALL: Maximum user level. Includes information displayed with the TYPICAL level with additional information (PROJECTION, ALIAS and information about REMOTE SQL if the operation is distributed). For finer control on the display output, the following keywords can be added to the above three standard format options to customize their default behavior. Each keyword either represents a logical group of plan table columns (such as PARTITION) or logical additions to the base plan table output (such as PREDICATE). Format keywords must be separated by either a comma or a space: • ROWS - if relevant, shows the number of rows estimated by the optimizer • BYTES - if relevant, shows the number of bytes estimated by the optimizer • COST - if relevant, shows optimizer cost information • PARTITION - if relevant, shows partition pruning information • PARALLEL - if relevant, shows PX information (distribution method and table queue information) • PREDICATE - if relevant, shows the predicate section • PROJECTION -if relevant, shows the projection section • ALIAS - if relevant, shows the "Query Block Name / Object Alias" section • REMOTE - if relevant, shows the information for distributed query (for example, remote from serial distribution and remote SQL) • NOTE - if relevant, shows the note section of the explain plan Format keywords can be prefixed by the sign '-' to exclude the specified information. For example, '-PROJECTION' excludes projection information. If the target plan table (see table_name parameter) also stores plan statistics columns (for example, it is a table used to capture the content of the fixed view V$SQL_PLAN_STATISTICS_ALL), additional format keywords can be used to specify which class of statistics to display when using the DISPLAY Function. These additional format keywords are IOSTATS, MEMSTATS, ALLSTATS, and LAST (see the DISPLAY_CURSOR Function or the DISPLAY_SQLSET Function for a full description of these four keywords).
filter_preds	SQL filter predicate(s) to restrict the set of rows selected from the table where the plan is stored. When value is NULL (the default), the plan displayed corresponds to the last executed explain plan. For example: filter_preds=>'plan_id = 10' Can reference any column of the table where the plan is stored and can contain any SQL construct (for example, sub-query, function calls (see **WARNING** under Usage Notes)

Usage Notes

Here are some ways you might use variations on the format parameter:

- Use 'ALL -PROJECTION -NOTE' to display everything except the projection and note sections.
- Use 'TYPICAL PROJECTION' to display using the typical format with the additional projection section (which is normally excluded under the typical format). Since typical is default, using simply 'PROJECTION' is equivalent.
- Use '-BYTES -COST -PREDICATE' to display using the typical format but excluding optimizer cost and byte estimates as well as the predicate section.
- Use 'BASIC ROWS' to display basic information with the additional number of rows estimated by the optimizer.

> **WARNING:**
>
> Application developers should expose the `filter_preds` parameter to end-users only after careful consideration because this could expose the application to SQL injection. Indeed, `filter_preds` can potentially reference any table or execute any server function for which the database user invoking the table function has privileges.

Examples

To display the result of the last `EXPLAIN PLAN` command stored in the plan table:

```
SELECT * FROM TABLE (DBMS_XPLAN.DISPLAY);
```

To display from other than the default plan table, "`my_plan_table`":

```
SELECT * FROM table (DBMS_XPLAN.DISPLAY('my_plan_table'));
```

To display the minimum plan information:

```
SELECT * FROM table (DBMS_XPLAN.DISPLAY('plan_table', null, 'basic'));
```

To display the plan for a statement identified by '`foo`', such as `statement_id='sales_query'`:

```
SELECT * FROM table (DBMS_XPLAN.DISPLAY('plan_table', 'sales_query'));
```

237.5.4 DISPLAY_AWR Function

This table function displays the contents of an execution plan stored in AWR.

> **Note:**
>
> This function is deprecated. Use `DISPLAY_WORKLOAD_REPOSITORY` instead. `DISPLAY_AWR` only works with snapshots for the local DBID, whereas `DISPLAY_WORKLOAD_REPOSITORY` supports all snapshots inside AWR, including remote and imported snapshots.

Syntax

```
DBMS_XPLAN.DISPLAY_AWR(
   sql_id            IN      VARCHAR2,
   plan_hash_value   IN      NUMBER DEFAULT NULL,
   db_id             IN      NUMBER DEFAULT NULL,
   format            IN      VARCHAR2 DEFAULT TYPICAL);
```

Parameters

Table 237-6 DISPLAY_AWR Table Function Parameters

Parameter	Description
sql_id	Specifies the SQL_ID of the SQL statement. You can retrieve the appropriate value for the SQL statement of interest by querying the column SQL_ID in DBA_HIST_SQLTEXT.
plan_hash_value	Specifies the PLAN_HASH_VALUE of a SQL statement. This parameter is optional. If omitted, the table function returns all stored execution plans for a given SQL_ID.
db_id	Specifies the database_id for which the plan of the SQL statement, identified by SQL_ID should be displayed. If not supplied, the database_id of the local database is used, as shown in V$DATABASE.

Table 237-6 (Cont.) DISPLAY_AWR Table Function Parameters

Parameter	Description
format	Controls the level of details for the plan. It accepts four values: • BASIC: Displays the minimum information in the plan—the operation ID, the operation name and its option. • TYPICAL: This is the default. Displays the most relevant information in the plan (operation id, name and option, #rows, #bytes and optimizer cost). Pruning, parallel and predicate information are only displayed when applicable. Excludes only PROJECTION, ALIAS and REMOTE SQL information (see below). • SERIAL: Like TYPICAL except that the parallel information is not displayed, even if the plan executes in parallel. • ALL: Maximum user level. Includes information displayed with the TYPICAL level with additional information (PROJECTION, ALIAS and information about REMOTE SQL if the operation is distributed). For finer control on the display output, the following keywords can be added to the above four standard format options to customize their default behavior. Each keyword either represents a logical group of plan table columns (such as PARTITION) or logical additions to the base plan table output (such as PREDICATE). Format keywords must be separated by either a comma or a space: • ROWS - if relevant, shows the number of rows estimated by the optimizer • BYTES - if relevant, shows the number of bytes estimated by the optimizer • COST - if relevant, shows optimizer cost information • PARTITION - if relevant, shows partition pruning information • PARALLEL - if relevant, shows PX information (distribution method and table queue information) • PREDICATE - if relevant, shows the predicate section • PROJECTION -if relevant, shows the projection section • ALIAS - if relevant, shows the "Query Block Name / Object Alias" section • REMOTE - if relevant, shows the information for distributed query (for example, remote from serial distribution and remote SQL) • NOTE - if relevant, shows the note section of the explain plan Format keywords can be prefixed by the sign '-' to exclude the specified information. For example, '-PROJECTION' excludes projection information.

Usage Notes

- To use the DISPLAY_AWR functionality, the calling user must have SELECT or READ privilege on DBA_HIST_SQL_PLAN, AWR_ROOT_SQL_PLAN, AWR_PDB_SQL_PLAN, DBA_HIST_SQLTEXT, AWR_ROOT_SQLTEXT, AWR_PDB_SQLTEXT, and V$DATABASE, otherwise it shows an appropriate error message. By default, select privilege for these views is granted to the select_catalog role.
- The following examples show different ways of using the format parameter:
  - Use 'BASIC ROWS' to display basic information with the additional number of rows estimated by the optimizer.
  - Use 'ALL -PROJECTION -NOTE' to display everything except the projection and note sections.

- Use 'TYPICAL PROJECTION' to display using the typical format with the additional projection section (which is normally excluded under the typical format). Since typical is default, using simply 'PROJECTION' is equivalent.
- Use '-BYTES -COST -PREDICATE' to display using the typical format but excluding optimizer cost and byte estimates and the predicate section.

Examples

To display the different execution plans associated with the SQL ID 'atfwcg8anrykp':

```
SELECT * FROM table(DBMS_XPLAN.DISPLAY_AWR('atfwcg8anrykp'));
```

To display all execution plans of all stored SQL statements containing the string 'TOTO':

```
SELECT tf.*
FROM   DBA_HIST_SQLTEXT ht, table(DBMS_XPLAN.DISPLAY_AWR(ht.sql_id,null,
null,  'ALL' )) tf
WHERE  ht.sql_text like '%TOTO%';
```

237.5.5 DISPLAY_CURSOR Function

This table function displays the explain plan of any cursor loaded in the cursor cache. In addition to the explain plan, various plan statistics (such as. I/O, memory and timing) can be reported (based on the V$SQL_PLAN_STATISTICS_ALL VIEWS).

Syntax

```
DBMS_XPLAN.DISPLAY_CURSOR(
   sql_id            IN  VARCHAR2  DEFAULT  NULL,
   cursor_child_no   IN  NUMBER    DEFAULT  0,
   format            IN  VARCHAR2  DEFAULT  'TYPICAL');
```

Parameters

Table 237-7 DISPLAY_CURSOR Function Parameters

Parameter	Description
sql_id	Specifies the SQL_ID of the SQL statement in the cursor cache. You can retrieve the appropriate value by querying the column SQL_ID in V$SQL or V$SQLAREA. Alternatively, you could choose the column PREV_SQL_ID for a specific session out of V$SESSION. This parameter defaults to NULL in which case the plan of the last cursor executed by the session is displayed.
cursor_child_no	Child number of the cursor to display. If not supplied, the execution plan of the child_number=0 cursor matching the supplied sql_id parameter are displayed. The child_number can be specified only if sql_id is specified.

Table 237-7 (Cont.) DISPLAY_CURSOR Function Parameters

Parameter	Description
format	Controls the level of details for the plan. It accepts five values: • BASIC: Displays the minimum information in the plan—the operation ID, the operation name and its option. • TYPICAL: This is the default. Displays the most relevant information in the plan (operation id, name and option, #rows, #bytes and optimizer cost). Pruning, parallel and predicate information are only displayed when applicable. Excludes only PROJECTION, ALIAS and REMOTE SQL information (see below). • SERIAL: Like TYPICAL except that the parallel information is not displayed, even if the plan executes in parallel. • ALL: Maximum user level. Includes information displayed with the TYPICAL level with additional information (PROJECTION, ALIAS and information about REMOTE SQL if the operation is distributed). • ADAPTIVE: – Displays the final plan, or the current plan if the execution has not completed. This section includes notes about runtime optimizations that affect the plan, such as switching from a Nested Loops join to a Hash join. – Plan lineage. This section shows the plans that were run previously due to automatic reoptimization. It also shows the default plan, if the plan changed due to dynamic plans. – Recommended plan. In reporting mode, the plan is chosen based on execution statistics displayed. Note that displaying the recommended plan for automatic reoptimization requires re-compiling the query with the optimizer adjustments collected in the child cursor. Displaying the recommended plan for a dynamic plan does not require this. – Dynamic plans. This summarizes the portions of the plan that differ from the default plan chosen by the optimizer. For finer control on the display output, you can add the following keywords to the preceding format options to customize their default behavior. Each keyword either represents a logical group of plan table columns (such as PARTITION) or logical additions to the base plan table output (such as PREDICATE). Format keywords must be separated by either a comma or a space: • ROWS - if relevant, shows the number of rows estimated by the optimizer • BYTES - if relevant, shows the number of bytes estimated by the optimizer • COST - if relevant, shows optimizer cost information • PARTITION - if relevant, shows partition pruning information • PARALLEL - if relevant, shows PX information (distribution method and table queue information) • PREDICATE - if relevant, shows the predicate section • PROJECTION -if relevant, shows the projection section • ALIAS - if relevant, shows the "Query Block Name / Object Alias" section • REMOTE - if relevant, shows the information for distributed query (for example, remote from serial distribution and remote SQL) • NOTE - if relevant, shows the note section of the explain plan • IOSTATS - assuming that basic plan statistics are collected when SQL statements are executed (either by using the gather_plan_statistics hint or by setting the parameter

Table 237-7 (Cont.) DISPLAY_CURSOR Function Parameters

Parameter	Description
	statistics_level to ALL), this format shows IO statistics for `ALL` (or only for the `LAST` as shown below) executions of the cursor. • `MEMSTATS` - Assuming that PGA memory management is enabled (that is, `pga_aggregate_target` parameter is set to a non 0 value), this format allows to display memory management statistics (for example, execution mode of the operator, how much memory was used, number of bytes spilled to disk, and so on). These statistics only apply to memory intensive operations like hash-joins, sort or some bitmap operators. • `ALLSTATS` - A shortcut for `'IOSTATS MEMSTATS'` • `LAST` - By default, plan statistics are shown for all executions of the cursor. The keyword `LAST` can be specified to see only the statistics for the last execution. The following formats are deprecated but supported for backward compatibility: • `RUNSTATS_TOT` - Same as `IOSTATS`, that is, displays IO statistics for all executions of the specified cursor. • `RUNSTATS_LAST` - Same as `IOSTATS LAST`, that is, displays the runtime statistics for the last execution of the cursor You can prefix format keywords with the sign '`-`' to exclude the specified information. For example, '`-PROJECTION`' excludes projection information.

Usage Notes

- To use the `DISPLAY_CURSOR` functionality, the calling user must have `SELECT` or `READ` privilege on the fixed views `V$SQL_PLAN_STATISTICS_ALL`, `V$SQL` and `V$SQL_PLAN`, otherwise it shows an appropriate error message.
- Here are some ways you might use variations on the `format` parameter:
  - Use '`ALL -PROJECTION -NOTE`' to display everything except the projection and note sections.
  - Use '`TYPICAL PROJECTION`' to display using the typical format with the additional projection section (which is normally excluded under the typical format). Since typical is default, using simply '`PROJECTION`' is equivalent.
  - Use '`-BYTES -COST -PREDICATE`' to display using the typical format but excluding optimizer cost and byte estimates as well as the predicate section.
  - Use '`BASIC ROWS`' to display basic information with the additional number of rows estimated by the optimizer.

Examples

To display the execution plan of the last `SQL` statement executed by the current session:

```
SELECT * FROM table (
   DBMS_XPLAN.DISPLAY_CURSOR);
```

To display the execution plan of all children associated with the SQL ID 'atfwcg8anrykp':

```
SELECT * FROM table (
   DBMS_XPLAN.DISPLAY_CURSOR('atfwcg8anrykp'));
```

To display runtime statistics for the cursor included in the preceding statement:

```
SELECT * FROM table (
   DBMS_XPLAN.DISPLAY_CURSOR('atfwcg8anrykp', NULL, 'ALLSTATS LAST');
```

237.5.6 DISPLAY_PLAN Function

This table function displays the contents of the plan table in a variety of formats with CLOB output type.

Syntax

```
DBMS_XPLAN.DISPLAY_PLAN (
   table_name       IN    VARCHAR2   DEFAULT 'PLAN_TABLE',
   statement_id     IN    VARCHAR2   DEFAULT NULL,
   format           IN    VARCHAR2   DEFAULT 'TYPICAL',
   filter_preds     IN    VARCHAR2   DEFAULT NULL,
   type             IN    VARCHAR2   DEFAULT 'TEXT')
  RETURN CLOB;
```

Parameters

Table 237-8 DISPLAY_PLAN Function Parameters

Parameter	Description
table_name	Specifies the table name where the plan is stored. This parameter defaults to PLAN_TABLE, which is the default plan table for the EXPLAIN PLAN command. If NULL is specified it also defaults to PLAN_TABLE.
statement_id	Specifies the statement_id of the plan to be displayed. This parameter defaults to NULL, which is the default when the EXPLAIN PLAN command is executed without a set statement_id clause.If no statement_id is specified, the function shows you the plan of the most recent explained statement.
filter_preds	SQL filter predicate(s) to restrict the set of rows selected from the table where the plan is stored. When value is NULL (the default), the plan displayed corresponds to the last executed explain plan. For example: filter_preds=>'plan_id = 10' Can reference any column of the table where the plan is stored and can contain any SQL construct (for example, sub-query, function calls (see **WARNING** under Usage Notes)

Table 237-8 (Cont.) DISPLAY_PLAN Function Parameters

Parameter	Description
format	Controls the level of details for the plan. It accepts five values: • BASIC: Displays the minimum information in the plan—the operation ID, the operation name and its option. • TYPICAL: This is the default. Displays the most relevant information in the plan (operation id, name and option, #rows, #bytes and optimizer cost). Pruning, parallel and predicate information are only displayed when applicable. Excludes only PROJECTION, ALIAS and REMOTE SQL information (see below). • SERIAL: Like TYPICAL except that the parallel information is not displayed, even if the plan executes in parallel. • ALL: Maximum user level. Includes information displayed with the TYPICAL level with additional information (PROJECTION, ALIAS and information about REMOTE SQL if the operation is distributed). • ADAPTIVE: Displays the default plan, and for each dynamic subplan (if stipulated): - A list of the rowsources from the original which may be replaced, and the rowsources to replace them - If outline display is specified in the format argument, the hints for each option in the dynamic subplan are displayed For finer control on the display output, the following keywords can be added to the above three standard format options to customize their default behavior. Each keyword either represents a logical group of plan table columns (such as PARTITION) or logical additions to the base plan table output (such as PREDICATE). Format keywords must be separated by either a comma or a space: • ROWS - if relevant, shows the number of rows estimated by the optimizer • BYTES - if relevant, shows the number of bytes estimated by the optimizer • COST - if relevant, shows optimizer cost information • PARTITION - if relevant, shows partition pruning information • PARALLEL - if relevant, shows PX information (distribution method and table queue information) • PREDICATE - if relevant, shows the predicate section • PROJECTION -if relevant, shows the projection section • ALIAS - if relevant, shows the "Query Block Name / Object Alias" section • REMOTE - if relevant, shows the information for distributed query (for example, remote from serial distribution and remote SQL) • NOTE - if relevant, shows the note section of the explain plan Format keywords can be prefixed by the sign '-' to exclude the specified information. For example, '-PROJECTION' excludes projection information. If the target plan table (see table_name parameter) also stores plan statistics columns (for example, it is a table used to capture the content of the fixed view V$SQL_PLAN_STATISTICS_ALL), additional format keywords can be used to specify which class of statistics to display when using the DISPLAY Function. These additional format keywords are IOSTATS, MEMSTATS, ALLSTATS and LAST (see the DISPLAY_CURSOR Function or the DISPLAY_SQLSET Function for a full description of these four keywords).
type	Output type, one of: 'TEXT', 'ACTIVE', 'HTML', or 'XML' (see Usage Notes regarding type ACTIVE). '

Return Values

Returns the requested report as `CLOB`

Usage Notes

Active reports have a rich, interactive user interface akin to that found in Enterprise Manager while not requiring any EM installation. The report file built is in HTML format, so it can be interpreted by most modern browsers. The code powering the active report is downloaded transparently by the web browser when the report is first viewed, hence viewing it requires outside connectivity.

> **WARNING:**
>
> Application developers should expose the `filter_preds` parameter to end-users only after careful consideration because this could expose the application to SQL injection. Indeed, `filter_preds` can potentially reference any table or execute any server function for which the database user invoking the table function has privileges.

237.5.7 DISPLAY_SQL_PLAN_BASELINE Function

This table function displays one or more execution plans for the specified SQL handle of a SQL plan baseline.

Syntax

```
DBMS_XPLAN.DISPLAY_SQL_PLAN_BASELINE (
   sql_handle      IN VARCHAR2 := NULL,
   plan_name       IN VARCHAR2 := NULL,
   format          IN VARCHAR2 := 'TYPICAL')
 RETURN dbms_xplan_type_table;
```

Parameters

Table 237-9 DISPLAY_SQL_PLAN_BASELINE Function Parameters

Parameter	Description
`sql_handle`	SQL statement handle. It identifies a SQL statement whose plans are to be displayed.
`plan_name`	Plan name. It identifies a specific plan. Default NULL means all plans associated with identified SQL statement are explained and displayed.
`format`	Format string determines what information stored in the plan displayed. The following format values are possible, each representing a common use case: `BASIC`, `TYPICAL`, and `ALL`.

Return Values

A PL/SQL type table

Usage Notes

This function uses plan information stored in the plan baseline to explain and display the plans. The `plan_id` stored in the SQL management base may not match the `plan_id` of the generated plan. A mismatch between the stored `plan_id` and generated `plan_id` means that it is a non-reproducible plan. Such a plan is deemed invalid and is bypassed by the optimizer during SQL compilation.

Examples

Display all plans of a SQL statement identified by the SQL handle `SYS_SQL_b1d49f6074ab95af` using `TYPICAL` format

```
SET LINESIZE 150
SET PAGESIZE 2000
SELECT t.*
FROM TABLE(DBMS_XPLAN.DISPLAY_SQL_PLAN_BASELINE('SYS_SQL_b1d49f6074ab95af'))
t;
```

Display all plans of one or more SQL statements containing the string `HR2` using `BASIC` format:

```
SET LINESIZE 150
SET PAGESIZE 2000
SELECT t.*
FROM (SELECT DISTINCT sql_handle FROM dba_sql_plan_baselines WHERE sql_text
LIKE '%HR2%') pb,
     TABLE(DBMS_XPLAN.DISPLAY_SQL_PLAN_BASELINE(pb.sql_handle, NULL,
'BASIC')) t;
```

237.5.8 DISPLAY_SQLSET Function

This table function displays the execution plan of a given statement stored in a SQL tuning set.

Syntax

```
DBMS_XPLAN.DISPLAY_SQLSET(
   sqlset_name      IN  VARCHAR2,
   sql_id           IN  VARCHAR2,
   plan_hash_value  IN NUMBER := NULL,
   format           IN  VARCHAR2  := 'TYPICAL',
   sqlset_owner     IN  VARCHAR2  := NULL)
  RETURN DBMS_XPLAN_TYPE_TABLE PIPELINED;
```

Parameters

Table 237-10 DISPLAY_SQLSET Function Parameters

Parameter	Description
`sqlset_name`	Name of the SQL Tuning Set

Table 237-10 (Cont.) DISPLAY_SQLSET Function Parameters

Parameter	Description
`sql_id`	Specifies the sql_id value for a SQL statement having its plan stored in the SQL tuning set. You can find all stored SQL statements by querying table function `DBMS_SQLTUNE.SELECT_SQLSET`
`plan_hash_value`	Optional parameter. Identifies a specific stored execution plan for a SQL statement. If suppressed, all stored execution plans are shown.

Table 237-10 (Cont.) DISPLAY_SQLSET Function Parameters

Parameter	Description
format	Controls the level of details for the plan. It accepts four values: • BASIC: Displays the minimum information in the plan—the operation ID, the operation name and its option. • TYPICAL: This is the default. Displays the most relevant information in the plan (operation id, name and option, #rows, #bytes and optimizer cost). Pruning, parallel and predicate information are only displayed when applicable. Excludes only PROJECTION, ALIAS and REMOTE SQL information (see below). • SERIAL: Like TYPICAL except that the parallel information is not displayed, even if the plan executes in parallel. • ALL: Maximum user level. Includes information displayed with the TYPICAL level with additional information (PROJECTION, ALIAS and information about REMOTE SQL if the operation is distributed). For finer control on the display output, the following keywords can be added to the above three standard format options to customize their default behavior. Each keyword either represents a logical group of plan table columns (such as PARTITION) or logical additions to the base plan table output (such as PREDICATE). Format keywords must be separated by either a comma or a space: • ROWS - if relevant, shows the number of rows estimated by the optimizer • BYTES - if relevant, shows the number of bytes estimated by the optimizer • COST - if relevant, shows optimizer cost information • PARTITION - if relevant, shows partition pruning information • PARALLEL - if relevant, shows PX information (distribution method and table queue information) • PREDICATE - if relevant, shows the predicate section • PROJECTION -if relevant, shows the projection section • ALIAS - if relevant, shows the "Query Block Name / Object Alias" section • REMOTE - if relevant, shows the information for distributed query (for example, remote from serial distribution and remote SQL) • NOTE - if relevant, shows the note section of the explain plan • IOSTATS - assuming that basic plan statistics are collected when SQL statements are executed (either by using the gather_plan_statistics hint or by setting the parameter STATISTICS_LEVEL to ALL), this format shows IO statistics for ALL (or only for the LAST as shown below) executions of the cursor. • MEMSTATS - Assuming that PGA memory management is enabled (that is, pga_aggregate_target parameter is set to a non 0 value), this format allows to display memory management statistics (for example, execution mode of the operator, how much memory was used, number of bytes spilled to disk, and so on). These statistics only apply to memory intensive operations like hash-joins, sort or some bitmap operators. • ALLSTATS - A shortcut for 'IOSTATS MEMSTATS' • LAST - By default, plan statistics are shown for all executions of the cursor. The keyword LAST can be specified to see only the statistics for the last execution. The following two formats are deprecated but supported for backward compatibility: • RUNSTATS_TOT - Same as IOSTATS, that is, displays IO statistics for all executions of the specified cursor.

Table 237-10 (Cont.) DISPLAY_SQLSET Function Parameters

Parameter	Description
	• RUNSTATS_LAST - Same as IOSTATS LAST, that is, displays the runtime statistics for the last execution of the cursor Format keywords can be prefixed by the sign '-' to exclude the specified information. For example, '-PROJECTION' excludes projection information.
sqlset_owner	The owner of the SQL tuning set. The default is the current user.

Usage Notes

Here are some ways you might use variations on the format parameter:

- Use 'ALL -PROJECTION -NOTE' to display everything except the projection and note sections.
- Use 'TYPICAL PROJECTION' to display using the typical format with the additional projection section (which is normally excluded under the typical format). Since typical is default, using simply 'PROJECTION' is equivalent.
- Use '-BYTES -COST -PREDICATE' to display using the typical format but excluding optimizer cost and byte estimates as well as the predicate section.
- Use 'BASIC ROWS' to display basic information with the additional number of rows estimated by the optimizer.

Examples

To display the execution plan for the SQL statement associated with SQL ID 'gwp663cqh5qbf' and PLAN HASH 3693697075 in the SQL Tuning Set called 'OLTP_optimization_0405':

```
SELECT * FROM table
(DBMS_XPLAN.DISPLAY_SQLSET('OLTP_optimization_0405','gwp663cqh5qbf',
3693697075));
```

To display all execution plans of the SQL ID 'atfwcg8anrykp' stored in the SQL tuning set:

```
SELECT * FROM table
(DBMS_XPLAN.DISPLAY_SQLSET('OLTP_optimization_0405','gwp663cqh5qbf'));
```

To display runtime statistics for the SQL statement included in the preceding statement:

```
SELECT * FROM table (
   DBMS_XPLAN.DISPLAY_SQLSET(
      'OLTP_optimization_0405', 'gwp663cqh5qbf', NULL, 'ALLSTATS LAST');
```

237.5.9 DISPLAY_WORKLOAD_REPOSITORY Function

This table function displays the contents of an execution plan stored in AWR.

> **Note:**
>
> This function replaces `DISPLAY_AWR`, which is deprecated.

Syntax

```
DBMS_XPLAN.DISPLAY_WORKLOAD_REPOSITORY(
   sql_id            IN      VARCHAR2,
   plan_hash_value   IN      NUMBER   DEFAULT NULL,
   format            IN      VARCHAR2 DEFAULT 'TYPICAL'
   dbid              IN      NUMBER   DEFAULT NULL,
   con_dbid          IN      NUMBER   DEFAULT NULL,
   awr_location      IN      VARCHAR2 DEFAULT 'AWR_ROOT'
);
```

Parameters

Table 237-11 DISPLAY_WORKLOAD_REPOSITORY Table Function Parameters

Parameter	Description
`sql_id`	Specifies the `SQL_ID` of the SQL statement. You can retrieve the appropriate value for the SQL statement of interest by querying the column `SQL_ID` in `DBA_HIST_SQLTEXT`.
`plan_hash_value`	Specifies the `PLAN_HASH_VALUE` of a SQL statement. This parameter is optional. If omitted, the table function returns all stored execution plans for a given `SQL_ID`.

Table 237-11 (Cont.) DISPLAY_WORKLOAD_REPOSITORY Table Function Parameters

Parameter	Description
format	Controls the level of details for the plan. It accepts four values: • BASIC: Displays the minimum information in the plan—the operation ID, the operation name and its option. • TYPICAL: This is the default. Displays the most relevant information in the plan (operation id, name and option, #rows, #bytes and optimizer cost). Pruning, parallel and predicate information are only displayed when applicable. Excludes only PROJECTION, ALIAS and REMOTE SQL information (see below). • SERIAL: Like TYPICAL except that the parallel information is not displayed, even if the plan executes in parallel. • ALL: Maximum user level. Includes information displayed with the TYPICAL level with additional information (PROJECTION, ALIAS and information about REMOTE SQL if the operation is distributed). For finer control on the display output, the following keywords can be added to the above four standard format options to customize their default behavior. Each keyword either represents a logical group of plan table columns (such as PARTITION) or logical additions to the base plan table output (such as PREDICATE). Format keywords must be separated by either a comma or a space: • ROWS - if relevant, shows the number of rows estimated by the optimizer • BYTES - if relevant, shows the number of bytes estimated by the optimizer • COST - if relevant, shows optimizer cost information • PARTITION - if relevant, shows partition pruning information • PARALLEL - if relevant, shows PX information (distribution method and table queue information) • PREDICATE - if relevant, shows the predicate section • PROJECTION -if relevant, shows the projection section • ALIAS - if relevant, shows the "Query Block Name / Object Alias" section • REMOTE - if relevant, shows the information for distributed query (for example, remote from serial distribution and remote SQL) • NOTE - if relevant, shows the note section of the explain plan Format keywords can be prefixed by the sign '-' to exclude the specified information. For example, '-PROJECTION' excludes projection information.
dbid	Identifies the plans for a specific database. If this parameter is omitted, then the value defaults to the DBID of the AWR repository pointed to by the initialization parameter AWR_LOCATION. In a CDB, if AWR_LOCATION is set to AWR_ROOT, then the value is set to the DBID of the CDB root. If it is set to AWR_PDB, then the value is set to the DBID of the container.
con_dbid	Identifies the plans for a specific container. If this parameter is omitted, then the value defaults to SYS_CONTEXT('userenv', 'con_id').
awr_location	Specifies the location of the AWR repository. Supported values are: • AWR_ROOT when the AWR to be accessed is in the root container. This is the default. • 'AWR_PDB', if the AWR to be accessed is in the local container.

Example 237-2 Querying an AWR Plan

Assume that you log in as an administrator and issue the following query:

```
select count(*) from sh.sAleS
```

You create an AWR snapshot as follows:

```
EXEC DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT;
```

You query joint `DBA_HIST_SQLTEST` to the function output as follows:

```
SET LINESIZE 150
SET PAGESIZE 5000

SELECT t.*
FROM   DBA_HIST_SQLTEXT ht,
       TABLE(DBMS_XPLAN.DISPLAY_WORKLOAD_REPOSITORY
         (ht.sql_id, null, '-PREDICATE +ALIAS',null,null,'AWR_ROOT')) t
WHERE  ht.SQL_TEXT LIKE '%sAleS%';

SQL_ID 2f4cx9qjnqd70
--------------------
select count(*) from sh.sAleS

Plan hash value: 1123225294

-------------------------------------------------------------------------------
----------------------------
| Id  | Operation                      | Name            | Rows  | Cost
(%CPU)| Time     | Pstart| Pstop |
-------------------------------------------------------------------------------
----------------------------
|   0 | SELECT STATEMENT               |                 |       |    27
(100)|          |       |       |
|   1 |  SORT AGGREGATE                |                 |     1 |
|          |       |       |
|   2 |   PARTITION RANGE ALL          |                 |   918K|    27
(0)| 00:00:01 |     1 |    28 |
|   3 |    BITMAP CONVERSION COUNT     |                 |   918K|    27
(0)| 00:00:01 |       |       |
|   4 |     BITMAP INDEX FAST FULL SCAN| SALES_PROMO_BIX |       |
|          |     1 |    28 |
-------------------------------------------------------------------------------
----------------------------

Query Block Name / Object Alias (identified by operation id):
-------------------------------------------------------------

   1 - SEL$1
   3 - SEL$1 / "SALES"@"SEL$1"
```

238
DBMS_XSLPROCESSOR

The DBMS_XSLPROCESSOR package provides an interface to manage the contents and structure of XML documents.

This chapter contains the following topics:

- DBMS_XSLPROCESSOR Overview
- DBMS_XSLPROCESSOR Security Model
- Summary of DBMS_XSLPROCESSOR Subprograms

> **See Also:**
>
> – *Oracle XML DB Developer's Guide*

238.1 DBMS_XSLPROCESSOR Overview

The DBMS_XSLPROCESSOR package provides an interface to manage the contents and structure of XML documents.

Standards

This PL/SQL implementation of the XSL processor follows the W3C XSLT working draft rev WD-xslt-19990813 and includes the required behavior of an XSL processor in terms of how it must read XSLT stylesheets and the transformation it must effect.

Concepts

The Extensible Stylesheet Language Transformation (XSLT) describes rules for transforming a source tree into a result tree. A transformation expressed in XSLT is called a stylesheet. The transformation specified is achieved by associating patterns with templates defined in the stylesheet. A template is instantiated to create part of the result tree.

Implementation

The following is the default behavior for this PL/SQL XSL Processor:

- A result tree which can be accessed by DOM programmatic interface
- Errors are not recorded unless an error log is specified; however, an application error will be raised if parsing fails

238.2 DBMS_XSLPROCESSOR Security Model

Owned by XDB, the DBMS_XSLPROCESSOR package must be created by SYS or XDB. The EXECUTE privilege is granted to PUBLIC. Subprograms in this package are executed using the privileges of the current user.

238.3 Summary of DBMS_XSLPROCESSOR Subprograms

This table lists the `DBMS_XSLPROCESSOR` subprograms and briefly describes them.

Table 238-1 *DBMS_XSLPROCESSOR Package Subprograms*

Method	Description
FREEPROCESSOR Procedure	Frees a processor object
FREESTYLESHEET Procedure	Frees a stylesheet object
NEWPROCESSOR Function	Returns a new processor instance
NEWSTYLESHEET Functions	Creates a new stylesheet from input and reference URLs
PROCESSXSL Functions and Procedures	Transforms an input XML document
REMOVEPARAM Procedure	Removes a top-level stylesheet parameter
RESETPARAMS Procedure	Resets the top-level stylesheet parameters
SELECTNODES Function	Selects nodes from a DOM tree that match a pattern
SELECTSINGLENODE Function	Selects the first node from the tree that matches a pattern
SETERRORLOG Procedure	Sets errors to be sent to the specified file
SETPARAM Procedure	Sets a top-level parameter in the stylesheet
SHOWWARNINGS Procedure	Turns warnings on or off
TRANSFORMNODE Function	Transforms a node in a DOM tree using a stylesheet
VALUEOF Function and Procedure	Gets the value of the first node that matches a pattern

238.3.1 FREEPROCESSOR Procedure

This procedure frees a `Processor` object.

Syntax

```
DBMS_XSLPROCESSOR.FREEPROCESSOR(
  p  IN  Processor);
```

Parameters

Table 238-2 FREEPROCESSOR Procedure Parameters

Parameter	Description
p	Processor

238.3.2 FREESTYLESHEET Procedure

This procedure frees a `Stylesheet` object.

Syntax

```
DBMS_XSLPROCESSOR.FREESTYLESHEET(
  ss  IN  Stylesheet);
```

Parameters

Table 238-3 FREESTYLESHEET Procedure Parameters

Parameter	Description
ss	Stylesheet

238.3.3 NEWPROCESSOR Function

This function returns a new `Processor` instance.

The function must be called before the default behavior of `Processor` can be changed and if other processor methods need to be used.

Syntax

```
DBMS_XSLPROCESSOR.NEWPROCESSOR
 RETURN Processor;
```

238.3.4 NEWSTYLESHEET Functions

This function creates and returns a new `Stylesheet` instance.

The options are described in the following table.

Syntax

Creates and returns a new stylesheet instance using the given `DOMDOCUMENT` and reference URLs:

```
DBMS_XSLPROCESSOR.NEWSTYLESHEET(
   xmldoc  IN   DOMDOCUMENT,
   ref     IN   VARCHAR2)
 RETURN Stylesheet;
```

Creates and returns a new `Stylesheet` instance using the given input and reference URLs:

```
DBMS_XSLPROCESSOR.NEWSTYLESHEET(
   inp    IN   VARCHAR2,
   ref    IN   VARCHAR2)
 RETURN Stylesheet;
```

Parameters

Table 238-4 NEWSTYLESHEET Function Parameters

Parameter	Description
xmldoc	`DOMDocument` to use for construction
inp	Input URL to use for construction
ref	Reference URL

238.3.5 PROCESSXSL Functions and Procedures

This function transforms input `XMLDocument`.

Any changes to the default processor behavior should be effected before calling this procedure. An application error is raised if processing fails.

Syntax

Transforms input `XMLDocument` using given `DOMDocument` and stylesheet, and returns the resultant document fragment:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p      IN   Processor,
   ss     IN   Stylesheet,
   xmldoc IN   DOMDOCUMENT),
 RETURN DOMDOCUMENTFRAGMENT;
```

Transforms input `XMLDocument` using given document as URL and the `Stylesheet`, and returns the resultant document fragment:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p     IN   Processor,
   ss    IN   Stylesheet,
   url   IN   VARCHAR2,
 RETURN DOMDOCUMENTFRAGMENT;
```

Transforms input `XMLDocument` using given document as `CLOB` and the `Stylesheet`, and returns the resultant document fragment:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p     IN   Processor,
   ss    IN  Stylesheet,
   clb   IN  CLOB)
 RETURN DOMDOCUMENTFRAGMENT;
```

Transforms input `XMLDocument` using given `DOMDOCUMENT` and the stylesheet, and writes the output to the specified file:

```
DBMS_XSLPROCESSOR.DBMS_XSLPROCESSOR.(
   p           IN    Processor,
   ss          IN    Stylesheet,
   xmldoc      IN    DOMDOCUMENT,
   dir         IN    VARCHAR2,
   fileName    IN    VARCHAR2);
```

Transforms input `XMLDocument` using given URL and the stylesheet, and writes the output to the specified file in a specified directory:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p          IN    Processor,
   ss         IN    Stylesheet,
   url        IN    VARCHAR2,
   dir        IN    VARCHAR2,
   fileName   IN    VARCHAR2);
```

Transforms input `XMLDocument` using given `DOMDOCUMENT` and the stylesheet, and writes the output to a `CLOB`:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p          IN      Processor,
   ss         IN      Stylesheet,
   xmldoc     IN      DOMDOCUMENT,
   cl         IN OUT  CLOB);
```

Transforms input `XMLDocument` using given `DOMDOCUMENTFRAGMENT` and the stylesheet, and returns the resultant document fragment:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p          IN    Processor,
   ss         IN    Stylesheet,
   xmldf      IN    DOMDOCUMENTFRAGMENT)
 RETURN DOMDOCUMENTFRAGMENT;
```

Transforms input `XMLDocumentFragment` using given `DOMDocumentFragment` and the stylesheet, and writes the output to the specified file in a specified directory:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p          IN    Processor,
   ss         IN    Stylesheet,
   xmldf      IN    DOMDOCUMENTFRAGMENT,
   dir        IN    VARCHAR2,
   filename   IN    VARCHAR2);
```

Transforms input `XMLDocumentFragment` using given `DOMDOCUMENTFRAGMENT` and the stylesheet, and writes the output to a buffer:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p          IN       Processor,
   ss         IN       Stylesheet,
   xmldf      IN       DOMDOCUMENTFRAGMENT,
   buf        IN OUT   VARCHAR2);
```

Transforms input `XMLDocumentFragment` using given `DOMDOCUMENTFRAGMENT` and the stylesheet, and writes the output to a `CLOB`:

```
DBMS_XSLPROCESSOR.PROCESSXSL(
   p          IN      Processor,
   ss         IN      Stylesheet,
   xmldf      IN      DOMDOCUMENTFRAGMENT,
   cl         IN OUT  CLOB);
```

Parameters

Table 238-5 PROCESSXSL Function and Procedure Parameters

Parameter	Description
`p`	`Processor` instance
`ss`	`Stylesheet` instance
`xmldoc`	XML document being transformed
`url`	URL for the information being transformed
`clb`	`CLOB` containing information to be transformed
`dir`	Directory where processing output file is saved
`filename`	Processing output file
`cl`	`CLOB` to which the processing output is saved

Table 238-5 (Cont.) PROCESSXSL Function and Procedure Parameters

Parameter	Description
xmldf	XMLDocumentFragment being transformed

238.3.6 REMOVEPARAM Procedure

This procedure removes a top level stylesheet parameter.

Syntax

```
DBMS_XSLPROCESSOR.REMOVEPARAM(
   ss      IN  Stylesheet,
   name    IN  VARCHAR2);
```

Parameters

Table 238-6 REMOVEPARAM Procedure Parameters

Parameter	Description
ss	Stylesheet instance
name	Name of the parameter

238.3.7 RESETPARAMS Procedure

This procedure resets the top-level stylesheet parameters.

Syntax

```
DBMS_XSLPROCESSOR.RESETPARAMS(
   ss  IN   Stylesheet);
```

Parameters

Table 238-7 RESETPARAMS Procedure Parameters

Parameter	Description
ss	Stylesheet instance

238.3.8 SELECTNODES Function

This function selects nodes which match the supplied path expression from a DOM tree, and returns the result of the selection.

Syntax

```
DBMS_XSLPROCESSOR.SELECTNODES(
   n           IN   DBMS_XMLDOM.DOMNODE,
   pattern     IN   VARCHAR2,
   namespace   IN VARCHAR2 := NULL)
 RETURN DBMS_XMLDOM.DOMNODELIST;
```

Parameters

Table 238-8 SELECTNODES Function Parameters

Parameter	Description
n	Root DOMNode of the tree
pattern	Pattern to use
namespace	Namespace declared

238.3.9 SELECTSINGLENODE Function

This function selects the first node from the tree that match the supplied path expression, and returns that node.

Syntax

```
DBMS_XSLPROCESSOR.SELECTSINGLENODE(
   n           IN   DBMS_XMLDOM.DOMNODE,
   pattern     IN   VARCHAR2,
   namespace   IN VARCHAR2 := NULL)
 RETURN DBMS_XMLDOM.DOMNODE;
```

Parameters

Table 238-9 SELECTSINGLENODE Function Parameters

Parameter	Description
n	Root DOMNode of the tree
pattern	Pattern to use
namespace	Namespace declared

238.3.10 SETERRORLOG Procedure

This deprecated procedure sets errors to be sent to the specified file.

> **Note:**
>
> This subprogram has been deprecated, and is included only for reasons of backward compatibility.

Syntax

```
DBMS_XSLPROCESSOR.SETERRORLOG(
   p         IN   Processor,
   fileName  IN   VARCHAR2);
```

Parameters

Table 238-10 SETERRORLOG Procedure Parameters

Parameter	Description
p	Processor instance
fileName	Complete path of the file to use as the error log

238.3.11 SETPARAM Procedure

This procedure sets a top level parameter in the stylesheet.

The parameter value must be a valid XPath expression. Literal string values must be quoted.

Syntax

```
DBMS_XSLPROCESSOR.SETPARAM(
   ss      IN   Stylesheet,
   name    IN   VARCHAR2,
   value   IN   VARCHAR2);
```

Parameters

Table 238-11 SETPARAM Procedure Parameters

Parameter	Description
ss	Stylesheet instance
name	Name of the parameter
value	Value of the parameter

238.3.12 SHOWWARNINGS Procedure

This procedure turns warnings on (TRUE) or off (FALSE).

Syntax

```
DBMS_XSLPROCESSOR.SHOWWARNINGS(
   p     IN   Processor,
   yes   IN   BOOLEAN);
```

Parameters

Table 238-12 SHOWWARNINGS Procedure Parameters

Parameter	Description
p	Processor instance
yes	Mode to set: TRUE to show warnings, FALSE otherwise

238.3.13 TRANSFORMNODE Function

This function transforms a node in a DOM tree using the given stylesheet, and returns the result of the transformation as a `DOMDocumentFragment`.

Syntax

```
DBMS_XSLPROCESSOR.TRANSFORMNODE(
   n    IN  DOMNODE,
   ss   IN  Stylesheet)
 RETURN DOMDocumentFragment;
```

Parameters

Table 238-13 TRANSFORMNODE Function Parameters

Parameter	Description
`n`	`DOMNode` to transform
`ss`	Stylesheet to use

238.3.14 VALUEOF Function and Procedure

This subprogram retrieves the value of the first node from the tree that matches the given pattern. You can use either a function or a procedure.

Syntax

```
DBMS_XSLPROCESSOR.VALUEOF(
  n           IN    DBMS_XMLDOM.DOMNODE,
  pattern     IN    VARCHAR2,
  namespace   IN    VARCHAR2 := NULL)
 RETURN VARCHAR2;

DBMS_XSLPROCESSOR.VALUEOF(
  n           IN    DBMS_XMLDOM.DOMNODE,
  pattern     IN    VARCHAR2,
  val         OUT   VARCHAR2,
  namespace   IN    VARCHAR2 := NULL);
```

Parameters

Table 238-14 VALUEOF Function and Procedure Parameters

Parameter	Description
`n`	Node whose value is being retrieved
`pattern`	Pattern to use
`val`	Retrieved value
namespace	Namespace to use

239
DBMS_XSTREAM_ADM

This `DBMS_XSTREAM_ADM` package provides interfaces for streaming database changes between an Oracle database and other systems. XStream enables applications to stream out or stream in database changes.

This chapter contains the following topic:

- Overview
- Security Model
- Operational Notes
- Summary of DBMS_XSTREAM_ADM Subprograms

> **See Also:**
>
> - *Oracle Database XStream Guide*
> - *Oracle Call Interface Programmer's Guide*
> - *Oracle Database XStream Java API Reference*

239.1 DBMS_XSTREAM_ADM Overview

The package provides interfaces for configuring outbound servers that stream database changes from an Oracle database to other systems. The package also provides interfaces for configuring inbound servers that stream database changes from other systems to an Oracle database.

In both cases, the database changes are encapsulated in logical change records (LCRs). Also, the other systems can be Oracle systems or a non-Oracle systems, such as non-Oracle databases or file systems.

XStream outbound servers can stream out LCRs from an Oracle database programmatically using C or Java. After receiving the LCRs, the other system can process them in any customized way. For example, the other system can save the contents of the LCRs to a file, send the LCRs to an Oracle database through an XStream inbound server, or generate SQL statements and execute them on any Oracle or non-Oracle databases.

XStream inbound servers accept LCRs from another system and either apply them to an Oracle database or process them in a customized way using apply handlers.

XStream can be used in a multitenant container database (CDB). A CDB is an Oracle database that includes zero, one, or many user-created pluggable databases (PDBs).

See Also:

- *Oracle Database XStream Guide*
- *Oracle Database Concepts* for more information about CDBs and PDBs

239.2 DBMS_XSTREAM_ADM Security Model

To ensure that the user who runs the subprograms in this package has the necessary privileges, configure an XStream administrator and connect as the XStream administrator when using this package.

An administrator must be granted the `DBA` role when the administrator is performing any of the following actions:

- Running the `ADD_OUTBOUND` procedure while connected as a user that is different from the configured connect user for an outbound server
- Running the `ALTER_OUTBOUND` procedure to change the capture user for a capture process or the connect user for an outbound server
- Running the `CREATE_OUTBOUND` procedure, because this procedure creates a capture process
- Running the `ALTER_INBOUND` procedure to change the apply user for an inbound server
- Running the `ADD_INBOUND` procedure while connected as a user that is different from the configured apply user for an inbound server

When the administrator does not need to perform the preceding tasks, the `DBA` role is not required.

See Also:

- GRANT_ADMIN_PRIVILEGE Procedure
- *Oracle Database XStream Guide, Chapter 4, "XStream Out and Security"* for more information about XStream and security.

239.3 DBMS_XSTREAM_ADM Operational Notes

Some subprograms in the `DBMS_APPLY_ADM` package can manage XStream outbound servers, and some subprograms in the `DBMS_APPLY_ADM` package can manage XStream inbound servers.

See Also:

DBMS_APPLY_ADM for details about which subprograms can manage outbound servers and inbound servers

239.4 Summary of DBMS_XSTREAM_ADM Subprograms

This table lists the `DBMS_XSTREAM_ADM` subprograms and briefly describes them.

Table 239-1 DBMS_XSTREAM_ADM Package Subprograms

Subprogram	Description
ADD_COLUMN Procedure	Either adds or removes a declarative rule-based transformation which adds a column to a row logical change record (row LCR) that satisfies the specified rule
ADD_GLOBAL_PROPAGATION_RULES Procedure	Either adds global rules to the positive rule set for a propagation, or adds global rules to the negative rule set for a propagation, and creates the specified propagation if it does not exist
ADD_GLOBAL_RULES Procedure	Adds global rules to either the positive or negative rule set of a capture process or apply process, and creates the specified capture process or apply process if it does not exist
ADD_OUTBOUND Procedure	Creates an XStream outbound server that dequeues LCRs from the specified queue
ADD_SCHEMA_PROPAGATION_RULES Procedure	Either adds schema rules to the positive rule set for a propagation, or adds schema rules to the negative rule set for a propagation, and creates the specified propagation if it does not exist
ADD_SCHEMA_RULES Procedure	Adds rules to a rule set of XStream clients.
ADD_SUBSET_OUTBOUND_RULES Procedure	Adds subset rules to an outbound server configuration
ADD_SUBSET_PROPAGATION_RULES Procedure	Adds subset rules to the positive rule set for a propagation, and creates the specified propagation if it does not exist
ADD_SUBSET_RULES Procedure	Adds subset rules to the positive rule set of a capture process or apply process, and creates the specified capture process or apply process if it does not exist
ADD_TABLE_PROPAGATION_RULES Procedure	Either adds table rules to the positive rule set for a propagation, or adds table rules to the negative rule set for a propagation, and creates the specified propagation if it does not exist
ADD_TABLE_RULES Procedure	This procedure adds rules to a rule set of an XStream client.
ALTER_INBOUND Procedure	Modifies an XStream inbound server
ALTER_OUTBOUND Procedure	Modifies an XStream outbound server
CREATE_INBOUND Procedure	Creates an XStream inbound server and its queue
CREATE_OUTBOUND Procedure	Creates an XStream outbound server, queue, and capture process to enable XStream client applications to stream out Oracle database changes encapsulated in LCRs
DELETE_COLUMN Procedure	Either adds or removes a declarative rule-based transformation which deletes a column from a row LCR that satisfies the specified rule
DROP_INBOUND Procedure	Removes an inbound server configuration

Table 239-1 (Cont.) DBMS_XSTREAM_ADM Package Subprograms

Subprogram	Description
DROP_OUTBOUND Procedure	Removes an outbound server configuration
ENABLE_GG_XSTREAM_FOR_STREAMS Procedure	Enables XStream performance optimizations for Oracle Replication components
GET_MESSAGE_TRACKING Function	Returns the tracking label for the current session
GET_TAG Function	Gets the binary tag for all redo entries generated by the current session
IS_GG_XSTREAM_FOR_STREAMS Function	Returns `TRUE` if XStream performance optimizations are enabled for Oracle Replication components, or returns `FALSE` if XStream performance optimizations are disabled for Oracle Replication components
KEEP_COLUMNS Procedure	Either adds or removes a declarative rule-based transformation which keeps a list of columns in a row LCR that satisfies the specified rule
MERGE_STREAMS Procedure	Merges a stream flowing from one capture process with a stream flowing from another capture process
MERGE_STREAMS_JOB Procedure	Determines whether the original capture process and the cloned capture are within the specified merge threshold and, if they are, runs the `MERGE_STREAMS` procedure to merge the two streams
PURGE_SOURCE_CATALOG Procedure	Removes all Oracle Replication data dictionary information at the local database for the specified object
RECOVER_OPERATION Procedure	Provides options for a split and merge operation that stopped because it encountered an error. This procedure either rolls forward the operation, rolls back the operation, or purges all of the metadata about the operation.
REMOVE_QUEUE Procedure	Removes the specified `ANYDATA` queue
REMOVE_RULE Procedure	Removes the specified rule or all rules from the rule set associated with the specified capture process, apply process, or propagation.
REMOVE_SUBSET_OUTBOUND_RULES Procedure	Removes subset rules from an outbound server configuration
REMOVE_XSTREAM_CONFIGURATION Procedure	Removes the XStream configuration at the local database
RENAME_COLUMN Procedure	Either adds or removes a declarative rule-based transformation which renames a column in a row LCR that satisfies the specified rule
RENAME_SCHEMA Procedure	Either adds or removes a declarative rule-based transformation which renames a schema in a row LCR that satisfies the specified rule
RENAME_TABLE Procedure	Either adds or removes a declarative rule-based transformation which renames a table in a row LCR that satisfies the specified rule
SET_MESSAGE_TRACKING Procedure	Sets the tracking label for logical change records (LCRs) produced by the current session
SET_PARAMETER Procedure	Sets a parameter for an outbound server, an inbound server, or an outbound server's capture process

Table 239-1 (Cont.) DBMS_XSTREAM_ADM Package Subprograms

Subprogram	Description
SET_TAG Procedure	Sets the binary tag for all redo entries subsequently generated by the current session
SET_UP_QUEUE Procedure	Creates a queue table and a queue for use with the capture, propagate, and apply functionality of XStream
SPLIT_STREAMS Procedure	Splits one stream flowing from a capture process off from all of the other streams flowing from the capture process
START_OUTBOUND Procedure	Starts an XStream outbound server
STOP_OUTBOUND Procedure	Stops an XStream outbound server

Note:

All subprograms commit unless specified otherwise.

239.4.1 ADD_COLUMN Procedure

This procedure either adds or removes a declarative rule-based transformation which adds a column to a row logical change record (row LCR) that satisfies the specified rule.

For the transformation to be performed when the specified rule evaluates to `TRUE`, the rule must be in the positive rule set of an XStream client. XStream clients include capture processes, propagations, and apply processes.

This procedure is overloaded. The `column_value` and `column_function` parameters are mutually exclusive.

Note:

- `ADD_COLUMN` transformations cannot add columns of the following data types: `BLOB`, `CLOB`, `NCLOB`, `BFILE`, `LONG`, `LONG RAW`, `ROWID`, user-defined types (including object types, `REF`s, varrays, nested tables), and Oracle-supplied types (including any types, XML types, spatial types, and media types).
- Declarative transformations can transform row LCRs only. Therefore, a DML rule must be specified when you run this procedure. If a DDL rule is specified, then the procedure raises an error.

See Also:

Oracle Database XStream Guide for more information about declarative rule-based transformations

Syntax

```
DBMS_XSTREAM_ADM.ADD_COLUMN(
   rule_name     IN  VARCHAR2,
   table_name    IN  VARCHAR2,
   column_name   IN  VARCHAR2,
   column_value  IN  ANYDATA,
   value_type    IN  VARCHAR2     DEFAULT 'NEW',
   step_number   IN  NUMBER       DEFAULT 0,
   operation     IN  VARCHAR2     DEFAULT 'ADD');

DBMS_XSTREAM_ADM.ADD_COLUMN(
   rule_name        IN  VARCHAR2,
   table_name       IN  VARCHAR2,
   column_name      IN  VARCHAR2,
   column_function  IN  VARCHAR2,
   value_type       IN  VARCHAR2     DEFAULT 'NEW',
   step_number      IN  NUMBER       DEFAULT 0,
   operation        IN  VARCHAR2     DEFAULT 'ADD');
```

Parameters

Table 239-2 ADD_COLUMN Procedure Parameters

Parameter	Description
rule_name	The name of the rule, specified as [*schema_name.*]*rule_name*. If NULL, then the procedure raises an error. For example, to specify a rule in the hr schema named employees12, enter hr.employees12. If the schema is not specified, then the current user is the default.
table_name	The name of the table to which the column is added in the row LCR, specified as [*schema_name.*]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default.
column_name	The name of the column added to each row LCR that satisfies the rule.
column_value	The value of the added column. Specify the appropriate ANYDATA function for the column datatype and the column value. For example, if the datatype of the column being added is NUMBER and the value is NULL, then specify the ANYDATA.ConvertNumber(NULL) function. This parameter cannot be specified if the column_function parameter is specified.

Table 239-2 (Cont.) ADD_COLUMN Procedure Parameters

Parameter	Description
column_function	Either the 'SYSDATE' or the 'SYSTIMESTAMP' SQL function. The 'SYSDATE' SQL function returns the current date and time. 'SYSDATE' uses the time zone of either the database host system or the database, depending on the setting of the TIME_AT_DBTIMEZONE initialization parameter. The datatype of the returned value is DATE, and the format returned depends on the value of the NLS_DATE_FORMAT initialization parameter. The 'SYSTIMESTAMP' SQL function returns the system date, including fractional seconds. 'SYSTIMESTAMP' uses the time zone of either the database host system or the database, depending on the setting of the TIME_AT_DBTIMEZONE initialization parameter. The return type is TIMESTAMP WITH TIME ZONE. The function executes when the rule evaluates to TRUE. This parameter cannot be specified if the column_value parameter is specified. **See Also:** *Oracle Database Reference* for more information about the TIME_AT_DBTIMEZONE parameter.
value_type	Specify 'NEW' to add the column to the new values in the row LCR. Specify 'OLD' to add the column to the old values in the row LCR.
step_number	The order of execution of the transformation. **See Also:** *Oracle Database XStream Guide* for more information about transformation ordering
operation	Specify 'ADD' to add the transformation to the rule. Specify 'REMOVE' to remove the transformation from the rule.

Usage Notes

When 'REMOVE' is specified for the operation parameter, all of the add column declarative rule-based transformations for the specified rule are removed that match the specified table_name, column_name, and step_number parameters. Nulls specified for these parameters act as wildcards. The following table lists the behavior of the ADD_COLUMN procedures when one or more of these parameters is NULL:

table_name	column_name	step_number	Result
NULL	NULL	NULL	Remove all add column transformations for the specified rule.
NULL	NULL	non-NULL	Remove all add column transformations with the specified step_number for the specified rule.
NULL	non-NULL	non-NULL	Remove all add column transformations with the specified column_name and step_number for the specified rule.

table_name	column_name	step_number	Result
non-NULL	NULL	non-NULL	Remove all add column transformations with the specified table_name and step_number for the specified rule.
NULL	non-NULL	NULL	Remove all add column transformations with the specified column_name for the specified rule.
non-NULL	non-NULL	NULL	Remove all add column transformations with the specified table_name and column_name for the specified rule.
non-NULL	NULL	NULL	Remove all add column transformations with the specified table_name for the specified rule.
non-NULL	non-NULL	non-NULL	Remove all add column transformations with the specified table_name, column_name, and step_number for the specified rule.

239.4.2 ADD_GLOBAL_PROPAGATION_RULES Procedure

This procedure either adds global rules to the positive rule set for a propagation, or adds global rules to the negative rule set for a propagation, and creates the specified propagation if it does not exist.

This procedure is overloaded. One version of this procedure contains two OUT parameters, and the other does not.

Syntax

```
DBMS_XSTREAM_ADM.ADD_GLOBAL_PROPAGATION_RULES(
   streams_name             IN   VARCHAR2  DEFAULT NULL,
   source_queue_name        IN   VARCHAR2,
   destination_queue_name   IN   VARCHAR2,
   include_dml              IN   BOOLEAN   DEFAULT TRUE,
   include_ddl              IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr       IN   BOOLEAN   DEFAULT FALSE,
   source_database          IN   VARCHAR2  DEFAULT NULL,
   dml_rule_name            OUT  VARCHAR2,
   ddl_rule_name            OUT  VARCHAR2,
   inclusion_rule           IN   BOOLEAN   DEFAULT TRUE,
   and_condition            IN   VARCHAR2  DEFAULT NULL,
   queue_to_queue           IN   BOOLEAN   DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_GLOBAL_PROPAGATION_RULES(
   streams_name             IN   VARCHAR2  DEFAULT NULL,
   source_queue_name        IN   VARCHAR2,
   destination_queue_name   IN   VARCHAR2,
   include_dml              IN   BOOLEAN   DEFAULT TRUE,
   include_ddl              IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr       IN   BOOLEAN   DEFAULT FALSE,
   source_database          IN   VARCHAR2  DEFAULT NULL,
   inclusion_rule           IN   BOOLEAN   DEFAULT TRUE,
   and_condition            IN   VARCHAR2  DEFAULT NULL,
   queue_to_queue           IN   BOOLEAN   DEFAULT NULL);
```

Parameters

Table 239-3 ADD_GLOBAL_PROPAGATION_RULES Procedure Parameters

Parameter	Description
streams_name	The name of the propagation. Do not specify an owner. If the specified propagation does not exist, then the procedure creates it automatically. If NULL and a propagation exists for the same source queue and destination queue (including database link), then the procedure uses this propagation. If NULL and no propagation exists for the same source queue and destination queue (including database link), then the procedure creates a propagation automatically with a system-generated name.
source_queue_name	The name of the source queue, specified as [*schema_name.*]*queue_name*. The current database must contain the source queue, and the queue must be ANYDATA type. For example, to specify a source queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default.
destination_queue_name	The name of the destination queue, including a database link, specified as [*schema_name.*]*queue_name*[*@dblink_name*], if the destination queue is in a remote database. The queue must be ANYDATA type. For example, to specify a destination queue named streams_queue in the strmadmin schema and use a database link named dbs2.net, enter strmadmin.streams_queue@dbs2.net for this parameter. If the schema is not specified, then the current user is the default. If the database link is omitted, then the procedure uses the global name of the current database, and the source queue and destination queue must be in the same database. **Note:** Connection qualifiers are not allowed.
include_dml	If TRUE, then the procedure creates a rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.

Table 239-3 (Cont.) ADD_GLOBAL_PROPAGATION_RULES Procedure Parameters

Parameter	Description
include_tagged_lcr	If TRUE, then the procedure does not add a condition regarding Oracle Replication tags to the generated rules. Therefore, these rules can evaluate to TRUE regardless of whether a logical change record (LCR) has a non-NULL tag. If the rules are added to the positive rule set for the propagation, then an LCR is always considered for propagation, regardless of whether it has a non-NULL tag. If the rules are added to a positive rule set, then setting this parameter to TRUE is appropriate for a full (for example, standby) copy of a database. If the rules are added to the negative rule set for the propagation, then whether an LCR is discarded does not depend on the tag for the LCR. If FALSE, then the procedure adds a condition to each generated rule that causes the rule to evaluate to TRUE only if an LCR has a NULL Oracle Replication tag. If the rules are added to the positive rule set for the propagation, then an LCR is considered for propagation only when the LCR contains a NULL tag. If the rules are added to a positive rule set, then setting this parameter to FALSE might be appropriate in update-anywhere configurations to avoid sending a change back to its source database. If the rules are added to the negative rule set for the propagation, then an LCR can be discarded only if it has a NULL tag. Usually, specify TRUE for this parameter if the inclusion_rule parameter is set to FALSE.
source_database	The global name of the source database. The source database is where the changes originated. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically. Oracle recommends that you specify a source database for propagation rules.
dml_rule_name	If include_dml is TRUE, then this parameter contains the DML rule name. If include_dml is FALSE, then this parameter contains a NULL.
ddl_rule_name	If include_ddl is TRUE, then this parameter contains the DDL rule name. If include_ddl is FALSE, then this parameter contains a NULL.
inclusion_rule	If inclusion_rule is TRUE, then the procedure adds the rules to the positive rule set for the propagation. If inclusion_rule is FALSE, then the procedure adds the rules to the negative rule set for the propagation. In either case, the system creates the rule set if it does not exist.

Table 239-3 (Cont.) ADD_GLOBAL_PROPAGATION_RULES Procedure Parameters

Parameter	Description
and_condition	If non-NULL, appends the specified condition to the system-generated rule condition using an AND clause in the following way: (*system_condition*) AND (*and_condition*) The variable in the specified condition must be :lcr. For example, to specify that the global rules generated by the procedure evaluate to TRUE only if the Oracle Replication tag is the hexadecimal equivalent of '02', specify the following condition: :lcr.get_tag() = HEXTORAW(''02'') The :lcr in the specified condition is converted to :dml or :ddl, depending on the rule that is being generated. If you are specifying an LCR member subprogram that is dependent on the LCR type (row or DDL), then make sure the procedure only generates the appropriate rule. Specifically, if you specify an LCR member subprogram that is valid only for row LCRs, then specify TRUE for the include_dml parameter and FALSE for the include_ddl parameter. If you specify an LCR member subprogram that is valid only for DDL LCRs, then specify FALSE for the include_dml parameter and TRUE for the include_ddl parameter. **See Also:** Logical Change Record TYPEs
queue_to_queue	If TRUE or NULL, then a new propagation created by this procedure is a queue to queue propagation. A queue-to-queue propagation always has its own propagation job and uses a service for automatic failover when the destination queue is a buffered queue in an Oracle Real Application Clusters (Oracle RAC) database. If FALSE, then a new propagation created by this procedure is a queue-to-dblink propagation. A queue-to-dblink propagation can share a propagation job with other propagations that use the same database link and does not support automatic failover in an Oracle RAC environment. The procedure cannot change the queue to queue property of an exiting propagation. If the specified propagation exists, then the procedure behaves in the following way for each setting: • If TRUE and the specified propagation is not a queue to queue propagation, then the procedure raises an error. • If FALSE and the specified propagation is a queue to queue propagation, then the procedure raises an error. • If NULL, then the procedure does not change the queue to queue property of the propagation.

Usage Notes

This procedure configures propagation using the current user. Only one propagation is allowed between a particular source queue and destination queue.

This procedure creates DML and DDL rules automatically based on include_dml and include_ddl parameter values, respectively. Each rule has a system-generated rule name that consists of the database name with a sequence number appended to it. The sequence number

is used to avoid naming conflicts. If the database name plus the sequence number is too long, then the database name is truncated. A propagation uses the rules for filtering.

Examples

The following is an example of a global rule condition created for DML changes:

```
(:dml.is_null_tag() = 'Y' and :dml.get_source_database_name() = 'DBS1.NET' )
```

239.4.3 ADD_GLOBAL_RULES Procedure

This procedure adds rules to an XStream clients rule set.

It adds rules to a rule set of one of the following types of XStream clients:

- When the `streams_type` parameter is set to `capture`, this procedure adds capture process rules for capturing changes to an entire database.

  This procedure creates the specified capture process if it does not exist.
- When the `streams_type` parameter is set to `apply` and the `streams_name` parameter specifies the name of an apply process, outbound server, or inbound server, this procedure adds apply rules for applying all logical change records (LCRs) it receives. The rules can specify that the LCRs must be from a particular source database.

  This procedure creates an apply process if no apply process, outbound server, or inbound server exists with the specified `streams_name`. This procedure can add rules to an outbound server or inbound server, but it cannot create an outbound server or inbound server.

This procedure is overloaded. One version of this procedure contains two `OUT` parameters, and the other does not.

> **Note:**
>
> If you add global rules to the positive rule set for a capture process, then make sure you add rules to the negative capture process rule set to exclude database objects that are not supported by Oracle Replication. Query the `DBA_XSTREAM_OUT_SUPPORT_MODE` data dictionary view to determine which database objects are not supported by Oracle Replication. If unsupported database objects are not excluded, then capture errors will result.

Syntax

```
DBMS_XSTREAM_ADM.ADD_GLOBAL_RULES(
   streams_type           IN   VARCHAR2,
   streams_name           IN   VARCHAR2  DEFAULT NULL,
   queue_name             IN   VARCHAR2  DEFAULT 'streams_queue',
   include_dml            IN   BOOLEAN   DEFAULT TRUE,
   include_ddl            IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr     IN   BOOLEAN   DEFAULT FALSE,
   source_database        IN   VARCHAR2  DEFAULT NULL,
   dml_rule_name          OUT  VARCHAR2,
   ddl_rule_name          OUT  VARCHAR2,
   inclusion_rule         IN   BOOLEAN   DEFAULT TRUE,
   and_condition          IN   VARCHAR2  DEFAULT NULL,
   source_root_name       IN   VARCHAR2  DEFAULT NULL,
   source_container_name  IN   VARCHAR2  DEFAULT NULL);
```

```
DBMS_XSTREAM_ADM.ADD_GLOBAL_RULES(
   streams_type           IN   VARCHAR2,
   streams_name           IN   VARCHAR2  DEFAULT NULL,
   queue_name             IN   VARCHAR2  DEFAULT 'streams_queue',
   include_dml            IN   BOOLEAN   DEFAULT TRUE,
   include_ddl            IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr     IN   BOOLEAN   DEFAULT FALSE,
   source_database        IN   VARCHAR2  DEFAULT NULL,
   inclusion_rule         IN   BOOLEAN   DEFAULT TRUE,
   and_condition          IN   VARCHAR2  DEFAULT NULL,
   source_root_name       IN   VARCHAR2  DEFAULT NULL,
   source_container_name  IN   VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-4 ADD_GLOBAL_RULES Procedure Parameters

Parameter	Description
streams_type	The type of XStream client: • Specify capture for a capture process. • Specify apply for an apply process.
streams_name	The name of the capture process or apply process. Do not specify an owner. If NULL, if streams_type is capture, and if one relevant capture process for the queue exists, then the relevant XStream client is used. If no relevant XStream client exists for the queue, then an XStream client is created automatically with a system-generated name. If NULL and multiple XStream clients of the specified streams_type for the queue exist, then the procedure raises an error. If NULL, if streams_type is apply, and if one relevant apply process exists, then the procedure uses the relevant apply process. The relevant apply process is identified in one of the following ways: • If one existing apply process has the source database specified in source_database and uses the queue specified in queue_name, then the procedure uses this apply process. • If source_database is NULL and one existing apply process is using the queue specified in queue_name, then the procedure uses this apply process. If NULL and no relevant apply process exists, then the procedure creates an apply process automatically with a system-generated name. If NULL and multiple relevant apply processes exist, then the procedure raises an error. Each apply process must have a unique name.
queue_name	The name of the local queue, specified as [*schema_name.*]*queue_name*. The current database must contain the queue, and the queue must be ANYDATA type. For example, to specify a queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default. For capture process rules, this is the queue into which a capture process enqueues LCRs. For outbound server rules, this is the queue from which the outbound server dequeues LCRs. For inbound server rules, this is the queue into which an inbound server enqueues error transactions.

Table 239-4 (Cont.) ADD_GLOBAL_RULES Procedure Parameters

Parameter	Description
include_dml	If TRUE, then the procedure creates a rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.
include_tagged_lcr	If TRUE, then the procedure does not add a condition regarding Oracle Replication tags to the generated rules. Therefore, these rules can evaluate to TRUE regardless of whether a redo entry or LCR has a non-NULL tag. If the rules are added to the positive rule set for the process, then a redo entry is always considered for capture, and an LCR is always considered for apply, regardless of whether the redo entry or LCR has a non-NULL tag. If the rules are added to a positive rule set, then setting this parameter to TRUE is appropriate for a full (for example, standby) copy of a database. If the rules are added to the negative rule set for the process, then whether a redo entry or LCR is discarded does not depend on the tag. If FALSE, then the procedure adds a condition to each generated rule that causes the rule to evaluate to TRUE only if a redo entry or LCR has a NULL Oracle Replication tag. If the rules are added to the positive rule set for the process, then a redo entry is considered for capture, and an LCR is considered for apply, only when the redo entry or LCR contains a NULL tag. If the rules are added to a positive rule set, then setting this parameter to FALSE might be appropriate in update-anywhere configurations to avoid sending a change back to its source database. If the rules are added to the negative rule set for the process, then a redo entry or LCR can be discarded only if it has a NULL tag. Usually, specify TRUE for this parameter if the inclusion_rule parameter is set to FALSE.
source_database	The global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. For capture process rules, specify NULL or the global name of the local database if you are creating a capture process locally at the source database. If you are adding rules to a downstream capture process rule set at a downstream database, then specify the source database of the changes that will be captured. For apply process rules, specify the source database of the changes that will be applied by the apply process. The source database is the database where the changes originated. If an apply process applies captured LCRs, then the apply process can apply LCRs from only one capture process at one source database. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically.
dml_rule_name	If include_dml is TRUE, then this parameter contains the DML rule name. If include_dml is FALSE, then this parameter contains a NULL.
ddl_rule_name	If include_ddl is TRUE, then this parameter contains the DDL rule name. If include_ddl is FALSE, then this parameter contains a NULL.

Table 239-4 (Cont.) ADD_GLOBAL_RULES Procedure Parameters

Parameter	Description
inclusion_rule	If inclusion_rule is TRUE, then the procedure adds the rules to the positive rule set for the XStream client. If inclusion_rule is FALSE, then the procedure adds the rules to the negative rule set for the XStream client. In either case, the system creates the rule set if it does not exist.
and_condition	If non-NULL, appends the specified condition to the system-generated rule condition using an AND clause in the following way: (*system_condition*) AND (*and_condition*) The variable in the specified condition must be :lcr. For example, to specify that the global rules generated by the procedure evaluate to TRUE only if the Oracle Replication tag is the hexadecimal equivalent of '02', specify the following condition: :lcr.get_tag() = HEXTORAW(''02'') The :lcr in the specified condition is converted to :dml or :ddl, depending on the rule that is being generated. If you are specifying an LCR member subprogram that is dependent on the LCR type (row or DDL), then make sure this procedure only generates the appropriate rule. Specifically, if you specify an LCR member subprogram that is valid only for row LCRs, then specify TRUE for the include_dml parameter and FALSE for the include_ddl parameter. If you specify an LCR member subprogram that is valid only for DDL LCRs, then specify FALSE for the include_dml parameter and TRUE for the include_ddl parameter. **See Also:** Logical Change Record TYPEs
source_root_name	The global name of the root in the source CDB. For example, mycdb.example.com. If this parameter is NULL, then the global name of the root in the local CDB is used. If you are configuring downstream capture, then this parameter must be a non-NULL value, and it must specify the global name of the root in the remote source CDB. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. **Note:** This parameter only applies to a CDB.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.

Usage Notes

This procedure creates DML and DDL rules automatically based on include_dml and include_ddl parameter values, respectively. Each rule has a system-generated rule name that consists of the database name with a sequence number appended to it. The sequence number is used to avoid naming conflicts. If the database name plus the sequence number is too long,

then the database name is truncated. A capture process or apply process uses the rules for filtering.

> **See Also:**
>
> - "Operational Notes"
> - "Security Model"

Examples

The following is an example of a global rule condition created for DML changes:

```
(:dml.is_null_tag() = 'Y' and :dml.get_source_database_name() = 'DBS1.NET' )
```

239.4.4 ADD_OUTBOUND Procedure

This procedure creates an XStream outbound server that dequeues LCRs from the specified queue. The outbound server streams out the LCRs to an XStream client application.

This procedure creates neither a capture process nor a queue. To create an outbound server, a capture process, and a queue with one procedure call, use the CREATE_OUTBOUND Procedure.

To create the capture process individually, use one of the following packages:

- `DBMS_XSTREAM_ADM`
- `DBMS_CAPTURE_ADM`

To create a queue individually, use the `SET_UP_QUEUE` procedure in the `DBMS_XSTREAM_ADM` package.

This procedure is overloaded. One `table_names` parameter is type `VARCHAR2` and the other `table_names` parameter is type `DBMS_UTILITY.UNCL_ARRAY`. Also, one `schema_names` parameter is type `VARCHAR2` and the other `schema_names` parameter is type `DBMS_UTILITY.UNCL_ARRAY`. These parameters enable you to enter the lists of tables and schemas in different ways and are mutually exclusive.

> **Note:**
>
> - A client application can create multiple sessions. Each session can attach to only one outbound server, and each outbound server can serve only one session at a time. However, different client application sessions can connect to different outbound servers. See *Oracle Call Interface Programmer's Guide* and *Oracle Database XStream Java API Reference* for information about attaching to an outbound server.
> - This procedure enables the outbound server that it creates.
> - Starting with Oracle Database 11*g* Release 2 (11.2.0.2), the `capture_name`, `start_scn`, and `start_time` parameters are included in this procedure.

Syntax

```
DBMS_XSTREAM_ADM.ADD_OUTBOUND(
   server_name             IN  VARCHAR2,
   queue_name              IN  VARCHAR2   DEFAULT NULL,
   source_database         IN  VARCHAR2   DEFAULT NULL,
   table_names             IN  DBMS_UTILITY.UNCL_ARRAY,
   schema_names            IN  DBMS_UTILITY.UNCL_ARRAY,
   connect_user            IN  VARCHAR2   DEFAULT NULL,
   comment                 IN  VARCHAR2   DEFAULT NULL,
   capture_name            IN  VARCHAR2   DEFAULT NULL,
   start_scn               IN  NUMBER     DEFAULT NULL,
   start_time              IN  TIMESTAMP  DEFAULT NULL,
   include_dml             IN  BOOLEAN    DEFAULT TRUE,
   include_ddl             IN  BOOLEAN    DEFAULT FALSE,
   source_root_name        IN  VARCHAR2   DEFAULT NULL,
   source_container_name   IN  VARCHAR2   DEFAULT NULL,
   lcrid_version           IN  NUMBER     DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_OUTBOUND(
   server_name             IN  VARCHAR2,
   queue_name              IN  VARCHAR2   DEFAULT NULL,
   source_database         IN  VARCHAR2   DEFAULT NULL,
   table_names             IN  VARCHAR2   DEFAULT NULL,
   schema_names            IN  VARCHAR2   DEFAULT NULL,
   connect_user            IN  VARCHAR2   DEFAULT NULL,
   comment                 IN  VARCHAR2   DEFAULT NULL,
   capture_name            IN  VARCHAR2   DEFAULT NULL,
   start_scn               IN  NUMBER     DEFAULT NULL,
   start_time              IN  TIMESTAMP  DEFAULT NULL,
   include_dml             IN  BOOLEAN    DEFAULT TRUE,
   include_ddl             IN  BOOLEAN    DEFAULT FALSE,
   source_root_name        IN  VARCHAR2   DEFAULT NULL,
   source_container_name   IN  VARCHAR2   DEFAULT NULL,
   lcrid_version           IN  NUMBER     DEFAULT NULL);
```

Parameters

Table 239-5 ADD_OUTBOUND Procedure Parameters

Parameter	Description
server_name	The name of the outbound server being created. A NULL specification is not allowed. Do not specify an owner. The specified name must not match the name of an existing outbound server, inbound server, apply process, or messaging client. **Note:** The server_name setting cannot exceed 30 bytes, and it cannot be altered after the outbound server is created.
queue_name	The name of the local queue from which the outbound server dequeues LCRs, specified as [*schema_name.*]*queue_name*. The current database must contain the queue, and the queue must be ANYDATA type. For example, to specify a queue named xstream_queue in the xstrmadmin schema, enter xstrmadmin.xstream_queue for this parameter. If the schema is not specified, then the current user is the default. If NULL, the procedure raises an error.

Table 239-5 (Cont.) ADD_OUTBOUND Procedure Parameters

Parameter	Description
source_database	The global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically.
table_names	The tables for which data manipulation language (DML) and data definition language (DDL) changes are streamed out to the XStream client application. The tables can be specified in the following ways: • Comma-delimited list of type VARCHAR2. • A PL/SQL associative array of type DBMS_UTILITY.UNCL_ARRAY, where each element is the name of a table. Specify the first table in position 1. The last position must be NULL. Each table should be specified as [*schema_name*.]*table_name*. For example, you can specify hr.employees. If the schema is not specified, then the current user is the default. **See Also:** "Usage Notes" for more information about this parameter
schema_names	The schemas for which DML and DDL changes are streamed out to the XStream client application. The schemas can be specified in the following ways: • Comma-delimited list of type VARCHAR2. • A PL/SQL associative array of type DBMS_UTILITY.UNCL_ARRAY, where each element is the name of a schema. Specify the first schema in position 1. The last position must be NULL. **Note:** This procedure does not concatenate the schema_names parameter with the table_names parameter. To specify tables, enter fully qualified table names in the table_names parameter (*schema_name.table_name*). **See Also:** "Usage Notes" for more information about this parameter
connect_user	The user who can attach to the specified outbound server to retrieve the LCR stream. The client application must attach to the outbound server as the specified connect user. See "CREATE_OUTBOUND Procedure" for information about the privileges required by a connect user. If NULL, then the current user is the default.
comment	An optional comment associated with the outbound server.
capture_name	The name of the capture process configured to capture changes for the outbound server. Do not specify an owner. If the specified name matches the name of an existing capture process for another outbound server, then the procedure uses the existing capture process and adds the rules for capturing changes to the database to the positive capture process rule set. If the specified name matches the name of an existing capture process for an apply process, then an error is raised. If the specified name does not match the name of an existing capture process, then an error is raised. If NULL, then the outbound server is created without a capture process.

Table 239-5 (Cont.) ADD_OUTBOUND Procedure Parameters

Parameter	Description
start_scn	A valid system change number (SCN) for the database from which the capture process starts capturing changes. If the capture_name parameter is NULL, then this parameter is ignored. If NULL and the capture_name parameter is non-NULL, then the start SCN of the capture process is not changed. An error is returned if an invalid SCN is specified. The start_scn and start_time parameters are mutually exclusive.
start_time	A valid time from which the capture process starts capturing changes. If the capture_name parameter is NULL, then this parameter is ignored. If NULL and the capture_name parameter is non-NULL, then the start SCN of the capture process is not changed. The start_scn and start_time parameters are mutually exclusive.
include_dml	If TRUE, then the procedure creates a rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.
source_root_name	The global name of the root in the source CDB. For example, mycdb.example.com. If this parameter is NULL, then the global name of the root in the local CDB is used. If you are configuring downstream capture, then this parameter must be a non-NULL value, and it must specify the global name of the root in the remote source CDB. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. **Note:** This parameter only applies to a CDB.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.
lcrid_version	The LCRID version for captured LCRs, either 1 or 2. If 2, then the LCRs are compatible with a database with its compatibility level at 12.2.0 or higher. If 1, then the LCRs are compatible with a database with its compatibility level at 12.1.0 or lower. If NULL, the default, and the database compatibility level is 12.2.0 or higher, then the lcrid_version is set to 2 internally. If the database compatibility level is 12.1.0 or lower, then the lcrid_version is set to 1 internally.

Usage Notes

The following list describes the behavior of the outbound server for various combinations of the table_names and schema_names parameters:

- If both the table_names and schema_names parameters are NULL or empty, then the outbound server streams all DML and DDL changes to the client application.

This procedure is overloaded. The `table_names` and `schema_names` parameters are defaulted to `NULL`. Do not specify `NULL` for both `table_names` and `schema_names` in the same call; otherwise, error `PLS-00307` is returned.

- If both the `table_names` and `schema_names` parameters are specified, then the outbound server streams DML and DDL changes for the specified tables and schemas.
- If the `table_names` parameter is specified and the `schema_names` parameter is `NULL` or empty, then the outbound server streams DML and DDL changes for the specified tables.
- If the `table_names` parameter is `NULL` or empty and the `schema_names` parameter is specified, then the outbound server streams DML and DDL changes for the specified schemas.

For the procedure that uses the `DBMS_UTILITY.UNCL_ARRAY` type for the `table_names` and `schema_names` parameters, both parameters must be specified. To specify only tables, the `schema_names` parameter must be specified and empty. To specify only schemas, the `table_names` parameter must be specified and empty.

> **Note:**
>
> An empty array includes one `NULL` entry.

239.4.5 ADD_SCHEMA_PROPAGATION_RULES Procedure

This procedure either adds schema rules to the positive rule set for a propagation, or adds schema rules to the negative rule set for a propagation, and creates the specified propagation if it does not exist.

This procedure is overloaded. One version of this procedure contains two `OUT` parameters, and the other does not.

Syntax

```
DBMS_XSTREAM_ADM.ADD_SCHEMA_PROPAGATION_RULES(
   schema_name             IN   VARCHAR2,
   streams_name            IN   VARCHAR2  DEFAULT NULL,
   source_queue_name       IN   VARCHAR2,
   destination_queue_name  IN   VARCHAR2,
   include_dml             IN   BOOLEAN   DEFAULT TRUE,
   include_ddl             IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr      IN   BOOLEAN   DEFAULT FALSE,
   source_database         IN   VARCHAR2  DEFAULT NULL,
   dml_rule_name           OUT  VARCHAR2,
   ddl_rule_name           OUT  VARCHAR2,
   inclusion_rule          IN   BOOLEAN   DEFAULT TRUE,
   and_condition           IN   VARCHAR2  DEFAULT NULL,
   queue_to_queue          IN   BOOLEAN   DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_SCHEMA_PROPAGATION_RULES(
   schema_name             IN   VARCHAR2,
   streams_name            IN   VARCHAR2  DEFAULT NULL,
   source_queue_name       IN   VARCHAR2,
   destination_queue_name  IN   VARCHAR2,
   include_dml             IN   BOOLEAN   DEFAULT TRUE,
   include_ddl             IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr      IN   BOOLEAN   DEFAULT FALSE,
```

```
  source_database         IN   VARCHAR2  DEFAULT NULL,
  inclusion_rule          IN   BOOLEAN   DEFAULT TRUE,
  and_condition           IN   VARCHAR2  DEFAULT NULL,
  queue_to_queue          IN   BOOLEAN   DEFAULT NULL);
```

Parameters

Table 239-6 ADD_SCHEMA_PROPAGATION_RULES Procedure Parameters

Parameter	Description
schema_name	The name of the schema. For example, hr.
streams_name	The name of the propagation. Do not specify an owner. If the specified propagation does not exist, then the procedure creates it automatically. If NULL and a propagation exists for the same source queue and destination queue (including database link), then the procedure uses this propagation. If NULL and no propagation exists for the same source queue and destination queue (including database link), then the procedure creates a propagation automatically with a system-generated name.
source_queue_name	The name of the source queue, specified as [*schema_name*.]*queue_name*. The current database must contain the source queue, and the queue must be ANYDATA type. For example, to specify a source queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default.
destination_queue_name	The name of the destination queue, including a database link, specified as [*schema_name*.]*queue_name*[*@dblink_name*], if the destination queue is in a remote database. The queue must be ANYDATA type. For example, to specify a destination queue named streams_queue in the strmadmin schema and use a database link named dbs2.net, enter strmadmin.streams_queue@dbs2.net for this parameter. If the schema is not specified, then the current user is the default. If the database link is omitted, then the procedure uses the global name of the current database, and the source queue and destination queue must be in the same database. **Note:** Connection qualifiers are not allowed.
include_dml	If TRUE, then the procedure creates a rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.

Table 239-6 (Cont.) ADD_SCHEMA_PROPAGATION_RULES Procedure Parameters

Parameter	Description
include_tagged_lcr	If TRUE, then the procedure does not add a condition regarding Oracle Replication tags to the generated rules. Therefore, these rules can evaluate to TRUE regardless of whether a logical change record (LCR) has a non-NULL tag. If the rules are added to the positive rule set for the propagation, then an LCR is always considered for propagation, regardless of whether it has a non-NULL tag. If the rules are added to a positive rule set, then setting this parameter to TRUE is appropriate for a full (for example, standby) copy of a database. If the rules are added to the negative rule set for the propagation, then whether an LCR is discarded does not depend on the tag for the LCR. If FALSE, then the procedure adds a condition to each generated rule that causes the rule to evaluate to TRUE only if an LCR has a NULL Oracle Replication tag. If the rules are added to the positive rule set for the propagation, then an LCR is considered for propagation only when the LCR contains a NULL tag. If the rules are added to a positive rule set, then setting this parameter to FALSE might be appropriate in update-anywhere configurations to avoid sending a change back to its source database. If the rules are added to the negative rule set for the propagation, then an LCR can be discarded only if it has a NULL tag. Usually, specify TRUE for this parameter if the inclusion_rule parameter is set to FALSE.
source_database	The global name of the source database. The source database is where the change originated. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically. Oracle recommends that you specify a source database for propagation rules.
dml_rule_name	If include_dml is TRUE, then this parameter contains the DML rule name. If include_dml is FALSE, then this parameter contains a NULL.
ddl_rule_name	If include_ddl is TRUE, then this parameter contains the DDL rule name. If include_ddl is FALSE, then this parameter contains a NULL.
inclusion_rule	If inclusion_rule is TRUE, then the procedure adds the rules to the positive rule set for the propagation. If inclusion_rule is FALSE, then the procedure adds the rules to the negative rule set for the propagation. In either case, the system creates the rule set if it does not exist.

Table 239-6 (Cont.) ADD_SCHEMA_PROPAGATION_RULES Procedure Parameters

Parameter	Description
and_condition	If non-NULL, appends the specified condition to the system-generated rule condition using an AND clause in the following way: (*system_condition*) AND (*and_condition*) The variable in the specified condition must be :lcr. For example, to specify that the schema rules generated by the procedure evaluate to TRUE only if the Oracle Replication tag is the hexadecimal equivalent of '02', specify the following condition: :lcr.get_tag() = HEXTORAW(''02'') The :lcr in the specified condition is converted to :dml or :ddl, depending on the rule that is being generated. If you are specifying an LCR member subprogram that is dependent on the LCR type (row or DDL), then make sure this procedure only generates the appropriate rule. Specifically, if you specify an LCR member subprogram that is valid only for row LCRs, then specify TRUE for the include_dml parameter and FALSE for the include_ddl parameter. If you specify an LCR member subprogram that is valid only for DDL LCRs, then specify FALSE for the include_dml parameter and TRUE for the include_ddl parameter. **See Also:** Logical Change Record TYPEs
queue_to_queue	If TRUE or NULL, then a new propagation created by this procedure is a queue to queue propagation. A queue-to-queue propagation always has its own propagation job and uses a service for automatic failover when the destination queue is a buffered queue in an Oracle Real Application Clusters (Oracle RAC) database. If FALSE, then a new propagation created by this procedure is a queue-to-dblink propagation. A queue-to-dblink propagation can share a propagation job with other propagations that use the same database link and does not support automatic failover in an Oracle RAC environment. This procedure cannot change the queue to queue property of an exiting propagation. If the specified propagation exists, then the procedure behaves in the following way for each setting: • If TRUE and the specified propagation is not a queue to queue propagation, then the procedure raises an error. • If FALSE and the specified propagation is a queue to queue propagation, then the procedure raises an error. • If NULL, then the procedure does not change the queue to queue property of the propagation.

Usage Notes

This procedure configures propagation using the current user. Only one propagation is allowed between a particular source queue and destination queue.

This procedure creates DML and DDL rules automatically based on include_dml and include_ddl parameter values, respectively. Each rule has a system-generated rule name that consists of the schema name with a sequence number appended to it. The sequence number

is used to avoid naming conflicts. If the schema name plus the sequence number is too long, then the schema name is truncated. A propagation uses the rules for filtering.

Examples

The following is an example of a schema rule condition created for DML changes:

```
((:dml.get_object_owner() = 'HR') and :dml.is_null_tag() = 'Y'
and :dml.get_source_database_name() = 'DBS1.NET' )
```

239.4.6 ADD_SCHEMA_RULES Procedure

This procedures adds rules to a rule set of one of the following types of XStream clients:

- When the `streams_type` parameter is set to `capture`, this procedure adds capture process rules for capturing changes to a specified schema.

  This procedure creates the specified capture process if it does not exist.

- When the `streams_type` parameter is set to `apply` and the `streams_name` parameter specifies the name of an apply process, outbound server, or inbound server, this procedure adds apply rules for applying logical change records (LCRs) that contain changes to a specified schema. The rules can specify that the LCRs must be from a particular source database.

  This procedure creates an apply process if no apply process, outbound server, or inbound server exists with the specified `streams_name`. This procedure can add rules to an outbound server or inbound server, but it cannot create an outbound server or inbound server.

This procedure is overloaded. One version of this procedure contains two `OUT` parameters, and the other does not.

> **Note:**
>
> If you add schema rules to the positive rule set for a capture process, then make sure you add rules to the negative capture process rule set to exclude database objects in the schema that are not supported by Oracle Replication. Query the `DBA_XSTREAM_OUT_SUPPORT_MODE` data dictionary view to determine which database objects are not supported by Oracle Replication. If unsupported database objects are not excluded, then capture errors will result.

Syntax

```
DBMS_XSTREAM_ADM.ADD_SCHEMA_RULES(
   schema_name             IN   VARCHAR2,
   streams_type            IN   VARCHAR2,
   streams_name            IN   VARCHAR2  DEFAULT NULL,
   queue_name              IN   VARCHAR2  DEFAULT 'streams_queue',
   include_dml             IN   BOOLEAN   DEFAULT TRUE,
   include_ddl             IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr      IN   BOOLEAN   DEFAULT FALSE,
   source_database         IN   VARCHAR2  DEFAULT NULL,
   dml_rule_name           OUT  VARCHAR2,
   ddl_rule_name           OUT  VARCHAR2,
   inclusion_rule          IN   BOOLEAN   DEFAULT TRUE,
   and_condition           IN   VARCHAR2  DEFAULT NULL,
   source_root_name        IN  VARCHAR2  DEFAULT NULL,
```

```
   source_container_name  IN  VARCHAR2  DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_SCHEMA_RULES(
   schema_name             IN   VARCHAR2,
   streams_type            IN   VARCHAR2,
   streams_name            IN   VARCHAR2  DEFAULT NULL,
   queue_name              IN   VARCHAR2  DEFAULT 'streams_queue',
   include_dml             IN   BOOLEAN   DEFAULT TRUE,
   include_ddl             IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr      IN   BOOLEAN   DEFAULT FALSE,
   source_database         IN   VARCHAR2  DEFAULT NULL,
   inclusion_rule          IN   BOOLEAN   DEFAULT TRUE,
   and_condition           IN   VARCHAR2  DEFAULT NULL,
   source_root_name        IN  VARCHAR2  DEFAULT NULL,
   source_container_name   IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-7 ADD_SCHEMA_RULES Procedure Parameters

Parameter	Description
schema_name	The name of the schema. For example, hr. You can specify a schema that does not yet exist, because Oracle Replication does not validate the existence of the schema.
streams_type	The type of XStream client: • Specify capture for a capture process. • Specify apply for an apply process.
streams_name	The name of the capture process or apply process. Do not specify an owner. If NULL, if streams_type is capture, and if one relevant capture process for the queue exists, then the relevant XStream client is used. If no relevant XStream client exists for the queue, then an XStream client is created automatically with a system-generated name. If NULL and multiple XStream clients of the specified streams_type for the queue exist, then the procedure raises an error. If NULL, if streams_type is apply, and if one relevant apply process exists, then the procedure uses the relevant apply process. The relevant apply process is identified in one of the following ways: • If one existing apply process has the source database specified in source_database and uses the queue specified in queue_name, then the procedure uses this apply process. • If source_database is NULL and one existing apply process is using the queue specified in queue_name, then the procedure uses this apply process. If NULL and no relevant apply process exists, then the procedure creates an apply process automatically with a system-generated name. If NULL and multiple relevant apply processes exist, then the procedure raises an error. Each apply process must have a unique name.

Table 239-7 (Cont.) ADD_SCHEMA_RULES Procedure Parameters

Parameter	Description
queue_name	The name of the local queue, specified as [*schema_name.*]*queue_name*. The current database must contain the queue, and the queue must be ANYDATA type. For example, to specify a queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default. For capture process rules, this is the queue into which a capture process enqueues LCRs. For outbound server rules, this is the queue from which the outbound server dequeues LCRs. For inbound server rules, this is the queue into which an inbound server enqueues error transactions.
include_dml	If TRUE, then the procedure creates a rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.
include_tagged_lcr	If TRUE, then the procedure does not add a condition regarding Oracle Replication tags to the generated rules. Therefore, these rules can evaluate to TRUE regardless of whether a redo entry or LCR has a non-NULL tag. If the rules are added to the positive rule set for the process, then a redo entry is always considered for capture, and an LCR is always considered for apply, regardless of whether the redo entry or LCR has a non-NULL tag. If the rules are added to a positive rule set, then setting this parameter to TRUE is appropriate for a full (for example, standby) copy of a database. If the rules are added to the negative rule set for the process, then whether a redo entry or LCR is discarded does not depend on the tag. If FALSE, then the procedure adds a condition to each generated rule that causes the rule to evaluate to TRUE only if a redo entry or LCR has a NULL Oracle Replication tag. If the rules are added to the positive rule set for the process, then a redo entry is considered for capture, and an LCR is considered for apply, only when the redo entry or LCR contains a NULL tag. If the rules are added to a positive rule set, then setting this parameter to FALSE might be appropriate in update-anywhere configurations to avoid sending a change back to its source database. If the rules are added to the negative rule set for the process, then a redo entry or LCR can be discarded only if it has a NULL tag. Usually, specify TRUE for this parameter if the inclusion_rule parameter is set to FALSE.

Table 239-7 (Cont.) ADD_SCHEMA_RULES Procedure Parameters

Parameter	Description
source_database	The global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. For capture process rules, specify NULL or the global name of the local database if you are creating a capture process locally at the source database. If you are adding rules to a downstream capture process rule set at a downstream database, then specify the source database of the changes that will be captured. For apply process rules, specify the source database of the changes that will be applied by the apply process. The source database is the database where the changes originated. If an apply process applies captured LCRs, then the apply process can apply LCRs from only one capture process at one source database. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically.
dml_rule_name	If include_dml is TRUE, then this parameter contains the DML rule name. If include_dml is FALSE, then this parameter contains a NULL.
ddl_rule_name	If include_ddl is TRUE, then this parameter contains the DDL rule name. If include_ddl is FALSE, then this parameter contains a NULL.
inclusion_rule	If inclusion_rule is TRUE, then the procedure adds the rules to the positive rule set for the XStream client. If inclusion_rule is FALSE, then the procedure adds the rules to the negative rule set for the XStream client. In either case, the system creates the rule set if it does not exist.

and_condition	If non-NULL, appends the specified condition to the system-generated rule condition using an AND clause in the following way:
	`(system_condition) AND (and_condition)`
	The variable in the specified condition must be :lcr. For example, to specify that the schema rules generated by the procedure evaluate to TRUE only if the Oracle Replication tag is the hexadecimal equivalent of '02', specify the following condition:
	`:lcr.get_tag() = HEXTORAW(''02'')`
	The :lcr in the specified condition is converted to :dml or :ddl, depending on the rule that is being generated. If you are specifying an LCR member subprogram that is dependent on the LCR type (row or DDL), then make sure this procedure only generates the appropriate rule.
	Specifically, if you specify an LCR member subprogram that is valid only for row LCRs, then specify TRUE for the include_dml parameter and FALSE for the include_ddl parameter. If you specify an LCR member subprogram that is valid only for DDL LCRs, then specify FALSE for the include_dml parameter and TRUE for the include_ddl parameter.
	See Also: Logical Change Record TYPEs

Table 239-7 (Cont.) ADD_SCHEMA_RULES Procedure Parameters

Parameter	Description
source_root_name	The global name of the root in the source CDB. For example, mycdb.example.com. If this parameter is NULL, then the global name of the root in the local CDB is used. If you are configuring downstream capture, then this parameter must be a non-NULL value, and it must specify the global name of the root in the remote source CDB. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. **Note:** This parameter only applies to a CDB.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.

Usage Notes

This procedure creates DML and DDL rules automatically based on include_dml and include_ddl parameter values, respectively. Each rule has a system-generated rule name that consists of the schema name with a sequence number appended to it. The sequence number is used to avoid naming conflicts. If the schema name plus the sequence number is too long, then the schema name is truncated. A capture process or apply process uses the rules for filtering.

See Also:

- "Operational Notes"
- "Security Model"

Examples

The following is an example of a schema rule condition created for DML changes:

```
((:dml.get_object_owner() = 'HR') and :dml.is_null_tag() = 'Y'
and :dml.get_source_database_name() = 'DBS1.NET' )
```

239.4.7 ADD_SUBSET_OUTBOUND_RULES Procedure

This procedure adds subset rules to an outbound server configuration. Subset rules instruct the outbound server to stream out a subset of the changes to the specified tables. Outbound servers can stream out a subset of both rows and columns.

This procedure is overloaded. One column_list parameter is type VARCHAR2 and the other column_list parameter is type DBMS_UTILITY.LNAME_ARRAY. These parameters enable you to enter the list of columns in different ways and are mutually exclusive.

> **Note:**
>
> This procedure does not add rules to the outbound server's capture process.

Syntax

```
DBMS_XSTREAM_ADM.ADD_SUBSET_OUTBOUND_RULES(
   server_name      IN VARCHAR2,
   table_name       IN VARCHAR2,
   condition        IN VARCHAR2  DEFAULT NULL,
   column_list      IN DBMS_UTILITY.LNAME_ARRAY,
   keep             IN BOOLEAN   DEFAULT TRUE,
   source_database  IN VARCHAR2  DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_SUBSET_OUTBOUND_RULES(
   server_name      IN VARCHAR2,
   table_name       IN VARCHAR2,
   condition        IN VARCHAR2  DEFAULT NULL,
   column_list      IN VARCHAR2  DEFAULT NULL,
   keep             IN BOOLEAN   DEFAULT TRUE,
   source_database  IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-8 ADD_SUBSET_OUTBOUND_RULES Procedure Parameters

Parameter	Description
`server_name`	The name of the outbound server to which rules are being added. Specify an existing outbound server. Do not specify an owner.
`table_name`	The name of the table specified as `[schema_name.]object_name`. For example, you can specify `hr.employees`. If the schema is not specified, then the current user is the default. If the outbound server configuration uses a local capture process, then the table must exist at the local source database. If the outbound server configuration uses a downstream capture process, then the table must exist at both the source database and at the downstream capture database. The specified table cannot have any LOB, `LONG`, or `LONG RAW` columns currently or in the future.
`condition`	The subset condition. Specify this condition similar to the way you specify conditions in a `WHERE` clause in SQL. For example, to specify rows in the `hr.employees` table where the `salary` is greater than `4000` and the `job_id` is `SA_MAN`, enter the following as the condition: `' salary > 4000 and job_id = ''SA_MAN'' '` If `NULL`, then the procedure raises an error. **Note:** The quotation marks in the preceding example are all single quotation marks.

Table 239-8 (Cont.) ADD_SUBSET_OUTBOUND_RULES Procedure Parameters

Parameter	Description
column_list	The list of columns either to include in the outbound server configuration or to exclude from the outbound server configuration. Whether the columns are included or excluded depends on the setting for the keep parameter. The columns can be specified in the following ways: • Comma-delimited list of type VARCHAR2. • A PL/SQL associative array of type DBMS_UTILITY.LNAME_ARRAY, where each element is the name of a column. Specify the first column in position 1. The last position must be NULL. To include or exclude all of the columns in a table, specify each column in the table in the list or array. If NULL, then the procedure raises an error.
keep	If TRUE, then the columns specified in the column_list parameter are kept as part of the outbound server configuration. Therefore, changes to these columns that satisfy the condition in the condition parameter are streamed to the outbound server's client application. If FALSE, then the columns specified in the column_list parameter are excluded from the outbound server configuration. Therefore, changes to these columns are not streamed to the outbound server's client application. **See Also:** "Usage Notes"
source_database	The global name of the container where the specified table_names and schema_names are located. If non-NULL, then a condition is added to the outbound server's rules to filter the LCRs based on the global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically.

Usage Notes

When the keep parameter is set to TRUE, this procedure creates a keep columns declarative rule-based transformation for the columns listed in column_list.

When the keep parameter is set to FALSE, this procedure creates a delete column declarative rule-based transformation for each column listed in column_list.

> **See Also:**
>
> *Oracle Database XStream Guide* for information about declarative rule-based transformations

239.4.8 ADD_SUBSET_PROPAGATION_RULES Procedure

This procedures adds propagation rules that propagate the logical change records (LCRs) related to a subset of the rows in the specified table in a source queue to a destination queue, and creates the specified propagation if it does not exist.

This procedure is overloaded. One version of this procedure contains three `OUT` parameters, and the other does not.

Syntax

```
DBMS_XSTREAM_ADM.ADD_SUBSET_PROPAGATION_RULES(
   table_name                IN   VARCHAR2,
   dml_condition             IN   VARCHAR2,
   streams_name              IN   VARCHAR2  DEFAULT NULL,
   source_queue_name         IN   VARCHAR2,
   destination_queue_name    IN   VARCHAR2,
   include_tagged_lcr        IN   BOOLEAN   DEFAULT FALSE,
   source_database           IN   VARCHAR2  DEFAULT NULL,
   insert_rule_name          OUT  VARCHAR2,
   update_rule_name          OUT  VARCHAR2,
   delete_rule_name          OUT  VARCHAR2,
   queue_to_queue            IN   BOOLEAN   DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_SUBSET_PROPAGATION_RULES(
   table_name                IN   VARCHAR2,
   dml_condition             IN   VARCHAR2,
   streams_name              IN   VARCHAR2  DEFAULT NULL,
   source_queue_name         IN   VARCHAR2,
   destination_queue_name    IN   VARCHAR2,
   include_tagged_lcr        IN   BOOLEAN   DEFAULT FALSE,
   source_database           IN   VARCHAR2  DEFAULT NULL,
   queue_to_queue            IN   BOOLEAN   DEFAULT NULL);
```

Parameters

Table 239-9 ADD_SUBSET_PROPAGATION_RULES Procedure Parameters

Parameter	Description
`table_name`	The name of the table specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default. The specified table must exist in the same database as the propagation. Also, the specified table cannot have any LOB, `LONG`, `LONG RAW`, or `XMLType` columns currently or in the future.
`dml_condition`	The subset condition. Specify this condition similar to the way you specify conditions in a `WHERE` clause in SQL. For example, to specify rows in the `hr.employees` table where the `salary` is greater than `4000` and the `job_id` is `SA_MAN`, enter the following as the condition: `' salary > 4000 and job_id = ''SA_MAN'' '` **Note:** The quotation marks in the preceding example are all single quotation marks.

Table 239-9 (Cont.) ADD_SUBSET_PROPAGATION_RULES Procedure Parameters

Parameter	Description
streams_name	The name of the propagation. Do not specify an owner. If the specified propagation does not exist, then the procedure creates it automatically. If NULL and a propagation exists for the same source queue and destination queue (including database link), then the procedure uses this propagation. If NULL and no propagation exists for the same source queue and destination queue (including database link), then the procedure creates a propagation automatically with a system-generated name.
source_queue_name	The name of the source queue, specified as [*schema_name.*]*queue_name*. The current database must contain the source queue, and the queue must be ANYDATA type. For example, to specify a source queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default.
destination_queue_name	The name of the destination queue, including a database link, specified as [*schema_name.*]*queue_name*[*@dblink_name*], if the destination queue is in a remote database. The queue must be ANYDATA type. For example, to specify a destination queue named streams_queue in the strmadmin schema and use a database link named dbs2.net, enter strmadmin.streams_queue@dbs2.net for this parameter. If the schema is not specified, then the current user is the default. If the database link is omitted, then the procedure uses the global name of the current database, and the source queue and destination queue must be in the same database. **Note:** Connection qualifiers are not allowed.
include_tagged_lcr	If TRUE, then an LCR is always considered for propagation, regardless of whether it has a non-NULL tag. This setting is appropriate for a full (for example, standby) copy of a database. If FALSE, then an LCR is considered for propagation only when the LCR contains a NULL tag. A setting of FALSE is often specified in update-anywhere configurations to avoid sending a change back to its source database.
source_database	The global name of the source database. The source database is where the change originated. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically. Oracle recommends that you specify a source database for propagation rules.
insert_rule_name	Contains the system-generated INSERT rule name. This rule handles inserts and updates that must be converted into inserts.

Table 239-9 (Cont.) ADD_SUBSET_PROPAGATION_RULES Procedure Parameters

Parameter	Description
update_rule_name	Contains the system-generated UPDATE rule name. This rule handles updates that remain updates.
delete_rule_name	Contains the system-generated DELETE rule name. This rule handles deletes and updates that must be converted into deletes
queue_to_queue	If TRUE or NULL, then a new propagation created by this procedure is a queue to queue propagation. A queue-to-queue propagation always has its own propagation job and uses a service for automatic failover when the destination queue is a buffered queue in an Oracle Real Application Clusters (Oracle RAC) database. If FALSE, then a new propagation created by this procedure is a queue-to-dblink propagation. A queue-to-dblink propagation can share a propagation job with other propagations that use the same database link and does not support automatic failover in an Oracle RAC environment. This procedure cannot change the queue to queue property of an exiting propagation. If the specified propagation exists, then the procedure behaves in the following way for each setting: • If TRUE and the specified propagation is not a queue to queue propagation, then the procedure raises an error. • If FALSE and the specified propagation is a queue to queue propagation, then the procedure raises an error. • If NULL, then the procedure does not change the queue to queue property of the propagation.

Usage Notes

This procedure configures propagation using the current user. Only one propagation is allowed between a particular source queue and destination queue.

Running this procedure generates three rules for the specified propagation: one for INSERT statements, one for UPDATE statements, and one for DELETE statements. For INSERT and DELETE statements, only row LCRs that satisfy the condition specified for the dml_condition parameter are propagated. For UPDATE statements, the following variations are possible:

- If both the new and old values in a row LCR satisfy the specified dml_condition, then the row LCR is propagated without any changes.
- If neither the new or old values in a row LCR satisfy the specified dml_condition, then the row LCR is not propagated.
- If the old values for a row LCR satisfy the specified dml_condition, but the new values do not, then the update row LCR is converted into a delete row LCR.
- If the new values for a row LCR satisfy the specified dml_condition, but the old values do not, then the update row LCR is converted to an insert row LCR.

When an update is converted into an insert or a delete, it is called row migration.

A propagation uses the rules for filtering. If the propagation does not have a positive rule set, then the procedure creates a positive rule set automatically, and the rules for propagating changes to the table are added to the positive rule set. A subset rule can be added to positive rule set only, not to a negative rule set. Other rules in an existing positive rule set for the

propagation are not affected. Additional rules can be added using either the DBMS_XSTREAM_ADM package or the DBMS_RULE_ADM package.

Rules for INSERT, UPDATE, and DELETE statements are created automatically when you run this procedure, and these rules are given a system-generated rule name. Each rule has a system-generated rule name that consists of the table name with a sequence number appended to it. The sequence number is used to avoid naming conflicts. If the table name plus the sequence number is too long, then the table name is truncated. The ADD_SUBSET_RULES procedure is overloaded, and the system-generated rule names for INSERT, UPDATE, and DELETE statements are returned.

When you create propagation subset rules for a table, you should create an unconditional supplemental log group at the source database with all the columns in the table. Supplemental logging is required if an update must be converted to an insert. The propagation rule must have all the column values to be able to perform this conversion correctly.

> **Note:**
>
> Subset rules should only reside in positive rule sets. You should not add subset rules to negative rule sets. Doing so might have unpredictable results because row migration would not be performed on LCRs that are not discarded by the negative rule set.

Examples

The following is an example of a rule condition created for filtering a row LCR containing an update operation when the dml_condition is region_id = 2, the table_name is hr.regions, and the source_database is dbs1.net:

```
:dml.get_object_owner()='HR' AND :dml.get_object_name()='REGIONS'
AND :dml.is_null_tag()='Y' AND :dml.get_source_database_name()='DBS1.NET'
AND :dml.get_command_type()='UPDATE'
AND (:dml.get_value('NEW','"REGION_ID"') IS NOT NULL)
AND (:dml.get_value('OLD','"REGION_ID"') IS NOT NULL)
AND (:dml.get_value('OLD','"REGION_ID"').AccessNumber()=2)
AND (:dml.get_value('NEW','"REGION_ID"').AccessNumber()=2)
```

239.4.9 ADD_SUBSET_RULES Procedure

This procedure adds rules to an XStream client.

It adds rules to a rule set of one of the following types of XStream clients:

- When the streams_type parameter is set to capture, this procedure adds capture process rules for capturing changes to a subset of rows in a specified table.

  This procedure creates the specified capture process if it does not exist.

- When the streams_type parameter is set to apply and the streams_name parameter specifies the name of an apply process, outbound server, or inbound server, this procedure adds apply rules for applying logical change records (LCRs) that contain changes to a subset of rows in a specified table. The rules can specify that the LCRs must be from a particular source database.

  This procedure creates an apply process if no apply process, outbound server, or inbound server exists with the specified streams_name. This procedure can add rules to an

outbound server or inbound server, but it cannot create an outbound server or inbound server.

This procedure is overloaded. One version of this procedure contains three `OUT` parameters, and the other does not.

Syntax

```
DBMS_XSTREAM_ADM.ADD_SUBSET_RULES(
   table_name              IN   VARCHAR2,
   dml_condition           IN   VARCHAR2,
   streams_type            IN   VARCHAR2 DEFAULT 'apply',
   streams_name            IN   VARCHAR2 DEFAULT NULL,
   queue_name              IN   VARCHAR2 DEFAULT 'streams_queue',
   include_tagged_lcr      IN   BOOLEAN  DEFAULT FALSE,
   source_database         IN   VARCHAR2 DEFAULT NULL,
   insert_rule_name        OUT  VARCHAR2,
   update_rule_name        OUT  VARCHAR2,
   delete_rule_name        OUT  VARCHAR2,
   source_root_name        IN  VARCHAR2  DEFAULT NULL,
   source_container_name   IN  VARCHAR2  DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_SUBSET_RULES(
   table_name              IN   VARCHAR2,
   dml_condition           IN   VARCHAR2,
   streams_type            IN   VARCHAR2 DEFAULT 'apply',
   streams_name            IN   VARCHAR2 DEFAULT NULL,
   queue_name              IN   VARCHAR2 DEFAULT 'streams_queue',
   include_tagged_lcr      IN   BOOLEAN  DEFAULT FALSE,
   source_database         IN   VARCHAR2 DEFAULT NULL,
   source_root_name        IN  VARCHAR2  DEFAULT NULL,
   source_container_name   IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-10 ADD_SUBSET_RULES Procedure Parameters

Parameter	Description
`table_name`	The name of the table specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default. The specified table must exist in the same database as the capture process or apply process. Also, the specified table cannot have any LOB, `LONG`, `LONG RAW`, or `XMLType` columns currently or in the future.
`dml_condition`	The subset condition. Specify this condition similar to the way you specify conditions in a `WHERE` clause in SQL. For example, to specify rows in the `hr.employees` table where the `salary` is greater than `4000` and the `job_id` is `SA_MAN`, enter the following as the condition: `' salary > 4000 and job_id = ''SA_MAN'' '` **Note:** The quotation marks in the preceding example are all single quotation marks.
`streams_type`	The type of XStream client: • Specify `capture` for a capture process. • Specify `apply` for an apply process.

Table 239-10 (Cont.) ADD_SUBSET_RULES Procedure Parameters

Parameter	Description
streams_name	The name of the capture process or apply process. Do not specify an owner. If NULL, if streams_type is capture, and if one relevant capture process for the queue exists, then the procedure uses the relevant XStream client. If no relevant XStream client exists for the queue, then the procedure creates an XStream client automatically with a system-generated name. If NULL and multiple XStream clients of the specified streams_type for the queue exist, then the procedure raises an error. If NULL, if streams_type is apply, and if one relevant apply process exists, then the procedure uses the relevant apply process. The relevant apply process is identified in one of the following ways: • If one existing apply process has the source database specified in source_database and uses the queue specified in queue_name, then the procedure uses this apply process. • If source_database is NULL and one existing apply process is using the queue specified in queue_name, then the procedure uses this apply process. If NULL and no relevant apply process exists, then the procedure creates an apply process automatically with a system-generated name. If NULL and multiple relevant apply processes exist, then the procedure raises an error. Each apply process must have a unique name.
queue_name	The name of the local queue, specified as [*schema_name.*]*queue_name*. The current database must contain the queue, and the queue must be ANYDATA type. For example, to specify a queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default. For capture process rules, this is the queue into which a capture process enqueues LCRs. For outbound server rules, this is the queue from which the outbound server dequeues LCRs. For inbound server rules, this is the queue into which an inbound server enqueues error transactions.
include_tagged_lcr	If TRUE, then the XStream client performs its action regardless of the tag: • A redo entry is always considered for capture by a capture process, regardless of whether the redo entry has a non-NULL tag. • An LCR is always considered for apply by an apply process, regardless of whether redo entry or LCR has a non-NULL tag. If FALSE, then an XStream client performs its action only when the tag is NULL: • A redo entry is considered for capture by a capture process only when the redo entry contains a NULL tag. • An LCR is considered for apply by an apply process only if the LCR contains a NULL tag. A setting of FALSE is often specified in update-anywhere configurations to avoid sending a change back to its source database.

Table 239-10 (Cont.) ADD_SUBSET_RULES Procedure Parameters

Parameter	Description
source_database	The global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. For capture process rules, specify NULL or the global name of the local database if you are creating a capture process locally at the source database. If you are adding rules to a downstream capture process rule set at a downstream database, then specify the source database of the changes that will be captured. For apply process rules, specify the source database of the changes that will be applied by the apply process. The source database is the database where the changes originated. If an apply process applies captured LCRs, then the apply process can apply LCRs from only one capture process at one source database. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically.
insert_rule_name	Contains the system-generated INSERT rule name. This rule handles inserts and updates that must be converted into inserts.
update_rule_name	Contains the system-generated UPDATE rule name. This rule handles updates that remain updates.
delete_rule_name	Contains the system-generated DELETE rule name. This rule handles deletes and updates that must be converted into deletes
source_root_name	The global name of the root in the source CDB. For example, mycdb.example.com. If this parameter is NULL, then the global name of the root in the local CDB is used. If you are configuring downstream capture, then this parameter must be a non-NULL value, and it must specify the global name of the root in the remote source CDB. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. **Note:** This parameter only applies to a CDB.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.

Usage Notes

Running this procedure generates three rules for the specified capture process or apply process: one for INSERT statements, one for UPDATE statements, and one for DELETE statements. For INSERT and DELETE statements, only DML changes that satisfy the condition specified for the dml_condition parameter are captured or applied. For UPDATE statements, the following variations are possible:

- If both the new and old values in a DML change satisfy the specified `dml_condition`, then the DML change is captured or applied without any changes.
- If neither the new or old values in a DML change satisfy the specified `dml_condition`, then the DML change is not captured or applied.
- If the old values for a DML change satisfy the specified `dml_condition`, but the new values do not, then the DML change is converted into a delete.
- If the new values for a DML change satisfy the specified `dml_condition`, but the old values do not, then the DML change is converted to an insert.

When an update is converted into an insert or a delete, it is called row migration.

A capture process or apply process uses the rules for filtering. If the XStream client does not have a positive rule set, then this procedure creates a positive rule set automatically, and adds the rules for the table to the positive rule set. A subset rule can be added to positive rule set only, not to a negative rule set. Other rules in an existing rule set for the process are not affected. Additional rules can be added using either the `DBMS_XSTREAM_ADM` package or the `DBMS_RULE_ADM` package.

Rules for `INSERT`, `UPDATE`, and `DELETE` statements are created automatically when you run this procedure, and these rules are given a system-generated rule name. Each rule has a system-generated rule name that consists of the table name with a sequence number appended to it. The sequence number is used to avoid naming conflicts. If the table name plus the sequence number is too long, then the table name is truncated. The `ADD_SUBSET_RULES` procedure is overloaded, and the system-generated rule names for `INSERT`, `UPDATE`, and `DELETE` statements are returned.

> **Note:**
>
> Subset rules should only reside in positive rule sets. You should not add subset rules to negative rule sets. Doing so might have unpredictable results because row migration would not be performed on LCRs that are not discarded by the negative rule set.

Examples

The following is an example of a rule condition created for filtering DML changes containing an update operation when the `dml_condition` is `region_id = 2`, the `table_name` is `hr.regions`, and the `source_database` is `dbs1.net`:

```
:dml.get_object_owner()='HR' AND :dml.get_object_name()='REGIONS'
AND :dml.is_null_tag()='Y' AND :dml.get_source_database_name()='DBS1.NET'
AND :dml.get_command_type()='UPDATE'
AND (:dml.get_value('NEW','"REGION_ID"') IS NOT NULL)
AND (:dml.get_value('OLD','"REGION_ID"') IS NOT NULL)
AND (:dml.get_value('OLD','"REGION_ID"').AccessNumber()=2)
AND (:dml.get_value('NEW','"REGION_ID"').AccessNumber()=2)
```

239.4.10 ADD_TABLE_PROPAGATION_RULES Procedure

This procedures adds table rules to the positive rule set for a propagation, or adds table rules to the negative rule set for a propagation, and creates the specified propagation if it does not exist.

This procedure is overloaded. One version of this procedure contains two `OUT` parameters, and the other does not.

Syntax

```
DBMS_XSTREAM_ADM.ADD_TABLE_PROPAGATION_RULES(
   table_name              IN   VARCHAR2,
   streams_name            IN   VARCHAR2  DEFAULT NULL,
   source_queue_name       IN   VARCHAR2,
   destination_queue_name  IN   VARCHAR2,
   include_dml             IN   BOOLEAN   DEFAULT TRUE,
   include_ddl             IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr      IN   BOOLEAN   DEFAULT FALSE,
   source_database         IN   VARCHAR2  DEFAULT NULL,
   dml_rule_name           OUT  VARCHAR2,
   ddl_rule_name           OUT  VARCHAR2,
   inclusion_rule          IN   BOOLEAN   DEFAULT TRUE,
   and_condition           IN   VARCHAR2  DEFAULT NULL,
   queue_to_queue          IN   BOOLEAN   DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_TABLE_PROPAGATION_RULES(
   table_name              IN   VARCHAR2,
   streams_name            IN   VARCHAR2  DEFAULT NULL,
   source_queue_name       IN   VARCHAR2,
   destination_queue_name  IN   VARCHAR2,
   include_dml             IN   BOOLEAN   DEFAULT TRUE,
   include_ddl             IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr      IN   BOOLEAN   DEFAULT FALSE,
   source_database         IN   VARCHAR2  DEFAULT NULL,
   inclusion_rule          IN   BOOLEAN   DEFAULT TRUE,
   and_condition           IN   VARCHAR2  DEFAULT NULL,
   queue_to_queue          IN   BOOLEAN   DEFAULT NULL);
```

Parameters

Table 239-11 ADD_TABLE_PROPAGATION_RULES Procedure Parameters

Parameter	Description
`table_name`	The name of the table specified as `[`*`schema_name`*`.]`*`table_name`*. For example, `hr.employees`. If the schema is not specified, then the current user is the default.
`streams_name`	The name of the propagation. Do not specify an owner. If the specified propagation does not exist, then the procedure creates it automatically. If `NULL` and a propagation exists for the same source queue and destination queue (including database link), then the procedure uses this propagation. If `NULL` and no propagation exists for the same source queue and destination queue (including database link), then the procedure creates a propagation automatically with a system-generated name.

Table 239-11 (Cont.) ADD_TABLE_PROPAGATION_RULES Procedure Parameters

Parameter	Description
source_queue_name	The name of the source queue, specified as [*schema_name.*]*queue_name*. The current database must contain the source queue, and the queue must be ANYDATA type. For example, to specify a source queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default.
destination_queue_name	The name of the destination queue, including a database link, specified as [*schema_name.*]*queue_name*[*@dblink_name*], if the destination queue is in a remote database. The queue must be ANYDATA type. For example, to specify a destination queue named streams_queue in the strmadmin schema and use a database link named dbs2.net, enter strmadmin.streams_queue@dbs2.net for this parameter. If the schema is not specified, then the current user is the default. If the database link is omitted, then the procedure uses the global name of the current database, and the source queue and destination queue must be in the same database. **Note:** Connection qualifiers are not allowed.
include_dml	If TRUE, then the procedure creates a rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted. The generated rule evaluates to TRUE for any DDL change that operates on the table or on an object that is part of the table, such as an index or trigger on the table. The rule evaluates to FALSE for any DDL change that either does not refer to the table or refers to the table in a subordinate way. For example, the rule evaluates to FALSE for changes that create synonyms or views based on the table. The rule also evaluates to FALSE for a change to a PL/SQL subprogram that refers to the table.

Table 239-11 (Cont.) ADD_TABLE_PROPAGATION_RULES Procedure Parameters

Parameter	Description
include_tagged_lcr	If TRUE, then the procedure does not add a condition regarding Oracle Replication tags to the generated rules. Therefore, these rules can evaluate to TRUE regardless of whether a logical change record (LCR) has a non-NULL tag. If the rules are added to the positive rule set for the propagation, then an LCR is always considered for propagation, regardless of whether it has a non-NULL tag. If the rules are added to a positive rule set, then setting this parameter to TRUE is appropriate for a full (for example, standby) copy of a database. If the rules are added to the negative rule set for the propagation, then whether an LCR is discarded does not depend on the tag for the LCR. If FALSE, then the procedure adds a condition to each generated rule that causes the rule to evaluate to TRUE only if an LCR has a NULL Oracle Replication tag. If the rules are added to the positive rule set for the propagation, then an LCR is considered for propagation only when the LCR contains a NULL tag. If the rules are added to a positive rule set, then setting this parameter to FALSE might be appropriate in update-anywhere configurations to avoid sending a change back to its source database. If the rules are added to the negative rule set for the propagation, then an LCR can be discarded only if it has a NULL tag. Usually, specify TRUE for this parameter if the inclusion_rule parameter is set to FALSE.
source_database	The global name of the source database. The source database is where the change originated. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically. Oracle recommends that you specify a source database for propagation rules.
dml_rule_name	If include_dml is TRUE, then this parameter contains the DML rule name. If include_dml is FALSE, then this parameter contains a NULL.
ddl_rule_name	If include_ddl is TRUE, then this parameter contains the DDL rule name. If include_ddl is FALSE, then this parameter contains a NULL.
inclusion_rule	If inclusion_rule is TRUE, then the procedure adds the rules to the positive rule set for the propagation. If inclusion_rule is FALSE, then the procedure adds the rules to the negative rule set for the propagation. In either case, the system creates the rule set if it does not exist.

Table 239-11 (Cont.) ADD_TABLE_PROPAGATION_RULES Procedure Parameters

Parameter	Description
and_condition	If non-NULL, appends the specified condition to the system-generated rule condition using an AND clause in the following way: (*system_condition*) AND (*and_condition*) The variable in the specified condition must be :lcr. For example, to specify that the table rules generated by the procedure evaluate to TRUE only if the Oracle Replication tag is the hexadecimal equivalent of '02', specify the following condition: :lcr.get_tag() = HEXTORAW(''02'') The :lcr in the specified condition is converted to :dml or :ddl, depending on the rule that is being generated. If you are specifying an LCR member subprogram that is dependent on the LCR type (row or DDL), then make sure this procedure only generates the appropriate rule. Specifically, if you specify an LCR member subprogram that is valid only for row LCRs, then specify TRUE for the include_dml parameter and FALSE for the include_ddl parameter. If you specify an LCR member subprogram that is valid only for DDL LCRs, then specify FALSE for the include_dml parameter and TRUE for the include_ddl parameter. **See Also:** Logical Change Record TYPEs
queue_to_queue	If TRUE or NULL, then a new propagation created by this procedure is a queue to queue propagation. A queue-to-queue propagation always has its own propagation job and uses a service for automatic failover when the destination queue is a buffered queue in an Oracle Real Application Clusters (Oracle RAC) database. If FALSE, then a new propagation created by this procedure is a queue-to-dblink propagation. A queue-to-dblink propagation can share a propagation job with other propagations that use the same database link and does not support automatic failover in an Oracle RAC environment. This procedure cannot change the queue to queue property of an exiting propagation. If the specified propagation exists, then the procedure behaves in the following way for each setting: • If TRUE and the specified propagation is not a queue to queue propagation, then the procedure raises an error. • If FALSE and the specified propagation is a queue to queue propagation, then the procedure raises an error. • If NULL, then the procedure does not change the queue to queue property of the propagation.

Usage Notes

This procedure configures propagation using the current user. Only one propagation is allowed between a particular source queue and destination queue.

This procedure creates DML and DDL rules automatically based on include_dml and include_ddl parameter values, respectively. Each rule has a system-generated rule name that consists of the table name with a sequence number appended to it. The sequence number is

used to avoid naming conflicts. If the table name plus the sequence number is too long, then the table name is truncated. A propagation uses the rules for filtering.

Examples

The following is an example of a table rule condition created for filtering DML statements:

```
(((:dml.get_object_owner() = 'HR' and :dml.get_object_name() = 'LOCATIONS'))
and :dml.is_null_tag() = 'Y' and :dml.get_source_database_name() = 'DBS1.NET' )
```

239.4.11 ADD_TABLE_RULES Procedure

This procedure adds rules to an XStream client rule set.

It adds rules to a rule set of one of the following types of XStream clients:

- When the `streams_type` parameter is set to `capture`, this procedure adds capture process rules for capturing changes to a specified table.

  This procedure creates the specified capture process if it does not exist.

- When the `streams_type` parameter is set to `apply` and the `streams_name` parameter specifies the name of an apply process, outbound server, or inbound server, this procedure adds apply rules for applying logical change records (LCRs) that contain changes to a specified table. The rules can specify that the LCRs must be from a particular source database.

  This procedure creates an apply process if no apply process, outbound server, or inbound server exists with the specified `streams_name`. This procedure can add rules to an outbound server or inbound server, but it cannot create an outbound server or inbound server.

This procedure is overloaded. One version of this procedure contains two `OUT` parameters, and the other does not.

Syntax

```
DBMS_XSTREAM_ADM.ADD_TABLE_RULES(
   table_name             IN   VARCHAR2,
   streams_type           IN   VARCHAR2,
   streams_name           IN   VARCHAR2  DEFAULT NULL,
   queue_name             IN   VARCHAR2  DEFAULT 'streams_queue',
   include_dml            IN   BOOLEAN   DEFAULT TRUE,
   include_ddl            IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr     IN   BOOLEAN   DEFAULT FALSE,
   source_database        IN   VARCHAR2  DEFAULT NULL,
   dml_rule_name          OUT  VARCHAR2,
   ddl_rule_name          OUT  VARCHAR2,
   inclusion_rule         IN   BOOLEAN   DEFAULT TRUE,
   and_condition          IN   VARCHAR2  DEFAULT NULL,
   source_root_name       IN  VARCHAR2   DEFAULT NULL,
   source_container_name  IN  VARCHAR2   DEFAULT NULL);

DBMS_XSTREAM_ADM.ADD_TABLE_RULES(
   table_name             IN   VARCHAR2,
   streams_type           IN   VARCHAR2,
   streams_name           IN   VARCHAR2  DEFAULT NULL,
   queue_name             IN   VARCHAR2  DEFAULT 'streams_queue',
   include_dml            IN   BOOLEAN   DEFAULT TRUE,
   include_ddl            IN   BOOLEAN   DEFAULT FALSE,
   include_tagged_lcr     IN   BOOLEAN   DEFAULT FALSE,
   source_database        IN   VARCHAR2  DEFAULT NULL,
```

```
  inclusion_rule         IN   BOOLEAN   DEFAULT TRUE,
  and_condition          IN   VARCHAR2  DEFAULT NULL,
  source_root_name       IN  VARCHAR2   DEFAULT NULL,
  source_container_name  IN  VARCHAR2   DEFAULT NULL);
```

Parameters

Table 239-12 ADD_TABLE_RULES Procedure Parameters

Parameter	Description
table_name	The name of the table specified as [*schema_name.*]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default. You can specify a table that does not yet exist, because Oracle Replication does not validate the existence of the table.
streams_type	The type of XStream client: • Specify capture for a capture process. • Specify apply for an apply process.
streams_name	The name of the capture process or apply process. Do not specify an owner. If NULL, if streams_type is capture, and if one relevant capture process for the queue exists, then the procedure uses the relevant XStream client. If no relevant XStream client exists for the queue, then the procedure creates an XStream client automatically with a system-generated name. If NULL and multiple XStream clients of the specified streams_type for the queue exist, then the procedure raises an error. If NULL, if streams_type is apply, and if one relevant apply process exists, then the procedure uses the relevant apply process. The relevant apply process is identified in one of the following ways: • If one existing apply process has the source database specified in source_database and uses the queue specified in queue_name, then the procedure uses this apply process. • If source_database is NULL and one existing apply process is using the queue specified in queue_name, then the procedure uses this apply process. If NULL and no relevant apply process exists, then the procedure creates an apply process automatically with a system-generated name. If NULL and multiple relevant apply processes exist, then the procedure raises an error. Each apply process must have a unique name.
queue_name	The name of the local queue, specified as [*schema_name.*]*queue_name*. The current database must contain the queue, and the queue must be ANYDATA type. For example, to specify a queue named streams_queue in the strmadmin schema, enter strmadmin.streams_queue for this parameter. If the schema is not specified, then the current user is the default. For capture process rules, this is the queue into which a capture process enqueues LCRs. For outbound server rules, this is the queue from which the outbound server dequeues LCRs. For inbound server rules, this is the queue into which an inbound server enqueues error transactions.
include_dml	If TRUE, then the procedure creates a DML rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.

Table 239-12 (Cont.) ADD_TABLE_RULES Procedure Parameters

Parameter	Description
include_ddl	If TRUE, then the procedure creates a DDL rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted. The generated rule evaluates to TRUE for any DDL change that operates on the table or on an object that is part of the table, such as an index or trigger on the table. The rule evaluates to FALSE for any DDL change that either does not refer to the table or refers to the table in a subordinate way. For example, the rule evaluates to FALSE for changes that create synonyms or views based on the table. The rule also evaluates to FALSE for a change to a PL/SQL subprogram that refers to the table.
include_tagged_lcr	If TRUE, then the procedure does not add a condition regarding Oracle Replication tags to the generated rules. Therefore, these rules can evaluate to TRUE regardless of whether a redo entry, session, or LCR has a non-NULL tag. If the rules are added to the positive rule set for the XStream client, then the XStream client performs its action regardless of the tag: • A redo entry is always considered for capture by a capture process, regardless of whether the redo entry has a non-NULL tag. • An LCR is always considered for apply by an apply process, regardless of whether redo entry or LCR has a non-NULL tag. If the rules are added to a positive rule set, then setting this parameter to TRUE is appropriate for a full (for example, standby) copy of a database. If the rules are added to the negative rule set for the XStream client, then whether a database change is discarded does not depend on the tag. If FALSE, then the procedure adds a condition to each generated rule that causes the rule to evaluate to TRUE only if a redo entry, session, or LCR has a NULL Oracle Replication tag. If the rules are added to the positive rule set for an XStream client, then the XStream client performs its action only when the tag is NULL: • A redo entry is considered for capture by a capture process only when the redo entry contains a NULL tag. • An LCR is considered for apply by an apply process only if the LCR contains a NULL tag. If the rules are added to a positive rule set, then setting this parameter to FALSE might be appropriate in update-anywhere configurations to avoid sending a change back to its source database. If the rules are added to the negative rule set for the XStream client, then a database change can be discarded only if it has a NULL tag. A setting of FALSE is often specified in update-anywhere configurations to avoid sending a change back to its source database. Usually, specify TRUE for this parameter if the inclusion_rule parameter is set to FALSE.

Table 239-12 (Cont.) ADD_TABLE_RULES Procedure Parameters

Parameter	Description
source_database	The global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. For capture process rules, specify NULL or the global name of the local database if you are creating a capture process locally at the source database. If you are adding rules to a downstream capture process rule set at a downstream database, then specify the source database of the changes that will be captured. For apply process rules, specify the source database of the changes that will be applied by the apply process. The source database is the database where the changes originated. If an apply process applies captured LCRs, then the apply process can apply LCRs from only one capture process at one source database. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is .NET, then the procedure specifies DBS1.NET automatically.
dml_rule_name	If include_dml is TRUE, then this parameter contains the DML rule name. If include_dml is FALSE, then this parameter contains a NULL.
ddl_rule_name	If include_ddl is TRUE, then this parameter contains the DDL rule name. If include_ddl is FALSE, then this parameter contains a NULL.
inclusion_rule	If inclusion_rule is TRUE, then the procedure adds the rules to the positive rule set for the XStream client. If inclusion_rule is FALSE, then the procedure adds the rules to the negative rule set for the XStream client. In either case, the system creates the rule set if it does not exist.
and_condition	If non-NULL, appends the specified condition to the system-generated rule condition using an AND clause in the following way: (*system_condition*) AND (*and_condition*) The variable in the specified condition must be :lcr. For example, to specify that the table rules generated by the procedure evaluate to TRUE only if the Oracle Replication tag is the hexadecimal equivalent of '02', specify the following condition: :lcr.get_tag() = HEXTORAW(''02'') The :lcr in the specified condition is converted to :dml or :ddl, depending on the rule that is being generated. If you are specifying an LCR member subprogram that is dependent on the LCR type (row or DDL), then make sure this procedure only generates the appropriate rule. Specifically, if you specify an LCR member subprogram that is valid only for row LCRs, then specify TRUE for the include_dml parameter and FALSE for the include_ddl parameter. If you specify an LCR member subprogram that is valid only for DDL LCRs, then specify FALSE for the include_dml parameter and TRUE for the include_ddl parameter. **See Also:** Logical Change Record TYPEs

Table 239-12 (Cont.) ADD_TABLE_RULES Procedure Parameters

Parameter	Description
source_root_name	The global name of the root in the source CDB. For example, mycdb.example.com. If this parameter is NULL, then the global name of the root in the local CDB is used. If you are configuring downstream capture, then this parameter must be a non-NULL value, and it must specify the global name of the root in the remote source CDB. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. **Note:** This parameter only applies to a CDB.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.

Usage Notes

This procedure creates DML and DDL rules automatically based on include_dml and include_ddl parameter values, respectively. Each rule has a system-generated rule name that consists of the table name with a sequence number appended to it. The sequence number is used to avoid naming conflicts. If the table name plus the sequence number is too long, then the table name is truncated. A capture process or apply process uses the rules for filtering.

See Also:

- "Operational Notes"
- "Security Model"

Examples

The following is an example of a table rule condition created for DML changes:

```
(((:dml.get_object_owner() = 'HR' and :dml.get_object_name() = 'LOCATIONS'))
and :dml.is_null_tag() = 'Y' and :dml.get_source_database_name() = 'DBS1.NET' )
```

239.4.12 ALTER_INBOUND Procedure

This procedure modifies an XStream inbound server.

Syntax

```
DBMS_XSTREAM_ADM.ALTER_INBOUND(
   server_name IN VARCHAR2,
   apply_user  IN VARCHAR2  DEFAULT NULL,
   comment     IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-13 ALTER_INBOUND Procedure Parameters

Parameter	Description
server_name	The name of the inbound server being altered. Specify an existing inbound server. Do not specify an owner.
apply_user	The user who applies all DML and DDL changes that satisfy the inbound server rule sets, who runs user-defined apply handlers, and who runs custom rule-based transformations configured for inbound server rules. The client application must attach to the inbound server as the apply user. Specify a user to change the apply user. In this case, the user who invokes the ALTER_INBOUND procedure must be granted the DBA role. Only the SYS user can set the apply_user to SYS. If NULL, then the apply user is not changed. See "CREATE_INBOUND Procedure" for information about the required privileges for an apply user.
comment	An optional comment associated with the inbound server. If non-NULL, then the specified comment replaces the existing comment. If NULL, then the existing comment is not changed.

239.4.13 ALTER_OUTBOUND Procedure

This procedure modifies an XStream outbound server configuration.

This procedure always alters the specified outbound server. This procedure can also alter the outbound server's capture process when either of the following conditions is met:

- The capture process was created by the CREATE_OUTBOUND procedure in this package.
- The queue used by the capture process was created by the CREATE_OUTBOUND procedure.

To check whether this procedure can alter the outbound server's capture process, query the CAPTURE_NAME column in the ALL_XSTREAM_OUTBOUND view. When the name of the capture process appears in the CAPTURE_NAME column of this view, the ALTER_OUTBOUND procedure can manage the capture process's rules or change the capture user for the capture process. When the CAPTURE_NAME column of this view is NULL, the ALTER_OUTBOUND procedure cannot manage the capture process.

This procedure is overloaded. One table_names parameter is type VARCHAR2 and the other table_names parameter is type DBMS_UTILITY.UNCL_ARRAY. Also, one schema_names parameter is type VARCHAR2 and the other schema_names parameter is type DBMS_UTILITY.UNCL_ARRAY. These parameters enable you to enter the list of tables and schemas in different ways and are mutually exclusive.

> **Note:**
>
> Starting with Oracle Database 11*g* Release 2 (11.2.0.2), the start_scn and start_time parameters are included in this procedure.

Syntax

```
DBMS_XSTREAM_ADM.ALTER_OUTBOUND(
   server_name           IN VARCHAR2,
   table_names           IN DBMS_UTILITY.UNCL_ARRAY,
   schema_names          IN DBMS_UTILITY.UNCL_ARRAY,
   add                   IN BOOLEAN    DEFAULT TRUE,
   capture_user          IN VARCHAR2   DEFAULT NULL,
   connect_user          IN VARCHR2    DEFAULT NULL,
   comment               IN VARCHAR2   DEFAULT NULL,
   inclusion_rule        IN BOOLEAN    DEFAULT TRUE,
   start_scn             IN NUMBER     DEFAULT NULL,
   start_time            IN TIMESTAMP  DEFAULT NULL,
   include_dml           IN BOOLEAN    DEFAULT TRUE,
   include_ddl           IN BOOLEAN    DEFAULT FALSE,
   source_database       IN VARCHAR2   DEFAULT NULL,
   source_container_name IN VARCHAR2   DEFAULT NULL);

DBMS_XSTREAM_ADM.ALTER_OUTBOUND(
   server_name           IN VARCHAR2,
   table_names           IN VARCHAR2   DEFAULT NULL,
   schema_names          IN VARCHAR2   DEFAULT NULL,
   add                   IN BOOLEAN    DEFAULT TRUE,
   capture_user          IN VARCHAR2   DEFAULT NULL,
   connect_user          IN VARCHAR2   DEFAULT NULL,
   comment               IN VARCHAR2   DEFAULT NULL,
   inclusion_rule        IN BOOLEAN    DEFAULT TRUE,
   start_scn             IN NUMBER     DEFAULT NULL,
   start_time            IN TIMESTAMP  DEFAULT NULL,
   include_dml           IN BOOLEAN    DEFAULT TRUE,
   include_ddl           IN BOOLEAN    DEFAULT FALSE,
   source_database       IN VARCHAR2   DEFAULT NULL,
   source_container_name IN VARCHAR2   DEFAULT NULL);
```

Parameters

Table 239-14 ALTER_OUTBOUND Procedure Parameters

Parameter	Description
`server_name`	The name of the outbound server being altered. Specify an existing outbound server. Do not specify an owner.
`table_names`	The tables that are either added to or removed from the XStream Out configuration. Whether the tables are added or removed depends on the setting for the `add` parameter. The tables can be specified in the following ways: • Comma-delimited list of type `VARCHAR2`. • A PL/SQL associative array of type `DBMS_UTILITY.UNCL_ARRAY`, where each element is the name of a table. Specify the first table in position 1. The last position must be `NULL`. Each table should be specified as `[`*`schema_name`*`.]`*`table_name`*. For example, `hr.employees`. If the schema is not specified, then the current user is the default. **See Also:** "Usage Notes" for more information about this parameter

Table 239-14 (Cont.) ALTER_OUTBOUND Procedure Parameters

Parameter	Description
schema_names	The schemas that are either added to or removed from the XStream Out configuration. Whether the schemas are added or removed depends on the setting for the add parameter. The schemas can be specified in the following ways: • Comma-delimited list of type VARCHAR2. • A PL/SQL associative array of type DBMS_UTILITY.UNCL_ARRAY, where each element is the name of a schema. Specify the first schema in position 1. The last position must be NULL. **Note:** This procedure does not concatenate the schema_names parameter with the table_names parameter. To specify tables, enter fully qualified table names in the table_names parameter (*schema_name.table_name*). **See Also:** "Usage Notes" for more information about this parameter
add	If TRUE, then the procedure adds to the XStream Out configuration the tables specified in the table_names parameter and the schemas specified in the schema_names parameter. If FALSE, then the procedure removes from the XStream Out configuration the tables specified in the table_names parameter and the schemas specified in the schema_names parameter.
capture_user	The user in whose security domain a capture process captures changes that satisfy its rule sets and runs custom rule-based transformations configured for capture process rules. Specify a user to change the capture user. In this case, the user who invokes the ALTER_OUTBOUND procedure must be granted the DBA role. Only the SYS user can set the capture_user to SYS. If NULL, then the capture user is not changed. If you change the capture user, then this procedure grants the new capture user enqueue privilege on the queue used by the capture process and configures the user as a secure queue user. Ensure that the capture user is granted the other required privileges. See "CREATE_OUTBOUND Procedure" for information about the privileges required by a capture user. The capture process is stopped and restarted automatically when you change the value of this parameter. **Note:** If the capture user for a capture process is dropped using DROP USER . . . CASCADE, then the capture process is also dropped automatically.
connect_user	The user who can attach to the specified outbound server to retrieve the change stream. The XStream client application must attach to the outbound server as the specified connect user. Specify a user to change the connect user. In this case, the user who invokes the ALTER_OUTBOUND procedure must be granted the DBA role. Only the SYS user can set the connect_user to SYS. If NULL, then the connect user is not changed. If you change the connect user, then this procedure grants the new connect user dequeue privileges on the queue used by the outbound server and configures the user as a secure queue user. Ensure that the connect user is granted the other required privileges. See "CREATE_OUTBOUND Procedure" for information about the privileges required by a connect user.

Table 239-14 (Cont.) ALTER_OUTBOUND Procedure Parameters

Parameter	Description
comment	An optional comment associated with the outbound server. If non-NULL, then the specified comment replaces the existing comment. If NULL, then the existing comment is not changed.
inclusion_rule	If TRUE and the add parameter is set to TRUE, then the procedure adds rules for the tables specified in the table_names parameter and the schemas specified in the schema_names parameter to the positive rule sets in the XStream Out configuration. When rules for tables and schemas are in positive rule sets, the XStream Out configuration streams DML and DDL changes to the tables and schemas out to the client application. If TRUE and the add parameter is set to FALSE, then the procedure removes rules for the tables specified in the table_names parameter and the schemas specified in the schema_names parameter from the positive rule sets in the XStream Out configuration. If FALSE and the add parameter is set to TRUE, then the procedure adds rules for the tables specified in the table_names parameter and the schemas specified in the schema_names parameter to the negative rule sets in the XStream Out configuration. When rules for tables and schemas are in negative rule sets, the XStream Out configuration does not stream changes to the tables and schemas out to the client application. If FALSE and the add parameter is set to FALSE, then the procedure removes rules for the tables specified in the table_names parameter and the schemas specified in the schema_names parameter from the negative rule sets in the XStream Out configuration.
start_scn	A valid SCN for the database from which the capture process starts capturing changes. To be valid, the SCN value must be greater than or equal to the first SCN for the capture process. If a valid SCN is specified, then the capture process captures changes from the specified SCN when it is restarted. An error is returned if an invalid SCN is specified. If NULL and the start_time parameter is NULL, then the start SCN is not changed. If NULL and the start_time parameter is non-NULL, then the start SCN is changed to match the specified start time. The start_scn and start_time parameters are mutually exclusive. **Note:** If the capture process is enabled, then the ALTER_OUTBOUND procedure automatically stops and restarts the capture process when the start_scn parameter is non-NULL. If the capture process is disabled, then the ALTER_OUTBOUND procedure automatically starts the capture process when the start_scn parameter is non-NULL.

Table 239-14 (Cont.) ALTER_OUTBOUND Procedure Parameters

Parameter	Description
start_time	A valid time from which the capture process starts capturing changes. To be valid, the time must correspond to an SCN value that is greater than or equal to the first SCN for the capture process. If a valid time is specified, then the capture process captures changes from the specified time when it is restarted. An error is returned if an invalid time is specified. If NULL and the start_scn parameter is NULL, then the start time is not changed. If NULL and the start_scn parameter is non-NULL, then the start time is changed to match the specified start SCN. The start_scn and start_time parameters are mutually exclusive. **Note:** If the capture process is enabled, then the ALTER_OUTBOUND procedure automatically stops and restarts the capture process when the start_time parameter is non-NULL. If the capture process is disabled, then the ALTER_OUTBOUND procedure automatically starts the capture process when the start_time parameter is non-NULL.
include_dml	If TRUE, then the procedure creates a DML rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a DDL rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.
source_database	The global name of the container where the specified table_names and schema_names are located. If source_database is non-NULL, then a condition is added to the outbound server's rules to filter the LCRs based on the global name of the source database.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.

Usage Notes

The following list describes the behavior of the outbound server for various combinations of the table_names and schema_names parameters:

- If both the table_names and schema_names parameters are NULL or empty, then no rules are changed for the XStream Out configuration.

  This procedure is overloaded. The table_names and schema_names parameters are defaulted to NULL. Do not specify NULL for both table_names and schema_names in the same call; otherwise, error PLS-00307 is returned.

- If both the table_names and schema_names parameters are specified, then the rules for the tables and schemas are added to or removed from the XStream Out configuration, depending on the setting of the add parameter.

- If the table_names parameter is specified and the schema_names parameter is NULL or empty, then the rules for the tables are added to or removed from the XStream Out configuration, depending on the setting of the add parameter. The existing rules for schemas are not changed for the XStream Out configuration.

- If the `table_names` parameter is `NULL` or empty and the `schema_names` parameter is specified, then the rules for the schemas are added to or removed from the XStream Out configuration, depending on the setting of the `add` parameter. The existing rules for tables are not changed for the XStream Out configuration.

For the procedure that uses the `DBMS_UTILITY.UNCL_ARRAY` type for the `table_names` and `schema_names` parameters, both parameters must be specified. To specify only tables, the `schema_names` parameter must be specified and empty. To specify only schemas, the `table_names` parameter must be specified and empty.

> **Note:**
>
> An empty array includes one `NULL` entry.

239.4.14 CREATE_INBOUND Procedure

This procedure creates an XStream inbound server and its queue.

> **Note:**
>
> A client application can create multiple sessions. Each session can attach to only one inbound server, and each inbound server can serve only one session at a time. However, different client application sessions can connect to different inbound servers. See *Oracle Call Interface Programmer's Guide* and *Oracle Database XStream Java API Reference* for information about attaching to an inbound server.

Syntax

```
DBMS_XSTREAM_ADM.CREATE_INBOUND(
   server_name IN VARCHAR2,
   queue_name  IN VARCHAR2,
   apply_user  IN VARCHAR2  DEFAULT NULL,
   comment     IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-15 CREATE_INBOUND Procedure Parameters

Parameter	Description
`server_name`	The name of the inbound server being created. A `NULL` specification is not allowed. Do not specify an owner. The specified name must not match the name of an existing outbound server, inbound server, apply process, or messaging client. **Note:** The `server_name` setting cannot exceed 30 bytes, and it cannot be altered after the inbound server is created.

Table 239-15 (Cont.) CREATE_INBOUND Procedure Parameters

Parameter	Description
queue_name	The name of the local queue used by the inbound server, specified as [*schema_name.*]*queue_name*. If the specified queue exists, then it is used. If the specified queue does not exist, then the procedure creates it. For example, to specify a queue named xstream_queue in the xstrmadmin schema, enter xstrmadmin.xstream_queue for this parameter. If the schema is not specified, then the current user is the default. **Note:** An inbound server's queue is used only to store error transactions.
apply_user	The apply user. If NULL, then the current user is the default. The client application must attach to the inbound server as the apply user. The apply user is the user in whose security domain an inbound server evaluates whether LCRs satisfy its rule sets, applies DML and DDL changes directly to database objects, runs custom rule-based transformations configured for inbound server rules, and runs apply handlers configured for the inbound server. This user must have the necessary privileges to perform these actions. This procedure grants the apply user dequeue privileges on the queue used by the inbound server and configures the user as a secure queue user. In addition to the privileges granted by this procedure, you must grant the following privileges to the apply user: • The necessary privileges to perform DML and DDL changes on the apply objects • EXECUTE privilege on the rule sets used by the inbound server • EXECUTE privilege on all rule-based transformation functions used in the rule set • EXECUTE privilege on all apply handler procedures You can grant these privileges directly to the apply user, or you can grant them through roles. In addition, the apply user must be granted EXECUTE privilege on all packages, including Oracle supplied packages, that are invoked in subprograms run by the inbound server. These privileges must be granted directly to the apply user. They cannot be granted through roles. **Note:** If the apply user for an inbound server is dropped using DROP USER . . . CASCADE, then the inbound server is also dropped automatically.
comment	An optional comment associated with the inbound server.

Usage Notes

By default, an inbound server does not use rules or rule sets. Therefore, an inbound server applies all of the LCRs sent to it by an XStream client application. However, to filter the LCRs sent to an inbound server, you can add rules and rule sets to an inbound server using the DBMS_XSTREAM_ADM and DBMS_RULE_ADM packages.

In a CDB, you can execute the CREATE_INBOUND procedure from either the root or a PDB. The inbound server is restricted to receiving LCRs from one source database and only applying the changes to its local container. If the inbound server is at the root level, then the apply user must be a common user.

> **See Also:**
>
> *Oracle Database XStream Guide*

239.4.15 CREATE_OUTBOUND Procedure

This procedure creates an XStream outbound server, queue, and capture process to enable client applications to stream out Oracle database changes.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

This procedure is overloaded. One `table_names` parameter is type `VARCHAR2` and the other `table_names` parameter is type `DBMS_UTILITY.UNCL_ARRAY`. Also, one `schema_names` parameter is type `VARCHAR2` and the other `schema_names` parameter is type `DBMS_UTILITY.UNCL_ARRAY`. These parameters enable you to enter the list of tables and schemas in different ways and are mutually exclusive.

> **Note:**
>
> - A client application can create multiple sessions. Each session can attach to only one outbound server, and each outbound server can serve only one session at a time. However, different client application sessions can connect to different outbound servers. See "OCIXStreamOutAttach()" in the *Oracle Call Interface Programmer's Guide* and *Oracle Database XStream Java API Reference* for information about attaching to an outbound server.
> - If the `capture_name` parameter is `NULL`, then this procedure automatically generates a name for the capture process that it creates.
> - This procedure automatically generates a name for the queue that it creates.
> - This procedure enables both the capture process and outbound server that it creates.
> - Starting with Oracle Database 11*g* Release 2 (11.2.0.2), the `capture_name` parameter is included in this procedure.

Syntax

```
DBMS_XSTREAM_ADM.CREATE_OUTBOUND(
   server_name            IN VARCHAR2,
   source_database        IN VARCHAR2  DEFAULT NULL,
   table_names            IN DBMS_UTILITY.UNCL_ARRAY,
   schema_names           IN DBMS_UTILITY.UNCL_ARRAY,
```

```
   capture_user           IN VARCHAR2  DEFAULT NULL,
   connect_user           IN VARCHAR2  DEFAULT NULL,
   comment                IN VARCHAR2  DEFAULT NULL,
   capture_name           IN VARCHAR2  DEFAULT NULL,
   include_dml            IN BOOLEAN   DEFAULT TRUE,
   include_ddl            IN BOOLEAN   DEFAULT FALSE,
   source_root_name       IN VARCHAR2  DEFAULT NULL,
   source_container_name  IN VARCHAR2  DEFAULT NULL,
   lcrid_version          IN NUMBER    DEFAULT NULL);

DBMS_XSTREAM_ADM.CREATE_OUTBOUND(
   server_name            IN VARCHAR2,
   source_database        IN VARCHAR2  DEFAULT NULL,
   table_names            IN VARCHAR2  DEFAULT NULL,
   schema_names           IN VARCHAR2  DEFAULT NULL,
   capture_user           IN VARCHAR2  DEFAULT NULL,
   connect_user           IN VARCHAR2  DEFAULT NULL,
   comment                IN VARCHAR2  DEFAULT NULL,
   capture_name           IN VARCHAR2  DEFAULT NULL,
   include_dml            IN BOOLEAN   DEFAULT TRUE,
   include_ddl            IN BOOLEAN   DEFAULT FALSE,
   source_root_name       IN VARCHAR2  DEFAULT NULL,
   source_container_name  IN VARCHAR2  DEFAULT NULL,
   lcrid_version          IN NUMBER    DEFAULT NULL);
```

Parameters

Table 239-16 CREATE_OUTBOUND Procedure Parameters

Parameter	Description
server_name	The name of the outbound server being created. A NULL specification is not allowed. Do not specify an owner. The specified name must not match the name of an existing outbound server, inbound server, apply process, or messaging client. **Note:** The server_name setting cannot exceed 30 bytes, and it cannot be altered after the outbound server is created.
source_database	The global name of the source database. The source database is where the changes to be captured originated. If non-NULL, then a condition is added to the outbound server's rules to filter the LCRs based on the global name of the source database. If NULL, then the procedure does not add a condition regarding the source database to the generated rules. In a CDB, specify the global name of the container to which the rules pertain. The container can be the root or a PDB. For example, mycdb.example.com or hrpdb.example.com. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. In a non-CDB, if non-NULL and the specified name is different from the global name of the current database, then downstream capture is assumed. In this case, configure the transmission of redo data from the source database to the downstream database before running the CREATE_OUTBOUND procedure. See *Oracle Database XStream Guide* for instructions. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically.

Table 239-16 (Cont.) CREATE_OUTBOUND Procedure Parameters

Parameter	Description
table_names	The tables for which DML and DDL changes are streamed out to the XStream client application. The tables can be specified in the following ways: • Comma-delimited list of type VARCHAR2. • A PL/SQL associative array of type DBMS_UTILITY.UNCL_ARRAY, where each element is the name of a table. Specify the first table in position 1. The last position must be NULL. Each table should be specified as [*schema_name.*]*table_name*. For example, hr.employees. If the schema is not specified, then the current user is the default. **See Also:** "Usage Notes" for more information about this parameter
schema_names	The schemas for which DML and DDL changes are streamed out to the XStream client application. The schemas can be specified in the following ways: • Comma-delimited list of type VARCHAR2. • A PL/SQL associative array of type DBMS_UTILITY.UNCL_ARRAY, where each element is the name of a schema. Specify the first schema in position 1. The last position must be NULL. **Note:** This procedure does not concatenate the schema_names parameter with the table_names parameter. To specify tables, enter fully qualified table names in the table_names parameter (*schema_name.table_name*). **See Also:** "Usage Notes" for more information about this parameter
capture_user	The user in whose security domain a capture process captures changes that satisfy its rule sets and runs custom rule-based transformations configured for capture process rules. If NULL, then the current user is the default. This procedure grants the capture user enqueue privilege on the queue used by the capture process and configures the user as a secure queue user. In addition, ensure that the capture user has the following privileges: • EXECUTE privilege on the rule sets used by the capture process • EXECUTE privilege on all rule-based transformation functions used in the positive rule set You can grant these privileges directly to the apply user, or you can grant them through roles. In addition, the capture user must be granted EXECUTE privilege on all packages, including Oracle supplied packages, that are invoked in rule-based transformations run by the capture process. These privileges must be granted directly to the capture user. They cannot be granted through roles. Only a user who is granted the DBA role can set a capture user. Only the SYS user can set the capture_user to SYS. A capture user does not require privileges on a database object to capture changes made to it. The capture process can pass these changes to a custom rule-based transformation function. Therefore, ensure that you consider security implications when you configure a capture process.

Table 239-16 (Cont.) CREATE_OUTBOUND Procedure Parameters

Parameter	Description
connect_user	The user who can attach to the specified outbound server to retrieve the change stream. The client application must attach to the outbound server as the specified connect user. If NULL, then the current user is the default. The connect user is the user in whose security domain an outbound server evaluates LCRs against its rule sets and runs custom rule-based transformations configured for outbound server rules. This user must have the necessary privileges to perform these actions. This procedure grants the connect user dequeue privileges on the queue used by the outbound server and configures the user as a secure queue user. In addition to the privileges granted by this procedure, grant the following privileges to the connect user: • EXECUTE privilege on the rule sets used by the outbound server • EXECUTE privilege on all rule-based transformation functions used in the rule set You can grant these privileges directly to the connect user, or you can grant them through roles. In addition, the connect user must be granted EXECUTE privilege on all packages, including Oracle supplied packages, that are invoked in subprograms run by the outbound server. These privileges must be granted directly to the apply user. They cannot be granted through roles.
comment	An optional comment associated with the outbound server.
capture_name	The name of the capture process configured to capture changes for the outbound server. Do not specify an owner. The capture process must not exist. If the specified name matches the name of an existing capture process, then an error is raised. If the name does not match the name of an existing capture process, then the procedure creates a new capture process with the specified name. If NULL, then the system creates a new capture process with a system-generated name. **Note:** The capture process name cannot be altered after the capture process is created.
include_dml	If TRUE, then the procedure creates a DML rule for DML changes. If FALSE, then the procedure does not create a DML rule. NULL is not permitted.
include_ddl	If TRUE, then the procedure creates a DDL rule for DDL changes. If FALSE, then the procedure does not create a DDL rule. NULL is not permitted.

Table 239-16 (Cont.) CREATE_OUTBOUND Procedure Parameters

Parameter	Description
source_root_name	The global name of the root in the source CDB. For example, mycdb.example.com. If this parameter is NULL, then the global name of the root in the local CDB is used. If you are configuring downstream capture, then this parameter must be a non-NULL value, and it must specify the global name of the root in the remote source CDB. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify DBS1 and the domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. **Note:** This parameter only applies to a CDB.
source_container_name	The short name of the source container. The container can be the root or a PDB. For example, CDB$ROOT or hrpdb. See *Oracle Database XStream Guide* for more information about setting this parameter in a CDB. **Note:** This parameter only applies to a CDB.
lcrid_version	The LCRID version for captured LCRs, either 1 or 2. If 2, then the LCRs are compatible with a database with its compatibility level at 12.2.0 or higher. If 1, then the LCRs are compatible with a database with its compatibility level at 12.1.0 or lower. If NULL, the default, and the database compatibility level is 12.2.0 or higher, then the lcrid_version is set to 2 internally. If the database compatibility level is 12.1.0 or lower, then the lcrid_version is set to 1 internally.

Usage Notes

The following list describes the behavior of the outbound server for various combinations of the table_names and schema_names parameters:

- If both the table_names and schema_names parameters are NULL or empty, then the outbound server streams all DML and DDL changes to the client application.

  This procedure is overloaded. The table_names and schema_names parameters are defaulted to NULL. Do not specify NULL for both table_names and schema_names in the same call; otherwise, error PLS-00307 is returned.
- If both the table_names and schema_names parameters are specified, then the outbound server streams DML and DDL changes for the specified tables and schemas.
- If the table_names parameter is specified and the schema_names parameter is NULL or empty, then the outbound server streams DML and DDL changes for the specified tables.
- If the table_names parameter is NULL or empty and the schema_names parameter is specified, then the outbound server streams DML and DDL changes for the specified schema.

For the procedure that uses the DBMS_UTILITY.UNCL_ARRAY type for the table_names and schema_names parameters, both parameters must be specified. To specify only tables, the

schema_names parameter must be specified and empty. To specify only schemas, the table_names parameter must be specified and empty.

> **Note:**
>
> An empty array includes one NULL entry.

239.4.16 DELETE_COLUMN Procedure

This procedure either adds or removes a declarative rule-based transformation which deletes a column from a row logical change record (LCR) that satisfies the specified rule.

For the transformation to be performed when the specified rule evaluates to TRUE, the rule must be in the positive rule set of an XStream client. XStream clients include capture processes, propagations, and apply processes.

> **Note:**
>
> - The DELETE_COLUMN procedure supports the same data types supported by Oracle Replication capture processes.
> - The DELETE_COLUMN procedure is useful when you want to delete a relatively small number of columns in a row LCR. To delete most of the columns in a row LCR and keep a relatively small number of columns, consider using the KEEP_COLUMNS procedure in this package.
> - Declarative transformations can transform row LCRs only. Therefore, a DML rule must be specified when you run this procedure. If a DDL rule is specified, then the procedure raises an error.

> **See Also:**
>
> - *Oracle Database XStream Guide* for more information about declarative rule-based transformations and about the data types supported by capture processes
> - KEEP_COLUMNS Procedure

Syntax

```
DBMS_XSTREAM_ADM.DELETE_COLUMN(
   rule_name      IN  VARCHAR2,
   table_name     IN  VARCHAR2,
   column_name    IN  VARCHAR2,
   value_type     IN  VARCHAR2   DEFAULT '*',
   step_number    IN  NUMBER     DEFAULT 0,
   operation      IN  VARCHAR2   DEFAULT 'ADD');
```

Parameters

Table 239-17 DELETE_COLUMN Procedure Parameters

Parameter	Description
rule_name	The name of the rule, specified as [*schema_name.*]*rule_name*. If NULL, then the procedure raises an error. For example, to specify a rule in the hr schema named employees12, enter hr.employees12. If the schema is not specified, then the current user is the default.
table_name	The name of the table from which the column is deleted in the row LCR, specified as [*schema_name.*]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default.
column_name	The name of the column deleted from each row LCR that satisfies the rule.
value_type	Specify 'NEW' to delete the column from the new values in the row LCR. Specify 'OLD' to delete the column from the old values in the row LCR. Specify '*' to delete the column from both the old and new values in the row LCR.
step_number	The order of execution of the transformation. **See Also:** *Oracle Database XStream Guide* for more information about transformation ordering
operation	Specify 'ADD' to add the transformation to the rule. Specify 'REMOVE' to remove the transformation from the rule.

Usage Notes

When 'REMOVE' is specified for the operation parameter, all of the delete column declarative rule-based transformations for the specified rule are removed that match the specified table_name, column_name, and step_number parameters. Nulls specified for these parameters act as wildcards. The following table lists the behavior of the DELETE_COLUMN procedure when one or more of these parameters is NULL:

table_name	column_name	step_number	Result
NULL	NULL	NULL	Remove all delete column transformations for the specified rule.
NULL	NULL	non-NULL	Remove all delete column transformations with the specified step_number for the specified rule.
NULL	non-NULL	non-NULL	Remove all delete column transformations with the specified column_name and step_number for the specified rule.
non-NULL	NULL	non-NULL	Remove all delete column transformations with the specified table_name and step_number for the specified rule.
NULL	non-NULL	NULL	Remove all delete column transformations with the specified column_name for the specified rule.

table_name	column_name	step_number	Result
non-NULL	non-NULL	NULL	Remove all delete column transformations with the specified table_name and column_name for the specified rule.
non-NULL	NULL	NULL	Remove all delete column transformations with the specified table_name for the specified rule.
non-NULL	non-NULL	non-NULL	Remove all delete column transformations with the specified table_name, column_name, and step_number for the specified rule.

239.4.17 DROP_INBOUND Procedure

This procedure removes an inbound server configuration.

This procedure always removes the specified inbound server. This procedure also removes the queue for the inbound server if all of the following conditions are met:

- One call to the CREATE_INBOUND procedure created the queue.
- The inbound server is the only subscriber to the queue.

> **See Also:**
>
> "CREATE_INBOUND Procedure"

Syntax

```
DBMS_XSTREAM_ADM.DROP_INBOUND(
   server_name IN VARCHAR2);
```

Parameters

Table 239-18 DROP_INBOUND Procedure Parameters

Parameter	Description
server_name	The name of the inbound server being removed. Specify an existing inbound server. Do not specify an owner.

239.4.18 DROP_OUTBOUND Procedure

This procedure removes an outbound server configuration.

This procedure always drops the specified outbound server. This procedure also drops the queue used by the outbound server if both of the following conditions are met:

- The queue was created by the CREATE_OUTBOUND procedure in this package.
- The outbound server is the only subscriber to the queue.

If either one of the preceding conditions is not met, then the `DROP_OUTBOUND` procedure only drops the outbound server. It does not drop the queue.

This procedure also drops the capture process for the outbound server if both of the following conditions are met:

- The procedure can drop the outbound server's queue.
- The capture process was created by the `CREATE_OUTBOUND` procedure.

If the procedure can drop the queue but cannot manage the capture process, then it drops the queue without dropping the capture process.

> **See Also:**
>
> - "ADD_OUTBOUND Procedure"
> - "CREATE_OUTBOUND Procedure"

Syntax

```
DBMS_XSTREAM_ADM.DROP_OUTBOUND(
   server_name IN VARCHAR2);
```

Parameters

Table 239-19 DROP_OUTBOUND Procedure Parameters

Parameter	Description
`server_name`	The name of the outbound server being removed. Specify an existing outbound server. Do not specify an owner.

239.4.19 ENABLE_GG_XSTREAM_FOR_STREAMS Procedure

This procedure enables XStream optimizations and performance optimizations for Oracle Replication components.

This procedure is intended for users of Oracle Replication who want to enable XStream optimizations and optimizations. For example, you can enable the optimizations for an Oracle Replication configuration that uses capture processes and apply processes to replicate changes between Oracle databases.

These capabilities and optimizations are enabled automatically for XStream components, such as outbound servers, inbound servers, and capture processes that send changes to outbound servers. It is not necessary to run this procedure for XStream components.

When XStream optimizations are enabled, Oracle Replication components can stream ID key LCRs and sequence LCRs. The XStream performance optimizations improve efficiency in various areas, including:

- LCR processing
- Handling large transactions
- DML execution during apply
- Dependency computation and scheduling

- Capture process parallelism

Syntax

```
DBMS_XSTREAM_ADM.ENABLE_GG_XSTREAM_FOR_STREAMS(
   enable IN BOOLEAN  TRUE);
```

Parameters

Table 239-20 ENABLE_GG_XSTREAM_FOR_STREAMS Procedure Parameters

Parameter	Description
`enable`	If `TRUE`, then enable XStream performance optimizations for Oracle Replication components. If `FALSE`, then disable XStream performance optimizations for Oracle Replication components.

Usage Notes

The following usage notes apply to this procedure:

- When you run this procedure, all capture processes and apply processes are restarted.
- After you run this procedure, the `PURPOSE` column in the following views displays `XStream Streams`:
  - `ALL_APPLY`
  - `DBA_APPLY`
  - `ALL_CAPTURE`
  - `DBA_CAPTURE`
- A license for the Oracle GoldenGate product is required to enable XStream performance optimizations for Oracle Replication components.

See Also:

- IS_GG_XSTREAM_FOR_STREAMS Function
- *Oracle Database XStream Guide, Chapter 1, Prerequisites for XStream"*

239.4.20 GET_MESSAGE_TRACKING Function

The `GET_MESSAGE_TRACKING` Function returns the tracking label for the current session.

See Also:

SET_MESSAGE_TRACKING Procedure

Syntax

```
DBMS_XSTREAM_ADM.GET_MESSAGE_TRACKING
RETURN VARCHAR2;
```

239.4.21 GET_TAG Function

This function gets the binary tag for all redo entries generated by the current session.

> **See Also:**
>
> "SET_TAG Procedure"

Syntax

```
DBMS_XSTREAM_ADM.GET_TAG
RETURN RAW;
```

Examples

The following example illustrates how to display the current logical change record (LCR) tag as output:

```
SET SERVEROUTPUT ON
DECLARE
   raw_tag RAW(2000);
BEGIN
   raw_tag := DBMS_XSTREAM_ADM.GET_TAG();
   DBMS_OUTPUT.PUT_LINE('Tag Value = ' || RAWTOHEX(raw_tag));
END;
/
```

You can also display the value by querying the `DUAL` view:

```
SELECT DBMS_XSTREAM_ADM.GET_TAG FROM DUAL;
```

239.4.22 IS_GG_XSTREAM_FOR_STREAMS Function

This function returns `TRUE` if XStream performance optimizations are enabled for Oracle Replication components, or this function returns `FALSE` if XStream performance optimizations are disabled for Oracle Replication components.

> **See Also:**
>
> "ENABLE_GG_XSTREAM_FOR_STREAMS Procedure"

Syntax

```
DBMS_XSTREAM_ADM.IS_GG_XSTREAM_FOR_STREAMS
RETURN BOOLEAN;
```

239.4.23 KEEP_COLUMNS Procedure

This procedure either adds or removes a declarative rule-based transformation which keeps a list of columns in a row logical change record (LCR) that satisfies the specified rule. The transformation deletes columns that are not in the list from the row LCR.

For the transformation to be performed when the specified rule evaluates to `TRUE`, the rule must be in the positive rule set of an XStream client. XStream clients include capture processes, propagations, and apply processes.

This procedure is overloaded. The `column_list` parameter is type `VARCHAR2` and the `column_table` parameter is type `DBMS_UTILITY.LNAME_ARRAY`. These parameters enable you to enter the list of columns in different ways and are mutually exclusive.

Note:

- The `KEEP_COLUMNS` procedure supports the same data types supported by Oracle Replication capture processes.
- The `KEEP_COLUMNS` procedure is useful when you want to keep a relatively small number of columns in a row LCR. To keep most of the columns in a row LCR and delete a relatively small number of columns, consider using the `DELETE_COLUMN` procedure in this package.
- Declarative transformations can transform row LCRs only. Therefore, a DML rule must be specified when you run this procedure. If a DDL rule is specified, then the procedure raises an error.

See Also:

- *Oracle Database XStream Guide* for more information about declarative rule-based transformations and about the data types supported by Oracle Replication capture processes
- DELETE_COLUMN Procedure

Syntax

```
DBMS_XSTREAM_ADM.KEEP_COLUMNS(
   rule_name     IN  VARCHAR2,
   table_name    IN  VARCHAR2,
   column_list   IN  VARCHAR2,
   value_type    IN  VARCHAR2 DEFAULT '*',
   step_number   IN  NUMBER DEFAULT 0,
   operation     IN  VARCHAR2 DEFAULT 'ADD');

DBMS_XSTREAM_ADM.KEEP_COLUMNS(
   rule_name     IN  VARCHAR2,
   table_name    IN  VARCHAR2,
   column_table  IN  DBMS_UTILITY.LNAME_ARRAY,
   value_type    IN  VARCHAR2 DEFAULT '*',
```

```
  step_number    IN  NUMBER DEFAULT 0,
  operation      IN  VARCHAR2 DEFAULT 'ADD');
```

Parameters

Table 239-21 KEEP_COLUMNS Procedure Parameters

Parameter	Description
rule_name	The name of the rule, specified as [*schema_name.*]*rule_name*. If NULL, then the procedure raises an error. For example, to specify a rule in the hr schema named employees12, enter hr.employees12. If the schema is not specified, then the current user is the default.
table_name	The name of the table for which the columns are kept in the row LCR, specified as [*schema_name.*]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default.
column_list	The names of the columns kept for each row LCR that satisfies the rule. Specify a comma-delimited list of type VARCHAR2. The transformation removes columns that are not in the list from the row LCR. If this parameter is set to NULL, and the column_table parameter is also set to NULL, then the procedure raises an error.
column_table	The names of the columns kept for each row LCR that satisfies the rule. Specify a PL/SQL associative array of type DBMS_UTILITY.LNAME_ARRAY, where each element is the name of a column. The first schema should be in position 1. The last position must be NULL. The transformation removes columns that are not in the table from the row LCR. If this parameter is set to NULL, and the column_list parameter is also set to NULL, then the procedure raises an error.
value_type	Specify 'NEW' to keep the columns in the new values in the row LCR. Specify 'OLD' to keep the columns in the old values in the row LCR. Specify '*' to keep the columns in both the old and new values in the row LCR.
step_number	The order of execution of the transformation. **See Also:** *Oracle Database XStream Guide* for more information about transformation ordering
operation	Specify 'ADD' to add the transformation to the rule. Specify 'REMOVE' to remove the transformation from the rule.

Usage Notes

When 'REMOVE' is specified for the operation parameter, all of the keep columns declarative rule-based transformations for the specified rule are removed that match the specified table_name, column_list, column_table, and step_number parameters. Nulls specified for these parameters act as wildcards. The following table lists the behavior of the KEEP_COLUMNS procedure when one or more of these parameters is NULL:

table_name	column_list/ column_table	step_number	Result
NULL	NULL	NULL	Remove all keep columns transformations for the specified rule.

table_name	column_list/ column_table	step_number	Result
NULL	NULL	non-NULL	Remove all keep columns transformations with the specified step_number for the specified rule.
NULL	non-NULL	non-NULL	Remove all keep columns transformations with the specified column_list/ column_table and step_number for the specified rule.
non-NULL	NULL	non-NULL	Remove all keep columns transformations with the specified table_name and step_number for the specified rule.
NULL	non-NULL	NULL	Remove all keep columns transformations with the specified column_list/ column_table for the specified rule.
non-NULL	non-NULL	NULL	Remove all keep columns transformations with the specified table_name and column_list/column_table for the specified rule.
non-NULL	NULL	NULL	Remove all keep columns transformations with the specified table_name for the specified rule.
non-NULL	non-NULL	non-NULL	Remove all keep columns transformations with the specified table_name, column_list/column_table, and step_number for the specified rule.

239.4.24 MERGE_STREAMS Procedure

This procedure merges a stream that is flowing from one capture process with a stream that is flowing from another capture process.

Typically, this procedure is used to merge two streams that were split using the SPLIT_STREAMS procedure in this package. The SPLIT_STREAMS procedure clones components of the original stream when it splits the streams. Therefore, the information in this section uses the following terminology:

- The stream before it was split off has the original queue, original capture process, and original propagation.
- The stream that was split off by the SPLIT_STREAMS procedure has a cloned queue, cloned capture process, and cloned propagation.

This procedure is called by the MERGE_STREAMS_JOB procedure. The MERGE_STREAMS_JOB procedure determines whether the streams are within a user-specified merge threshold so that the streams can be merged safely. If the streams are not within the merge threshold, then the MERGE_STREAMS_JOB procedure does nothing. Typically, it is best to run the MERGE_STREAMS_JOB procedure instead of running the MERGE_STREAMS procedure directly.

However, you can choose to run the MERGE_STREAMS procedure directly when the following conditions are met:

- The problem at the destination of the split stream has been corrected, and the destination queue can accept changes.

- The cloned capture process used by the split stream is started and is capturing changes.
- The apply process at the destination database is applying the changes captured by the cloned capture process.
- The `CAPTURE_MESSAGE_CREATE_TIME` in the `GV$XSTREAM_CAPTURE` view of the cloned capture process has caught up to, or nearly caught up to, the `CAPTURE_MESSAGE_CREATE_TIME` of the original capture process. The cloned capture process might never completely catch up to the original capture process. Therefore, you can merge the split stream when the cloned capture process has nearly caught up to the original capture process.

The `MERGE_STREAMS` procedure performs the following actions:

1. Stops the cloned capture process.
2. Stops the original capture process.
3. Copies the cloned propagation back to the original propagation. The propagation has the same name as the original propagation after it is copied back.
4. Starts the original capture process from the lower SCN value of these two SCN values:
   - The acknowledged SCN of the cloned propagation.
   - The lowest acknowledged SCN of the other propagations that propagate changes captured by the original capture process.

   When the original capture process is started, it might recapture changes that it already captured, or it might capture changes that were already captured by the cloned capture process. In either case, the relevant apply processes will discard any duplicate changes they receive.
5. Drops the cloned propagation.
6. Drops the cloned capture process.
7. Drops the cloned queue.

> **See Also:**
>
> SPLIT_STREAMS Procedure

Syntax

```
DBMS_XSTREAM_ADM.MERGE_STREAMS(
   cloned_propagation_name  IN  VARCHAR2,
   propagation_name         IN  VARCHAR2  DEFAULT NULL,
   queue_name               IN  VARCHAR2  DEFAULT NULL,
   perform_actions          IN  BOOLEAN   DEFAULT TRUE,
   script_name              IN  VARCHAR2  DEFAULT NULL,
   script_directory_object  IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-22 MERGE_STREAMS Procedure Parameters

Parameter	Description
cloned_propagation_name	The name of the cloned propagation used by the stream that was split off from the original stream using the SPLIT_STREAMS procedure. The name of the cloned propagation also identifies the cloned queue and capture process used by the cloned propagation. You must specify an existing propagation name. Do not specify an owner.
propagation_name	The name of the propagation that is merged back to the original stream. If NULL, then the name of the original propagation in the original stream is used. Specify NULL only if the streams were split using the SPLIT_STREAMS procedure. Specify a non-NULL value to use a name that is different from the original propagation name or if you are merging two streams that were not split by the SPLIT_STREAMS procedure. If a non-NULL value is specified, then an error is raised under either of the following conditions: • The queue specified in the queue_name parameter does not exist. • The queue specified in the queue_name parameter exists but is not used by a capture process.
queue_name	The name of the queue that is the source queue for the propagation that is merged back. If NULL, then the existing, original queue is the source queue for the propagation that is merged back. Specify NULL only if the streams were split using the SPLIT_STREAMS procedure. Specify a non-NULL value if you are merging two streams that were not split by the SPLIT_STREAMS procedure. Specify the name of the existing queue used by the capture process that will capture changes in the merged stream.
perform_actions	If TRUE, then the procedure performs the necessary actions to merge the streams directly. If FALSE, then the procedure does not perform the necessary actions to merge the streams directly. Specify FALSE when this procedure is generating a script that you can edit and then run. The procedure raises an error if you specify FALSE and either of the following parameters is NULL: • script_name • script_directory_object

Table 239-22 (Cont.) MERGE_STREAMS Procedure Parameters

Parameter	Description
script_name	If non-NULL and the perform_actions parameter is FALSE, then specify the name of the script generated by this procedure. The script contains all of the statements used to merge the streams. If a file with the specified script name exists in the specified directory for the script_directory_object parameter, then the procedure appends the statements to the existing file. If non-NULL and the perform_actions parameter is TRUE, then the procedure generates the specified script and performs the actions to split the stream directly. If NULL and the perform_actions parameter is TRUE, then the procedure performs the actions to merge the streams directly and does not generate a script. If NULL and the perform_actions parameter is FALSE, then the procedure raises an error.
script_directory_object	The directory object for the directory on the local computer system into which the generated script is placed. If the script_name parameter is NULL, then the procedure ignores this parameter and does not generate a script. If NULL and the script_name parameter is non-NULL, then the procedure raises an error. **Note:** The specified directory object cannot point to an Oracle Automatic Storage Management (ASM) disk group.

Usage Notes

You can use the MERGE_STREAMS procedure to merge two streams that were not split using the SPLIT_STREAMS procedure. Merging streams in this way can save resources and improve performance when a single database is running two or more capture processes.

239.4.25 MERGE_STREAMS_JOB Procedure

This procedure determines whether the original capture process and the cloned capture process are within the specified merge threshold. If they are within the merge threshold, then this procedure runs the MERGE_STREAMS procedure to merge the two streams.

Typically, this procedure is used to merge two streams that were split using the SPLIT_STREAMS procedure in this package. The SPLIT_STREAMS procedure clones components of the original stream when it splits the streams. Therefore, the information in this section uses the following terminology:

- The stream before it was split off has the original queue, original capture process, and original propagation.
- The stream that was split off by the SPLIT_STREAMS procedure has a cloned queue, cloned capture process, and cloned propagation.

If the auto_merge_threshold parameter was set to a positive number in the SPLIT_STREAMS procedure that split the streams, then a merge job runs the MERGE_STREAMS_JOB procedure automatically according to its schedule. The schedule name is specified for the schedule_name parameter, and the merge job name is specified for the merge_job_name parameter when the

MERGE_STREAMS_JOB procedure is run automatically. The merge job and its schedule were created by the SPLIT_STREAMS procedure.

If the auto_merge_threshold parameter was set to NULL or 0 (zero) in the SPLIT_STREAMS procedure that split the streams, then you can run the MERGE_STREAMS_JOB procedure manually. In this case, it is not run automatically.

See Also:

- MERGE_STREAMS Procedure
- SPLIT_STREAMS Procedure

Syntax

```
DBMS_XSTREAM_ADM.MERGE_STREAMS_JOB(
   cloned_propagation_name IN VARCHAR2,
   propagation_name        IN VARCHAR2 DEFAULT NULL,
   queue_name              IN VARCHAR2 DEFAULT NULL,
   merge_threshold         IN NUMBER,
   schedule_name           IN VARCHAR2 DEFAULT NULL,
   merge_job_name          IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 239-23 MERGE_STREAMS_JOB Procedure Parameters

Parameter	Description
cloned_propagation_name	The name of the cloned propagation used by the stream that was split off from the original stream using the SPLIT_STREAMS procedure. The name of the cloned propagation also identifies the cloned queue and capture process used by the cloned propagation. You must specify an existing propagation name. Do not specify an owner.
propagation_name	The name of the propagation that is merged back to the original stream. If NULL, then the name of the original propagation in the original stream is used. Specify NULL only if the streams were split using the SPLIT_STREAMS procedure. Specify a non-NULL value to use a name that is different from the original propagation name or if you are merging two streams that were not split by the SPLIT_STREAMS procedure. If a non-NULL value is specified, then an error is raised under either of the following conditions: • The queue specified in the queue_name parameter does not exist. • The queue specified in the queue_name parameter exists but is not used by a capture process.

Table 239-23 (Cont.) MERGE_STREAMS_JOB Procedure Parameters

Parameter	Description
queue_name	The name of the queue that is the source queue for the propagation that is merged back. If NULL, then the existing, original queue is the source queue for the propagation that is merged back. Specify NULL only if the streams were split using the SPLIT_STREAMS procedure. Specify a non-NULL value if you are merging two streams that were not split by the SPLIT_STREAMS procedure. Specify the name of the existing queue used by the capture process that will capture changes in the merged stream.
merge_threshold	The merge threshold in seconds. The value of the CAPTURE_MESSAGE_CREATE_TIME column for each capture process in the GV$XSTREAM_CAPTURE dynamic performance view determines whether the streams are merged. Specifically, if the difference, in seconds, between the CAPTURE_MESSAGE_CREATE_TIME of the cloned capture process and the original capture process is less than or equal to the value specified for this parameter, then this procedure runs the MERGE_STREAMS procedure to merge the streams. If the difference is greater than the value specified by this parameter, then this procedure does nothing.
schedule_name	The name of the schedule for the merge job. If NULL, then no schedule name is specified. Typically, you set this parameter to NULL when the auto_merge_threshold parameter was set to NULL or 0 (zero) in the SPLIT_STREAMS procedure that split the streams. Specify NULL if you run this procedure manually.
merge_job_name	The name of the job that merges the streams. If NULL, then no merge job name is specified. Typically, you set this parameter to NULL when the auto_merge_threshold parameter was set to NULL or 0 (zero) in the SPLIT_STREAMS procedure that split the streams. Specify NULL if you run this procedure manually.

Usage Notes

You can use the MERGE_STREAMS_JOB procedure to merge two streams that were not split using the SPLIT_STREAMS procedure. Merging streams in this way can save resources and improve performance when a single database is running two or more capture processes.

After the MERGE_STREAMS_JOB procedure completes, you can query the DBA_CAPTURE and DBA_PROPAGATION views to determine whether the streams were merged. If the streams were merged, then the cloned capture process and cloned propagation do not appear in these views.

If the streams were merged and the schedule_name and merge_job_name parameters were non-NULL, then the specified schedule and merge job are deleted automatically.

239.4.26 PURGE_SOURCE_CATALOG Procedure

This procedure removes all Oracle Replication data dictionary information at the local database for the specified object.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

You can use this procedure to remove Oracle Replication metadata that is not needed currently and will not be needed in the future.

Syntax

```
DBMS_XSTREAM_ADM.PURGE_SOURCE_CATALOG(
   source_database     IN  VARCHAR2,
   source_object_name  IN  VARCHAR2,
   source_object_type  IN  VARCHAR2,
   source_root_name    IN  VARCHAR2);
```

Parameters

Table 239-24 PURGE_SOURCE_CATALOG Procedure Parameters

Parameter	Description
`source_database`	In a non-CDB, specify the global name of the source database containing the database object. In a CDB, specify the global name of the container containing the database object. The container can be the root or a PDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify `DBS1` and the domain is `.NET`, then the procedure specifies `DBS1.NET` automatically.
`source_object_name`	The name of the object specified as `[`*`schema_name`*`.]`*`object_name`*. For example, `hr.employees`. If the schema is not specified, then the current user is the default.
`source_object_type`	Type of the object. Currently, `TABLE` is the only possible object type.
`source_root_name`	The global name of the source root containing the object in a CDB. The source root is where the changes being captured originated in a CDB. If you do not include the domain name, then the procedure appends it to the database name automatically. For example, if you specify `DBS1` and the domain is `EXAMPLE.COM`, then the procedure specifies `DBS1.EXAMPLE.COM` automatically. If the `source_root_name` parameter is `NULL`, then the global name of the local root is the default. **Note:** This parameter only applies to a CDB.

Usage Notes

The global name of the source database containing the object must be specified for the `source_database` parameter. If the current database is not the source database for the object, then the procedure removes data dictionary information about the object from the current database, not the source database.

For example, suppose changes to the `hr.employees` table at the `dbs1.net` source database are being applied to the `hr.employees` table at the `dbs2.net` destination database. Also, suppose `hr.employees` at `dbs2.net` is not a source at all. In this case, specifying `dbs2.net` as the `source_database` for this table results in an error. However, specifying `dbs1.net` as the `source_database` for this table while running the `PURGE_SOURCE_CATALOG` procedure at the `dbs2.net` database removes data dictionary information about the table at `dbs2.net`.

Do not run this procedure at a database if either of the following conditions is true:

- Logical change records (LCRs) captured by the capture process for the object are or might be applied locally without reinstantiating the object.
- LCRs captured by the capture process for the object are or might be forwarded by the database without reinstantiating the object.

> **Note:**
>
> These conditions do not apply to LCRs that were not created by the capture process. That is, these conditions do not apply to user-created LCRs.

239.4.27 RECOVER_OPERATION Procedure

This procedure provides options for split and merge operations that stopped because they encountered an errors.

This procedure either rolls forward the operation, rolls back the operation, or purges all of the metadata about the operation. Split and merge operations might be run in an XStream Out environment in which multiple outbound servers use the same capture process.

This procedure only can perform these actions for split and merge operations using the `split_threshold` and `merge_threshold` capture process parameters set to non-`NULL` values to enable automatic split and merge.

Information about the operation is stored in the following data dictionary views when the operation is in process:

- `DBA_RECOVERABLE_SCRIPT`
- `DBA_RECOVERABLE_SCRIPT_PARAMS`
- `DBA_RECOVERABLE_SCRIPT_BLOCKS`
- `DBA_RECOVERABLE_SCRIPT_ERRORS`

The data dictionary views are populated at the database that contains the capture process.

When the operation completes successfully, metadata about the operation is moved from the `DBA_RECOVERABLE_SCRIPT` view to the `DBA_RECOVERABLE_SCRIPT_HIST` view. The other views, `DBA_RECOVERABLE_SCRIPT_PARAMS`, `DBA_RECOVERABLE_SCRIPT_BLOCKS`, and

DBA_RECOVERABLE_SCRIPT_ERRORS, retain information about the operation until it is purged automatically after 30 days.

When one of these operations encounters an error and stops, metadata about the operation remains in these views. In this case, you can either roll forward, roll back, or purge the metadata about the operation using the RECOVER_OPERATION procedure. If you choose to roll forward the operation, then correct conditions that caused the errors reported in DBA_RECOVERABLE_SCRIPT_ERRORS before proceeding.

Run the RECOVER_OPERATION procedure at the database that contains the capture process.

Note:

To run the RECOVER_OPERATION procedure, both databases must be Oracle Database 10*g* Release 2 or later databases.

See Also:

- "SPLIT_STREAMS Procedure"
- "MERGE_STREAMS Procedure"
- "MERGE_STREAMS_JOB Procedure"

Syntax

```
DBMS_XSTREAM_ADM.RECOVER_OPERATION(
   script_id       IN  RAW,
   operation_mode  IN  VARCHAR2  DEFAULT 'FORWARD');
```

Parameters

Table 239-25 RECOVER_OPERATION Procedure Parameters

Parameter	Description
script_id	The operation id of the operation that is being rolled forward, rolled back, or purged. Query the SCRIPT_ID column of the DBA_RECOVERABLE_SCRIPT data dictionary view to determine the operation id.

Table 239-25 (Cont.) RECOVER_OPERATION Procedure Parameters

Parameter	Description
operation_mode	If FORWARD, then the procedure rolls forward the operation. Specify FORWARD to try to complete the operation. If ROLLBACK, then the procedure rolls back all of the actions performed in the operation. If the rollback is successful, then this option also moves the metadata about the operation from the DBA_RECOVERABLE_SCRIPT view to the DBA_RECOVERABLE_SCRIPT_HIST view. The other views retain information about the operation for 30 days. If PURGE, then the procedure moves the metadata about the operation from the DBA_RECOVERABLE_SCRIPT view to the DBA_RECOVERABLE_SCRIPT_HIST view without rolling the operation back. The other views retain information about the operation for 30 days.

239.4.28 REMOVE_QUEUE Procedure

This procedure removes the specified ANYDATA queue.

Specifically, this procedure performs the following actions:

1. Waits until all current enqueue and dequeue transactions commit.
2. Stops the queue, which means that no further enqueues into the queue or dequeues from the queue are allowed.
3. Drops the queue.
4. If the drop_unused_queue_table parameter is set to TRUE, then drops the queue table if it is empty and no other queues are using it.
5. If the cascade parameter is set to TRUE, then drops all of the XStream clients that are using the queue.

> **Note:**
>
> The specified queue must be a ANYDATA queue.

Syntax

```
DBMS_XSTREAM_ADM.REMOVE_QUEUE(
   queue_name                IN  VARCHAR2,
   cascade                   IN  BOOLEAN  DEFAULT FALSE,
   drop_unused_queue_table   IN  BOOLEAN  DEFAULT TRUE);
```

Parameters

Table 239-26 REMOVE_QUEUE Procedure Parameters

Parameter	Description
queue_name	The name of the queue to remove, specified as [*schema_name*.]*queue_name*. For example, strmadmin.streams_queue. If the schema is not specified, then the current user is the default.
cascade	If TRUE, then the procedure drops any XStream clients that use the queue. If FALSE, then the procedure raises an error if there are any XStream clients that use the queue. Before you run this procedure with the cascade parameter set to FALSE, make sure no XStream clients are using the queue currently.
drop_unused_queue_table	If TRUE and the queue table for the queue is empty, then the procedure drops the queue table. The queue table is not dropped if it contains any messages or if it is used by another queue. If FALSE, then the procedure does not drop the queue table.

239.4.29 REMOVE_RULE Procedure

This procedure removes the specified rule or all rules from the rule set associated with the specified capture process, apply process, or propagation.

> **Note:**
>
> If a rule was automatically created by the system, and you want to drop the rule, then you should use this procedure to remove the rule instead of the DBMS_RULE_ADM.DROP_RULE procedure. If you use the DBMS_RULE_ADM.DROP_RULE procedure, then some metadata about the rule might remain.

Syntax

```
DBMS_XSTREAM_ADM.REMOVE_RULE(
   rule_name         IN  VARCHAR2,
   streams_type      IN  VARCHAR2,
   streams_name      IN  VARCHAR2,
   drop_unused_rule  IN  BOOLEAN  DEFAULT TRUE,
   inclusion_rule    IN  BOOLEAN  DEFAULT TRUE);
```

Parameters

Table 239-27 REMOVE_RULE Procedure Parameters

Parameter	Description
rule_name	The name of the rule to remove, specified as [*schema_name*.]*rule_name*. If NULL, then the procedure removes all rules from the specified capture process, apply process, or propagation rule set. For example, to specify a rule in the hr schema named prop_rule1, enter hr.prop_rule1. If the schema is not specified, then the current user is the default.
streams_type	The type of XStream client: • Specify capture for a capture process. • Specify propagation for a propagation. • Specify apply for an apply process.
streams_name	The name of the XStream client, which can be a capture process, propagation, or apply process. Do not specify an owner. If the specified XStream client does not exist, but there is metadata in the data dictionary that associates the rule with this client, then the procedure removes the metadata. If the specified XStream client does not exist, and there is no metadata in the data dictionary that associates the rule with this client, then the procedure raises an error.
drop_unused_rule	If TRUE and the rule is not in any rule set, then the procedure drops the rule from the database. If TRUE and the rule exists in any rule set, then the procedure does not drop the rule from the database. If FALSE, then the procedure does not drop the rule from the database.
inclusion_rule	If inclusion_rule is TRUE, then the procedure removes the rule from the positive rule set for the XStream client. If inclusion_rule is FALSE, then the procedure removes the rule from the negative rule set for the XStream client.

239.4.30 REMOVE_SUBSET_OUTBOUND_RULES Procedure

This procedure removes subset rules from an outbound server configuration.

The names of the specified insert, update, and delete rules must match those generated by the ADD_SUBSET_OUTBOUND_RULES procedure. To view the rule names for subset rules, run the following query:

```
SELECT RULE_OWNER, SUBSETTING_OPERATION, RULE_NAME
   FROM ALL_XSTREAM_RULES
   WHERE SUBSETTING_OPERATION IS NOT NULL;
```

Note:

- This procedure removes the declarative rule-based transformation associated with each rule it removes.
- This procedure does not remove rules from the outbound server's capture process.

See Also:

"ADD_SUBSET_OUTBOUND_RULES Procedure"

Syntax

```
DBMS_XSTREAM_ADM.REMOVE_SUBSET_OUTBOUND_RULES(
   server_name      IN VARCHAR2,
   insert_rule_name IN VARCHAR2,
   update_rule_name IN VARCHAR2,
   delete_rule_name IN VARCHAR2);
```

239.4.31 REMOVE_XSTREAM_CONFIGURATION Procedure

This procedure removes the XStream configuration at the local database.

Syntax

```
DBMS_XSTREAM_ADM.REMOVE_XSTREAM_CONFIGURATION(
   container IN VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-28 REMOVE_XSTREAM_CONFIGURATION Procedure Parameters

Parameter	Description
container	If CURRENT, then the XStream configuration is removed from the current container. CURRENT can be specified while connected to the root or to a PDB in a CDB. If ALL, then the XStream configuration is removed from all of the containers in the CDB. To specify ALL, the procedure must be invoked in the root. If a container name, then the XStream configuration is removed from the specified container. To specify root, use CDB$ROOT while connected to the root. To specify a PDB, the procedure must be invoked in the root. **Note:** This parameter only applies to a CDB.

Usage Notes

Specifically, this procedure performs the following actions at the local database:

- Drops all capture processes
- If any tables have been prepared for instantiation, then aborts preparation for instantiation for the table using the ABORT_TABLE_INSTANTIATION procedure in the DBMS_CAPTURE_ADM package

- If any schemas have been prepared for instantiation, then aborts preparation for instantiation for the schema using the ABORT_SCHEMA_INSTANTIATION procedure in the DBMS_CAPTURE_ADM package
- If the database has been prepared for instantiation, then aborts preparation for instantiation for the database using the ABORT_GLOBAL_INSTANTIATION procedure in the DBMS_CAPTURE_ADM package
- Drops propagations that were created using either the DBMS_XSTREAM_ADM package or the DBMS_PROPAGATION_ADM package. Before a propagation is dropped, its propagation job is disabled. Does not drop propagations that were created using the DBMS_AQADM package.
- Disables all propagation jobs used by propagations
- Drops all apply processes. If there are apply errors in the error queue for an apply process, then this procedure deletes these apply errors before it drops the apply process.
- Removes specifications for DDL handlers used by apply processes, but does not delete the PL/SQL procedures used by these handlers
- Removes specifications for message handlers used by apply processes, but does not delete the PL/SQL procedures used by these handlers
- Removes specifications for precommit handlers used by apply processes, but does not delete the PL/SQL procedures used by these handlers
- Removes the instantiation SCN and ignore SCN for each apply object and schema and for the entire database
- Removes messaging clients
- Unsets message notification specifications that were set using the SET_MESSAGE_NOTIFICATION procedure in the DBMS_XSTREAM_ADM package
- Removes specifications for procedure DML handlers and error handlers, but does not delete the PL/SQL procedures used by these handlers
- Removes update conflict handlers
- Removes specifications for substitute key columns for apply tables
- Drops rule sets and rules that were created using the DBMS_XSTREAM_ADM package.
- Drops unused rule sets that were used by capture processes, propagations, apply processes, and messaging clients, and removes the rules in these rule sets. These rules and rule sets are removed regardless of whether they were created using the DBMS_XSTREAM_ADM package or the DBMS_RULE_ADM package.

This procedure stops capture processes and apply processes before it drops them.

This procedure does not drop rule sets or rules if they meet both of the following conditions:

- The rule sets or rules were created using the DBMS_RULE_ADM package.
- The rule sets or rules were not used by a capture process, propagation, apply process, or messaging client.

> **Note:**
>
> Running this procedure is dangerous. You should run this procedure only if you are sure you want to remove the entire XStream configuration at a database. If an Oracle Replication configuration exists at the database, then this procedure also removes the entire Oracle Replication configuration.

> **Note:**
>
> - Running this procedure repeatedly does not cause errors. If the procedure fails to complete, then you can run it again.
> - This procedure commits multiple times.

239.4.32 RENAME_COLUMN Procedure

This procedure either adds or removes a declarative rule-based transformation which renames a column in a row logical change record (LCR) that satisfies the specified rule.

For the transformation to be performed when the specified rule evaluates to `TRUE`, the rule must be in the positive rule set of an XStream client. XStream clients include capture processes, propagations, and apply processes.

> **Note:**
>
> - The `RENAME_COLUMN` procedure supports the same data types supported by Oracle Replication capture processes.
> - Declarative transformations can transform row LCRs only. Therefore, a DML rule must be specified when you run this procedure. If a DDL rule is specified, then the procedure raises an error.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about declarative rule-based transformations and about the data types supported by Oracle Replication capture processes

Syntax

```
DBMS_XSTREAM_ADM.RENAME_COLUMN(
   rule_name         IN  VARCHAR2,
   table_name        IN  VARCHAR2,
   from_column_name  IN  VARCHAR2,
   to_column_name    IN  VARCHAR2,
   value_type        IN  VARCHAR2  DEFAULT '*',
```

```
  step_number       IN  NUMBER    DEFAULT 0,
  operation         IN  VARCHAR2  DEFAULT 'ADD');
```

Parameters

Table 239-29 RENAME_COLUMN Procedure Parameters

Parameter	Description
rule_name	The name of the rule, specified as [*schema_name.*]*rule_name*. If NULL, then the procedure raises an error. For example, to specify a rule in the hr schema named employees12, enter hr.employees12. If the schema is not specified, then the current user is the default.
table_name	The name of the table in which the column is renamed in the row LCR, specified as [*schema_name.*]*object_name*. For example, hr.employees. If the schema is not specified, then the current user is the default.
from_column_name	The name of the column to be renamed in each row LCR that satisfies the rule.
to_column_name	The new name of the column in each row LCR that satisfies the rule.
value_type	Specify 'NEW' to rename the column in the new values in the row LCR. Specify 'OLD' to rename the column in the old values in the row LCR. Specify '*' to rename the column in both the old and new values in the row LCR.
step_number	The order of execution of the transformation. **See Also:** *Oracle Database XStream Guide* for more information about transformation ordering
operation	Specify 'ADD' to add the transformation to the rule. Specify 'REMOVE' to remove the transformation from the rule.

Usage Notes

When 'REMOVE' is specified for the operation parameter, all of the rename column declarative rule-based transformations for the specified rule are removed that match the specified table_name, column_name, and step_number parameters. Nulls specified for these parameters act as wildcards. The following table lists the behavior of the RENAME_COLUMN procedure when one or more of these parameters is NULL:

table_name	from_column_name	to_column_name	step_number	Result
NULL	NULL	NULL	NULL	Remove all rename column transformations for the specified rule.
NULL	NULL	NULL	non-NULL	Remove all rename column transformations with the specified step_number for the specified rule.
NULL	NULL	non-NULL	non-NULL	Remove all rename column transformations with the specified to_column_name and step_number for the specified rule.

table_name	from_column_name	to_column_name	step_number	Result
NULL	non-NULL	non-NULL	non-NULL	Remove all rename column transformations with the specified table_name and step_number for the specified rule.
NULL	NULL	non-NULL	NULL	Remove all rename column transformations with the specified column_name for the specified rule.
non-NULL	NULL	non-NULL	NULL	Remove all rename column transformations with the specified table_name and column_name for the specified rule.
NULL	non-NULL	NULL	NULL	Remove all rename column transformations with the specified table_name for the specified rule.
NULL	non-NULL	non-NULL	NULL	Remove all rename column transformations with the specified table_name, column_name, and step_number for the specified rule.
non-NULL	NULL	non-NULL	NULL	Remove all rename column transformations with the specified table_name, column_name, and step_number for the specified rule.
non-NULL	non-NULL	non-NULL	NULL	Remove all rename column transformations with the specified table_name, column_name, and step_number for the specified rule.
non-NULL	non-NULL	non-NULL	non-NULL	Remove all rename column transformations with the specified table_name, column_name, and step_number for the specified rule.

239.4.33 RENAME_SCHEMA Procedure

This procedure either adds or removes a declarative rule-based transformation which renames a schema in a row logical change record (LCR) that satisfies the specified rule.

For the transformation to be performed when the specified rule evaluates to TRUE, the rule must be in the positive rule set of an XStream client. XStream clients include capture processes, propagations, and apply processes.

> **Note:**
>
> Declarative transformations can transform row LCRs only. Therefore, a DML rule must be specified when you run this procedure. If a DDL rule is specified, then the procedure raises an error.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about declarative rule-based transformations

Syntax

```
DBMS_XSTREAM_ADM.RENAME_SCHEMA(
   rule_name         IN  VARCHAR2,
   from_schema_name  IN  VARCHAR2,
   to_schema_name    IN  VARCHAR2,
   step_number       IN  NUMBER    DEFAULT 0,
   operation         IN  VARCHAR2  DEFAULT 'ADD');
```

Parameters

Table 239-30 RENAME_SCHEMA Procedure Parameters

Parameter	Description
`rule_name`	The name of the rule, specified as `[schema_name.]rule_name`. If `NULL`, then the procedure raises an error. For example, to specify a rule in the `hr` schema named `employees12`, enter `hr.employees12`. If the schema is not specified, then the current user is the default.
`from_schema_name`	The name of the schema to be renamed in each row LCR that satisfies the rule.
`to_schema_name`	The new name of the schema in each row LCR that satisfies the rule.
`step_number`	The order of execution of the transformation. **See Also:** *Oracle Database XStream Guide* for more information about transformation ordering
`operation`	Specify `'ADD'` to add the transformation to the rule. Specify `'REMOVE'` to remove the transformation from the rule.

Usage Notes

When `'REMOVE'` is specified for the `operation` parameter, all of the rename schema declarative rule-based transformations for the specified rule are removed that match the specified `from_schema_name`, `to_schema_name`, and `step_number` parameters. Nulls specified for these parameters act as wildcards. The following table lists the behavior of the `RENAME_SCHEMA` procedure when one or more of these parameters is `NULL`:

from_schema_name	to_schema_name	step_number	Result
`NULL`	`NULL`	`NULL`	Remove all rename schema transformations for the specified rule.
`NULL`	`NULL`	non-`NULL`	Remove all rename schema transformations with the specified `step_number` for the specified rule.

from_schema_name	to_schema_name	step_number	Result
NULL	non-NULL	non-NULL	Remove all rename schema transformations with the specified to_schema_name and step_number for the specified rule.
non-NULL	NULL	non-NULL	Remove all rename schema transformations with the specified from_schema_name and step_number for the specified rule.
NULL	non-NULL	NULL	Remove all rename schema transformations with the specified to_schema_name for the specified rule.
non-NULL	non-NULL	NULL	Remove all rename schema transformations with the specified from_schema_name and to_schema_name for the specified rule.
non-NULL	NULL	NULL	Remove all rename schema transformations with the specified from_schema_name for the specified rule.
non-NULL	non-NULL	non-NULL	Remove all rename schema transformations with the specified from_schema_name, to_schema_name, and step_number for the specified rule.

239.4.34 RENAME_TABLE Procedure

This procedure either adds or removes a declarative rule-based transformation which renames a table in a row logical change record (row LCR) that satisfies the specified rule.

For the transformation to be performed when the specified rule evaluates to TRUE, the rule must be in the positive rule set of an XStream client. XStream clients include capture processes, propagations, and apply processes.

> **Note:**
>
> Declarative transformations can transform row LCRs only. Therefore, a DML rule must be specified when you run this procedure. If a DDL rule is specified, then the procedure raises an error.

> **See Also:**
>
> *Oracle Database XStream Guide* for more information about declarative rule-based transformations

Syntax

```
DBMS_XSTREAM_ADM.RENAME_TABLE(
   rule_name         IN  VARCHAR2,
   from_table_name   IN  VARCHAR2,
   to_table_name     IN  VARCHAR2,
   step_number       IN  NUMBER    DEFAULT 0,
   operation         IN  VARCHAR2  DEFAULT 'ADD');
```

Parameters

Table 239-31 RENAME_TABLE Procedure Parameters

Parameter	Description
`rule_name`	The name of the rule, specified as `[schema_name.]rule_name`. If `NULL`, then the procedure raises an error. For example, to specify a rule in the `hr` schema named `employees12`, enter `hr.employees12`. If the schema is not specified, then the current user is the default.
`from_table_name`	The name of the table to be renamed in each row LCR that satisfies the rule, specified as `[schema_name.]object_name`. For example, `hr.employees`. If the schema is not specified, then the current user is the default.
`to_table_name`	The new name of the table in each row LCR that satisfies the rule, specified as `[schema_name.]object_name`. For example, `humres.staff`. The transformation can rename the table only, the schema only, or the table and the schema. If the schema is not specified, then the current user is the default.
`step_number`	The order of execution of the transformation. **See Also:** *Oracle Database XStream Guide* for more information about transformation ordering
`operation`	Specify `'ADD'` to add the transformation to the rule. Specify `'REMOVE'` to remove the transformation from the rule.

Usage Notes

When `'REMOVE'` is specified for the `operation` parameter, all of the rename table declarative rule-based transformations for the specified rule are removed that match the specified `from_table_name`, `to_table_name`, and `step_number` parameters. Nulls specified for these parameters act as wildcards. The following table lists the behavior of the `RENAME_TABLE` procedure when one or more of these parameters is `NULL`:

from_table_name	to_table_name	step_number	Result
`NULL`	`NULL`	`NULL`	Remove all rename table transformations for the specified rule.

from_table_name	to_table_name	step_number	Result
NULL	NULL	non-NULL	Remove all rename table transformations with the specified step_number for the specified rule.
NULL	non-NULL	non-NULL	Remove all rename table transformations with the specified to_table_name and step_number for the specified rule.
non-NULL	NULL	non-NULL	Remove all rename table transformations with the specified from_table_name and step_number for the specified rule.
NULL	non-NULL	NULL	Remove all rename table transformations with the specified to_table_name for the specified rule.
non-NULL	non-NULL	NULL	Remove all rename table transformations with the specified from_table_name and to_table_name for the specified rule.
non-NULL	NULL	NULL	Remove all rename table transformations with the specified from_table_name for the specified rule.
non-NULL	non-NULL	non-NULL	Remove all rename table transformations with the specified from_table_name, to_table_name, and step_number for the specified rule.

239.4.35 SET_MESSAGE_TRACKING Procedure

This procedure sets the tracking label for logical change records (LCRs) produced by the current session.

This procedure affects only the current session. Any LCRs produced by the current session are tracked, including captured LCRs and persistent LCRs.

> **Note:**
>
> The tracking label set by this procedure does not track non-LCR messages.

> **See Also:**
>
> KAWGET_MESSAGE_TRACKING

Syntax

```
DBMS_XSTREAM_ADM.SET_MESSAGE_TRACKING(
   tracking_label IN VARCHAR2  DEFAULT 'Streams_tracking',
   actions        IN NUMBER    DEFAULT DBMS_XSTREAM_ADM.ACTION_MEMORY);
```

Parameters

Table 239-32 SET_MESSAGE_TRACKING Procedure Parameters

Parameter	Description
`tracking_label`	The label used to track the LCRs produced by the session. Set this parameter to `NULL` to stop message tracking in the current session. The size limit for a label is 4,000 bytes.
`actions`	When `DBMS_XSTREAM_ADM.ACTION_MEMORY` is specified, the LCRs are tracked in memory. Currently, `DBMS_XSTREAM_ADM.ACTION_MEMORY` is the only valid setting for this parameter. The value specified for this parameter is an enumerated constant. Enumerated constants must be prefixed with the package name.

239.4.36 SET_PARAMETER Procedure

This procedure sets a parameter for an outbound server, an inbound server, or an outbound server's capture process.

Syntax

```
DBMS_XSTREAM_ADM.SET_PARAMETER(
   streams_name    IN  VARCHAR2,
   streams_type    IN  VARCHAR2,
   parameter       IN  VARCHAR2,
   value           IN  VARCHAR2  DEFAULT NULL,
   no_wait         IN  BOOLEAN   DEFAULT FALSE,
   source_database IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-33 SET_PARAMETER Procedure Parameters

Parameter	Description
`streams_type`	The type of XStream client: • Specify `capture` for a capture process. • Specify `apply` for an outbound server or inbound server.
`streams_name`	The name of the capture process, outbound server, or inbound server. Do not specify an owner.

Table 239-33 (Cont.) SET_PARAMETER Procedure Parameters

Parameter	Description
parameter	The name of the parameter you are setting. See "Capture Process Parameters" for information about capture process parameters. See "Apply Component Parameters" for information about outbound server and inbound server parameters.
value	The value to which the parameter is set. If NULL, then the parameter is set to its default value.
no_wait	If TRUE, then the parameter is set immediately. If FALSE, then the parameter is set after synchronizing with the running capture process, inbound server, or outbound server. When you modify multiple parameters for the same process consecutively, setting this parameter to TRUE speeds up each call. However, if the process is currently running, you must set this parameter to FALSE in the last to the procedure to ensure that the process uses the modified parameter values. If the no_wait parameter is set to TRUE for the last call to the procedure, the running process might not detect the parameter changes.
source_database	If CURRENT, then the parameter is set only in the container where the procedure is invoked. CURRENT can be specified while connected to the root or to a PDB. If ALL, then the parameter is set in all containers in the CDB and all PDBs created after the procedure is invoked. To specify ALL, the procedure must be invoked in the root. If a container name, then the parameter is set in the specified container. To specify root, use CDB$ROOT while connected to the root. To specify a PDB, the procedure must be invoked in the root. **Note:** This parameter only applies to CDBs. Also, a non-null value can be specified only for the following parameters: • include_objects capture parameter • excludetag capture or apply parameter • excludetrans capture or apply parameter • excludeuser capture or apply parameter • excludeuserid capture or apply parameter • getreplicates capture or apply parameter • getapplops capture or apply parameter

239.4.37 SET_TAG Procedure

This procedure sets the binary tag for all redo entries subsequently generated by the current session.

Each redo entry generated by DML or DDL statements in the current session will have this tag. This procedure affects only the current session.

Syntax

```
DBMS_XSTREAM_ADM.SET_TAG(
   tag  IN RAW  DEFAULT NULL);
```

Parameters

Table 239-34 SET_TAG Procedure Parameters

Parameter	Description
tag	The binary tag for all subsequent redo entries generated by the current session. A raw value is a sequence of bytes, and a byte is a sequence of bits. By default, the tag for a session is NULL. The size limit for a tag value is 2000 bytes.

Usage Notes

To set the tag to the hexadecimal value of '17' in the current session, run the following procedure:

```
EXEC DBMS_XSTREAM_ADM.SET_TAG(tag => HEXTORAW('17'));
```

The following are considerations for the SET_TAG procedure:

- This procedure is not transactional. That is, the effects of SET_TAG cannot be rolled back.
- If the SET_TAG procedure is run to set a non-NULL session tag before a data dictionary build has been performed on the database, then the redo entries for a transaction that started before the dictionary build might not include the specified tag value for the session. Therefore, perform a data dictionary build before using the SET_TAG procedure in a session. A data dictionary build happens when the DBMS_CAPTURE_ADM.BUILD procedure is run. The BUILD procedure can be run automatically when a capture process is created.

> **See Also:**
>
> BUILD Procedure

239.4.38 SET_UP_QUEUE Procedure

This procedure creates a queue table and a ANYDATA queue.

Syntax

```
DBMS_XSTREAM_ADM.SET_UP_QUEUE(
   queue_table     IN  VARCHAR2  DEFAULT 'streams_queue_table',
   storage_clause  IN  VARCHAR2  DEFAULT NULL,
   queue_name      IN  VARCHAR2  DEFAULT 'streams_queue',
   queue_user      IN  VARCHAR2  DEFAULT NULL,
   comment         IN  VARCHAR2  DEFAULT NULL);
```

Parameters

Table 239-35 SET_UP_QUEUE Procedure Parameters

Parameter	Description
queue_table	The name of the queue table specified as [*schema_name*.]*queue_table_name*. For example, strmadmin.streams_queue_table. If the schema is not specified, then the current user is the default. If the queue table owner is not specified, then the procedure specifies the user who runs this procedure automatically as the queue table owner. Queue table names can be a maximum of 24 bytes.
storage_clause	The storage clause for queue table The storage parameter is included in the CREATE TABLE statement when the queue table is created. You can specify any valid table storage clause. If a tablespace is not specified here, then the procedure creates the queue table and all its related objects in the default user tablespace of the user who runs this procedure. If a tablespace is specified here, then the procedure creates the queue table and all its related objects in the tablespace specified in the storage clause. If NULL, then the procedure uses the storage characteristics of the tablespace in which the queue table is created. **See Also:** *Oracle Database SQL Language Reference*for more information about storage clauses
queue_name	The name of the queue that will function as the ANYDATA queue, specified as [*schema_name*.]*queue_name*. For example, strmadmin.streams_queue. If the schema is not specified, then the procedure uses the queue table owner. The owner of the queue table must also be the owner of the queue. The queue owner automatically has privileges to perform all queue operations on the queue. If the schema is not specified for this parameter, and the queue table owner is not specified in queue_table, then the current user is the default. Queue names can be a maximum of 24 bytes.
queue_user	The name of the user who requires ENQUEUE and DEQUEUE privileges for the queue. This user also is configured as a secure queue user of the queue. The queue user cannot grant these privileges to other users because they are not granted with the GRANT option. If NULL, then the procedure does not grant any privileges. You can also grant queue privileges to the appropriate users using the DBMS_AQADM package.
comment	The comment for the queue

Usage Notes

Set up includes the following actions:

- If the specified queue table does not exist, then this procedure runs the CREATE_QUEUE_TABLE procedure in the DBMS_AQADM package to create the queue table with the specified storage clause. If this procedure creates the queue table, then it creates a multiple consumer ANYDATA queue that is both a secure queue and a transactional queue.

Also, if the database is Oracle Database 10*g* release 2 or later, the `sort_list` setting in `CREATE_QUEUE_TABLE` is set to `commit_time`. If the database is a release before Oracle Database 10*g* release 2, the `sort_list` setting in `CREATE_QUEUE_TABLE` is set to `enq_time`.

- If the specified queue table exists, then the queue uses the properties of the existing queue table.
- If the specified queue name does not exist, then this procedure runs the `CREATE_QUEUE` procedure in the `DBMS_AQADM` package to create the queue.
- This procedure starts the queue.
- If a queue user is specified, then this procedure configures this user as a secure queue user of the queue and grants `ENQUEUE` and `DEQUEUE` privileges on the queue to the specified queue user.

  To configure the queue user as a secure queue user, this procedure creates an Advanced Queuing agent with the same name as the user name, if one does not exist. If an agent with this name exists and is associated with the queue user only, then it is used. `SET_UP_QUEUE` then runs the `ENABLE_DB_ACCESS` procedure in the `DBMS_AQADM` package, specifying the agent and the user.

> **Note:**
>
> If the agent that `SET_UP_QUEUE` tries to create exists and is associated with a user other than the user specified by `queue_user`, then the procedure raises an error. In this case, rename or remove the existing agent, and retry `SET_UP_QUEUE`.

239.4.39 SPLIT_STREAMS Procedure

This procedure splits one stream flowing from a capture process off from all of the other streams flowing from the capture process.

This procedure is intended for an Oracle Replication environment in which a capture process captures changes that are propagated to two or more destination databases. When one destination of a propagation stops accepting the captured changes, the changes remain in the capture process's queue. The queue can grow and begin to spill LCRs to the hard disk, degrading the performance of the Oracle Replication environment. A destination might stop accepting changes for several reasons. For example, the destination database might be down.

Specifically, this procedure performs the following actions:

1. Creates a new queue at the database running the capture process. The new queue is called the cloned queue because it is a clone of the queue used by the original stream. The new queue will be used by the new, cloned capture process, and it will be the source queue for the new, cloned propagation.
2. Creates a new propagation that propagates LCRs from the source queue created in Step 1 to the existing destination queue. The new propagation is called the cloned propagation because it is a clone of the propagation used by the original stream. The cloned propagation uses the same rule set as the original propagation.
3. Stops the capture process.
4. Queries the acknowledge SCN for the original propagation. The acknowledged SCN is the last SCN acknowledged by the apply process that applies the changes sent by the propagation.

5. Creates a new capture process. The new capture process is called the cloned capture process because it is a clone of the capture process used by the original stream. The procedure sets the start SCN for the cloned capture process to the value of the queried acknowledged SCN. The cloned capture process uses the same rule set as the original capture process.
6. Drops the original propagation.
7. Starts the original capture process with the start SCN set to the acknowledged SCN queried in Step 4.
8. If the `auto_merge_threshold` parameter is set to a positive number, then creates an Oracle Scheduler job to run the `MERGE_STREAMS_JOB` procedure at set intervals according to its schedule. When the two streams are within the specified merge threshold, the `MERGE_STREAMS_JOB` procedure runs the `MERGE_STREAMS` procedure to merge the streams automatically.

After the `SPLIT_STREAMS` procedure has finished running, the cloned capture process is disabled. When the problem at the destination database is solved, and the destination queue can accept changes, you should start the cloned capture process using the `START_CAPTURE` procedure in the `DBMS_CAPTURE_ADM` package.

Note:

If the original capture process is a downstream capture process, then you must configure the cloned capture process to read the redo log from the source database before you start the cloned capture process.

See Also:

- "MERGE_STREAMS Procedure"
- "MERGE_STREAMS_JOB Procedure"

Syntax

```
DBMS_XSTREAM_ADM.SPLIT_STREAMS(
   propagation_name          IN      VARCHAR2,
   cloned_propagation_name   IN      VARCHAR2  DEFAULT NULL,
   cloned_queue_name         IN      VARCHAR2  DEFAULT NULL,
   cloned_capture_name       IN      VARCHAR2  DEFAULT NULL,
   perform_actions           IN      BOOLEAN   DEFAULT TRUE,
   script_name               IN      VARCHAR2  DEFAULT NULL,
   script_directory_object   IN      VARCHAR2  DEFAULT NULL,
   auto_merge_threshold      IN      NUMBER    DEFAULT NULL,
   schedule_name             IN OUT  VARCHAR2,
   merge_job_name            IN OUT  VARCHAR2);
```

Parameters

Table 239-36 SPLIT_STREAMS Procedure Parameters

Parameter	Description
propagation_name	The name of the propagation that cannot send LCRs to its destination queue. The specified propagation is the propagation for the stream that is being split off from the other streams. You must specify an existing propagation name. Do not specify an owner.
cloned_propagation_name	The name of the new propagation created by this procedure for the stream that is split off. If NULL, then the system generates a propagation name.
cloned_queue_name	The name of the new queue created by this procedure for the stream that is split off. If NULL, then the system generates a queue name.
cloned_capture_name	The name of the new capture process created by this procedure for the stream that is split off. If NULL, then the system generates a capture process name.
perform_actions	If TRUE, then the procedure performs the necessary actions to split the stream directly. If FALSE, then the procedure does not perform the necessary actions to split the stream directly. Specify FALSE when this procedure is generating a script that you can edit and then run. The procedure raises an error if you specify FALSE and either of the following parameters is NULL: • script_name • script_directory_object
script_name	If non-NULL and the perform_actions parameter is FALSE, then specify the name of the script generated by this procedure. The script contains all of the statements used to split the stream. If a file with the specified script name exists in the specified directory for the script_directory_object parameter, then the procedure appends the statements to the existing file. If non-NULL and the perform_actions parameter is TRUE, then the procedure generates the specified script and performs the actions to split the stream directly. If NULL and the perform_actions parameter is TRUE, then the procedure performs the actions to split the stream directly and does not generate a script. If NULL and the perform_actions parameter is FALSE, then the procedure raises an error.
script_directory_object	The directory object for the directory on the local computer system into which the generated script is placed. If the script_name parameter is NULL, then the procedure ignores this parameter and does not generate a script. If NULL and the script_name parameter is non-NULL, then the procedure raises an error. **Note:** The specified directory object cannot point to an Oracle Automatic Storage Management (ASM) disk group.

Table 239-36 (Cont.) SPLIT_STREAMS Procedure Parameters

Parameter	Description
auto_merge_threshold	If a positive number is specified, then the stream that was split off is automatically merged back into all of the other streams flowing from the capture process by an Oracle Scheduler job. The job runs the MERGE_STREAMS_JOB procedure at set intervals according to its schedule. The value of the CAPTURE_MESSAGE_CREATE_TIME column for each capture process in the GV$XSTREAM_CAPTURE dynamic performance view determines when the streams are merged. Specifically, if the difference, in seconds, between CAPTURE_MESSAGE_CREATE_TIME of the cloned capture process and the original capture process is less than or equal to the value specified for the auto_merge_threshold parameter, then the two streams are merged automatically. The cloned capture process must be started before the split stream can be merged back with the original stream. If NULL or 0 (zero) is specified, then the split stream is not merged back with the original stream automatically. To merge the split stream with the original stream, run the MERGE_STREAM procedure manually when the CAPTURE_MESSAGE_CREATE_TIME of the cloned capture process catches up to, or nearly catches up to, the CAPTURE_MESSAGE_CREATE_TIME of the original capture process. The CAPTURE_MESSAGE_CREATE_TIME records the time when a captured change was recorded in the redo log.
schedule_name	The Oracle Scheduler schedule name, specified as [*schema_name.*]*schedule_name*. For example, strmadmin.merge_schedule. If the schema is not specified, then the current user is the default. If auto_merge_threshold is a non-NULL positive number, then the schedule is used by the job that will automatically merge the streams at the appropriate time. You can specify a schedule name to adhere to naming conventions or to track the schedule more easily. If NULL and auto_merge_threshold is a non-NULL positive number, then the system generates a schedule name. If auto_merge_threshold is NULL or 0 (zero), then this parameter must be NULL. If this procedure creates a schedule, the schedule starts when the procedure completes. You can modify the schedule to control how often the merge job is run. If an existing schedule name is specified, an error is raised.

Table 239-36 (Cont.) SPLIT_STREAMS Procedure Parameters

Parameter	Description
merge_job_name	The Oracle Scheduler job name, specified as [*schema_name.*]*merge_job_name*. For example, strmadmin.merge_job. If the schema is not specified, then the current user is the default. If auto_merge_threshold is a non-NULL positive number, then the job will automatically merge the streams at the appropriate time. Specify a merge job name to adhere to naming conventions or to track the job more easily. If NULL and auto_merge_threshold is a non-NULL positive number, then the system generates a job name. If auto_merge_threshold is NULL or 0 (zero), then this parameter must be NULL. If an existing job name is specified, an error is raised.

See Also:

Oracle Database Administrator's Guide for information about Oracle Scheduler

239.4.40 START_OUTBOUND Procedure

This procedure starts an XStream outbound server. The outbound server streams out the LCRs to an XStream client application.

Syntax

```
DBMS_XSTREAM_ADM.START_OUTBOUND(
   server_name      IN  VARCHAR2);
```

Parameters

Table 239-37 START_OUTBOUND Procedure Parameters

Parameter	Description
server_name	The name of the outbound server being started. A NULL specification is not allowed. Do not specify an owner.

239.4.41 STOP_OUTBOUND Procedure

This procedure stops an XStream outbound server. The outbound server streams out the LCRs to an XStream client application.

Syntax

```
DBMS_XSTREAM_ADM.STOP_OUTBOUND(
   server_name  IN  VARCHAR2,
   force        IN  BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 239-38 STOP_OUTBOUND Procedure Parameters

Parameter	Description
`server_name`	The name of the outbound server being stopped. A `NULL` specification is not allowed. Do not specify an owner.
`force`	If `TRUE`, then the procedure stops the outbound server and its capture process as soon as possible. If `FALSE`, then the procedure stops the outbound server after ensuring that there are no gaps in the set of applied transactions. The behavior of the apply component depends on the setting specified for the `force` parameter and the setting specified for the `commit_serialization` apply component parameter.

240
XSTREAM_CAPTURE

The `XSTREAM_CAPTURE` package provides subprograms for granting privileges to and revoking privileges from XStream administrators.

This chapter contains the following topic:

- Overview
- Security Model

240.1 XSTREAM_CAPTURE Overview

This package provides subprograms for granting privileges to XStream administrators and revoking privileges from XStream administrators.

240.2 XSTREAM_CAPTURE Security Model

Security on this package can be controlled by either granting `EXECUTE` on this package to selected users or roles, or by granting `EXECUTE_CATALOG_ROLE` to selected users or roles.

The user executing the subprograms in the `XSTREAM_CAPTURE` package must have `SYSDBA` administrative privilege, and the user must exercise the privilege using `AS SYSDBA` at connect time.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

To ensure that the user who runs the subprograms in this package has the necessary privileges, connect as an administrative user who can create users, grant privileges, and create tablespaces when using this package.

240.3 Summary of XSTREAM_CAPTURE Subprograms

This table lists the `XSTREAM_CAPTURE` subprograms and briefly describes them.

Table 240-1 XSTREAM_CAPTURE Package Subprograms

Subprogram	Description
GRANT_ADMIN_PRIVILEGE Procedure	Either grants the privileges needed by a user to be an XStream administrator directly, or generates a script that grants these privileges
GRANT_REMOTE_ADMIN_ACCESS Procedure	Enables a remote XStream administrator to perform administrative actions at the local database by connecting to the grantee using a database link
REVOKE_ADMIN_PRIVILEGE Procedure	Either revokes XStream administrator privileges from a user directly, or generates a script that revokes these privileges

Table 240-1 (Cont.) XSTREAM_CAPTURE Package Subprograms

Subprogram	Description
REVOKE_REMOTE_ADMIN_ACCESS Procedure	Disables a remote XStream administrator from performing administrative actions by connecting to the grantee using a database link

> **Note:**
>
> All subprograms commit unless specified otherwise.

240.3.1 XSTREAM_APPLY Procedure

This procedure either grants the privileges needed by a user to be an XStream administrator directly, or generates a script that grants these privileges.

> **See Also:**
>
> "XSTREAM_APPLY Procedure"

Syntax

```
XSTREAM_CAPTURE.XSTREAM_APPLY(
   grantee                    IN  VARCHAR2,
   privilege_type             IN  VARCHAR2  DEFAULT '*',
   grant_select_privileges    IN  BOOLEAN   DEFAULT FALSE,
   do_grants                  IN  BOOLEAN   DEFAULT TRUE,
   file_name                  IN  VARCHAR2  DEFAULT NULL,
   directory_name             IN  VARCHAR2  DEFAULT NULL
   grant_optional_privileges  IN  VARCHAR2  DEFAULT NULL,
   container                  IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameters

Table 240-2 XSTREAM_APPLY Procedure Parameters

Parameter	Description
grantee	The user to whom privileges are granted

Table 240-2 (Cont.) XSTREAM_APPLY Procedure Parameters

Parameter	Description
privilege_type	Specify one of the following values: • CAPTURE Specifying CAPTURE grants the minimum privileges required by the user to administer capture processes. • APPLY Specifying APPLY grants the minimum privileges required by the user to administer outbound servers, inbound servers, and apply processes. • * Specifying * grants the minimum privileges required by the user to administer capture processes, outbound servers, inbound servers, and apply processes.
grant_select_privileges	If TRUE, then the procedure grants a set of privileges, including SELECT_CATALOG_ROLE, to the user. If FALSE, then the procedure does not grant the set of privileges to the user. SELECT_CATALOG_ROLE enables the user to select from the data dictionary. Set this parameter to TRUE for the XStream trusted user model. Set this parameter to FALSE for the XStream untrusted user model.
do_grants	If TRUE, then the procedure grants the privileges to the specified grantee directly, and adds the grantee to the DBA_XSTREAM_ADMINISTRATOR data dictionary view with YES for both the LOCAL_PRIVILEGES column and the ACCESS_FROM_REMOTE column. If the user already has an entry in this data dictionary view, then the procedure does not make another entry, and no error is raised. If TRUE and any of the grant statements fails, then the procedure raises an error. If FALSE, then the procedure does not grant the privileges to the specified grantee directly, and does not add the grantee to the DBA_XSTREAM_ADMINISTRATOR data dictionary view. You specify FALSE when the procedure is generating a file that you will run later. If you specify FALSE and either the file_name or directory_name parameter is NULL, then the procedure raises an error.
file_name	The name of the file generated by the procedure. The file contains all of the statements that grant the privileges. If a file with the specified file name exists in the specified directory name, then the grant statements are appended to the existing file. If NULL, then the procedure does not generate a file.
directory_name	The directory into which the generated file is placed. The specified directory must be a directory object created using the SQL statement CREATE DIRECTORY. If you specify a directory, then the user who invokes the procedure must have the WRITE privilege on the directory object. If the file_name parameter is NULL, then this parameter is ignored, and the procedure does not generate a file. If NULL and the file_name parameter is non-NULL, then the procedure raises an error.
grant_optional_privileges	A comma-separated list of optional privileges to grant to the grantee, such as the DV_XSTREAM_ADMIN and DV_GOLDENGATE_ADMIN privileges

Table 240-2 (Cont.) XSTREAM_APPLY Procedure Parameters

Parameter	Description
container	If CURRENT, then grants privileges to the grantee only in the container where the procedure is invoked. CURRENT can be specified while connected to the root or to a PDB. If ALL, then grants privileges to the grantee in all containers in the CDB and all PDBs created after the procedure is invoked. To specify ALL, the procedure must be invoked in the root. If a container name, then grants privileges to the grantee only in the specified container. To specify root, use CDB$ROOT while connected to the root. To specify a PDB, the procedure must be invoked in the root. **Note:** This parameter only applies to CDBs.

Usage Notes

The user who runs the procedure must be an administrative user who can grant privileges to other users.

Specifically, the procedure grants the following privileges to the specified user:

- The RESTRICTED SESSION system privilege
- EXECUTE on the following packages:
  - DBMS_APPLY_ADM
  - DBMS_AQ
  - DBMS_AQADM
  - DBMS_AQIN
  - DBMS_AQELM
  - DBMS_CAPTURE_ADM
  - DBMS_FLASHBACK
  - DBMS_LOCK
  - DBMS_PROPAGATION_ADM
  - DBMS_RULE_ADM
  - DBMS_TRANSFORM
  - DBMS_XSTREAM_ADM
- Privileges to enqueue messages into and dequeue messages from any queue
- Privileges to manage any queue
- Privileges to create, alter, and execute any of the following types of objects in the user's own schema and in other schemas:
  - Evaluation contexts
  - Rule sets
  - Rules

  In addition, the grantee can grant these privileges to other users.

- SELECT_CATALOG_ROLE
- SELECT or READ privilege on data dictionary views related to XStream and Oracle Replication
- The ability to allow a remote XStream administrator to perform administrative actions through a database link by connecting to the grantee

  This ability is enabled by running the GRANT_REMOTE_ADMIN_ACCESS procedure in this package.

> **Note:**
>
> – To view all of the statements run by the procedure in detail, you can use the procedure to generate a script and then view the script in a text editor.
>
> – This procedure grants only the privileges necessary to configure and administer an XStream environment. You can grant additional privileges to the grantee if necessary.

> **See Also:**
>
> – "GRANT_REMOTE_ADMIN_ACCESS Procedure"
>
> – *Oracle Database SQL Language Reference* for information about the CREATE DIRECTORY SQL statement

240.3.2 GRANT_REMOTE_ADMIN_ACCESS Procedure

This procedure enables a remote XStream administrator to perform administrative actions at the local database by connecting to the grantee using a database link.

Syntax

```
DBMS_XSTREAM_AUTH.GRANT_REMOTE_ADMIN_ACCESS(
   grantee  IN  VARCHAR2);
```

Parameters

Table 240-3 GRANT_REMOTE_ADMIN_ACCESS Procedure Parameter

Parameter	Description
grantee	The user who allows remote access. The procedure adds the grantee to the DBA_XSTREAM_ADMINISTRATOR data dictionary view with YES for the ACCESS_FROM_REMOTE column. If the user already has an entry in this data dictionary view, then the procedure does not make another entry. Instead, it updates the ACCESS_FROM_REMOTE column to YES.

Usage Notes

Typically, you run the procedure and specify a grantee at a local source database if a downstream capture process captures changes originating at the local source database. The

XStream administrator at a downstream capture database administers the source database using this connection.

> **Note:**
>
> The `XSTREAM_APPLY` procedure in this package runs this procedure.

> **See Also:**
>
> "XSTREAM_APPLY Procedure"

240.3.3 REVOKE_ADMIN_PRIVILEGE Procedure

This procedure either revokes XStream administrator privileges from a user directly, or generates a script that revokes these privileges.

Syntax

```
XSTREAM_CAPTURE.REVOKE_ADMIN_PRIVILEGE(
   grantee                     IN  VARCHAR2,
   privilege_type              IN  VARCHAR2  DEFAULT '*',
   revoke_select_privileges    IN  BOOLEAN   DEFAULT FALSE,
   do_revokes                  IN  BOOLEAN   DEFAULT TRUE,
   file_name                   IN  VARCHAR2  DEFAULT NULL,
   directory_name              IN  VARCHAR2  DEFAULT NULL
   revoke_optional_privileges  IN  VARCHAR2  DEFAULT NULL,
   container                   IN  VARCHAR2  DEFAULT 'CURRENT');
```

Parameters

Table 240-4 *REVOKE_ADMIN_PRIVILEGE Procedure Parameters*

Parameter	Description
`grantee`	The user from whom privileges are revoked
`privilege_type`	Specify one of the following values: • `CAPTURE` Specifying `CAPTURE` revokes the minimum privileges required by the user to administer capture processes. • `APPLY` Specifying `APPLY` revokes the minimum privileges required by the user to administer outbound servers, inbound servers, and apply processes. • `*` Specifying `*` revokes the minimum privileges required by the user to administer capture processes, outbound servers, inbound servers, and apply processes.

Table 240-4 (Cont.) *REVOKE_ADMIN_PRIVILEGE Procedure Parameters*

Parameter	Description
revoke_select_privileges	If TRUE, then the procedure revokes a set of privileges, including SELECT_CATALOG_ROLE, to the user. If FALSE, then the procedure does not revoke the set of privileges to the user. SELECT_CATALOG_ROLE enables the user to select from the data dictionary.
do_revokes	If TRUE, then the procedure revokes the privileges from the specified user directly, and removes the user from the DBA_XSTREAM_ADMINISTRATOR data dictionary view. If the user does not have a record in this data dictionary view, then the procedure does not remove a record from the view, and no error is raised. If TRUE and any of the revoke statements fails, then the procedure raises an error. A revoke statement fails if the user is not granted the privilege that is being revoked. If FALSE, then the procedure does not revoke the privileges from the specified user directly, and does not remove the user from the DBA_XSTREAM_ADMINISTRATOR data dictionary view. You specify FALSE when the procedure is generating a file that you will run later. If you specify FALSE and either the file_name or directory_name parameter is NULL, then the procedure does not raise an error.
file_name	The name of the file generated by this procedure. The file contains all of the statements that revoke the privileges. If a file with the specified file name exists in the specified directory name, then the revoke statements are appended to the existing file. If NULL, then the procedure does not generate a file.
directory_name	The directory into which the generated file is placed. The specified directory must be a directory object created using the SQL statement CREATE DIRECTORY. If you specify a directory, then the user who invokes the procedure must have the WRITE privilege on the directory object. If the file_name parameter is NULL, then this parameter is ignored, and the procedure does not generate a file. If NULL and the file_name parameter is non-NULL, then the procedure raises an error.
revoke_optional_privileges	A comma-separated list of optional privileges to revoke from the grantee, such as the DV_XSTREAM_ADMIN and DV_GOLDENGATE_ADMIN privileges
container	If CURRENT, then revokes privileges from the grantee only in the container where the procedure is invoked. CURRENT can be specified while connected to the root or to a PDB. If ALL, then revokes privileges from the grantee in all containers in the CDB. To specify ALL, the procedure must be invoked in the root. If a container name, then revokes privileges from the grantee only in the specified container. To specify root, use CDB$ROOT while connected to the root. To specify a PDB, the procedure must be invoked in the root. **Note:** This parameter only applies to CDBs.

Usage Notes

The user who runs this procedure must be an administrative user who can revoke privileges from other users. Specifically, this procedure revokes the privileges granted by running the XSTREAM_APPLY procedure in this package.

> **Note:**
>
> To view all of the statements run by this procedure in detail, you can use the procedure to generate a script and then view the script in a text editor.

> **See Also:**
>
> - "XSTREAM_APPLY Procedure"
> - *Oracle Database SQL Language Reference* for information about the `CREATE DIRECTORY` SQL statement

240.3.4 REVOKE_REMOTE_ADMIN_ACCESS Procedure

This procedure disables a remote XStream administrator from performing administrative actions by connecting to the grantee using a database link.

> **Note:**
>
> The `REVOKE_ADMIN_PRIVILEGE` procedure in this package runs this procedure.

> **See Also:**
>
> "REVOKE_ADMIN_PRIVILEGE Procedure"

Syntax

```
DBMS_XSTREAM_AUTH.REVOKE_REMOTE_ADMIN_ACCESS(
   grantee  IN  VARCHAR2);
```

Parameters

Table 240-5 REVOKE_REMOTE_ADMIN_ACCESS Procedure Parameter

Parameter	Description
`grantee`	The user for whom access from a remote XStream administrator is disabled. If a row for the grantee exists in the `DBA_XSTREAM_ADMINISTRATOR` data dictionary view, then the procedure updates the `ACCESS_FROM_REMOTE` column for the grantee to `NO`. If, after this update, both the `LOCAL_PRIVILEGES` column and the `ACCESS_FROM_REMOTE` column are `NO` for the grantee, then the procedure removes the grantee from the view. If no row for the grantee exists in the `DBA_XSTREAM_ADMINISTRATOR` data dictionary view, then the procedure does not update the view and does not raise an error.

241
DEBUG_EXTPROC

The DEBUG_EXTPROC package enables you to start up the extproc agent within a session. This utility package can help you debug external procedures.

This chapter contains the following topics:

- Security Model
- Operational Notes
- Rules and Limits
- Summary of DEBUG_EXTPROC Subprograms

241.1 DEBUG_EXTPROC Security Model

Your Oracle account must have EXECUTE privileges on the package and CREATE LIBRARY privileges.

241.2 DEBUG_EXTPROC Operational Notes

These operational notes apply to DEBUG_EXTPROC.

To install the package, run the script DBGEXTP.SQL.

- Install/load this package in the Oracle USER where you want to debug the 'extproc' process.
- Ensure that you have execute privileges on package DEBUG_EXTPROC

```
SELECT SUBSTR(OBJECT_NAME, 1, 20)
FROM USER_OBJECTS
WHERE OBJECT_NAME = 'DEBUG_EXTPROC';
```

- You can install this package as any other user, as long as you have EXECUTE privileges on the package.

> **Note:**
>
> These notes assumes that you built your shared library with debug symbols to aid in the debugging process. Please check the C compiler manual pages for the appropriate C compiler switches to build the shared library with debug symbols.

Having installed the package, proceed accordingly:

- Start a new Oracle session through SQL*Plus or OCI program by connecting to ORACLE.
- Execute procedure DEBUG_EXTPROC.STARTUP_EXTPROC_AGENT to startup the extproc agent in this session; for example, execute DEBUG_EXTPROC.STARTUP_EXTPROC_AGENT; Do not exit this session, because that terminates the extproc agent.

- Determine the PID of the extproc agent that was started up for this session.
- Using a debugger (for example, gdb, dbx, or the native system debugger), load the extproc executable and attach to the running process.
- Set a breakpoint on function 'pextproc' and let the debugger continue with its execution.
- Now execute your external procedure in the same session where you first executed `DEBUG_EXTPROC.STARTUP_EXTPROC_AGENT`
- Your debugger should now break in function 'pextproc'. At this point in time, the shared library referenced by your PL/SQL external function would have been loaded and the function resolved. Now set a breakpoint in your C function and let the debugger continue its execution.

Because PL/SQL loads the shared library at runtime, the debugger you use may or may not automatically be able to track the new symbols from the shared library. You may have to issue some debugger command to load the symbols (for example, 'share' in gdb)

- The debugger should now break in your C function. Its assumed that you had built the shared library with debugging symbols.
- Now proceed with your debugging.

241.3 Rules and Limits

`DEBUG_EXTPROC` works only on platforms with debuggers that can attach to a running process.

241.4 Summary of DEBUG_EXTPROC Subprograms

The `STARTUP_EXTPROC_AGENT` procedure is the only `DEBUG_EXTPROC` subprogram.

Table 241-1 DEBUG_EXTPROC Package Subprograms

Subprogram	Description
STARTUP_EXTPROC_AGENT Procedure	Starts up the extproc agent process in the session

241.4.1 STARTUP_EXTPROC_AGENT Procedure

This procedure starts up the extproc agent process in the session. This enables you to get the PID of the executing process. This PID is needed to be able to attach to the running process using a debugger.

Syntax

```
DEBUG_EXTPROC.STARTUP_EXTPROC_AGENT;
```

242
HTF

The `HTF` (hypertext functions) and `HTP` (hypertext procedures) packages generate HTML tags. For example, the `HTF.ANCHOR` function generates the HTML anchor tag, `<A>`.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

This chapter contains the following topics:

- Deprecated Subprograms
- Operational Notes
- Rules and Limits
- Examples
- Summary of Tags
- Summary of HTF Subprograms

242.1 HTF Deprecated Subprograms

The `ESCAPE_URL` function is deprecated with Oracle Database 10*g*.

> **Note:**
>
> Oracle recommends that you do not use deprecated procedures in new applications. Support for deprecated features is for backward compatibility only.

Related Topics

- ESCAPE_URL Function
  This deprecated function replaces characters that have special meaning in HTML and HTTP with their escape sequences.

242.2 HTF Operational Notes

These operational notes apply to the `HTF` package.

For every HTF function that generates one or more HTML tags, there is a corresponding HTP procedure with identical parameters with the following exception:

- The PRINTS Procedure and the PS Procedure do not have HTF function equivalents. Use the ESCAPE_SC Function or the ESCAPE_URL Function if you need a string conversion function. Note that while there is a ESCAPE_SC Procedure that performs the same operation as the PRINTS Procedure and the PS Procedure, there is no procedural equivalent for the ESCAPE_URL Function.
- The FORMAT_CELL Function does not have an HTP equivalent. The function formats column values inside an HTML table using TABLEDATA Function which does have an HTP equivalent in the TABLEDATA Procedure. The advantage of this using the FORMAT_CELL Function is that it allows for better control over the HTML tables.

The function versions do not directly generate output in your Web page. Instead, they pass their output as return values to the statements that invoked them. Use these functions when you need to nest calls. To print the output of HTF functions, call the functions from within the HTF.`PRINT` function. It then prints its parameters to the generated Web page.

242.3 Rules and Limits

If you use values of the LONG datatype in functions such as `HTF.PRINT`, `HTF.PRN`, `HTF.PA` or `OWA_UTIL.CELLSPRINT`, only the first 32 K of the `LONG` data is used. The `LONG` data is bound to a `VARCHAR2` datatype in the function.

242.4 HTF Examples

The commands in this example generate a simple HTML document.

```
CREATE OR REPLACE PROCEDURE hello AS
BEGIN
   HTP.P (HTF.HTMLOPEN); -- generates <HTML>
   HTP.P (HTF.HEADOPEN); -- generates <HEAD>
   HTP.P (HTF.TITLE('Hello')); -- generates <TITLE>Hello</TITLE>
   HTP.P (HTF.HEADCLOSE); -- generates </HEAD>
   HTP.P (HTF.BODYOPEN); -- generates <BODY>
   HTP.P (HTF.HEADER(1, 'Hello')); -- generates <H1>Hello</H1>
   HTP.P (HTF.BODYCLOSE); -- generates </BODY>
   HTP.P (HTF.HTMLCLOSE); -- generates </HTML>
END;
```

242.5 Summary of Tags

This list summarizes the HTML tags generated by the `HTF` package.

HTML, HEAD, and BODY Tags

HTMLOPEN Function, HTMLCLOSE Function - generate `<HTML>` and `</HTML>`

HEADOPEN Function, HEADCLOSE Function - generate `<HEAD>` and `</HEAD>`

BODYOPEN Function, BODYCLOSE Function - generate `<BODY>` and `</BODY>`

Comment Tag

COMMENT Function - generates <!-- and -->

http://www.w3.org.BASE Function - generates `<BASE>`

LINKREL Function - generates `<LINK>` with the `REL` attribute

LINKREV Function - generates <LINK> with the REV attribute

TITLE Function - generates <TITLE>

META Function - generates <META>

SCRIPT Function - generates <SCRIPT>

STYLE Function - generates <STYLE>

ISINDEX Function - generates <ISINDEX>

Applet Tags

APPLETOPEN Function, APPLETCLOSE Function - generate <APPLET> and </APPLET>

PARAM Function - generates <PARAM>

List Tags

OLISTOPEN Function, OLISTCLOSE Function - generate <OL> and </OL>

ULISTOPEN Function, ULISTCLOSE Function - generate <UL> and </UL>

DLISTOPEN Function, DLISTCLOSE Function- generate <DL> and </DL>

DLISTTERM Function - generates <DT>

DLISTDEF Function - generates <DD>

DIRLISTOPEN Function, DIRLISTCLOSE Function - generate <DIR> and </DIR>

LISTHEADER Function - generates <LH>

LISTINGOPEN Function, LISTINGCLOSE Function - generate <LISTING> and </LISTING>

MENULISTOPEN Function - generate <MENU> and </MENU>

LISTITEM Function - generates <LI>

Form Tags

FORMOPEN Function, FORMCLOSE Function - generate <FORM> and </FORM>

FORMCHECKBOX Function - generates <INPUT TYPE="CHECKBOX">

FORMHIDDEN Function - generates <INPUT TYPE="HIDDEN">

FORMIMAGE Function - generates <INPUT TYPE="IMAGE">

FORMPASSWORD Function - generates <INPUT TYPE="PASSWORD">

FORMRADIO Function - generates <INPUT TYPE="RADIO">

FORMSELECTOPEN Function, FORMSELECTCLOSE Function - generate <SELECT> and </SELECT>

FORMSELECTOPTION Function - generates <OPTION>

FORMTEXT Function - generates <INPUT TYPE="TEXT">

FORMTEXTAREA Function - generate <TEXTAREA>

FORMTEXTAREAOPEN Function, FORMTEXTAREACLOSE Function - generate <TEXTAREA> and </TEXTAREA>

FORMRESET Function - generates <INPUT TYPE="RESET">

FORMSUBMIT Function - generates <INPUT TYPE="SUBMIT">

Table Tags

TABLEOPEN Function, TABLECLOSE Function - generate <TABLE> and </TABLE>

TABLECAPTION Function - generates <CAPTION>

TABLEROWOPEN Function, TABLEROWCLOSE Function - generate <TR> and </TR>

TABLEHEADER Function - generates <TH>

TABLEDATA Function - generates <TD>

IMG, HR, and A Tags

HR Function, LINE Function - generate <HR>

IMG Function, IMG2 Function - generate <IMG>

ANCHOR Function, ANCHOR2 Function - generate <A>

MAPOPEN Function, MAPCLOSE Function - generate <MAP> and </MAP>

Paragraph Formatting Tags

HEADER Function - generates heading tags (<H1> to <H6>)

PARA Function, PARAGRAPH Function - generate <P>

PRN Functions, PRINT Functions - generate any text that is passed in

PRN Functions, S Function - generate any text that is passed in; special characters in HTML are escaped

PREOPEN Function, PRECLOSE Function - generate <PRE> and </PRE>

BLOCKQUOTEOPEN Function, BLOCKQUOTECLOSE Function - generate <BLOCKQUOTE> and </BLOCKQUOTE>

DIV Function - generates <DIV>

NL Function, BR Function - generate

NOBR Function - generates <NOBR>

WBR Function - generates <WBR>

PLAINTEXT Function - generates <PLAINTEXT>

ADDRESS Function - generates <ADDRESS>

MAILTO Function - generates <A> with the MAILTO attribute

AREA Function - generates <AREA>

BGSOUND Function - generates <BGSOUND>

Character Formatting Tags

BASEFONT Function - generates <BASEFONT>

BIG Function - generates <BIG>

BOLD Function - generates <B>

CENTER Function - generates <CENTER> and </CENTER>

CENTEROPEN Function, CENTERCLOSE Function - generate <CENTER> and </CENTER>

CITE Function - generates <CITE>

CODE Function - generates <CODE>

DFN Function - generates <DFN>

EM Function, EMPHASIS Function - generate <EM>

FONTOPEN Function, FONTCLOSE Function - generate <FONT> and </FONT>

ITALIC Function - generates <I>

KBD Function, KEYBOARD Function - generate <KBD> and </KBD>

S Function - generates <S>

SAMPLE Function - generates <SAMP>

SMALL Function - generates <SMALL>

STRIKE Function - generates <STRIKE>

STRONG Function - generates <STRONG>

SUB Function - generates <SUB>

SUP Function - generates <SUP>

TELETYPE Function - generates <TT>

UNDERLINE Function - generates <U>

VARIABLE Function - generates <VAR>

Frame Tags

FRAME Function - generates <FRAME>

FRAMESETOPEN Function, FRAMESETCLOSE Function - generate <FRAMESET> and </FRAMESET>

NOFRAMESOPEN Function, NOFRAMESCLOSE Function - generate <NOFRAMES> and </NOFRAMES>

242.6 Summary of HTF Subprograms

This table lists the HTF subprograms and briefly describes them.

Table 242-1 HTF Package Subprograms

Subprogram	Description
ADDRESS Function	Generates the <ADDRESS> and </ADDRESS> tags which specify the address, author and signature of a document
ANCHOR Function	Generates the <A> and </A> tags which specify the source or destination of a hypertext link
ANCHOR2 Function	Generates the <A> and </A> tags which specify the source or destination of a hypertext link
APPLETCLOSE Function	Closes the applet invocation with the </APPLET> tag
APPLETOPEN Function	Generates the <APPLET> tag which begins the invocation of a Java applet
AREA Function	Generates the <AREA> tag, which defines a client-side image map
BASE Function	Generates the <BASE> tag which records the URL of the document
BASEFONT Function	Generates the <BASEFONT> tag which specifies the base font size for a Web page
BGSOUND Function	Generates the <BGSOUND> tag which includes audio for a Web page
BIG Function	Generates the <BIG> and </BIG> tags which direct the browser to render the text in a bigger font
BLOCKQUOTECLOSE Function	Generates the </BLOCKQUOTE> tag which mark the end of a section of quoted text
BLOCKQUOTEOPEN Function	Generates the <BLOCKQUOTE> tag, which marks the beginning of a section of quoted text
BODYCLOSE Function	Generates the </BODY> tag which marks the end of a body section of an HTML document
BODYOPEN Function	Generates the <BODY> tag which marks the beginning of the body section of an HTML document
BOLD Function	Generates the <B> and </B> tags which direct the browser to display the text in boldface
BR Function	Generates the tag which begins a new line of text
CENTER Function	Generates the <CENTER> and </CENTER> tags which center a section of text within a Web page
CENTERCLOSE Function	Generates the </CENTER> tag which marks the end of a section of text to center
CENTEROPEN Function	Generates the <CENTER> tag which mark the beginning of a section of text to center
CITE Function	Generates the <CITE> and </CITE> tags which direct the browser to render the text as a citation
CODE Function	Generates the <CODE> and </CODE> tags which direct the browser to render the text in monospace font or however "code" is defined stylistically
COMMENT Function	Generates the comment tags <!-- ctext -->
DFN Function	Generates the <DFN> and </DFN> tags which direct the browser to mark the text as italics or however "definition" is defined stylistically
DIRLISTCLOSE Function	Generates the </DIR> tag which ends a directory list section
DIRLISTOPEN Function	Generates the <DIR> which starts a directory list section

Table 242-1 (Cont.) HTF Package Subprograms

Subprogram	Description
DIV Function	Generates the <DIV> tag which creates document divisions
DLISTCLOSE Function	Generates the </DL> tag which ends a definition list
DLISTDEF Function	Generates the <DD> tag, which inserts definitions of terms
DLISTOPEN Function	Generates the <DL> tag which starts a definition list
DLISTTERM Function	Generates the <DT> tag which defines a term in a definition list <DL>
EM Function	Generates the <EM> and </EM> tags, which define text to be emphasized
EMPHASIS Function	Generates the <EM> and </EM> tags, which define text to be emphasized
ESCAPE_SC Function	Replaces characters that have special meaning in HTML with their escape sequences
ESCAPE_URL Function	Replaces characters that have special meaning in HTML and HTTP with their escape sequences
FONTCLOSE Function	Generates the </FONT> tag which marks the end of a section of text with the specified font characteristics
FONTOPEN Function	Generates the <FONT> which marks the beginning of section of text with the specified font characteristics
FORMAT_CELL Function	formats column values inside an HTML table using the TABLEDATA Function
FORMCHECKBOX Function	Generates the <INPUT> tag with TYPE="checkbox" which inserts a checkbox element in a form
FORMCLOSE Function	Generates the </FORM> tag which marks the end of a form section in an HTML document
FORMFILE Function	Generates the <INPUT> tag with TYPE="file" which inserts a file form element, and is used for file uploading for a given page
FORMHIDDEN Function	Generates the <INPUT> tag with TYPE="hidden"which inserts a hidden form element
FORMIMAGE Function	Generates the <INPUT> tag with TYPE="image" which creates an image field that the user clicks to submit the form immediately
FORMOPEN Function	Generates the <FORM> tag which marks the beginning of a form section in an HTML document
FORMPASSWORD Function	Generates the <INPUT> tag with TYPE="password" which creates a single-line text entry field
FORMRADIO Function	Generates the <INPUT> tag with TYPE="radio", which creates a radio button on the HTML form
FORMRESET Function	Generates the <INPUT> tag with TYPE="reset" which creates a button that, when selected, resets the form fields to their initial values
FORMSELECTCLOSE Function	Generates the </SELECT> tag which marks the end of a Select form element
FORMSELECTOPEN Function	Generates the </SELECT> tag which marks the beginning of a Select form element

Table 242-1 (Cont.) HTF Package Subprograms

Subprogram	Description
FORMSELECTOPTION Function	Generates the <OPTION> tag which represents one choice in a Select element
FORMSUBMIT Function	Generates the <INPUT> tag with TYPE="submit" which creates a button that, when clicked, submits the form
FORMTEXT Function	Generates the <INPUT> tag with TYPE="text", which creates a field for a single line of text
FORMTEXTAREA Function	Generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area
FORMTEXTAREA2 Function	Generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area with the ability to specify a wrap style
FORMTEXTAREACLOSE Function	Generates the </TEXTAREA> tag which ends a text area form element
FORMTEXTAREAOPEN Function	Generates the <TEXTAREA> which marks the beginning of a text area form element
FORMTEXTAREAOPEN2 Function	Generates the <TEXTAREA> which marks the beginning of a text area form element with the ability to specify a wrap style
FRAME Function	Generates the <FRAME> tag which defines the characteristics of a frame created by a <FRAMESET> tag
FRAMESETCLOSE Function	Generates the </FRAMESET> tag which ends a frameset section
FRAMESETOPEN Function	Generates the </FRAMESET> tag which begins a frameset section
HEADCLOSE Function	Generates the </HEAD> tag which marks the end of an HTML document head section
HEADER Function	Generates opening heading tags (<H1> to <H6>) and their corresponding closing tags (</H1> to </H6>)
HEADOPEN Function	Generates the <HEAD> tag which marks the beginning of the HTML document head section
HR Function	Generates the <HR> tag, which generates a line in the HTML document
HTMLCLOSE Function	Generates the </HTML> tag which marks the end of an HTML document
HTMLOPEN Function	Generates the <HTML> tag which marks the beginning of an HTML document
IMG Function	Generates the <IMG> tag which directs the browser to load an image onto the HTML page
IMG2 Function	Generates the <IMG> tag which directs the browser to load an image onto the HTML page with the option of specifying values for the USEMAP attribute
ISINDEX Function	Creates a single entry field with a prompting text, such as "*enter value,*" then sends that value to the URL of the page or program
ITALIC Function	Generates the <I> and </I> tags which direct the browser to render the text in italics
KBD Function	Generates the <KBD> and </KBD> tags which direct the browser to render the text in monospace font

Table 242-1 (Cont.) HTF Package Subprograms

Subprogram	Description
KEYBOARD Function	Generates the <KBD> and </KBD> tags, which direct the browser to render the text in monospace font
LINE Function	Generates the <HR> tag, which generates a line in the HTML document
LINKREL Function	Generates the <LINK> tag with the REL attribute which delineates the relationship described by the hypertext link from the anchor to the target
LINKREV Function	Generates the <LINK> tag with the REV attribute which delineates the relationship described by the hypertext link from the target to the anchor
LISTHEADER Function	Generates the <LH> and </LH> tags which print an HTML tag at the beginning of the list
LISTINGCLOSE Function	Generates the </LISTING> tags which marks the end of a section of fixed-width text in the body of an HTML page
LISTINGOPEN Function	Generates the <LISTING> tag which marks the beginning of a section of fixed-width text in the body of an HTML page
LISTITEM Function	Generates the <LI> tag, which indicates a list item
MAILTO Function	Generates the <A> tag with the HREF set to 'mailto' prepended to the mail address argument
MAPCLOSE Function	Generates the </MAP> tag which marks the end of a set of regions in a client-side image map
MAPOPEN Function	Generates the <MAP> tag which mark the beginning of a set of regions in a client-side image map
MENULISTCLOSE Function	Generates the </MENU> tag which ends a list that presents one line for each item
MENULISTOPEN Function	Generates the <MENU> tag which begins a list that presents one line for each item
META Function	Generates the <META> tag, which embeds meta-information about the document and also specifies values for HTTP headers
NL Function	Generates the tag which begins a new line of text
NOBR Function	Generates the <NOBR> and </NOBR> tags which turn off line-breaking in a section of text
NOFRAMESCLOSE Function	Generates the </NOFRAMES> tag which marks the end of a no-frames section
NOFRAMESOPEN Function	Generates the <NOFRAMES> tag which mark the beginning of a no-frames section
OLISTCLOSE Function	Generates the </OL> tag which defines the end of an ordered list
OLISTOPEN Function	Generates the <OL> tag which marks the beginning of an ordered list
PARA Function	Generates the <P> tag which indicates that the text that comes after the tag is to be formatted as a paragraph
PARAGRAPH Function	Adds attributes to the <P> tag
PARAM Function	Generates the <PARAM> tag which specifies parameter values for Java applets

Table 242-1 (Cont.) HTF Package Subprograms

Subprogram	Description
PLAINTEXT Function	Generates the <PLAINTEXT> and </PLAINTEXT> tags which direct the browser to render the text they surround in fixed-width type
PRECLOSE Function	Generates the </PRE> tag which marks the end of a section of preformatted text in the body of the HTML page
PREOPEN Function	Generates the <PRE> tag which marks the beginning of a section of preformatted text in the body of the HTML page
PRINT Functions	Generates the specified parameter as a string terminated with the \n newline character
PRN Functions	Generates the specified parameter as a string
S Function	Generates the <S> and </S> tags which direct the browser to render the text they surround in strikethrough type
SAMPLE Function	Generates the <SAMP> and </SAMP> tags which direct the browser to render the text they surround in monospace font or however "sample" is defined stylistically
SCRIPT Function	Generates the <SCRIPT> and </SCRIPT> tags which contain a script written in languages such as JavaScript and VBscript
SMALL Function	Generates the <SMALL> and </SMALL> tags, which direct the browser to render the text they surround using a small font
STRIKE Function	Generates the <STRIKE> and </STRIKE> tags which direct the browser to render the text they surround in strikethrough type
STRONG Function	Generates the <STRONG> and </STRONG> tags which direct the browser to render the text they surround in bold or however "strong" is defined stylistically
STYLE Function	Generates the <STYLE> and </STYLE> tags which include a style sheet in a Web page
SUB Function	Generates the _{and} tags which direct the browser to render the text they surround as subscript
SUP Function	Generates the ^{and} tags which direct the browser to render the text they surround as superscript
TABLECAPTION Function	Generates the <CAPTION> and </CAPTION> tags which place a caption in an HTML table
TABLECLOSE Function	Generates the </TABLE> tag which marks the end of an HTML table
TABLEDATA Function	Generates the <TD> and </TD> tags which insert data into a cell of an HTML table
TABLEHEADER Function	Generates the <TH> and </TH> tags which insert a header cell in an HTML table.
TABLEOPEN Function	Generates the <TABLE> tag which marks the beginning of an HTML table
TABLEROWCLOSE Function	Generates the </TR> tag which marks the end of a new row in an HTML table
TABLEROWOPEN Function	Generates the <TR> tag which marks the beginning of a new row in an HTML table

Table 242-1 (Cont.) HTF Package Subprograms

Subprogram	Description
TELETYPE Function	Generates the <TT> and </TT> tags which direct the browser to render the text they surround in a fixed width typewriter font, for example, the courier font
TITLE Function	Generates the <TITLE> and </TITLE> tags which specify the text to display in the titlebar of the browser window
ULISTCLOSE Function	Generates the </UL> tag which marks the end of an unordered list
ULISTOPEN Function	Generates the <UL> tag which marks the beginning of an unordered list
UNDERLINE Function	Generates the <U> and </U> tags, which direct the browser to render the text they surround with an underline
VARIABLE Function	Generates the <VAR> and </VAR> tags which direct the browser to render the text they surround in italics or however "variable" is defined stylistically.
WBR Function	Generates the <WBR> tag, which inserts a soft line break within a section of NOBR text

242.6.1 ADDRESS Function

This function generates the <ADDRESS> and </ADDRESS> tags which specify the address, author and signature of a document.

Syntax

```
HTF.ADDRESS (
   cvalue          IN       VARCHAR2
   cnowrap         IN       VARCHAR2   DEFAULT NULL
   cclear          IN       VARCHAR2   DEFAULT NULL
   cattributes     IN       VARCHAR2   DEFAULT NULL)
 RETURN VARCHAR2;
```

Parameters

Table 242-2 ADDRESS Function Parameters

Parameter	Description
cvalue	The string that goes between the <ADDRESS> and </ADDRESS> tags.
cnowrap	If the value for this parameter is not NULL, the NOWRAP attribute is included in the tag
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag

Examples

This function generates

```
<ADDRESS CLEAR="cclear" NOWRAP cattributes>cvalue</ADDRESS>
```

242.6.2 ANCHOR Function

Like the ANCHOR2 function, this function generates the <A> and </A> HTML tags which specify the source or destination of a hypertext link.

The difference between these subprograms is that the ANCHOR2 Function provides a target and therefore can be used for a frame.

Syntax

```
HTF.ANCHOR (
   curl           IN       VARCHAR2,
   ctext          IN       VARCHAR2,
   cname          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
 RETURN VARCHAR2;
```

Parameters

Table 242-3 ANCHOR Function Parameters

Parameter	Description
curl	The value for the HREF attribute.
ctext	The string that goes between the <A> and </A> tags.
cname	The value for the NAME attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<A HREF="curl" NAME="cname" cattributes>ctext</A>
```

Usage Notes

This tag accepts several attributes, but either HREF or NAME is required. HREF specifies to where to link. NAME allows this tag to be a target of a hypertext link.

242.6.3 ANCHOR2 Function

Like the ANCHOR function, this function generates the <A> and </A> HTML tags which specify the source or destination of a hypertext link.

The difference between this and the ANCHOR function is that this function provides a target and therefore can be used for a frame.

Syntax

```
HTF.ANCHOR2 (
   curl           IN       VARCHAR2,
   ctext          IN       VARCHAR2,
   cname          IN       VARCHAR2   DEFAULT NULL,
   ctarget        in       varchar2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
 RETURN VARCHAR2;
```

Parameters

Table 242-4 ANCHOR2 Function Parameters

Parameter	Description
curl	The value for the HREF attribute.
ctext	The string that goes between the <A> and </A> tags.
cname	The value for the NAME attribute
ctarget	The value for the TARGET attribute.
cattributes	The other attributes to be included as-is in the tag

Examples

This function generates

```
<A HREF="curl" NAME="cname" TARGET = "ctarget" cattributes>ctext</A>
```

242.6.4 APPLETCLOSE Function

This function closes the applet invocation with the </APPLET> tag.

You must first invoke the a Java applet using APPLETOPEN Function

Syntax

```
HTF.APPLETCLOSE
 RETURN VARCHAR2;
```

242.6.5 APPLETOPEN Function

This function generates the <APPLET> tag which begins the invocation of a Java applet.

You close the applet invocation with APPLETCLOSE Function which generates the </APPLET> tag.

Syntax

```
HTF.APPLETOPEN (
   ccode          IN       VARCHAR2,
   cheight        IN       NUMBER,
   cwidth         IN       NUMBER,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-5 APPLETOPEN Function Parameters

Parameter	Description
ccode	The value for the CODE attribute which specifies the name of the applet class.
cheight	The value for the HEIGHT attribute.
cwidth	The value for the WIDTH attribute.

Table 242-5 (Cont.) APPLETOPEN Function Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<APPLET CODE=ccode HEIGHT=cheight WIDTH=cwidth cattributes>
```

so that, for example,

```
HTF.appletopen('testclass.class', 100, 200, 'CODEBASE="/ows-applets"')
```

generates

```
<APPLET CODE="testclass.class" height=100 width=200 CODEBASE="/ows-applets">
```

Usage Notes

- Specify parameters to the Java applet using the PARAM Function function.
- Use the `cattributes` parameter to specify the `CODEBASE` attribute since the PL/SQL cartridge does not know where to find the class files. The `CODEBASE` attribute specifies the virtual path containing the class files.

242.6.6 AREA Function

This function generates the `<AREA>` tag, which defines a client-side image map. The `<AREA>` tag defines areas within the image and destinations for the areas.

Syntax

```
HTF.AREA (
   ccoords        IN       VARCHAR2
   cshape         IN       VARCHAR2   DEFAULT NULL,
   chref          IN       VARCHAR2   DEFAULT NULL,
   cnohref        IN       VARCHAR2   DEFAULT NULL,
   ctarget        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-6 AREA Function Parameters

Parameter	Description
ccords	The value for the `COORDS` attribute.
cshape	The value for the `SHAPE` attribute.
chref	The value for the `HREF` attribute.
cnohref	If the value for this parameter is not `NULL`, the `NOHREF` attribute is added to the tag.
ctarget	The value for the `TARGET` attribute.

Table 242-6 (Cont.) AREA Function Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<AREA COORDS="ccoords" SHAPE="cshape" HREF="chref" NOHREF TARGET="ctarget" cattributes>
```

242.6.7 BASE Function

This function generates the <BASE> tag which records the URL of the document.

Syntax

```
HTF.BASE (
   ctarget        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-7 BASE Function Parameters

Parameter	Description
ctarget	The value for the TARGET attribute which establishes a window name to which all links in this document are targeted.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BASE HREF="<current URL>" TARGET="ctarget" cattributes>
```

242.6.8 BASEFONT Function

This function generates the <BASEFONT> tag which specifies the base font size for a Web page.

Syntax

```
HTF.BASEFONT (
   nsize    IN    INTEGER)
  RETURN VARCHAR2;
```

Parameters

Table 242-8 BASEFONT Function Parameters

Parameter	Description
nsize	The value for the SIZE attribute.

Examples

This function generates

```
<BASEFONT SIZE="nsize">
```

242.6.9 BGSOUND Function

This function generates the <BGSOUND> tag which includes audio for a Web page.

Syntax

```
HTF.BGSOUND (
   csrc           IN       VARCHAR2,
   cloop          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-9 BGSOUND Function Parameters

Parameter	Description
csrc	The value for the SRC attribute.
cloop	The value for the LOOP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BGSOUND SRC="csrc" LOOP="cloop" cattributes>
```

242.6.10 BIG Function

This function generates the <BIG> and </BIG> tags which direct the browser to render the text in a bigger font.

Syntax

```
HTF.BIG (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-10 BIG Function Parameters

Parameter	Description
ctext	The text that goes between the tags.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BIG cattributes>ctext</BIG>
```

242.6.11 BLOCKQUOTECLOSE Function

This function generates the `</BLOCKQUOTE>` tag which mark the end of a section of quoted text.

You mark the beginning of a section of text by means of the BLOCKQUOTEOPEN Function.

Syntax

```
HTF.BLOCKQUOTECLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</BLOCKQUOTE>
```

242.6.12 BLOCKQUOTEOPEN Function

This function generates the `<BLOCKQUOTE>` tag, which marks the beginning of a section of quoted text.

You mark the end of a section of text by means of the BLOCKQUOTECLOSE Function.

Syntax

```
HTF.BLOCKQUOTEOPEN (
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-11 BLOCKQUOTEOPEN Function Parameters

Parameter	Description
`cnowrap`	If the value for this parameter is not `NULL`, the `NOWRAP` attribute is added to the tag.
`cclear`	The value for the `CLEAR` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BLOCKQUOTE CLEAR="cclear" NOWRAP cattributes>
```

242.6.13 BODYCLOSE Function

This function generates the `</BODY>` tag which marks the end of a body section of an HTML document.

You mark the beginning of a body section by means of the BODYOPEN Function.

Syntax

```
HTF.BODYCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</BODY>
```

242.6.14 BODYOPEN Function

This function generates the `<BODY>` tag which marks the beginning of the body section of an HTML document.

You mark the end of a body section by means of the BODYCLOSE Function.

Syntax

```
HTF.BODYOPEN (
   cbackground    IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-12 BODYOPEN Function Parameters

Parameter	Description
`cbackground`	The value for the `BACKGROUND` attribute which specifies a graphic file to use for the background of the document.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BODY background="cbackground" cattributes>
```

so that

```
HTF.BODYOPEN('/img/background.gif')
  RETURN VARCHAR2;
```

generates:

```
<BODY background="/img/background.gif">
```

242.6.15 BOLD Function

This function generates the <B> and </B> tags which direct the browser to display the text in boldface.

Syntax

```
HTF.BOLD (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-13 BOLD Function Parameters

Parameter	Description
ctext	The text that goes between the tags.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<B cattributes>ctext</B>
```

242.6.16 BR Function

This function generates the <BR> tag which begins a new line of text.

It performs the same operation as the NL Function.

Syntax

```
HTF.BR(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-14 BR Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BR CLEAR="cclear" cattributes>
```

242.6.17 CENTER Function

This function generates the <CENTER> and </CENTER> tags which center a section of text within a Web page.

Syntax

```
HTF.CENTER (
   ctext          IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-15 CENTER Parameters

Parameter	Description
ctext	The text that goes between the tags.

Examples

This function generates

```
<CENTER>ctext</CENTER>
```

242.6.18 CENTERCLOSE Function

This function generates the </CENTER> tag which marks the end of a section of text to center.

You mark the beginning of a section of text to center with the CENTEROPEN Function.

Syntax

```
HTF.CENTERCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</CENTER>
```

242.6.19 CENTEROPEN Function

This function generates the <CENTER> tag which mark the beginning of a section of text to center.

You mark the beginning of a section of text to center with the CENTERCLOSE Function.

Syntax

```
HTF.CENTEROPEN
  RETURN VARCHAR2;
```

Examples

This function generates

```
<CENTER>
```

242.6.20 CITE Function

This function generates the <CITE> and </CITE> tags which direct the browser to render the text as a citation.

Syntax

```
HTF.CITE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-16 CITE Function Parameters

Parameter	Description
ctext	The text to render as citation.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<CITE cattributes>ctext</CITE>
```

242.6.21 CODE Function

This function generates the <CODE> and </CODE> tags which direct the browser to render the text in monospace font or however "code" is defined stylistically.

Syntax

```
HTF.CODE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-17 CODE Function Parameters

Parameter	Description
ctext	The text to render as code.
cattributes	The other attributes to be included as-is in the tag

Examples

This function generates

```
<CODE cattributes>ctext</CODE>
```

242.6.22 COMMENT Function

This function generates the comment tags.

Syntax

```
HTF.COMMENT (
   ctext          IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-18 COMMENT Function Parameters

Parameter	Description
ctext	The comment.

Examples

This function generates

```
<!-- ctext -->
```

242.6.23 DFN Function

This function generates the <DFN> and </DFN> tags which direct the browser to mark the text in italics or however "definition" is described stylistically.

Syntax

```
HTF.DFN (
   ctext          IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-19 DFN Function Parameters

Parameter	Description
ctext	The text to render in italics.

Examples

This function generates

```
<DFN>ctext</DFN>
```

242.6.24 DIRLISTCLOSE Function

This function generates the </DIR> tag which ends a directory list section. You start a directory list section with the DIRLISTOPEN Function.

Syntax

```
HTF.DIRLISTCLOSE
  RETURN VARCHAR2;
```

Usage Notes

A directory list presents a list of items that contains up to 20 characters. Items in this list are typically arranged in columns, 24 characters wide. Insert the <LI> tag directly or invoke the LISTITEM Function so that the <LI> tag appears directly after the </DIR> tag to define the items as a list.

Examples

This function generates

```
</DIR>
```

Related Topics

- DIRLISTOPEN Function
  This function generates the <DIR> which starts a directory list section. You end a directory list section with the DIRLISTCLOSE Function.

242.6.25 DIRLISTOPEN Function

This function generates the <DIR> which starts a directory list section. You end a directory list section with the DIRLISTCLOSE Function.

Syntax

```
HTF.DIRLISTOPEN
  RETURN VARCHAR2;
```

Usage Notes

A directory list presents a list of items that contains up to 20 characters. Items in this list are typically arranged in columns, 24 characters wide. Insert the <LI> tag directly or invoke the LISTITEM Function so that the <LI> tag appears directly after the </DIR> tag to define the items as a list.

Examples

This function generates

```
<DIR>
```

Related Topics

- DIRLISTCLOSE Function
  This function generates the </DIR> tag which ends a directory list section. You start a directory list section with the DIRLISTOPEN Function.

242.6.26 DIV Function

This function generates the <DIV> tag which creates document divisions.

Syntax

```
HTF.DIV (
   calign          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-20 DIV Function Parameters

Parameter	Description
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<DIV ALIGN="calign" cattributes>
```

242.6.27 DLISTCLOSE Function

This function generates the </DL> tag which ends a definition list. You start a definition list by means of the DLISTOPEN Function.

Syntax

```
HTF.DLISTCLOSE
  RETURN VARCHAR2;
```

Usage Notes

A definition list looks like a glossary: it contains terms and definitions. Terms are inserted using the DLISTTERM Function and definitions are inserted using the DLISTDEF Function.

Examples

This function generates

```
</DL>
```

Related Topics

- DLISTOPEN Function
  This function generates the <DL> tag which starts a definition list. You end a definition list by means of the DLISTCLOSE Function.

242.6.28 DLISTDEF Function

This function generates the <DD> tag, which inserts definitions of terms. Use this tag for a definition list <DL>. Terms are tagged <DT> and definitions are tagged <DD>.

Syntax

```
HTF.DLISTDEF (
   ctext          IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-21 DLISTDEF Function Parameters

Parameter	Description
ctext	The definition of the term.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<DD CLEAR="cclear" cattributes>ctext
```

242.6.29 DLISTOPEN Function

This function generates the <DL> tag which starts a definition list. You end a definition list by means of the DLISTCLOSE Function.

Syntax

```
HTF.DLISTOPEN (
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-22 DLISTOPEN Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Usage Notes

A definition list looks like a glossary: it contains terms and definitions. Terms are inserted using the DLISTTERM Function and definitions are inserted using the DLISTDEF Function.

Examples

This function generates

```
<DL CLEAR="cclear" cattributes>
```

Related Topics

- DLISTCLOSE Function
  This function generates the </DL> tag which ends a definition list. You start a definition list by means of the DLISTOPEN Function.

242.6.30 DLISTTERM Function

This function generates the <DT> tag which defines a term in a definition list <DL>.

Syntax

```
HTF.DLISTTERM (
   ctext          IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-23 DLISTTERM Function Parameters

Parameter	Description
ctext	The term.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<DT CLEAR="cclear" cattributes>ctext
```

242.6.31 EM Function

This function generates the <EM> and </EM> tags, which define text to be emphasized.

It performs the same task as the EMPHASIS Function.

Syntax

```
HTF.EM(
     ctext          IN       VARCHAR2,
     cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-24 EM Function Parameters

Parameter	Description
ctext	The text to emphasize.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<EM cattributes>ctext</EM>
```

242.6.32 EMPHASIS Function

This function generates the <EM> and </EM> tags, which define text to be emphasized.

It performs the same task as the EM Function.

Syntax

```
HTF.EMPHASIS(
     ctext          IN       VARCHAR2,
     cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-25 EMPHASIS Function Parameters

Parameter	Description
ctext	The text to emphasize.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<EM cattributes>ctext</EM>
```

242.6.33 ESCAPE_SC Function

This function replaces characters that have special meaning in HTML with their escape sequences.

The following characters are converted:

- & to &
- " to &quot:
- < to <
- > to >

This function performs the same operation as HTP. PRINTS Procedure and HTP. PS Procedure.

Syntax

```
HTF.ESCAPE_SC(
     ctext           IN       VARCHAR2);
```

Parameters

Table 242-26 ESCAPE_SC Function Parameters

Parameter	Description
ctext	The text string to convert.

242.6.34 ESCAPE_URL Function

This deprecated function replaces characters that have special meaning in HTML and HTTP with their escape sequences.

> **Note:**
>
> This procedure, deprecated in Release 10*g*, and provided here only for reasons of backward compatibility, does not comply with the Internet Engineering Task Force (IETF) Request for Comments (RFC) standards of URL encoding. If you need to encode URLs, it is recommended you use the ESCAPE Functionin the UTL_URL package.

The following characters are converted:

- & to &
- " to &quot:
- < to <
- > to >
- % to &25

Syntax

```
HTF.ESCAPE_URL(
     p_url           IN       VARCHAR2);
```

Parameters

Table 242-27 ESCAPE_URL Function Parameters

Parameter	Description
p_url	The string to convert.

242.6.35 FONTCLOSE Function

This function generates the</FONT> tag which marks the end of a section of text with the specified font characteristics.

You mark the beginning of the section text by means of the FONTOPEN Function.

Syntax

```
HTF.FONTCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

</FONT>

242.6.36 FONTOPEN Function

This function generates the <FONT> which marks the beginning of section of text with the specified font characteristics.

You mark the end of the section text by means of the FONTCLOSE Function.

Syntax

```
HTF.FONTOPEN(
   ccolor          IN       VARCHAR2   DEFAULT NULL,
   cface           IN       VARCHAR2   DEFAULT NULL,
   csize           IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-28 FONTOPEN Function Parameters

Parameter	Description
ccolor	The value for the COLOR attribute.
cface	The value for the FACE attribute
csize	The value for the SIZE attribute
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<FONT COLOR="ccolor" FACE="cface" SIZE="csize" cattributes>
```

242.6.37 FORMAT_CELL Function

This function formats column values inside an HTML table using the TABLEDATA Function. It allows for better control over the HTML tables.

Syntax

```
HTF.FORMAT_CELL(
   columnValue          IN       VARCHAR2
   format_numbers       IN       VARCHAR2   DEFAULT NULL
 RETURN VARCHAR2;
```

Parameters

Table 242-29 FORMAT_CELL Function Parameters

Parameter	Description
`columnValue`	The value that needs to be formatted in an HTML table.
`format_numbers`	The format that numeric data is displayed in. If the value of this parameter is not `NULL`, the number fields are right-justified and rounded to two decimal places.

Examples

This function generates

```
<TD >columnValue</TD>
```

Related Topics

- TABLEDATA Function
  This function generates the `<TD>` and `</TD>` tags which insert data into a cell of an HTML table.

242.6.38 FORMCHECKBOX Function

This function generates the `<INPUT>` tag with `TYPE="checkbox"` which inserts a checkbox element in a form.

A checkbox element is a button that the user toggles on or off.

Syntax

```
HTF.FORMCHECKBOX(
   cname          IN       VARCHAR2,
   cvalue         IN       VARCHAR2   DEFAULT 'ON',
   cchecked       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-30 FORMCHECKBOX Function Parameters

Parameter	Description
`cname`	The value for the `NAME` attribute.

Table 242-30 (Cont.) FORMCHECKBOX Function Parameters

Parameter	Description
cvalue	The value for the VALUE attribute.
cchecked	If the value for this parameter is not NULL, the CHECKED attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="checkbox" NAME="cname" VALUE="cvalue" CHECKED cattributes>
```

242.6.39 FORMCLOSE Function

This function generates the </FORM> tag which marks the end of a form section in an HTML document.

You mark the beginning of the form section by means of the FORMOPEN Function.

Syntax

```
HTF.FORMCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</FORM>
```

242.6.40 FORMFILE Function

This function generates the <INPUT> tag with TYPE="file" which inserts a file form element. This is used for file uploading for a given page.

Syntax

```
HTF.FORMFILE(
   cname          IN       VARCHAR2,
   caccept        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-31 FORMFILE Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
caccept	A comma-delimited list of MIME types for upload.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="file" NAME="cname" ACCEPT="caccept" cattributes>
```

242.6.41 FORMHIDDEN Function

This function generates the <INPUT> tag with TYPE="hidden", which inserts a hidden form element.

This element is not seen by the user. It submits additional values to the script.

Syntax

```
HTF.FORMHIDDEN(
   cname          IN       VARCHAR2,
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-32 FORMHIDDEN Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="hidden" NAME="cname" VALUE="cvalue" cattributes>
```

242.6.42 FORMIMAGE Function

This function generates the <INPUT> tag with TYPE="image" which creates an image field that the user clicks to submit the form immediately.

The coordinates of the selected point are measured in pixels, and returned (along with other contents of the form) in two name/value pairs. The x coordinate is submitted under the name of the field with .x appended, and the y coordinate with .y appended. Any VALUE attribute is ignored.

Syntax

```
HTF.FORMIMAGE(
   cname          IN       VARCHAR2,
   csrc           IN       VARCHAR2,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-33 FORMIMAGE Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
csrc	The value for the SRC attribute that specifies the image file.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="image" NAME="cname" SRC="csrc" ALIGN="calign" cattributes>
```

242.6.43 FORMOPEN Function

This function generates the <FORM> tag which marks the beginning of a form section in an HTML document.

You mark the end of the form section by means of the FORMCLOSE Function.

Syntax

```
HTF.FORMOPEN(
   curl           IN       VARCHAR2,
   cmethod        IN       VARCHAR2   DEFAULT 'POST',
   ctarget        IN       VARCHAR2   DEFAULT NULL,
   cenctype       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-34 FORMOPEN Function Parameters

Parameter	Description
curl	The URL of the Web Request Broker or CGI script where the contents of the form is sent. This parameter is required.
cmethod	The value for the METHOD attribute. The value can be "GET" or "POST".
ctarget	The value for the TARGET attribute.
cenctype	The value for the ENCTYPE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<FORM ACTION="curl" METHOD="cmethod" TARGET="ctarget" ENCTYPE="cenctype" cattributes>
```

242.6.44 FORMPASSWORD Function

This function generates the <INPUT> tag with TYPE="password" which creates a single-line text entry field.

When the user enters text in the field, each character is represented by one asterisk. This is used for entering passwords.

Syntax

```
HTF.FORMPASSWORD(
   cname          IN       VARCHAR2,
   csize          IN       VARCHAR2,
   cmaxlength     IN       VARCHAR2   DEFAULT NULL,
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-35 FORMPASSWORD Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
csize	The value for the SIZE attribute.
cmaxlength	The value for the MAXLENGTH attribute.
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="password" NAME="cname" SIZE="csize" MAXLENGTH="cmaxlength" VALUE="cvalue"
cattributes>
```

242.6.45 FORMRADIO Function

This function generates the <INPUT> tag with TYPE="radio", which creates a radio button on the HTML form. Within a set of radio buttons, the user selects only one. Each radio button in the same set has the same name, but different values. The selected radio button generates a name/value pair.

Syntax

```
HTF.FORMRADIO(
   cname          IN       VARCHAR2,
   cvalue         IN       VARCHAR2,
   cchecked       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-36 FORMRADIO Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.
cchecked	If the value for this parameter is not NULL, the CHECKED attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="radio" NAME="cname" VALUE="cvalue" CHECKED cattributes>
```

242.6.46 FORMRESET Function

This function generates the <INPUT> tag with TYPE="reset" which creates a button that, when selected, resets the form fields to their initial values.

Syntax

```
HTF.FORMRESET(
   cvalue          IN       VARCHAR2   DEFAULT 'Reset',
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-37 FORMRESET Function Parameters

Parameter	Description
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="reset" VALUE="cvalue" cattributes>
```

242.6.47 FORMSELECTCLOSE Function

This function generates the </SELECT> tag which marks the end of a Select form element.

A Select form element is a listbox where the user selects one or more values. You mark the beginning of Select form element by means of the FORMSELECTOPEN Function.The values are inserted using FORMSELECTOPTION Function.

Syntax

```
HTF.FORMSELECTCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</SELECT>
```

as shown under Examples of the FORMSELECTOPEN Function.

242.6.48 FORMSELECTOPEN Function

This function generates the <SELECT> tags which begins a Select form element.

A Select form element is a listbox where the user selects one or more values. You mark the end of Select form element by means of the FORMSELECTCLOSE Function.The values are inserted using FORMSELECTOPTION Function.

Syntax

```
HTF.FORMSELECTOPEN(
  cname        IN    VARCHAR2,
  cprompt      IN    VARCHAR2   DEFAULT NULL,
  nsize        IN    INTEGER    DEFAULT NULL,
  cattributes  IN    VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-38 FORMSELECTOPEN Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
cprompt	The string preceding the list box.
nsize	The value for the SIZE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
cprompt <SELECT NAME="cname" SIZE="nsize" cattributes>
</SELECT>
```

so that

```
HTF.FORMSELECTOPEN('greatest_player';
    'Pick the greatest player:');
HTF.FORMSELECTOPTION('Messier');
HTF.FORMSELECTOPTION('Howe');
HTF.FORMSELECTOPTION('Gretzky');.
HTF.FORMSELECTCLOSE;
```

generates

```
Pick the greatest player:
<SELECT NAME="greatest_player">
<OPTION>Messier
<OPTION>Howe
<OPTION>Gretzky
</SELECT>
```

242.6.49 FORMSELECTOPTION Function

This function generates the `<OPTION>` tag which represents one choice in a Select element.

Syntax

```
HTF.FORMSELECTOPTION(
   cvalue           IN       VARCHAR2,
   cselected       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-39 FORMSELECTOPTION Function Parameters

Parameter	Description
`cvalue`	The text for the option
`cvalue`	If the value for this parameter is not `NULL`, the `SELECTED` attribute is added to the tag.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates `<OPTION SELECTED cattributes>cvalue` as shown under the Examples section of the FORMSELECTOPEN Function.

242.6.50 FORMSUBMIT Function

This function generates the `<INPUT>` tag with `TYPE="submit"` which creates a button that, when clicked, submits the form. If the button has a `NAME` attribute, the button contributes a name/value pair to the submitted data.

Syntax

```
HTF.FORMSUBMIT(
   cname          IN       VARCHAR2   DEFAULT NULL,
   cvalue         IN       VARCHAR2   DEFAULT 'Submit',
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-40 FORMSUBMIT Function Parameters

Parameter	Description
`cname`	The value for the `NAME` attribute.

Table 242-40 (Cont.) FORMSUBMIT Function Parameters

Parameter	Description
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="submit" NAME="cname" VALUE="cvalue" cattributes>
```

242.6.51 FORMTEXT Function

This function generates the <INPUT> tag with TYPE="text", which creates a field for a single line of text.

Syntax

```
HTF.FORMTEXT(
   cname          IN       VARCHAR2,
   csize          IN       VARCHAR2   DEFAULT NULL,
   cmaxlength     IN       VARCHAR2   DEFAULT NULL,
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-41 FORMTEXT Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
csize	The value for the SIZE attribute.
cmaxlength	The value for the MAXLENGTH attribute.
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<INPUT TYPE="text" NAME="cname" SIZE="csize" MAXLENGTH="cmaxlength" VALUE="cvalue"
cattributes>
```

242.6.52 FORMTEXTAREA Function

This function generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area. This field enables entering several lines of text.

The same operation is performed by the FORMTEXTAREA2 Function which in addition has the cwrap parameter that lets you specify a wrap style.

Syntax

```
HTF.FORMTEXTAREA(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-42 FORMTEXTAREA Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" cattributes></
TEXTAREA>
```

242.6.53 FORMTEXTAREA2 Function

This function generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area. This field enables entering several lines of text.

The same operation is performed by the FORMTEXTAREA Function except that in that case you cannot specify a wrap style.

Syntax

```
HTF.FORMTEXTAREA2(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cwrap          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-43 FORMTEXTAREA2 Function Parameters

Parameter	Description
cname	The value for the NAME attribute.

Table 242-43 (Cont.) FORMTEXTAREA2 Function Parameters

Parameter	Description
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cwrap	The value for the WRAP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" WRAP="cwrap"
cattributes></TEXTAREA>
```

242.6.54 FORMTEXTAREACLOSE Function

This function generates the </TEXTAREA> tag which ends a text area form element.

You open a text area element by means of eitherFORMTEXTAREAOPEN Function or FORMTEXTAREAOPEN2 Function.

Syntax

```
HTF.FORMTEXTAREACLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</TEXTAREA>
```

242.6.55 FORMTEXTAREAOPEN Function

This function generates the <TEXTAREA> which marks the beginning of a text area form element.

The same operation is performed by the FORMTEXTAREAOPEN2 Function which in addition has the cwrap parameter that lets you specify a wrap style. You mark the end of a text area form element by means of the FORMTEXTAREACLOSE Function.

Syntax

```
HTF.FORMTEXTAREAOPEN(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-44 FORMTEXTAREAOPEN Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" cattributes>
```

242.6.56 FORMTEXTAREAOPEN2 Function

This function generates the <TEXTAREA> which marks the beginning of a text area form element.

The same operation is performed by the FORMTEXTAREAOPEN Function except that in that case you cannot specify a wrap style. You mark the end of a text area form element by means of the FORMTEXTAREACLOSE Function.

Syntax

```
HTF.FORMTEXTAREAOPEN2(
   cname           IN       VARCHAR2,
   nrows           IN       INTEGER,
   ncolumns        IN       INTEGER,
   calign          IN       VARCHAR2   DEFAULT NULL,
   cwrap           IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-45 FORMTEXTAREAOPEN2 Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cwrap	The value for the WRAP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" WRAP = "cwrap"
cattributes>
```

242.6.57 FRAME Function

This function generates the <FRAME> tag which begins the characteristics of a frame created by a <FRAMESET> tag.

Syntax

```
HTF.FRAME(
   csrc           IN       VARCHAR2,
   cname          IN       VARCHAR2   DEFAULT NULL,
   cmarginwidth   IN       VARCHAR2   DEFAULT NULL,
   cmarginheight  IN       VARCHAR2   DEFAULT NULL,
   cscrolling     IN       VARCHAR2   DEFAULT NULL,
   cnoresize      IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-46 FRAME Function Parameters

Parameter	Description
csrc	The URL to display in the frame.
cname	The value for the NAME attribute.
cmarginwidth	The value for the MARGINWIDTH attribute.
cscrolling	The value for the SCROLLING attribute.
cnoresize	If the value for this parameter is not NULL, the NORESIZE attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<FRAME SRC="csrc" NAME="cname" MARGINWIDTH="cmarginwidth" MARGINHEIGHT="cmarginheight"
SCROLLING="cscrolling" NORESIZE cattributes>
```

242.6.58 FRAMESETCLOSE Function

This function generates the </FRAMESET> tag which ends a frameset section.

You mark the beginning of a frameset section by means of the FRAMESETOPEN Function.

Syntax

```
HTF.FRAMESETCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</FRAMESET>
```

242.6.59 FRAMESETOPEN Function

This function generates the <FRAMESET> tag which define a frameset section.

You mark the end of a frameset section by means of the FRAMESETCLOSE Function.

Syntax

```
HTF.FRAMESETOPEN(
   crows          IN       VARCHAR2   DEFAULT NULL,
   ccols          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-47 FRAMESETOPEN Function Parameters

Parameter	Description
crows	The value for the ROWS attribute.
ccols	The value for the COLS attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<FRAMESET ROWS="crows" COLS="ccols" cattributes>
```

242.6.60 HEADCLOSE Function

This function generates the </HEAD> tag which marks the end of an HTML document head section.

You mark the beginning of an HTML document head section by means of the HEADOPEN Function.

Syntax

```
HTF.HEADCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</HEAD>
```

242.6.61 HEADER Function

This function generates opening heading tags (<H1> to <H6>) and their corresponding closing tags (</H1> to </H6>).

Syntax

```
HTF.HEADER(
   nsize          IN       INTEGER,
   cheader        IN       VARCHAR2,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-48 HEADER Function Parameters

Parameter	Description
nsize	The heading level. This is an integer between 1 and 6.
cheader	The text to display in the heading.
calign	The value for the ALIGN attribute.
cnowrap	The value for the NOWRAP attribute.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

```
HTF.header (1,'Overview')
  RETURN VARCHAR2;
```

produces:

```
<H1>Overview</H1>
```

242.6.62 HEADOPEN Function

This function generates the <HEAD> tag which marks the beginning of the HTML document head section.

You mark the end of an HTML document head section by means of the HEADCLOSE Function.

Syntax

```
HTF.HEADOPEN
  RETURN VARCHAR2;
```

Examples

This function generates

```
<HEAD>
```

242.6.63 HR Function

This function generates the <HR> tag, which generates a line in the HTML document.

This subprogram performs the same operation as the LINE Function.

Syntax

```
HTF.HR(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   csrc            IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-49 HR Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
csrc	The value for the SRC attribute which specifies a custom image as the source of the line.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<HR CLEAR="cclear" SRC="csrc" cattributes>
```

242.6.64 HTMLCLOSE Function

This function generates the </HTML> tag which marks the end of an HTML document.

You use the HTMLOPEN Function to mark the beginning of an HTML document.

Syntax

```
HTF.HTMLCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</HTML>
```

242.6.65 HTMLOPEN Function

This function generates the <HTML> tag which marks the beginning of an HTML document.

You use the HTMLCLOSE Function to mark the end of the an HTML document.

Syntax

```
HTF.HTMLOPEN
  RETURN VARCHAR2;
```

Examples

This function generates

```
<HTML>
```

242.6.66 IMG Function

This function generates the `<IMG>` tag which directs the browser to load an image onto the HTML page.

The IMG2 Function performs the same operation but additionally uses the `cusemap` parameter.

Syntax

```
HTF.IMG(
   curl           IN       VARCHAR2   DEFAULT NULL,
   calign         IN       VARCHAR2   DEFAULT NULL,
   calt           IN       VARCHAR2   DEFAULT NULL,
   cismap         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-50 IMG Function Parameters

Parameter	Description
`curl`	The value for the SRC attribute.
`calign`	The value for the `ALIGN` attribute.
`calt`	The value for the `ALT` attribute which specifies alternative text to display if the browser does not support images.
`cismap`	If the value for this parameter is not `NULL`, the `ISMAP` attribute is added to the tag. The attribute indicates that the image is an imagemap.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<IMG SRC="curl" ALIGN="calign" ALT="calt" ISMAP cattributes>
```

242.6.67 IMG2 Function

This function generates the <IMG> tag, which directs the browser to load an image onto the HTML page.

The IMG Function performs the same operation but does not use the `cusemap` parameter.

Syntax

```
HTF.IMG2(
   curl           IN       VARCHAR2   DEFAULT NULL,
   calign         IN       VARCHAR2   DEFAULT NULL,
   calt           IN       VARCHAR2   DEFAULT NULL,
   cismap         IN       VARCHAR2   DEFAULT NULL,
   cusemap        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-51 IMG2 Function Parameters

Parameter	Description
curl	The value for the SRC attribute.
calign	The value for the ALIGN attribute.
calt	The value for the ALT attribute which specifies alternative text to display if the browser does not support images.
cismap	If the value for this parameter is not NULL, the ISMAP attribute is added to the tag. The attribute indicates that the image is an imagemap.
cusemap	The value for the USEMAP attribute which specifies a client-side image map.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<IMG SRC="curl" ALIGN="calign" ALT="calt" ISMAP USEMAP="cusemap" cattributes>
```

242.6.68 ISINDEX Function

This function creates a single entry field with a prompting text, such as "*enter value*," then sends that value to the URL of the page or program.

Syntax

```
HTF.ISINDEX(
   cprompt        IN       VARCHAR2    DEFAULT NULL,
   curl           IN       VARCHAR2    DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-52 ISINDEX Function Parameters

Parameter	Description
cprompt	The value for the PROMPT attribute.
curl	The value for the HREF attribute.

Examples

This function generates

```
<ISINDEX PROMPT="cprompt" HREF="curl">
```

242.6.69 ITALIC Function

This function generates the `<I> and </I>` tags which direct the browser to render the text in italics.

Syntax

```
HTF.ITALIC(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-53 ITALIC Function Parameters

Parameter	Description
`ctext`	The text to be rendered in italics.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<I cattributes>ctext</I>
```

242.6.70 KBD Function

This function generates the `<KBD>` and `</KBD>` tags which direct the browser to render the text in monospace font.

This subprogram performs the same operation as the KEYBOARD Function.

Syntax

```
HTF.KBD(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-54 KBD Function Parameters

Parameter	Description
`ctext`	The text to be rendered in monospace.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<KBD cattributes>ctext</KBD>
```

242.6.71 KEYBOARD Function

This function generates the <KBD> and </KBD> tags, which direct the browser to render the text in monospace font.

This subprogram performs the same operation as the KBD Function.

Syntax

```
HTF.KEYBOARD(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-55 KEYBOARD Function Parameters

Parameter	Description
ctext	The text to be rendered in monospace.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<KBD cattributes>ctext</KBD>
```

242.6.72 LINE Function

This function generates the <HR> tag, which generates a line in the HTML document.

This subprogram performs the same operation as the HR Function.

Syntax

```
HTF.LINE(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   csrc           IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-56 LINE Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.

Table 242-56 (Cont.) LINE Function Parameters

Parameter	Description
csrc	The value for the SRC attribute which specifies a custom image as the source of the line.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<HR CLEAR="cclear" SRC="csrc" cattributes>
```

242.6.73 LINKREL Function

This function generates the <LINK> tag with the REL attribute which delineates the relationship described by the hypertext link from the anchor to the target. This is only used when the HREF attribute is present.

This is the opposite of LINKREV Function. This tag indicates a relationship between documents but does not create a link. To create a link, use the ANCHOR Function.

Syntax

```
HTF.LINKREL(
   crel           IN       VARCHAR2,
   curl           IN       VARCHAR2,
   ctitle         IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-57 LINKREL Function Parameters

Parameter	Description
crel	The value for the REL attribute.
curl	The value for the URL attribute.
ctitle	The value for the TITLE attribute.

Examples

This function generates

```
<LINK REL="crel" HREF="curl" TITLE="ctitle">
```

242.6.74 LINKREV Function

This function generates the <LINK> tag with the REV attribute which delineates the relationship described by the hypertext link from the target to the anchor.

This is the opposite of the LINKREL Function. This tag indicates a relationship between documents, but does not create a link. To create a link, use the ANCHOR Function.

Syntax

```
HTF.LINKREV(
   crev           IN       VARCHAR2,
   curl           IN       VARCHAR2,
   ctitle         IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-58 LINKREV Function Parameters

Parameter	Description
crev	The value for the REV attribute.
curl	The value for the URL attribute.
ctitle	The value for the TITLE attribute.

Examples

This function generates

```
<LINK REV="crev" HREF="curl" TITLE="ctitle">
```

242.6.75 LISTHEADER Function

This function generates the <LH> and </LH> tags which print an HTML tag at the beginning of the list.

Syntax

```
HTF.LISTHEADER(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-59 LISTHEADER Function Parameters

Parameter	Description
ctext	The text to place between <LH> and </LH>.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<LH cattributes>ctext</LH>
```

242.6.76 LISTINGCLOSE Function

This function generates the </LISTING> tags which marks the end of a section of fixed-width text in the body of an HTML page.

To mark the beginning of a section of fixed-width text in the body of an HTML page, use the LISTINGOPEN Function.

Syntax

```
HTF.LISTINGCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</LISTING>
```

242.6.77 LISTINGOPEN Function

This function generates the <LISTING> tag which marks the beginning of a section of fixed-width text in the body of an HTML page.

To mark the end of a section of fixed-width text in the body of an HTML page, use the LISTINGCLOSE Function.

Syntax

```
HTF.LISTINGOPEN
  RETURN VARCHAR2;
```

Examples

This function generates

```
<LISTING>
```

242.6.78 LISTITEM Function

This function generates the <LI> tag, which indicates a list item.

Syntax

```
HTF.LISTITEM(
   ctext          IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cdingbat       IN       VARCHAR2   DEFAULT NULL,
   csrc           IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-60 LISTITEM Function Parameters

Parameter	Description
ctext	The text for the list item.
cclear	The value for the CLEAR attribute.
cdingbat	The value for the DINGBAT attribute.
csrc	The value for the SRC attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<LI CLEAR="cclear" DINGBAT="cdingbat" SRC="csrc" cattributes>ctext
```

242.6.79 MAILTO Function

This function generates the <A> tag with the HREF set to 'mailto' prepended to the mail address argument.

Syntax

```
HTF.MAILTO(
   caddress       IN       VARCHAR2,
   ctext          IN       VARCHAR2,
   cname          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-61 MAILTO Function Parameters

Parameter	Description
caddress	The email address of the recipient.
ctext	The clickable portion of the link.
cname	The value for the NAME attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<A HREF="mailto:caddress" NAME="cname" cattributes>ctext</A>
```

so that

```
HTF.mailto('pres@white_house.gov','Send Email to the President');
```

generates:

```
<A HREF="mailto:pres@white_house.gov">Send Email to the President</A>
```

242.6.80 MAPCLOSE Function

This function generates the </MAP> tag which marks the end of a set of regions in a client-side image map

. To mark the beginning of a set of regions in a client-side image map, use the MAPOPEN Function.

Syntax

```
HTF.MAPCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</MAP>
```

242.6.81 MAPOPEN Function

This function generates the <MAP> tag which mark the beginning of a set of regions in a client-side image map.

To mark the end of a set of regions in a client-side image map, use the MAPCLOSE Function.

Syntax

```
HTF.MAPOPEN(
   cname          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-62 MAPOPEN Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<MAP NAME="cname" cattributes>
```

242.6.82 MENULISTCLOSE Function

This function generates the </MENU> tag which ends a list that presents one line for each item.

To begin a list of this kind, use the MENULISTOPEN Function. The items in the list appear more compact than an unordered list. The LISTITEM Function defines the list items in a menu list.

Syntax

```
HTF.MENULISTCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</MENU>
```

242.6.83 MENULISTOPEN Function

This function generates the <MENU> tag which begins a list that presents one line for each item.

To end a list of this kind, use the MENULISTCLOSE Function.The items in the list appear more compact than an unordered list. The LISTITEM Function defines the list items in a menu list.

Syntax

```
HTF.MENULISTOPEN
  RETURN VARCHAR2;
```

Examples

This function generates

```
<MENU>
```

242.6.84 META Function

This function generates the <META> tag, which embeds meta-information about the document and also specifies values for HTTP headers. For example, you can specify the expiration date, keywords, and author name.

Syntax

```
HTF.META(
   chttp_equiv    IN       VARCHAR2,
   cname          IN       VARCHAR2,
   ccontent       IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-63 META Function Parameters

Parameter	Description
chttp_equiv	The value for the CHTTP_EQUIV attribute.
cname	The value for the NAME attribute.
ccontent	The value for the CONTENT attribute.

Examples

This function generates

```
<META HTTP-EQUIV="chttp_equiv" NAME ="cname" CONTENT="ccontent">
```

so that

```
HTF.meta ('Refresh', NULL, 120);
```

generates

```
<META HTTP-EQUIV="Refresh" CONTENT=120>
```

On some Web browsers, this causes the current URL to be reloaded automatically every 120 seconds.

242.6.85 NL Function

This function generates the <BR> tag which begins a new line of text.

It performs the same operation as the BR Function.

Syntax

```
HTF.NL(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-64 NL Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<BR CLEAR="cclear" cattributes>
```

242.6.86 NOBR Function

This function generates the <NOBR> and </NOBR> tags which turn off line-breaking in a section of text.

Syntax

```
HTF.NOBR(
ctext         IN         VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-65 NOBR Function Parameters

Parameter	Description
ctext	The text that is to be rendered on one line.

Examples

This function generates

```
<NOBR>ctext</NOBR>
```

242.6.87 NOFRAMESCLOSE Function

This function generates the `</NOFRAMES>` tag which marks the end of a no-frames section.

To mark the beginning of a no-frames section, use the FRAMESETOPEN Function. See also FRAME Function, FRAMESETOPEN Function and FRAMESETCLOSE Function.

Syntax

```
HTF.NOFRAMESCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</NOFRAMES>
```

242.6.88 NOFRAMESOPEN Function

This function generates the `<NOFRAMES>` tag which mark the beginning of a no-frames section.

To mark the end of a no-frames section, use the FRAMESETCLOSE Function. See also FRAME Function, FRAMESETOPEN Function and FRAMESETCLOSE Function.

Syntax

```
HTF.NOFRAMESOPEN
  RETURN VARCHAR2;
```

Examples

This function generates

```
<NOFRAMES>
```

242.6.89 OLISTCLOSE Function

This function generates the `</OL>` tag which defines the end of an ordered list. An ordered list presents a list of numbered items.

To mark the beginning of a list of this kind, use the OLISTOPEN Function. Numbered items are added using LISTITEM Function.

Syntax

```
HTF.OLISTCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</OL>
```

242.6.90 OLISTOPEN Function

This function generates the <OL> tag which marks the beginning of an ordered list. An ordered list presents a list of numbered items.

To mark the end of a list of this kind, use the OLISTCLOSE Function. Numbered items are added using LISTITEM Function.

Syntax

```
HTF.OLISTOPEN(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cwrap           IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-66 OLISTOPEN Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cwrap	The value for the WRAP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<OL CLEAR="cclear" WRAP="cwrap" cattributes>
```

242.6.91 PARA Function

This function generates the <P> tag which indicates that the text that comes after the tag is to be formatted as a paragraph.

You can add attributes to the tag by means of the PARAGRAPH Function.

Syntax

```
HTF.PARA
  RETURN VARCHAR2;
```

Examples

This function generates

```
<P>
```

242.6.92 PARAGRAPH Function

You can use this function to add attributes to the <P> tag created by the PARA Function.

Syntax

```
HTF.PARAGRAPH(
   calign          IN       VARCHAR2   DEFAULT NULL,
   cnowrap         IN       VARCHAR2   DEFAULT NULL,
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-67 PARAGRAPH Function Parameters

Parameter	Description
calign	The value for the ALIGN attribute.
cnowrap	If the value for this parameter is not NULL, the NOWRAP attribute is added to the tag.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<P ALIGN="calign" NOWRAP CLEAR="cclear" cattributes>
```

Related Topics

- PARA Function
  This function generates the <P> tag which indicates that the text that comes after the tag is to be formatted as a paragraph.

242.6.93 PARAM Function

This function generates the <PARAM> tag which specifies parameter values for Java applets.

The values can reference HTML variables. To invoke a Java applet from a Web page, use APPLETOPEN Function to begin the invocation. Use one PARAM Function for each desired name-value pair, and use APPLETCLOSE Function to end the applet invocation.

Syntax

```
HTF.PARAM(
   cname           IN       VARCHAR2
   cvalue          IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-68 PARAM Function Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.

Examples

This function generates

```
<PARAM NAME=cname VALUE="cvalue">
```

242.6.94 PLAINTEXT Function

This function generates the <PLAINTEXT> and </PLAINTEXT> tags which direct the browser to render the text they surround in fixed-width type.

Syntax

```
HTF.PLAINTEXT(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-69 PLAINTEXT Function Parameters

Parameter	Description
ctext	The text to be rendered in fixed-width font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<PLAINTEXT cattributes>ctext</PLAINTEXT>
```

242.6.95 PRECLOSE Function

This function generates the </PRE> tag which marks the end of a section of preformatted text in the body of the HTML page.

To mark the beginning of a section of preformatted text in the body of the HTML page, use the PREOPEN Function.

Syntax

```
HTF.PRECLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</PRE>
```

242.6.96 PREOPEN Function

This function generates the `<PRE>` tag which marks the beginning of a section of preformatted text in the body of the HTML page.

To mark the end of a section of preformatted text in the body of the HTML page, use the PRECLOSE Function.

Syntax

```
HTF.PREOPEN(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cwidth          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-70 PREOPEN Function Parameters

Parameter	Description
`cclear`	The value for the `CLEAR` attribute.
`cwidth`	The value for the `WIDTH` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<PRE CLEAR="cclear" WIDTH="cwidth" cattributes>
```

242.6.97 PRINT Functions

These functions generate the specified parameter as a string terminated with the `\n` newline character.

The PRN Functions performs the same operation but does not terminate with a newline character.

Syntax

```
HTF.PRINT (
   cbuf      IN       VARCHAR2)
  RETURN VARCHAR2;

HTF.PRINT (
   dbuf      IN       DATE)
  RETURN VARCHAR2;

HTF.PRINT (
```

```
   nbuf      IN       NUMBER)
  RETURN VARCHAR2;
```

Parameters

Table 242-71 PRINT Function Parameters

Parameter	Description
cbuf	The string to generate terminated by a newline.
dbuf	The string to generate terminated by a newline.
nbuf	The string to generate terminated by a newline.

Usage Notes

- The \n character is not the same as <BR>. The \n character formats the HTML source but it does not affect how the browser renders the HTML source. Use <BR> to control how the browser renders the HTML source.
- These functions do not have function equivalents.

242.6.98 PRN Functions

These functions generate the specified parameter as a string.

Unlike the PRINT Functions the string is not terminated with the \n newline character.

Syntax

```
HTF.PRN (
   cbuf      IN       VARCHAR2)
  RETURN VARCHAR2;

HTF.PRN (
   dbuf      IN       DATE)
  RETURN VARCHAR2;

HTF.PRN (
   nbuf      IN       NUMBER)
  RETURN VARCHAR2;
```

Parameters

Table 242-72 PRN Function Parameters

Parameter	Description
cbuf	The string to generate (not terminated by a newline).
dbuf	The string to generate (not terminated by a newline).
nbuf	The string to generate (not terminated by a newline).

Usage Notes

These functions do not have function equivalents.

242.6.99 S Function

This function generates the <S> and </S> tags which direct the browser to render the text they surround in strikethrough type.

This performs the same operation as STRIKE Function.

Syntax

```
HTF.S (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-73 S Function Parameters

Parameter	Description
ctext	The text to be rendered in strikethrough type.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<S cattributes>ctext</S>
```

242.6.100 SAMPLE Function

This function generates the <SAMP> and </SAMP> tags which direct the browser to render the text they surround in monospace font or however "sample" is defined stylistically.

Syntax

```
HTF.SAMPLE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-74 SAMPLE Function Parameters

Parameter	Description
ctext	The text to be rendered in monospace font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<SAMP cattributes>ctext</SAMP>
```

242.6.101 SCRIPT Function

This function generates the <SCRIPT> and </SCRIPT> tags which contain a script written in languages such as JavaScript and VBscript.

Syntax

```
HTF.SCRIPT (
   cscript        IN       VARCHAR2,
   clanguage      IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-75 SCRIPT Function Parameters

Parameter	Description
cscript	The text of the script. This is the text that makes up the script itself, not the name of a file containing the script.
clanguage	The language in which the script is written. If this parameter is omitted, the user's browser determines the scripting language.

Examples

This function generates

```
<SCRIPT LANGUAGE=clanguage>cscript</SCRIPT>
```

so that

```
HTF.script ('Erupting_Volcano', 'Javascript');
```

generates

```
<SCRIPT LANGUAGE=Javascript>"script text here"</SCRIPT>
```

This causes the browser to run the script enclosed in the tags.

242.6.102 SMALL Function

This function generates the <SMALL> and </SMALL> tags, which direct the browser to render the text they surround using a small font.

Syntax

```
HTF.SMALL (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-76 SMALL Function Parameters

Parameter	Description
ctext	The text to be rendered in small font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<SMALL cattributes>ctext</SMALL>
```

242.6.103 STRIKE Function

This function generates the <STRIKE> and </STRIKE> tags which direct the browser to render the text they surround in strikethrough type.

This performs the same operation as S Function.

Syntax

```
STRIKE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-77 STRIKE Function Parameters

Parameter	Description
ctext	The text to be rendered in strikethrough type.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<STRIKE cattributes>ctext</STRIKE>
```

242.6.104 STRONG Function

This function generates the <STRONG> and </STRONG> tags which direct the browser to render the text they surround in bold, or however "strong" is defined.

Syntax

```
HTF.STRONG(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-78 STRONG Function Parameters

Parameter	Description
ctext	The text to be emphasized.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<STRONG cattributes>ctext</STRONG>
```

242.6.105 STYLE Function

This function generates the <STYLE> and </STYLE> tags which include a style sheet in a Web page.

You can get more information about style sheets at http://www.w3.org. This feature is not compatible with browsers that support only HTML versions 2.0 or earlier. Such browsers will ignore this tag.

Syntax

```
HTF.STYLE(
   cstyle          IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-79 STYLE Function Parameters

Parameter	Description
cstyle	The style information to include.

Examples

This function generates

```
<STYLE>cstyle</STYLE>
```

242.6.106 SUB Function

This function generates the _{and} tags which direct the browser to render the text they surround as subscript.

Syntax

```
HTF.SUB(
   ctext          IN       VARCHAR2,
   calign         in       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-80 SUB Function Parameters

Parameter	Description
ctext	The text to render in subscript.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<SUB ALIGN="calign" cattributes>ctext</SUB>
```

242.6.107 SUP Function

This function generates the ^{and} tags which direct the browser to render the text they surround as superscript.

Syntax

```
HTF.SUP(
   ctext          IN       VARCHAR2,
   calign         in       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-81 SUP Function Parameters

Parameter	Description
ctext	The text to render in superscript.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<SUP ALIGN="calign" cattributes>ctext</SUP>
```

242.6.108 TABLECAPTION Function

This function generates the <CAPTION> and </CAPTION> tags which place a caption in an HTML table.

Syntax

```
HTF.TABLECAPTION(
   ccaption       IN       VARCHAR2,
   calign         in       VARCHAR2   DEFAULT NULL,
```

```
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-82 TABLECAPTION Function Parameters

Parameter	Description
ctext	The text for the caption.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<CAPTION ALIGN="calign" cattributes>ccaption</CAPTION>
```

242.6.109 TABLECLOSE Function

This function generates the </TABLE> tag which marks the end of an HTML table.

To define the beginning of an HTML table, use the TABLEOPEN Function.

Syntax

```
HTF.TABLECLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</TABLE>
```

242.6.110 TABLEDATA Function

This function generates the <TD> and </TD> tags which insert data into a cell of an HTML table.

Syntax

```
HTF.TABLEDATA(
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cdp            IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   crowspan       IN       VARCHAR2   DEFAULT NULL,
   ccolspan       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-83 TABLEDATA Function Parameters

Parameter	Description
cvalue	The data for the cell in the table.
calign	The value for the ALIGN attribute.
cdp	The value for the DP attribute.
cnowrap	If the value of this parameter is not NULL, the NOWRAP attribute is added to the tag.
ccolspan	The value for the COLSPAN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TD ALIGN="calign" DP="cdp" ROWSPAN="crowspan" COLSPAN="ccolspan" NOWRAP
cattributes>cvalue</TD>
```

242.6.111 TABLEHEADER Function

This function generates the <TH> and </TH> tags which insert a header cell in an HTML table.

The <TH> tag is similar to the <TD> tag except that the text in this case the rows are usually rendered in bold type.

Syntax

```
HTF.TABLEHEADER(
   cvalue          IN       VARCHAR2   DEFAULT NULL,
   calign          IN       VARCHAR2   DEFAULT NULL,
   cdp             IN       VARCHAR2   DEFAULT NULL,
   cnowrap         IN       VARCHAR2   DEFAULT NULL,
   crowspan        IN       VARCHAR2   DEFAULT NULL,
   ccolspan        IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-84 TABLEHEADER Function Parameters

Parameter	Description
cvalue	The data for the cell in the table.
calign	The value for the ALIGN attribute.
cdp	The value for the DP attribute.
cnowrap	If the value of this parameter is not NULL, the NOWRAP attribute is added to the tag.
crispen	The value for the ROWSPAN attribute.
ccolspan	The value for the COLSPAN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TH ALIGN="calign" DP="cdp" ROWSPAN="crowspan" COLSPAN="ccolspan" NOWRAP
cattributes>cvalue</TH>
```

242.6.112 TABLEOPEN Function

This function generates the <TABLE> tag which marks the beginning of an HTML table.

To define the end of an HTML table, use the TABLECLOSE Function.

Syntax

```
HTF.TABLEOPEN(
   cborder        IN       VARCHAR2   DEFAULT NULL
   calign         IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-85 TABLEOPEN Function Parameters

Parameter	Description
`border`	The value for the `BORDER` attribute.
`calign`	The value for the `ALIGN` attribute.
`cnowrap`	If the value of this parameter is not `NULL`, the `NOWRAP` attribute is added to the tag.
`cclear`	The value for the `CLEAR` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TABLE "cborder" NOWRAP ALIGN="calign" CLEAR="cclear" cattributes>
```

242.6.113 TABLEROWCLOSE Function

This function generates the `</TR>` tag which marks the end of a new row in an HTML table.

To mark the beginning of a new row, use the TABLEROWOPEN Function.

Syntax

```
HTF.TABLEROWCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</TABLE>
```

242.6.114 TABLEROWOPEN Function

This function generates the <TR> tag which marks the beginning of a new row in an HTML table.

To mark the end of a new row, use the TABLEROWCLOSE Function.

Syntax

```
HTF.TABLEROWOPEN(
   calign          IN       VARCHAR2   DEFAULT NULL,
   cvalign         IN       VARCHAR2   DEFAULT NULL,
   cdp             IN       VARCHAR2   DEFAULT NULL,
   cnowrap         IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-86 TABLEROWOPEN Function Parameters

Parameter	Description
calign	The value for the ALIGN attribute.
cvalign	The value for the VALIGN attribute.
cdp	The value for the DP attribute.
cnowrap	If the value of this parameter is not NULL, the NOWRAP attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<<TR ALIGN="calign" VALIGN="cvalign" DP="cdp" NOWRAP cattributes>
```

242.6.115 TELETYPE Function

This function generates the <TT> and </TT> tags which direct the browser to render the text they surround in a fixed width typewriter font, for example, the courier font.

Syntax

```
HTF.TELETYPE(
   ctext           IN       VARCHAR2,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-87 TELETYPE Function Parameters

Parameter	Description
ctext	The text to render in a fixed width typewriter font.

Table 242-87 (Cont.) TELETYPE Function Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<TT cattributes>ctext</TT>
```

242.6.116 TITLE Function

This function generates the <TITLE> and </TITLE> tags which specify the text to display in the titlebar of the browser window.

Syntax

```
HTF.TITLE(
   ctitle          IN       VARCHAR2)
  RETURN VARCHAR2;
```

Parameters

Table 242-88 TITLE Function Parameters

Parameter	Description
ctitle	The text to display in the titlebar of the browser window.

Examples

This function generates

```
<TITLE>ctitle</TITLE>
```

242.6.117 ULISTCLOSE Function

This function generates the </UL> tag which marks the end of an unordered list. An unordered list presents items with bullets.

To mark the beginning of an unordered list, use the ULISTOPEN Function. Add list items with LISTITEM Function.

Syntax

```
HTF.ULISTCLOSE
  RETURN VARCHAR2;
```

Examples

This function generates

```
</UL>
```

242.6.118 ULISTOPEN Function

This function generates the <UL> tag which marks the beginning of an unordered list. An unordered list presents items with bullets.

To mark the end of an unordered list, use the ULISTCLOSE Function. Add list items with LISTITEM Function.

Syntax

```
HTF.ULISTOPEN(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cwrap           IN       VARCHAR2   DEFAULT NULL,
   cdingbat        IN       VARCHAR2   DEFAULT NULL,
   csrc            IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-89 ULISTOPEN Function Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cwrap	The value for the WRAP attribute.
cdingbat	The value for the DINGBAT attribute.
csrc	The value for the SRC attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<UL CLEAR="cclear" WRAP="cwrap" DINGBAT="cdingbat" SRC="csrc" cattributes>
```

242.6.119 UNDERLINE Function

This function generates the <U> and </U> tags, which direct the browser to render the text they surround with an underline.

Syntax

```
HTF.UNDERLINE(
   ctext           IN       VARCHAR2,
   cattributes     IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-90 UNDERLINE Function Parameters

Parameter	Description
ctext	The text to render with an underline.

Table 242-90 (Cont.) UNDERLINE Function Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<U cattributes>ctext</U>
```

242.6.120 VARIABLE Function

This function generates the <VAR> and </VAR> tags which direct the browser to render the text they surround in italics or however "variable" is defined stylistically.

Syntax

```
HTF.VARIABLE(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 242-91 VARIABLE Function Parameters

Parameter	Description
ctext	The text to render in italics.
cattributes	The other attributes to be included as-is in the tag.

Examples

This function generates

```
<VAR cattributes>ctext</VAR>
```

242.6.121 WBR Function

This function generates the <WBR> tag, which inserts a soft line break within a section of NOBR text.

Syntax

```
HTF.WBR
  RETURN VARCHAR2;
```

Examples

This function generates

```
<WBR>
```

243
HTP

The `HTP` (hypertext procedures) and `HTF` (hypertext functions) packages generate HTML tags. For example, the `HTP.ANCHOR` procedure generates the HTML anchor tag, `<A>`.

This chapter contains the following topics:

- Operational Notes
- Rules and Limits
- Examples
- Summary of Tags
- Summary of HTP Subprograms

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

243.1 HTP Operational Notes

Certain operational notes apply to the `HTP` package.

For every HTP procedure that generates one or more HTML tags, there is a corresponding HTF function with identical parameters with the following exception:

- The PRINTS Procedure and the PS Procedure do not have HTF function equivalents. Use the ESCAPE_SC Function or the ESCAPE_URL Function if you need a string conversion function. Note that while there is a ESCAPE_SC Procedure that performs the same operation as the PRINTS Procedure and the PS Procedure, there is no procedural equivalent for the ESCAPE_URL Function.
- The FORMAT_CELL Function does not have an HTP equivalent. The function formats column values inside an HTML table using TABLEDATA Function which does have an HTP equivalent in the TABLEDATA Procedure. The advantage of this using the FORMAT_CELL Function is that it allows for better control over the HTML tables.

The function versions do not directly generate output in your Web page. Instead, they pass their output as return values to the statements that invoked them. Use these functions when you need to nest calls. To print the output of HTF functions, call the functions from within the HTP.`PRINT` procedure. It then prints its parameters to the generated Web page.

243.2 HTP Rules and Limits

If you use values of the LONG datatype in procedures such as `HTP.PRINT`, `HTP.PRN`, `HTP.PRINTS`, `HTP.PA` or `OWA_UTIL.CELLSPRINT`, only the first 32 K of the `LONG` data is used. The `LONG` data is bound to a `VARCHAR2` datatype in the procedure.

243.3 HTP Examples

These commands generate a simple HTML document.

```
CREATE OR REPLACE PROCEDURE hello AS
BEGIN
    HTP.HTMLOPEN;           -- generates <HTML>
    HTP.HEADOPEN;           -- generates <HEAD>
    HTP.TITLE('Hello');     -- generates <TITLE>Hello</TITLE>
    HTP.HEADCLOSE;          -- generates </HEAD>
    HTP.BODYOPEN;           -- generates <BODY>
    HTP.HEADER(1, 'Hello'); -- generates <H1>Hello</H1>
    HTP.BODYCLOSE;          -- generates </BODY>
    HTP.HTMLCLOSE;          -- generates </HTML>
END;
```

243.4 HTP Summary of Tags

`HTP` tags can be grouped by type, to make understanding them easier.

HTML, HEAD, and BODY Tags

HTMLOPEN Procedure, HTMLCLOSE Procedure - generate `<HTML>` and `</HTML>`

HEADOPEN Procedure, HEADCLOSE Procedure - generate `<HEAD>` and `</HEAD>`

BODYOPEN Procedure, BODYCLOSE Procedure - generate `<BODY>` and `</BODY>`

Comment Tag

COMMENT Procedure - generates <!-- and -->

Tags in the <HEAD> Area

BASE Procedure - generates `<BASE>`

LINKREL Procedure - generates `<LINK>` with the `REL` attribute

LINKREV Procedure - generates `<LINK>` with the `REV` attribute

TITLE Procedure - generates `<TITLE>`

META Procedure - generates `<META>`

SCRIPT Procedure - generates `<SCRIPT>`

STYLE Procedure - generates `<STYLE>`

ISINDEX Procedure - generates `<ISINDEX>`

Applet Tags

APPLETOPEN Procedure, APPLETCLOSE Procedure - generate `<APPLET>` and `</APPLET>`

PARAM Procedure - generates `<PARAM>`

List Tags

OLISTOPEN Procedure, OLISTCLOSE Procedure - generate `<OL>` and `</OL>`

ULISTOPEN Procedure, ULISTCLOSE Procedure - generate `<UL>` and `</UL>`

DLISTOPEN Procedure, DLISTCLOSE Procedure- generate `<DL>` and `</DL>`

DLISTTERM Procedure - generates `<DT>`

DLISTDEF Procedure - generates `<DD>`

DIRLISTOPEN Procedure, DIRLISTCLOSE Procedure - generate `<DIR>` and `</DIR>`

LISTHEADER Procedure - generates `<LH>`

LISTINGOPEN Procedure, LISTINGCLOSE Procedure - generate `<LISTING>` and `</LISTING>`

MENULISTOPEN Procedure - generate `<MENU>` and `</MENU>`

LISTITEM Procedure - generates `<LI>`

Form Tags

FORMOPEN Procedure, FORMCLOSE Procedure - generate `<FORM>` and `</FORM>`

FORMCHECKBOX Procedure - generates `<INPUT TYPE="CHECKBOX">`

FORMHIDDEN Procedure - generates `<INPUT TYPE="HIDDEN">`

FORMIMAGE Procedure - generates `<INPUT TYPE="IMAGE">`

FORMPASSWORD Procedure - generates `<INPUT TYPE="PASSWORD">`

FORMRADIO Procedure - generates `<INPUT TYPE="RADIO">`

FORMSELECTOPEN Procedure, FORMSELECTCLOSE Procedure - generate `<SELECT>` and `</SELECT>`

FORMSELECTOPTION Procedure - generates `<OPTION>`

FORMTEXT Procedure - generates `<INPUT TYPE="TEXT">`

FORMTEXTAREA Procedure - generate `<TEXTAREA>`

FORMTEXTAREAOPEN Procedure, FORMTEXTAREACLOSE Procedure - generate `<TEXTAREA>` and `</TEXTAREA>`

FORMRESET Procedure - generates `<INPUT TYPE="RESET">`

FORMSUBMIT Procedure - generates `<INPUT TYPE="SUBMIT">`

Table Tags

TABLEOPEN Procedure, TABLECLOSE Procedure - generate `<TABLE>` and `</TABLE>`

TABLECAPTION Procedure - generates <CAPTION>

TABLEROWOPEN Procedure, TABLEROWCLOSE Procedure - generate <TR> and </TR>

TABLEHEADER Procedure - generates <TH>

TABLEDATA Procedure - generates <TD>

IMG, HR, and A Tags

HR Procedure, LINE Procedure - generate <HR>

IMG Procedure, IMG2 Procedure - generate <IMG>

ANCHOR Procedure, ANCHOR2 Procedure - generate <A>

MAPOPEN Procedure, MAPCLOSE Procedure - generate <MAP> and </MAP>

Paragraph Formatting Tags

HEADER Procedure - generates heading tags (<H1> to <H6>)

PARA Procedure, PARAGRAPH Procedure - generate <P>

PRN Procedures, PRINT Procedures - generate any text that is passed in

PRINTS Procedure, PS Procedure - generate any text that is passed in; special characters in HTML are escaped

PREOPEN Procedure, PRECLOSE Procedure - generate <PRE> and </PRE>

BLOCKQUOTEOPEN Procedure, BLOCKQUOTECLOSE Procedure - generate <BLOCKQUOTE> and </BLOCKQUOTE>

DIV Procedure - generates <DIV>

NL Procedure, BR Procedure - generate

NOBR Procedure - generates <NOBR>

WBR Procedure - generates <WBR>

PLAINTEXT Procedure - generates <PLAINTEXT>

ADDRESS Procedure - generates <ADDRESS>

MAILTO Procedure - generates <A> with the MAILTO attribute

AREA Procedure - generates <AREA>

BGSOUND Procedure - generates <BGSOUND>

Character Formatting Tags

BASEFONT Procedure - generates <BASEFONT>

BIG Procedure - generates <BIG>

BOLD Procedure - generates <B>

CENTER Procedure - generates <CENTER> and </CENTER>

CENTEROPEN Procedure, CENTERCLOSE Procedure - generate <CENTER> and </CENTER>

CITE Procedure - generates <CITE>

CODE Procedure - generates <CODE>

DFN Procedure - generates <DFN>

EM Procedure, EMPHASIS Procedure - generate <EM>

FONTOPEN Procedure, FONTCLOSE Procedure - generate <FONT> and </FONT>

ITALIC Procedure - generates <I>

KBD Procedure, KEYBOARD Procedure - generate <KBD> and </KBD>

S Procedure - generates <S>

SAMPLE Procedure - generates <SAMP>

SMALL Procedure - generates <SMALL>

STRIKE Procedure - generates <STRIKE>

STRONG Procedure - generates <STRONG>

SUB Procedure - generates <SUB>

SUP Procedure - generates <SUP>

TELETYPE Procedure - generates <TT>

UNDERLINE Procedure - generates <U>

VARIABLE Procedure - generates <VAR>

Frame Tags

FRAME Procedure - generates <FRAME>

FRAMESETOPEN Procedure, FRAMESETCLOSE Procedure - generate <FRAMESET> and </FRAMESET>

NOFRAMESOPEN Procedure, NOFRAMESCLOSE Procedure - generate <NOFRAMES> and </NOFRAMES>

243.5 Summary of HTP Subprograms

This table lists the HTP subprograms and briefly describes them.

Table 243-1 HTP Package Subprograms

Subprogram	Description
ADDRESS Procedure	Generates the <ADDRESS> and </ADDRESS> tags which specify the address, author and signature of a document
ANCHOR Procedure	Generates the <A> and </A> tags which specify the source or destination of a hypertext link
ANCHOR2 Procedure	Generates the <A> and </A> tags which specify the source or destination of a hypertext link
APPLETCLOSE Procedure	Closes the applet invocation with the </APPLET> tag

Table 243-1 (Cont.) HTP Package Subprograms

Subprogram	Description
APPLETOPEN Procedure	Generates the <APPLET> tag which begins the invocation of a Java applet
AREA Procedure	Generates the <AREA> tag, which defines a client-side image map
BASE Procedure	Generates the <BASE> tag which records the URL of the document
BASEFONT Procedure	Generates the <BASEFONT> tag which specifies the base font size for a Web page
BGSOUND Procedure	Generates the <BGSOUND> tag which includes audio for a Web page
BIG Procedure	Generates the <BIG> and </BIG> tags which direct the browser to render the text in a bigger font
BLOCKQUOTECLOSE Procedure	Generates the </BLOCKQUOTE> tag which mark the end of a section of quoted text
BLOCKQUOTEOPEN Procedure	Generates the <BLOCKQUOTE> tag, which marks the beginning of a section of quoted text
BODYCLOSE Procedure	Generates the </BODY> tag which marks the end of a body section of an HTML document
BODYOPEN Procedure	Generates the <BODY> tag which marks the beginning of the body section of an HTML document
BOLD Procedure	Generates the <B> and </B> tags which direct the browser to display the text in boldface
BR Procedure	Generates the tag which begins a new line of text
CENTER Procedure	Generates the <CENTER> and </CENTER> tags which center a section of text within a Web page
CENTERCLOSE Procedure	Generates the </CENTER> tag which marks the end of a section of text to center
CENTEROPEN Procedure	Generates the <CENTER> tag which mark the beginning of a section of text to center
CITE Procedure	Generates the <CITE> and </CITE> tags which direct the browser to render the text as a citation
CODE Procedure	Generates the <CODE> and </CODE> tags which direct the browser to render the text in monospace font or however "code" is defined stylistically
COMMENT Procedure	Generates procedure generates the comment tags <!-- ctext -->
DFN Procedure	Generates the <DFN> and </DFN> tags which direct the browser to mark the text as italics or however "definition" is defined stylistically
DIRLISTCLOSE Procedure	Generates the </DIR> tag which ends a directory list section
DIRLISTOPEN Procedure	Generates the <DIR> which starts a directory list section
DIV Procedure	Generates the <DIV> tag which creates document divisions
DLISTCLOSE Procedure	Generates the </DL> tag which ends a definition list
DLISTDEF Procedure	Generates the <DD> tag, which inserts definitions of terms
DLISTOPEN Procedure	Generates the <DL> tag which starts a definition list
DLISTTERM Procedure	Generates the <DT> tag which defines a term in a definition list <DL>

Table 243-1 (Cont.) HTP Package Subprograms

Subprogram	Description
EM Procedure	Generates the <EM> and </EM> tags, which define text to be emphasized
EMPHASIS Procedure	Generates the <EM> and </EM> tags, which define text to be emphasized
ESCAPE_SC Procedure	Replaces characters that have special meaning in HTML with their escape sequences
FONTCLOSE Procedure	Generates the </FONT> tag which marks the end of a section of text with the specified font characteristics
FONTOPEN Procedure	Generates the <FONT> which marks the beginning of section of text with the specified font characteristics
FORMCHECKBOX Procedure	Generates the <INPUT> tag with TYPE="checkbox" which inserts a checkbox element in a form
FORMCLOSE Procedure	Generates the </FORM> tag which marks the end of a form section in an HTML document
FORMOPEN Procedure	Generates the <FORM> tag which marks the beginning of a form section in an HTML document
FORMFILE Procedure	Generates the <INPUT> tag with TYPE="file" which inserts a file form element, and is used for file uploading for a given page
FORMHIDDEN Procedure	Generates the <INPUT> tag with TYPE="hidden"which inserts a hidden form element
FORMIMAGE Procedure	Generates the <INPUT> tag with TYPE="image" which creates an image field that the user clicks to submit the form immediately
FORMPASSWORD Procedure	Generates the <INPUT> tag with TYPE="password" which creates a single-line text entry field
FORMRADIO Procedure	Generates the <INPUT> tag with TYPE="radio", which creates a radio button on the HTML form
FORMRESET Procedure	Generates the <INPUT> tag with TYPE="reset" which creates a button that, when selected, resets the form fields to their initial values
FORMSELECTCLOSE Procedure	Generates the </SELECT> tag which marks the end of a Select form element
FORMSELECTOPEN Procedure	Generates the </SELECT> tag which marks the beginning of a Select form element
FORMSELECTOPTION Procedure	Generates the <OPTION> tag which represents one choice in a Select element
FORMSUBMIT Procedure	Generates the <INPUT> tag with TYPE="submit" which creates a button that, when clicked, submits the form
FORMTEXT Procedure	Generates the <INPUT> tag with TYPE="text", which creates a field for a single line of text
FORMTEXTAREA Procedure	Generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area
FORMTEXTAREA2 Procedure	Generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area with the ability to specify a wrap style

Table 243-1 (Cont.) HTP Package Subprograms

Subprogram	Description
FORMTEXTAREACLOSE Procedure	Generates the </TEXTAREA> tag which ends a text area form element
FORMTEXTAREAOPEN Procedure	Generates the <TEXTAREA> which marks the beginning of a text area form element
FORMTEXTAREAOPEN2 Procedure	Generates the <TEXTAREA> which marks the beginning of a text area form element with the ability to specify a wrap style
FRAME Procedure	Generates the <FRAME> tag which begins the characteristics of a frame created by a <FRAMESET> tag
FRAMESETCLOSE Procedure	Generates the </FRAMESET> tag which ends a frameset section
FRAMESETOPEN Procedure	Generates the </FRAMESET> tag which begins a frameset section
HEADCLOSE Procedure	Generates the </HEAD> tag which marks the end of an HTML document head section
HEADER Procedure	Generates opening heading tags (<H1> to <H6>) and their corresponding closing tags (</H1> to </H6>)
HEADOPEN Procedure	Generates the <HEAD> tag which marks the beginning of the HTML document head section
HR Procedure	Generates the <HR> tag, which generates a line in the HTML document
HTMLCLOSE Procedure	Generates the </HTML> tag which marks the end of an HTML document
HTMLOPEN Procedure	Generates the <HTML> tag which marks the beginning of an HTML document
IMG Procedure	Generates the <IMG> tag which directs the browser to load an image onto the HTML page
IMG2 Procedure	Generates the <IMG> tag which directs the browser to load an image onto the HTML page with the option of specifying values for the USEMAP attribute
ISINDEX Procedure	Creates a single entry field with a prompting text, such as "*enter value,*" then sends that value to the URL of the page or program
ITALIC Procedure	Generates the <I> and </I> tags which direct the browser to render the text in italics
KBD Procedure	Generates the <KBD> and </KBD> tags which direct the browser to render the text in monospace font
KEYBOARD Procedure	Generates the <KBD> and </KBD> tags, which direct the browser to render the text in monospace font
LINE Procedure	Generates the <HR> tag, which generates a line in the HTML document
LINKREL Procedure	Generates the <LINK> tag with the REL attribute which delineates the relationship described by the hypertext link from the anchor to the target
LINKREV Procedure	Generates the <LINK> tag with the REV attribute which delineates the relationship described by the hypertext link from the target to the anchor
LISTHEADER Procedure	Generates the <LH> and </LH> tags which print an HTML tag at the beginning of the list

Table 243-1 (Cont.) HTP Package Subprograms

Subprogram	Description
LISTINGCLOSE Procedure	Generates the </LISTING> tags which marks the end of a section of fixed-width text in the body of an HTML page
LISTINGOPEN Procedure	Generates the <LISTING> tag which marks the beginning of a section of fixed-width text in the body of an HTML page
LISTITEM Procedure	Generates the <LI> tag, which indicates a list item
MAILTO Procedure	Generates the <A> tag with the HREF set to 'mailto' prepended to the mail address argument
MAPCLOSE Procedure	Generates the </MAP> tag which marks the end of a set of regions in a client-side image map
MAPOPEN Procedure	Generates the <MAP> tag which mark the beginning of a set of regions in a client-side image map
MENULISTCLOSE Procedure	Generates the </MENU> tag which ends a list that presents one line for each item
MENULISTOPEN Procedure	Generates the <MENU> tag which begins a list that presents one line for each item
META Procedure	Generates the <META> tag, which embeds meta-information about the document and also specifies values for HTTP headers
NL Procedure	Generates the tag which begins a new line of text
NOBR Procedure	Generates the <NOBR> and </NOBR> tags which turn off line-breaking in a section of text
NOFRAMESCLOSE Procedure	Generates the </NOFRAMES> tag which marks the end of a no-frames section
NOFRAMESOPEN Procedure	Generates the <NOFRAMES> tag which mark the beginning of a no-frames section
OLISTCLOSE Procedure	Generates the </OL> tag which defines the end of an ordered list
OLISTOPEN Procedure	Generates the <OL> tag which marks the beginning of an ordered list
PARA Procedure	Generates the <P> tag which indicates that the text that comes after the tag is to be formatted as a paragraph
PARAGRAPH Procedure	Adds attributes to the <P> tag
PARAM Procedure	Generates the <PARAM> tag which specifies parameter values for Java applets
PLAINTEXT Procedure	Generates the <PLAINTEXT> and </PLAINTEXT> tags which direct the browser to render the text they surround in fixed-width type
PRECLOSE Procedure	Generates the </PRE> tag which marks the end of a section of preformatted text in the body of the HTML page
PREOPEN Procedure	Generates the <PRE> tag which marks the beginning of a section of preformatted text in the body of the HTML page
PRINT Procedures	Generates the specified parameter as a string terminated with the \n newline character
PRINTS Procedure	Generates a string and replaces the following characters with the corresponding escape sequence
PRN Procedures	Generates the specified parameter as a string

Table 243-1 (Cont.) HTP Package Subprograms

Subprogram	Description
PS Procedure	Generates a string and replaces the following characters with the corresponding escape sequence.
S Procedure	Generates the <S> and </S> tags which direct the browser to render the text they surround in strikethrough type
SAMPLE Procedure	Generates the <SAMP> and </SAMP> tags which direct the browser to render the text they surround in monospace font or however "sample" is defined stylistically
SCRIPT Procedure	Generates the <SCRIPT> and </SCRIPT> tags which contain a script written in languages such as JavaScript and VBscript
SMALL Procedure	Generates the <SMALL> and </SMALL> tags, which direct the browser to render the text they surround using a small font
STRIKE Procedure	Generates the <STRIKE> and </STRIKE> tags which direct the browser to render the text they surround in strikethrough type
STRONG Procedure	Generates the <STRONG> and </STRONG> tags which direct the browser to render the text they surround in bold or however "strong" is defined stylistically
STYLE Procedure	Generates the <STYLE> and </STYLE> tags which include a style sheet in a Web page
SUB Procedure	Generates the _{and} tags which direct the browser to render the text they surround as subscript
SUP Procedure	Generates the ^{and} tags which direct the browser to render the text they surround as superscript
TABLECAPTION Procedure	Generates the <CAPTION> and </CAPTION> tags which place a caption in an HTML table
TABLECLOSE Procedure	Generates the </TABLE> tag which marks the end of an HTML table
TABLEDATA Procedure	Generates the <TD> and </TD> tags which insert data into a cell of an HTML table
TABLEHEADER Procedure	Generates the <TH> and </TH> tags which insert a header cell in an HTML table.
TABLEOPEN Procedure	Generates the <TABLE> tag which marks the beginning of an HTML table
TABLEROWCLOSE Procedure	Generates the </TR> tag which marks the end of a new row in an HTML table
TABLEROWOPEN Procedure	Generates the <TR> tag which marks the beginning of a new row in an HTML table
TELETYPE Procedure	Generates the <TT> and </TT> tags which direct the browser to render the text they surround in a fixed width typewriter font, for example, the courier font
TITLE Procedure	Generates the <TITLE> and </TITLE> tags which specify the text to display in the titlebar of the browser window
ULISTCLOSE Procedure	Generates the </UL> tag which marks the end of an unordered list
ULISTOPEN Procedure	Generates the <UL> tag which marks the beginning of an unordered list

Table 243-1 (Cont.) HTP Package Subprograms

Subprogram	Description
UNDERLINE Procedure	Generates the <U> and </U> tags, which direct the browser to render the text they surround with an underline
VARIABLE Procedure	Generates the <VAR> and </VAR> tags which direct the browser to render the text they surround in italics or however "variable" is defined stylistically.
WBR Procedure	Generates the <WBR> tag, which inserts a soft line break within a section of NOBR text

243.5.1 ADDRESS Procedure

This procedure generates the <ADDRESS> and </ADDRESS> tags which specify the address, author and signature of a document.

Syntax

```
HTP.ADDRESS (
   cvalue          IN       VARCHAR2
   cnowrap         IN       VARCHAR2   DEFAULT NULL
   cclear          IN       VARCHAR2   DEFAULT NULL
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-2 ADDRESS Procedure Parameters

Parameter	Description
cvalue	The string that goes between the <ADDRESS> and </ADDRESS> tags.
cnowrap	If the value for this parameter is not NULL, the NOWRAP attribute is included in the tag
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag

Examples

This procedure generates

```
<ADDRESS CLEAR="cclear" NOWRAP cattributes>cvalue</ADDRESS>
```

243.5.2 ANCHOR Procedure

Like the ANCHOR2 procedure, this procedure generates the <A> and </A> HTML tags which specify the source or destination of a hypertext link.

The difference between these subprograms is that the ANCHOR2 Procedure provides a target and therefore can be used for a frame.

Syntax

```
HTP.ANCHOR (
   curl           IN       VARCHAR2,
   ctext          IN       VARCHAR2,
   cname          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-3 ANCHOR Procedure Parameters

Parameter	Description
curl	The value for the HREF attribute.
ctext	The string that goes between the <A> and </A> tags.
cname	The value for the NAME attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<A HREF="curl" NAME="cname" cattributes>ctext</A>
```

Usage Notes

This tag accepts several attributes, but either HREF or NAME is required. HREF specifies to where to link. NAME allows this tag to be a target of a hypertext link.

243.5.3 ANCHOR2 Procedure

Like the ANCHOR procedure, this procedure generates the <A> and </A> HTML tags which specify the source or destination of a hypertext link.

The difference between this procedure and the ANCHOR procedure is that this procedure provides a target and therefore can be used for a frame.

Syntax

```
HTP.ANCHOR2 (
   curl           IN       VARCHAR2,
   ctext          IN       VARCHAR2,
   cname          IN       VARCHAR2   DEFAULT NULL,
   ctarget        in       varchar2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-4 ANCHOR2 Procedure Parameters

Parameter	Description
curl	The value for the HREF attribute.
ctext	The string that goes between the <A> and </A> tags.

Table 243-4 (Cont.) ANCHOR2 Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute
ctarget	The value for the TARGET attribute.
cattributes	The other attributes to be included as-is in the tag

Examples

This procedure generates

```
<A HREF="curl" NAME="cname" TARGET = "ctarget" cattributes>ctext</A>
```

243.5.4 APPLETCLOSE Procedure

This procedure closes the applet invocation with the </APPLET> tag.

You must first invoke the a Java applet using APPLETOPEN Procedure.

Syntax

```
HTP.APPLETCLOSE;
```

243.5.5 APPLETOPEN Procedure

This procedure generates the <APPLET> tag which begins the invocation of a Java applet.

You close the applet invocation with APPLETCLOSE Procedure which generates the </APPLET> tag.

Syntax

```
HTP.APPLETOPEN (
   ccode          IN       VARCHAR2,
   cheight        IN       NUMBER,
   cwidth         IN       NUMBER,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-5 APPLETOPEN Procedure Parameters

Parameter	Description
ccode	The the value for the CODE attribute which specifies the name of the applet class.
cheight	The value for the HEIGHT attribute.
cwidth	The value for the WIDTH attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<APPLET CODE=ccode HEIGHT=cheight WIDTH=cwidth cattributes>
```

so that, for example,

```
HTP.appletopen('testclass.class', 100, 200, 'CODEBASE="/ows-applets"')
```

generates

```
<APPLET CODE="testclass.class" height=100 width=200 CODEBASE="/ows-applets">
```

Usage Notes

- Specify parameters to the Java applet using the PARAM Procedure.
- Use the `cattributes` parameter to specify the `CODEBASE` attribute since the PL/SQL cartridge does not know where to find the class files. The `CODEBASE` attribute specifies the virtual path containing the class files.

243.5.6 AREA Procedure

This procedure generates the `<AREA>` tag, which defines a client-side image map. The `<AREA>` tag defines areas within the image and destinations for the areas.

Syntax

```
HTP.AREA (
   ccoords        IN       VARCHAR2
   cshape         IN       VARCHAR2   DEFAULT NULL,
   chref          IN       VARCHAR2   DEFAULT NULL,
   cnohref        IN       VARCHAR2   DEFAULT NULL,
   ctarget        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-6 AREA Procedure Parameters

Parameter	Description
`ccords`	The the value for the `COORDS` attribute.
`cshape`	The value for the `SHAPE` attribute.
`chref`	The value for the `HREF` attribute.
`cnohref`	If the value for this parameter is not `NULL`, the `NOHREF` attribute is added to the tag.
`ctarget`	The value for the `TARGET` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<AREA COORDS="ccoords" SHAPE="cshape" HREF="chref" NOHREF TARGET="ctarget" cattributes>
```

243.5.7 BASE Procedure

This procedure generates the <BASE> tag which records the URL of the document.

Syntax

```
HTP.BASE (
   ctarget        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-7 BASE Procedure Parameters

Parameter	Description
ctarget	The value for the TARGET attribute which establishes a window name to which all links in this document are targeted.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BASE HREF="<current URL>" TARGET="ctarget" cattributes>
```

243.5.8 BASEFONT Procedure

This procedure generates the <BASEFONT> tag which specifies the base font size for a Web page.

Syntax

```
HTP.BASEFONT (
   nsize    IN    INTEGER);
```

Parameters

Table 243-8 BASEFONT Procedure Parameters

Parameter	Description
nsize	The value for the SIZE attribute.

Examples

This procedure generates

```
<BASEFONT SIZE="nsize">
```

243.5.9 BGSOUND Procedure

This procedure generates the <BGSOUND> tag which includes audio for a Web page.

Syntax

```
HTP.BGSOUND (
   csrc           IN       VARCHAR2,
   cloop          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-9 BGSOUND Procedure Parameters

Parameter	Description
csrc	The value for the SRC attribute.
cloop	The value for the LOOP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BGSOUND SRC="csrc" LOOP="cloop" cattributes>
```

243.5.10 BIG Procedure

This procedure generates the <BIG> and </BIG> tags which direct the browser to render the text in a bigger font.

Syntax

```
HTP.BIG (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-10 BIG Procedure Parameters

Parameter	Description
ctext	The the text that goes between the tags.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BIG cattributes>ctext</BIG>
```

243.5.11 BLOCKQUOTECLOSE Procedure

This procedure generates the `</BLOCKQUOTE>` tag which mark the end of a section of quoted text.

You mark the beginning of a section of text by means of the BLOCKQUOTEOPEN Procedure.

Syntax

```
HTP.BLOCKQUOTECLOSE;
```

Examples

This procedure generates

```
</BLOCKQUOTE>
```

243.5.12 BLOCKQUOTEOPEN Procedure

This procedure generates the `<BLOCKQUOTE>` tag, which marks the beginning of a section of quoted text.

You mark the end of a section of text by means of the BLOCKQUOTECLOSE Procedure.

Syntax

```
HTP.BLOCKQUOTEOPEN (
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-11 BLOCKQUOTEOPEN Procedure Parameters

Parameter	Description
`cnowrap`	If the value for this parameter is not `NULL`, the `NOWRAP` attribute is added to the tag.
`cclear`	The value for the `CLEAR` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BLOCKQUOTE CLEAR="cclear" NOWRAP cattributes>
```

243.5.13 BODYCLOSE Procedure

This procedure generates the `</BODY>` tag which marks the end of a body section of an HTML document.

You mark the beginning of a body section by means of the BODYOPEN Procedure.

Syntax

```
HTP.BODYCLOSE;
```

Examples

This procedure generates

```
</BODY>
```

243.5.14 BODYOPEN Procedure

This procedure generates the <BODY> tag which marks the beginning of the body section of an HTML document.

You mark the end of a body section by means of the BODYCLOSE Procedure.

Syntax

```
HTP.BODYOPEN (
   cbackground    IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-12 BODYOPEN Procedure Parameters

Parameter	Description
cbackground	The value for the BACKGROUND attribute which specifies a graphic file to use for the background of the document.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BODY background="cbackground" cattributes>
```

so that

```
HTP.BODYOPEN('/img/background.gif');
```

generates:

```
<BODY background="/img/background.gif">
```

243.5.15 BOLD Procedure

This procedure generates the <B> and </B> tags which direct the browser to display the text in boldface.

Syntax

```
HTP.BOLD (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-13 BOLD Procedure Parameters

Parameter	Description
ctext	The text that goes between the tags.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<B cattributes>ctext</B>
```

243.5.16 BR Procedure

This procedure generates the <BR> tag which begins a new line of text.

It performs the same operation as the NL Procedure.

Syntax

```
HTP.BR(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-14 BR Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BR CLEAR="cclear" cattributes>
```

243.5.17 CENTER Procedure

This procedure generates the <CENTER> and </CENTER> tags which center a section of text within a Web page.

Syntax

```
HTP.CENTER (
   ctext           IN       VARCHAR2);
```

Parameters

Table 243-15 CENTER Parameters

Parameter	Description
ctext	The text that goes between the tags.

Examples

This procedure generates

```
<CENTER>ctext</CENTER>
```

243.5.18 CENTERCLOSE Procedure

This procedure generates the </CENTER> tag which marks the end of a section of text to center.

You mark the beginning of a section of text to center with the CENTEROPEN Procedure.

Syntax

```
HTP.CENTERCLOSE;
```

Examples

This procedure generates

```
</CENTER>
```

243.5.19 CENTEROPEN Procedure

This procedure generates the <CENTER> tag which mark the beginning of a section of text to center.

You mark the beginning of a of a section of text to center by means of the CENTERCLOSE Procedure.

Syntax

```
HTP.CENTEROPEN;
```

Examples

This procedure generates

```
<CENTER>
```

243.5.20 CITE Procedure

This procedure generates the <CITE> and </CITE> tags which direct the browser to render the text as a citation.

Syntax

```
HTP.CITE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-16 CITE Procedure Parameters

Parameter	Description
ctext	The text to render as citation.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<CITE cattributes>ctext</CITE>
```

243.5.21 CODE Procedure

This procedure generates the <CODE> and </CODE> tags which direct the browser to render the text in monospace font or however "code" is defined stylistically.

Syntax

```
HTP.CODE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-17 CODE Procedure Parameters

Parameter	Description
ctext	The text to render as code.
cattributes	The other attributes to be included as-is in the tag

Examples

This procedure generates

```
<CODE cattributes>ctext</CODE>
```

243.5.22 COMMENT Procedure

This procedure generates the comment tags.

Syntax

```
HTP.COMMENT (
   ctext          IN       VARCHAR2);
```

Parameters

Table 243-18 COMMENT Procedure Parameters

Parameter	Description
ctext	The comment.

Examples

This procedure generates

```
<!-- ctext -->
```

243.5.23 DFN Procedure

This procedure generates the <DFN> and </DFN> tags which direct the browser to mark the text in italics or however "definition" is described stylistically.

Syntax

```
HTP.DFN (
   ctext          IN       VARCHAR2);
```

Parameters

Table 243-19 DFN Procedure Parameters

Parameter	Description
ctext	The text to render in italics.

Examples

This procedure generates

```
<DFN>ctext</DFN>
```

243.5.24 DIRLISTCLOSE Procedure

This procedure generates the </DIR> tag which ends a directory list section. You start a directory list section with the DIRLISTOPEN Procedure.

Syntax

```
HTP.DIRLISTCLOSE;
```

Usage Notes

A directory list presents a list of items that contains up to 20 characters. Items in this list are typically arranged in columns, 24 characters wide. Insert the <LI> tag directly or invoke the LISTITEM Procedure so that the <LI> tag appears directly after the </DIR> tag to define the items as a list.

Examples

This procedure generates

```
</DIR>
```

Related Topics

- DIRLISTOPEN Procedure
  This procedure generates the <DIR> which starts a directory list section. You end a directory list section with the DIRLISTCLOSE Procedure.

243.5.25 DIRLISTOPEN Procedure

This procedure generates the <DIR> which starts a directory list section. You end a directory list section with the DIRLISTCLOSE Procedure.

Syntax

```
HTP.DIRLISTOPEN;
```

Usage Notes

A directory list presents a list of items that contains up to 20 characters. Items in this list are typically arranged in columns, 24 characters wide. Insert the <LI> tag directly or invoke the LISTITEM Procedure so that the <LI> tag appears directly after the </DIR> tag to define the items as a list.

Examples

This procedure generates

```
<DIR>
```

Related Topics

- DIRLISTCLOSE Procedure
  This procedure generates the </DIR> tag which ends a directory list section. You start a directory list section with the DIRLISTOPEN Procedure.

243.5.26 DIV Procedure

This procedure generates the <DIV> tag which creates document divisions.

Syntax

```
HTP.DIV (
   calign         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-20 DIV Procedure Parameters

Parameter	Description
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<DIV ALIGN="calign" cattributes>
```

243.5.27 DLISTCLOSE Procedure

This procedure generates the </DL> tag which ends a definition list. You start a definition list by means of the DLISTOPEN Procedure.

Syntax

```
HTP.DLISTCLOSE;
```

Usage Notes

A definition list looks like a glossary: it contains terms and definitions. Terms are inserted using the DLISTTERM Procedure and definitions are inserted using the DLISTDEF Procedure.

Examples

This procedure generates

```
</DL>
```

Related Topics

- DLISTOPEN Procedure
  This procedure generates the <DL> tag which starts a definition list. You end a definition list by means of the DLISTCLOSE Procedure.

243.5.28 DLISTDEF Procedure

This procedure generates the <DD> tag, which inserts definitions of terms. Use this tag for a definition list <DL>. Terms are tagged <DT> and definitions are tagged <DD>.

Syntax

```
HTP.DLISTDEF (
   ctext          IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-21 DLISTDEF Procedure Parameters

Parameter	Description
ctext	The definition of the term.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<DD CLEAR="cclear" cattributes>ctext
```

243.5.29 DLISTOPEN Procedure

This procedure generates the <DL> tag which starts a definition list. You end a definition list by means of the DLISTCLOSE Procedure.

Syntax

```
HTP.DLISTOPEN (
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-22 DLISTOPEN Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Usage Notes

A definition list looks like a glossary: it contains terms and definitions. Terms are inserted using the DLISTTERM Procedure and definitions are inserted using the DLISTDEF Procedure.

Examples

This procedure generates

```
<DL CLEAR="cclear" cattributes>
```

Related Topics

- DLISTCLOSE Procedure
  This procedure generates the </DL> tag which ends a definition list. You start a definition list by means of the DLISTOPEN Procedure.

243.5.30 DLISTTERM Procedure

This procedure generates the <DT> tag which defines a term in a definition list <DL>.

Syntax

```
HTP.DLISTTERM (
   ctext           IN       VARCHAR2   DEFAULT NULL,
   cclear          IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-23 DLISTTERM Procedure Parameters

Parameter	Description
ctext	The term.
cclear	The value for the CLEAR attribute.

Table 243-23 (Cont.) DLISTTERM Procedure Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<DT CLEAR="cclear" cattributes>ctext
```

243.5.31 EM Procedure

This procedure generates the <EM> and </EM> tags, which define text to be emphasized.

It performs the same task as the EMPHASIS Procedure.

Syntax

```
HTP.EM(
     ctext          IN       VARCHAR2,
     cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-24 EM Procedure Parameters

Parameter	Description
ctext	The text to emphasize.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<EM cattributes>ctext</EM>
```

243.5.32 EMPHASIS Procedure

This procedure generates the <EM> and </EM> tags, which define text to be emphasized.

It performs the same task as the EM Procedure.

Syntax

```
HTP.EMPHASIS(
     ctext          IN       VARCHAR2,
     cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-25 EMPHASIS Procedure Parameters

Parameter	Description
ctext	The text to emphasize.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<EM cattributes>ctext</EM>
```

243.5.33 ESCAPE_SC Procedure

This procedure replaces characters that have special meaning in HTML with their escape sequences.

The following characters are converted:

- & to &
- " to &quot:
- < to <
- > to >

This procedure performs the same operation as PRINTS Procedures and PS Procedure.

Syntax

```
HTP.ESCAPE_SC(
     ctext           IN       VARCHAR2);
```

Parameters

Table 243-26 ESCAPE_SC Procedure Parameters

Parameter	Description
ctext	The text string to convert.

243.5.34 FONTCLOSE Procedure

This procedure generates the</FONT> tag which marks the end of a section of text with the specified font characteristics.

You mark the beginning of the section text by means of the FONTOPEN Procedure.

Syntax

```
HTP.FONTCLOSE;
```

Examples

This procedure generates

```
</FONT>
```

243.5.35 FONTOPEN Procedure

This procedure generates the `<FONT>` which marks the beginning of section of text with the specified font characteristics.

You mark the end of the section text by means of the FONTCLOSE Procedure.

Syntax

```
HTP.FONTOPEN(
   ccolor          IN       VARCHAR2   DEFAULT NULL,
   cface           IN       VARCHAR2   DEFAULT NULL,
   csize           IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-27 FONTOPEN Procedure Parameters

Parameter	Description
`ccolor`	The value for the `COLOR` attribute.
`cface`	The value for the `FACE` attribute
`csize`	The value for the `SIZE` attribute
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<FONT COLOR="ccolor" FACE="cface" SIZE="csize" cattributes>
```

243.5.36 FORMCHECKBOX Procedure

This procedure generates the `<INPUT>` tag with `TYPE="checkbox"` which inserts a checkbox element in a form.

A checkbox element is a button that the user toggles on or off.

Syntax

```
HTP.FORMCHECKBOX(
   cname           IN       VARCHAR2,
   cvalue          IN       VARCHAR2   DEFAULT 'ON',
   cchecked        IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-28 FORMCHECKBOX Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.
cchecked	If the value for this parameter is not NULL, the CHECKED attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="checkbox" NAME="cname" VALUE="cvalue" CHECKED cattributes>
```

243.5.37 FORMCLOSE Procedure

This procedure generates the </FORM> tag which marks the end of a form section in an HTML document.

You mark the beginning of the form section by means of the FORMOPEN Procedure.

Syntax

```
HTP.FORMCLOSE;
```

Examples

This procedure generates

```
</FORM>
```

243.5.38 FORMOPEN Procedure

This procedure generates the <FORM> tag which marks the beginning of a form section in an HTML document.

You mark the end of the form section by means of the FORMCLOSE Procedure.

Syntax

```
HTP.FORMOPEN(
   curl           IN       VARCHAR2,
   cmethod        IN       VARCHAR2   DEFAULT 'POST',
   ctarget        IN       VARCHAR2   DEFAULT NULL,
   cenctype       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-29 FORMOPEN Procedure Parameters

Parameter	Description
curl	The URL of the WRB or CGI script where the contents of the form is sent. This parameter is required.
cmethod	The value for the METHOD attribute. The value can be "GET" or "POST".
ctarget	The value for the TARGET attribute.
cenctype	The value for the ENCTYPE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<FORM ACTION="curl" METHOD="cmethod" TARGET="ctarget" ENCTYPE="cenctype" cattributes>
```

243.5.39 FORMFILE Procedure

This procedure generates the <INPUT> tag with TYPE="file" which inserts a file form element. This is used for file uploading for a given page.

Syntax

```
HTP.FORMFILE(
   cname          IN       VARCHAR2,
   caccept        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-30 FORMFILE Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
caccept	A comma-delimited list of MIME types for upload.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="file" NAME="cname" ACCEPT="caccept" cattributes>
```

243.5.40 FORMHIDDEN Procedure

This procedure generates the <INPUT> tag with TYPE="hidden", which inserts a hidden form element.

This element is not seen by the user. It submits additional values to the script.

Syntax

```
HTP.FORMHIDDEN(
   cname          IN       VARCHAR2,
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-31 FORMHIDDEN Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="hidden" NAME="cname" VALUE="cvalue" cattributes>
```

243.5.41 FORMIMAGE Procedure

This procedure generates the <INPUT> tag with TYPE="image" which creates an image field that the user clicks to submit the form immediately.

The coordinates of the selected point are measured in pixels, and returned (along with other contents of the form) in two name/value pairs. The x coordinate is submitted under the name of the field with .x appended, and the y coordinate with .y appended. Any VALUE attribute is ignored.

Syntax

```
HTP.FORMIMAGE(
   cname          IN       VARCHAR2,
   csrc           IN       VARCHAR2,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-32 FORMIMAGE Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
csrc	The value for the SRC attribute that specifies the image file.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="image" NAME="cname" SRC="csrc" ALIGN="calign" cattributes>
```

243.5.42 FORMPASSWORD Procedure

This procedure generates the <INPUT> tag with TYPE="password" which creates a single-line text entry field. When the user enters text in the field, each character is represented by one asterisk. This is used for entering passwords.

Syntax

```
HTP.FORMPASSWORD(
   cname          IN       VARCHAR2,
   csize          IN       VARCHAR2,
   cmaxlength     IN       VARCHAR2   DEFAULT NULL,
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-33 FORMPASSWORD Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
csize	The value for the SIZE attribute.
cmaxlength	The value for the MAXLENGTH attribute.
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="password" NAME="cname" SIZE="csize" MAXLENGTH="cmaxlength" VALUE="cvalue"
cattributes>
```

243.5.43 FORMRADIO Procedure

This procedure generates the <INPUT> tag with TYPE="radio", which creates a radio button on the HTML form. Within a set of radio buttons, the user selects only one. Each radio button in the same set has the same name, but different values. The selected radio button generates a name/value pair.

Syntax

```
HTP.FORMRADIO(
   cname          IN       VARCHAR2,
   cvalue         IN       VARCHAR2,
   cchecked       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-34 FORMRADIO Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.
cchecked	If the value for this parameter is not NULL, the CHECKED attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="radio" NAME="cname" VALUE="cvalue" CHECKED cattributes>
```

243.5.44 FORMRESET Procedure

This procedure generates the <INPUT> tag with TYPE="reset" which creates a button that, when selected, resets the form fields to their initial values.

Syntax

```
HTP.FORMRESET(
   cvalue         IN       VARCHAR2   DEFAULT 'Reset',
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-35 FORMRESET Procedure Parameters

Parameter	Description
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="reset" VALUE="cvalue" cattributes>
```

243.5.45 FORMSELECTCLOSE Procedure

This procedure generates the </SELECT> tag which marks the end of a Select form element.

A Select form element is a listbox where the user selects one or more values. You mark the beginning of Select form element by means of the FORMSELECTOPEN Procedure.The values are inserted using FORMSELECTOPTION Procedure.

Syntax

```
HTP.FORMSELECTCLOSE;
```

Examples

This procedure generates

```
</SELECT>
```

as shown under Examples of the FORMSELECTOPEN Procedure.

243.5.46 FORMSELECTOPEN Procedure

This procedure generates the <SELECT> tags which creates a Select form element.

A Select form element is a listbox where the user selects one or more values. You mark the end of Select form element by means of the FORMSELECTCLOSE Procedure.The values are inserted using FORMSELECTOPTION Procedure.

Syntax

```
FORMSELECTOPEN(
   cname          IN       VARCHAR2,
   cprompt        IN       VARCHAR2   DEFAULT NULL,
   nsize          IN       INTEGER    DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-36 FORMSELECTOPEN Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cprompt	The string preceding the list box.
nsize	The value for the SIZE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
cprompt <SELECT NAME="cname" SIZE="nsize" cattributes>
</SELECT>
```

so that

```
HTP.FORMSELECTOPEN('greatest_player';
    'Pick the greatest player:');
HTP.FORMSELECTOPTION('Messier');
HTP.FORMSELECTOPTION('Howe');
HTP.FORMSELECTOPTION('Gretzky');.
HTP.FORMSELECTCLOSE;
```

generates

```
Pick the greatest player:
<SELECT NAME="greatest_player">
<OPTION>Messier
<OPTION>Howe
```

```
<OPTION>Gretzky
</SELECT>
```

243.5.47 FORMSELECTOPTION Procedure

This procedure generates the <OPTION> tag which represents one choice in a Select element.

Syntax

```
HTP.FORMSELECTOPTION(
   cvalue          IN       VARCHAR2,
   cselected       IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-37 FORMSELECTOPTION Procedure Parameters

Parameter	Description
cvalue	The text for the option.
cvalue	If the value for this parameter is not NULL, the SELECTED attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<OPTION SELECTED cattributes>cvalue
```

as shown under Examples of the FORMSELECTOPEN Procedure.

243.5.48 FORMSUBMIT Procedure

This procedure generates the <INPUT> tag with TYPE="submit" which creates a button that, when clicked, submits the form. If the button has a NAME attribute, the button contributes a name/value pair to the submitted data.

Syntax

```
HTP.FORMSUBMIT(
   cname           IN       VARCHAR2   DEFAULT NULL,
   cvalue         IN       VARCHAR2   DEFAULT 'Submit',
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-38 FORMSUBMIT Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="submit" NAME="cname" VALUE="cvalue" cattributes>
```

243.5.49 FORMTEXT Procedure

This procedure generates the `<INPUT>` tag with `TYPE="text"`, which creates a field for a single line of text.

Syntax

```
HTP.FORMTEXT(
   cname          IN       VARCHAR2,
   csize          IN       VARCHAR2   DEFAULT NULL,
   cmaxlength     IN       VARCHAR2   DEFAULT NULL,
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-39 FORMTEXT Procedure Parameters

Parameter	Description
`cname`	The value for the `NAME` attribute.
`csize`	The value for the `SIZE` attribute.
`cmaxlength`	The value for the `MAXLENGTH` attribute.
`cvalue`	The value for the `VALUE` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<INPUT TYPE="text" NAME="cname" SIZE="csize" MAXLENGTH="cmaxlength" VALUE="cvalue"
cattributes>
```

243.5.50 FORMTEXTAREA Procedure

This procedure generates the `<TEXTAREA>` tag, which creates a text field that has no predefined text in the text area. This field enables entering several lines of text.

The same operation is performed by the FORMTEXTAREA2 Procedure which in addition has the `cwrap` parameter that lets you specify a wrap style.

Syntax

```
HTP.FORMTEXTAREA(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign       , IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-40 FORMTEXTAREA Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" cattributes></
TEXTAREA>
```

243.5.51 FORMTEXTAREA2 Procedure

This procedure generates the <TEXTAREA> tag, which creates a text field that has no predefined text in the text area. This field enables entering several lines of text.

The same operation is performed by the FORMTEXTAREA Procedure except that in that case you cannot specify a wrap style.

Syntax

```
HTP.FORMTEXTAREA2(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cwrap          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-41 FORMTEXTAREA2 Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cwrap	The value for the WRAP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" WRAP="cwrap"
cattributes></TEXTAREA>
```

243.5.52 FORMTEXTAREACLOSE Procedure

This procedure generates the `</TEXTAREA>` tag which ends a text area form element.

You open a text area element by means of either FORMTEXTAREAOPEN Procedure or FORMTEXTAREAOPEN2 Procedure.

Syntax

```
HTP.FORMTEXTAREACLOSE;
```

Examples

This procedure generates

```
</TEXTAREA>
```

243.5.53 FORMTEXTAREAOPEN Procedure

This procedure generates the `<TEXTAREA>` which marks the beginning of a text area form element.

The same operation is performed by the FORMTEXTAREAOPEN2 Procedure which in addition has the `cwrap` parameter that lets you specify a wrap style. You mark the end of a text area form element by means of the FORMTEXTAREACLOSE Procedure.

Syntax

```
HTP.FORMTEXTAREAOPEN(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-42 FORMTEXTAREAOPEN Procedure Parameters

Parameter	Description
`cname`	The value for the `NAME` attribute.
`nrows`	The value for the `ROWS` attribute.This is an integer.
`ncolumns`	The value for the `COLS` attribute.This is an integer.
`calign`	The value for the `ALIGN` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" cattributes>
```

243.5.54 FORMTEXTAREAOPEN2 Procedure

This procedure generates the <TEXTAREA> which marks the beginning of a text area form element.

The same operation is performed by the FORMTEXTAREAOPEN Procedure except that in that case you cannot specify a wrap style. You mark the end of a text area form element by means of the FORMTEXTAREACLOSE Procedure.

Syntax

```
HTP.FORMTEXTAREAOPEN2(
   cname          IN       VARCHAR2,
   nrows          IN       INTEGER,
   ncolumns       IN       INTEGER,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cwrap          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-43 FORMTEXTAREAOPEN2 Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
nrows	The value for the ROWS attribute.This is an integer.
ncolumns	The value for the COLS attribute.This is an integer.
calign	The value for the ALIGN attribute.
cwrap	The value for the WRAP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TEXTAREA NAME="cname" ROWS="nrows" COLS="ncolumns" ALIGN="calign" WRAP = "cwrap"
cattributes>
```

243.5.55 FRAME Procedure

This procedure generates the <FRAME> tag which begins the characteristics of a frame created by a <FRAMESET> tag.

Syntax

```
HTP.FRAME(
   csrc           IN       VARCHAR2,
   cname          IN       VARCHAR2   DEFAULT NULL,
   cmarginwidth   IN       VARCHAR2   DEFAULT NULL,
   cmarginheight  IN       VARCHAR2   DEFAULT NULL,
   cscrolling     IN       VARCHAR2   DEFAULT NULL,
   cnoresize      IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-44 FRAME Procedure Parameters

Parameter	Description
csrc	The URL to display in the frame.
cname	The value for the NAME attribute.
cmarginwidth	The value for the MARGINWIDTH attribute.
cscrolling	The value for the SCROLLING attribute.
cnoresize	If the value for this parameter is not NULL, the NORESIZE attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<FRAME SRC="csrc" NAME="cname" MARGINWIDTH="cmarginwidth" MARGINHEIGHT="cmarginheight"
SCROLLING="cscrolling" NORESIZE cattributes>
```

243.5.56 FRAMESETCLOSE Procedure

This procedure generates the </FRAMESET> tag which ends a frameset section.

You mark the beginning of a frameset section by means of the FRAMESETOPEN Procedure.

Syntax

```
HTP.FRAMESETCLOSE;
```

Examples

This procedure generates

```
</FRAMESET>
```

243.5.57 FRAMESETOPEN Procedure

This procedure generates the <FRAMESET> tag which define a frameset section.

You mark the end of a frameset section by means of the FRAMESETCLOSE Procedure.

Syntax

```
HTP.FRAMESETOPEN(
   crows          IN       VARCHAR2   DEFAULT NULL,
   ccols          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-45 FRAMESETOPEN Procedure Parameters

Parameter	Description
crows	The value for the ROWS attribute.
ccols	The value for the COLS attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<FRAMESET ROWS="crows" COLS="ccols" cattributes>
```

243.5.58 HEADCLOSE Procedure

This procedure generates the </HEAD> tag which marks the end of an HTML document head section.

You mark the beginning of an HTML document head section by means of the HEADOPEN Procedure.

Syntax

```
HTP.HEADCLOSE;
```

Examples

This procedure generates

```
</HEAD>
```

243.5.59 HEADER Procedure

This procedure generates opening heading tags (<H1> to <H6>) and their corresponding closing tags (</H1> to </H6>).

Syntax

```
HTP.HEADER(
   nsize          IN       INTEGER,
   cheader        IN       VARCHAR2,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-46 HEADER Procedure Parameters

Parameter	Description
nsize	The the heading level. This is an integer between 1 and 6.

Table 243-46 (Cont.) HEADER Procedure Parameters

Parameter	Description
cheader	The text to display in the heading.
calign	The value for the ALIGN attribute.
cnowrap	The value for the NOWRAP attribute.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

```
HTP.header (1,'Overview');
```

produces:

```
<H1>Overview</H1>
```

243.5.60 HEADOPEN Procedure

This procedure generates the <HEAD> tag which marks the beginning of the HTML document head section.

You mark the end of an HTML document head section by means of the HEADCLOSE Procedure.

Syntax

```
HTP.HEADOPEN;
```

Examples

This procedure generates

```
<HEAD>
```

243.5.61 HR Procedure

This procedure generates the <HR> tag, which generates a line in the HTML document.

This subprogram performs the same operation as the LINE Procedure.

Syntax

```
HTP.HR(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   csrc           IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-47 HR Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
csrc	The value for the SRC attribute which specifies a custom image as the source of the line.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<HR CLEAR="cclear" SRC="csrc" cattributes>
```

243.5.62 HTMLCLOSE Procedure

This procedure generates the </HTML> tag which marks the end of an HTML document.

You use the HTMLOPEN Procedure to mark the beginning of an HTML document.

Syntax

```
HTP.HTMLCLOSE;
```

Examples

This procedure generates

```
</HTML>
```

243.5.63 HTMLOPEN Procedure

This procedure generates the <HTML> tag which marks the beginning of an HTML document.

You use the HTMLCLOSE Procedure to mark the end of the an HTML document.

Syntax

```
HTP.HTMLOPEN;
```

Examples

This procedure generates

```
<HTML>
```

243.5.64 IMG Procedure

This procedure generates the <IMG> tag which directs the browser to load an image onto the HTML page.

The IMG2 Procedure performs the same operation but additionally uses the cusemap parameter.

Syntax

```
HTP.IMG(
   curl           IN       VARCHAR2   DEFAULT NULL,
   calign         IN       VARCHAR2   DEFAULT NULL,
   calt           IN       VARCHAR2   DEFAULT NULL,
   cismap         IN       VARCHAR2  DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-48 IMG Procedure Parameters

Parameter	Description
curl	The value for the SRC attribute.
calign	The value for the ALIGN attribute.
calt	The value for the ALT attribute which specifies alternative text to display if the browser does not support images.
cismap	If the value for this parameter is not NULL, the ISMAP attribute is added to the tag. The attribute indicates that the image is an imagemap.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<IMG SRC="curl" ALIGN="calign" ALT="calt" ISMAP cattributes>
```

243.5.65 IMG2 Procedure

This procedure generates the <IMG> tag, which directs the browser to load an image onto the HTML page.

The IMG Procedure performs the same operation but does not use the cusemap parameter.

Syntax

```
HTP.IMG2(
   curl           IN       VARCHAR2   DEFAULT NULL,
   calign         IN       VARCHAR2   DEFAULT NULL,
   calt           IN       VARCHAR2   DEFAULT NULL,
   cismap         IN       VARCHAR2  DEFAULT NULL,
   cusemap        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-49 IMG2 Procedure Parameters

Parameter	Description
curl	The value for the SRC attribute.
calign	The value for the ALIGN attribute.

Table 243-49 (Cont.) IMG2 Procedure Parameters

Parameter	Description
calt	The value for the ALT attribute which specifies alternative text to display if the browser does not support images.
cismap	If the value for this parameter is not NULL, the ISMAP attribute is added to the tag. The attribute indicates that the image is an imagemap.
cusemap	The value for the USEMAP attribute which specifies a client-side image map.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<IMG SRC="curl" ALIGN="calign" ALT="calt" ISMAP USEMAP="cusemap" cattributes>
```

243.5.66 ISINDEX Procedure

This procedure creates a single entry field with a prompting text, such as "*enter value,*" then sends that value to the URL of the page or program.

Syntax

```
HTP.ISINDEX(
   cprompt         IN       VARCHAR2    DEFAULT NULL,
   curl            IN       VARCHAR2    DEFAULT NULL);
```

Parameters

Table 243-50 ISINDEX Procedure Parameters

Parameter	Description
cprompt	The value for the PROMPT attribute.
curl	The value for the HREF attribute.

Examples

This procedure generates

```
<ISINDEX PROMPT="cprompt" HREF="curl">
```

243.5.67 ITALIC Procedure

This procedure generates the <I> and </I> tags which direct the browser to render the text in italics.

Syntax

```
HTP.ITALIC(
   ctext           IN       VARCHAR2,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-51 ITALIC Procedure Parameters

Parameter	Description
ctext	The text to be rendered in italics.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<I cattributes>ctext</I>
```

243.5.68 KBD Procedure

This procedure generates the <KBD> and </KBD> tags which direct the browser to render the text in monospace font.

This subprogram performs the same operation as the KEYBOARD Procedure.

Syntax

```
HTP.KBD(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-52 KBD Procedure Parameters

Parameter	Description
ctext	The text to be rendered in monospace.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<KBD cattributes>ctext</KBD>
```

243.5.69 KEYBOARD Procedure

This procedure generates the <KBD> and </KBD> tags, which direct the browser to render the text in monospace font.

This subprogram performs the same operation as the KBD Procedure.

Syntax

```
HTP.KEYBOARD(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-53 KEYBOARD Procedure Parameters

Parameter	Description
ctext	The text to be rendered in monospace.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<KBD cattributes>ctext</KBD>
```

243.5.70 LINE Procedure

This procedure generates the <HR> tag, which generates a line in the HTML document.

This subprogram performs the same operation as the HR Procedure.

Syntax

```
HTP.LINE(
   cclear          IN       VARCHAR2   DEFAULT NULL,
   csrc            IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-54 LINE Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
csrc	The value for the SRC attribute which specifies a custom image as the source of the line.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<HR CLEAR="cclear" SRC="csrc" cattributes>
```

243.5.71 LINKREL Procedure

This procedure generates the <LINK> tag with the REL attribute which delineates the relationship described by the hypertext link from the anchor to the target. This is only used when the HREF attribute is present.

This is the opposite of LINKREV Procedure. This tag indicates a relationship between documents but does not create a link. To create a link, use the ANCHOR Procedure.

Syntax

```
HTP.LINKREL(
   crel           IN       VARCHAR2,
   curl           IN       VARCHAR2,
   ctitle         IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-55 LINKREL Procedure Parameters

Parameter	Description
crel	The value for the REL attribute.
curl	The value for the URL attribute.
ctitle	The value for the TITLE attribute.

Examples

This procedure generates

```
<LINK REL="crel" HREF="curl" TITLE="ctitle">
```

243.5.72 LINKREV Procedure

This procedure generates the <LINK> tag with the REV attribute which delineates the relationship described by the hypertext link from the target to the anchor.

This is the opposite of the LINKREL Procedure. This tag indicates a relationship between documents, but does not create a link. To create a link, use the ANCHOR Procedure.

Syntax

```
HTP.LINKREV(
   crev           IN       VARCHAR2,
   curl           IN       VARCHAR2,
   ctitle         IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-56 LINKREV Procedure Parameters

Parameter	Description
crel	The value for the REV attribute.
curl	The value for the URL attribute.
ctitle	The value for the TITLE attribute.

Examples

This procedure generates

```
<LINK REV="crev" HREF="curl" TITLE="ctitle">
```

243.5.73 LISTHEADER Procedure

This procedure generates the <LH> and </LH> tags which print an HTML tag at the beginning of the list.

Syntax

```
HTP.LISTHEADER(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-57 LISTHEADER Procedure Parameters

Parameter	Description
ctext	The text to place between <LH> and </LH>.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<LH cattributes>ctext</LH>
```

243.5.74 LISTINGCLOSE Procedure

This procedure generates the </LISTING> tags which marks the end of a section of fixed-width text in the body of an HTML page.

To mark the beginning of a section of fixed-width text in the body of an HTML page, use the LISTINGOPEN Procedure.

Syntax

```
HTP.LISTINGCLOSE;
```

Examples

This procedure generates

```
</LISTING>
```

243.5.75 LISTINGOPEN Procedure

This procedure generates the <LISTING> tag which marks the beginning of a section of fixed-width text in the body of an HTML page.

To mark the end of a section of fixed-width text in the body of an HTML page, use the LISTINGCLOSE Procedure.

Syntax

```
HTP.LISTINGOPEN;
```

Examples

This procedure generates

```
<LISTING>
```

243.5.76 LISTITEM Procedure

This procedure generates the <LI> tag, which indicates a list item.

Syntax

```
HTP.LISTITEM(
   ctext          IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cdingbat       IN       VARCHAR2   DEFAULT NULL,
   csrc           IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-58 LISTITEM Procedure Parameters

Parameter	Description
ctext	The text for the list item.
cclear	The value for the CLEAR attribute.
cdingbat	The value for the DINGBAT attribute.
csrc	The value for the SRC attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<LI CLEAR="cclear" DINGBAT="cdingbat" SRC="csrc" cattributes>ctext
```

243.5.77 MAILTO Procedure

This procedure generates the <A> tag with the HREF set to 'mailto' prepended to the mail address argument.

Syntax

```
HTP.MAILTO(
   caddress       IN       VARCHAR2,
   ctext          IN       VARCHAR2,
   cname          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-59 MAILTO Procedure Parameters

Parameter	Description
caddress	The email address of the recipient.
ctext	The clickable portion of the link.
cname	The value for the NAME attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<A HREF="mailto:caddress" NAME="cname" cattributes>ctext</A>
```

so that

```
HTP.mailto('pres@white_house.gov','Send Email to the President');
```

generates:

```
<A HREF="mailto:pres@white_house.gov">Send Email to the President</A>
```

243.5.78 MAPCLOSE Procedure

This procedure generates the </MAP> tag which marks the end of a set of regions in a client-side image map.

To mark the beginning of a set of regions in a client-side image map, use the MAPOPEN Procedure.

Syntax

```
HTP.MAPCLOSE;
```

Examples

This procedure generates

```
</MAP>
```

243.5.79 MAPOPEN Procedure

This procedure generates the <MAP> tag which mark the beginning of a set of regions in a client-side image map.

To mark the end of a set of regions in a client-side image map, use the MAPCLOSE Procedure.

Syntax

```
HTP.MAPOPEN(
   cname          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-60 MAPOPEN Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<MAP NAME="cname" cattributes>
```

243.5.80 MENULISTCLOSE Procedure

This procedure generates the </MENU> tag which ends a list that presents one line for each item.

To begin a list of this kind, use the MENULISTOPEN Procedure. The items in the list appear more compact than an unordered list. The LISTITEM Procedure defines the list items in a menu list.

Syntax

```
HTP.MENULISTCLOSE;
```

Examples

This procedure generates

```
</MENU>
```

243.5.81 MENULISTOPEN Procedure

This procedure generates the <MENU> tag which begins a list that presents one line for each item.

To end a list of this kind, use the MENULISTCLOSE Procedure.The items in the list appear more compact than an unordered list. The LISTITEM Procedure defines the list items in a menu list.

Syntax

```
HTP.MENULISTOPEN;
```

Examples

This procedure generates

```
<MENU>
```

243.5.82 META Procedure

This procedure generates the <META> tag, which embeds meta-information about the document and also specifies values for HTTP headers. For example, you can specify the expiration date, keywords, and author name.

Syntax

```
HTP.META(
   chttp_equiv    IN       VARCHAR2,
   cname          IN       VARCHAR2,
   ccontent       IN       VARCHAR2);
```

Parameters

Table 243-61 META Procedure Parameters

Parameter	Description
chttp_equiv	The value for the CHTTP_EQUIV attribute.
cname	The value for the NAME attribute.
ccontent	The value for the CONTENT attribute.

Examples

This procedure generates

```
<META HTTP-EQUIV="chttp_equiv" NAME ="cname" CONTENT="ccontent">
```

so that

```
HTP.meta ('Refresh', NULL, 120);
```

generates

```
<META HTTP-EQUIV="Refresh" CONTENT=120>
```

On some Web browsers, this causes the current URL to be reloaded automatically every 120 seconds.

243.5.83 NL Procedure

This procedure generates the <BR> tag which begins a new line of text.

It performs the same operation as the BR Procedure.

Syntax

```
HTP.NL(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-62 NL Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<BR CLEAR="cclear" cattributes>
```

243.5.84 NOBR Procedure

This procedure generates the <NOBR> and </NOBR> tags which turn off line-breaking in a section of text.

Syntax

```
HTP.NOBR(
ctext        IN        VARCHAR2);
```

Parameters

Table 243-63 NOBR Procedure Parameters

Parameter	Description
ctext	The text that is to be rendered on one line.

Examples

This procedure generates

```
<NOBR>ctext</NOBR>
```

243.5.85 NOFRAMESCLOSE Procedure

This procedure generates the </NOFRAMES> tag which marks the end of a no-frames section.

To mark the beginning of a no-frames section, use the FRAMESETOPEN Procedure. See also FRAME Procedure, FRAMESETOPEN Procedure and FRAMESETCLOSE Procedure.

Syntax

```
HTP.NOFRAMESCLOSE;
```

Examples

This procedure generates

```
</NOFRAMES>
```

243.5.86 NOFRAMESOPEN Procedure

This procedure generates the <NOFRAMES> tag which mark the beginning of a no-frames section.

To mark the end of a no-frames section, use the FRAMESETCLOSE Procedure. See also FRAME Procedure, FRAMESETOPEN Procedure and FRAMESETCLOSE Procedure.

Syntax

```
HTP.NOFRAMESOPEN;
```

Examples

This procedure generates

```
<NOFRAMES>
```

243.5.87 OLISTCLOSE Procedure

This procedure generates the </OL> tag which defines the end of an ordered list. An ordered list presents a list of numbered items.

To mark the beginning of a list of this kind, use the OLISTOPEN Procedure. Numbered items are added using LISTITEM Procedure.

Syntax

```
HTP.OLISTCLOSE;
```

Examples

This procedure generates

```
</OL>
```

243.5.88 OLISTOPEN Procedure

This procedure generates the <OL> tag which marks the beginning of an ordered list. An ordered list presents a list of numbered items.

To mark the end of a list of this kind, use the OLISTCLOSE Procedure. Numbered items are added using LISTITEM Procedure.

Syntax

```
HTP.OLISTOPEN(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cwrap          IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-64 OLISTOPEN Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cwrap	The value for the WRAP attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<OL CLEAR="cclear" WRAP="cwrap" cattributes>
```

243.5.89 PARA Procedure

This procedure generates the <P> tag which indicates that the text that comes after the tag is to be formatted as a paragraph.

You can add attributes to the tag by means of the PARAGRAPH Procedure.

Syntax

```
HTP.PARA;
```

Examples

This procedure generates

```
<P>
```

243.5.90 PARAGRAPH Procedure

You can use this procedure to add attributes to the <P> tag created by the PARA Procedure.

Syntax

```
HTP.PARAGRAPH(
   calign         IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-65 PARAGRAPH Procedure Parameters

Parameter	Description
calign	The value for the ALIGN attribute.
cnowrap	If the value for this parameter is not NULL, the NOWRAP attribute is added to the tag.
cclear	The value for the CLEAR attribute.

Table 243-65 (Cont.) PARAGRAPH Procedure Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<P ALIGN="calign" NOWRAP CLEAR="cclear" cattributes>
```

Related Topics

- PARA Procedure
  This procedure generates the <P> tag which indicates that the text that comes after the tag is to be formatted as a paragraph.

243.5.91 PARAM Procedure

This procedure generates the <PARAM> tag which specifies parameter values for Java applets.

The values can reference HTML variables. To invoke a Java applet from a Web page, use APPLETOPEN Procedure to begin the invocation. Use one PARAM Procedure for each desired name-value pair, and use APPLETCLOSE Procedure to end the applet invocation.

Syntax

```
HTP.PARAM(
   cname          IN       VARCHAR2
   cvalue         IN       VARCHAR2);
```

Parameters

Table 243-66 PARAM Procedure Parameters

Parameter	Description
cname	The value for the NAME attribute.
cvalue	The value for the VALUE attribute.

Examples

This procedure generates

```
<PARAM NAME=cname VALUE="cvalue">
```

243.5.92 PLAINTEXT Procedure

This procedure generates the <PLAINTEXT> and </PLAINTEXT> tags which direct the browser to render the text they surround in fixed-width type.

Syntax

```
HTP.PLAINTEXT(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-67 PLAINTEXT Procedure Parameters

Parameter	Description
ctext	The text to be rendered in fixed-width font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<PLAINTEXT cattributes>ctext</PLAINTEXT>
```

243.5.93 PRECLOSE Procedure

This procedure generates the </PRE> tag which marks the end of a section of preformatted text in the body of the HTML page.

To mark the beginning of a section of preformatted text in the body of the HTML page, use the PREOPEN Procedure.

Syntax

```
HTP.PRECLOSE;
```

Examples

This procedure generates

```
</PRE>
```

243.5.94 PREOPEN Procedure

This procedure generates the <PRE> tag which marks the beginning of a section of preformatted text in the body of the HTML page.

To mark the end of a section of preformatted text in the body of the HTML page, use the PRECLOSE Procedure.

Syntax

```
HTP.PREOPEN(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cwidth         IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-68 PREOPEN Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cwidth	The value for the WIDTH attribute.

Table 243-68 (Cont.) PREOPEN Procedure Parameters

Parameter	Description
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<PRE CLEAR="cclear" WIDTH="cwidth" cattributes>
```

243.5.95 PRINT Procedures

These procedures generate the specified parameter as a string terminated with the \n newline character.

The PRN Procedures performs the same operation but does not terminate with a newline character.

Syntax

```
HTP.PRINT (
   cbuf      IN       VARCHAR2);

HTP.PRINT (
   dbuf      IN       DATE);

HTP.PRINT (
   nbuf      IN       NUMBER);
```

Parameters

Table 243-69 PRINT Procedure Parameters

Parameter	Description
cbuf	The string to generate terminated by a newline.
dbuf	The string to generate terminated by a newline.
nbuf	The string to generate terminated by a newline.

Usage Notes

- The \n character is not the same as <BR>. The \n character formats the HTML source but it does not affect how the browser renders the HTML source. Use <BR> to control how the browser renders the HTML source.
- These procedures do not have function equivalents.

243.5.96 PRINTS Procedure

This procedure generates a string and replaces certain characters with a corresponding escape sequence.

The following characters are replaced with the corresponding escape sequence.

- < to <
- > to >
- " to "
- & to &

If not replaced, the special characters are interpreted as HTML control characters and produce garbled output. This procedure an the PS Procedure perform the same operation as the PRN Procedures but with character substitution.

Syntax

```
HTP.PRINTS (
   ctext      IN       VARCHAR2);
```

Parameters

Table 243-70 PRINTS Procedure Parameters

Parameter	Description
ctext	The string where to perform character substitution.

Usage Notes

This procedure does not have an HTF function equivalent (see Operational Notes for the HTF implementation).

243.5.97 PRN Procedures

These procedures generate the specified parameter as a string.

Unlike the PRINT Procedures the string is not terminated with the \n newline character.

Syntax

```
HTP.PRN (
   cbuf       IN       VARCHAR2);

HTP.PRN (
   dbuf       IN       DATE);

HTP.PRN (
   nbuf       IN       NUMBER);
```

Parameters

Table 243-71 PRN Procedure Parameters

Parameter	Description
cbuf	The string to generate (not terminated by a newline).
dbuf	The string to generate (not terminated by a newline).
nbuf	The string to generate (not terminated by a newline).

Usage Notes

These procedures do not have function equivalents.

243.5.98 PS Procedure

This procedure generates a string and replaces certain characters with the corresponding escape sequence.

This procedure replaces the following characters with the corresponding escape sequence.

- < to `<`
- > to `>`
- " to `"`
- & to `&`

If not replaced, the special characters are interpreted as HTML control characters and produce garbled output. This procedure and the PRINTS Procedure perform the same operation as the PRN Procedures but with character substitution.

Syntax

```
HTP.PS (
   ctext      IN       VARCHAR2);
```

Parameters

Table 243-72 PS Procedure Parameters

Parameter	Description
ctext	The string where to perform character substitution.

Usage Notes

This procedure does not have an HTF function equivalent (see Operational Notes for the HTF implementation).

243.5.99 S Procedure

This procedure generates the <S> and </S> tags which direct the browser to render the text they surround in strikethrough type.

This performs the same operation as STRIKE Procedure.

Syntax

```
HTP.S (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-73 S Procedure Parameters

Parameter	Description
ctext	The text to be rendered in strikethrough type.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<S cattributes>ctext</S>
```

243.5.100 SAMPLE Procedure

This procedure generates the <SAMP> and </SAMP> tags which direct the browser to render the text they surround in monospace font or however "sample" is defined stylistically.

Syntax

```
HTP.SAMPLE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-74 SAMPLE Procedure Parameters

Parameter	Description
ctext	The text to be rendered in monospace font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<SAMP cattributes>ctext</SAMP>
```

243.5.101 SCRIPT Procedure

This procedure generates the <SCRIPT> and </SCRIPT> tags which contain a script written in languages such as JavaScript and VBscript.

Syntax

```
HTP.SCRIPT (
   cscript        IN       VARCHAR2,
   clanguage      IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-75 SCRIPT Procedure Parameters

Parameter	Description
cscript	The text of the script. This is the text that makes up the script itself, not the name of a file containing the script.
clanguage	The language in which the script is written. If this parameter is omitted, the user's browser determines the scripting language.

Examples

This procedure generates

```
<SCRIPT LANGUAGE=clanguage>cscript</SCRIPT>
```

so that

```
HTP.script ('Erupting_Volcano', 'Javascript');
```

generates

```
<SCRIPT LANGUAGE=Javascript>"script text here"</SCRIPT>
```

This causes the browser to run the script enclosed in the tags.

243.5.102 SMALL Procedure

This procedure generates the <SMALL> and </SMALL> tags, which direct the browser to render the text they surround using a small font.

Syntax

```
HTP.SMALL (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-76 SMALL Procedure Parameters

Parameter	Description
ctext	The text to be rendered in small font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<SMALL cattributes>ctext</SMALL>
```

243.5.103 STRIKE Procedure

This procedure generates the <STRIKE> and </STRIKE> tags which direct the browser to render the text they surround in strikethrough type.

This performs the same operation as S Procedure.

Syntax

```
HTP.STRIKE (
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-77 STRIKE Procedure Parameters

Parameter	Description
ctext	The text to be rendered in strikethrough type.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<STRIKE cattributes>ctext</STRIKE>
```

243.5.104 STRONG Procedure

This procedure generates the <STRONG> and </STRONG> tags which direct the browser to render the text they surround in bold, or however "strong" is defined.

Syntax

```
HTP.STRONG(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-78 STRONG Procedure Parameters

Parameter	Description
ctext	The text to be emphasized.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<STRONG cattributes>ctext</STRONG>
```

243.5.105 STYLE Procedure

This procedure generates the <STYLE> and </STYLE> tags which include a style sheet in a Web page.

You can get more information about style sheets at http://www.w3.org. This feature is not compatible with browsers that support only HTML versions 2.0 or earlier. Such browsers will ignore this tag.

Syntax

```
HTP.STYLE(
   cstyle          IN       VARCHAR2);
```

Parameters

Table 243-79 STYLE Procedure Parameters

Parameter	Description
cstyle	The the style information to include.

Examples

This procedure generates

```
<STYLE>cstyle</STYLE>
```

243.5.106 SUB Procedure

This procedure generates the _{and} tags which direct the browser to render the text they surround as subscript.

Syntax

```
HTP.SUB(
   ctext           IN       VARCHAR2,
   calign          in       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-80 SUB Procedure Parameters

Parameter	Description
ctext	The text to render in subscript.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<SUB ALIGN="calign" cattributes>ctext</SUB>
```

243.5.107 SUP Procedure

This procedure generates the ^{and} tags which direct the browser to render the text they surround as superscript.

Syntax

```
HTP.SUP(
   ctext          IN       VARCHAR2,
   calign         in       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-81 SUP Procedure Parameters

Parameter	Description
ctext	The text to render in superscript.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<SUP ALIGN="calign" cattributes>ctext</SUP>
```

243.5.108 TABLECAPTION Procedure

This procedure generates the <CAPTION> and </CAPTION> tags which place a caption in an HTML table.

Syntax

```
HTP.TABLECAPTION(
   ccaption       IN       VARCHAR2,
   calign         in       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-82 TABLECAPTION Procedure Parameters

Parameter	Description
ctext	The text for the caption.
calign	The value for the ALIGN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<CAPTION ALIGN="calign" cattributes>ccaption</CAPTION>
```

243.5.109 TABLECLOSE Procedure

This procedure generates the `</TABLE>` tag which marks the end of an HTML table.

To define the beginning of an HTML table, use the TABLEOPEN Procedure.

Syntax

```
HTP.TABLECLOSE;
```

Examples

This procedure generates

```
</TABLE>
```

243.5.110 TABLEDATA Procedure

This procedure generates the `<TD>` and `</TD>` tags which insert data into a cell of an HTML table.

Syntax

```
HTP.TABLEDATA(
   cvalue         IN       VARCHAR2   DEFAULT NULL,
   calign         IN       VARCHAR2   DEFAULT NULL,
   cdp            IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   crowspan       IN       VARCHAR2   DEFAULT NULL,
   ccolspan       IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-83 TABLEDATA Procedure Parameters

Parameter	Description
`cvalue`	The data for the cell in the table.
`calign`	The value for the `ALIGN` attribute.
`cdp`	The value for the `DP` attribute.
`cnowrap`	If the value of this parameter is not `NULL`, the `NOWRAP` attribute is added to the tag.
`ccolspan`	The value for the `COLSPAN` attribute.
`cattributes`	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TD ALIGN="calign" DP="cdp" ROWSPAN="crowspan" COLSPAN="ccolspan" NOWRAP
cattributes>cvalue</TD>
```

243.5.111 TABLEHEADER Procedure

This procedure generates the <TH> and </TH> tags which insert a header cell in an HTML table.

The <TH> tag is similar to the <TD> tag except that the text in this case the rows are usually rendered in bold type.

Syntax

```
HTP.TABLEHEADER(
   cvalue          IN       VARCHAR2   DEFAULT NULL,
   calign          IN       VARCHAR2   DEFAULT NULL,
   cdp             IN       VARCHAR2   DEFAULT NULL,
   cnowrap         IN       VARCHAR2   DEFAULT NULL,
   crowspan        IN       VARCHAR2   DEFAULT NULL,
   ccolspan        IN       VARCHAR2   DEFAULT NULL,
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-84 TABLEHEADER Procedure Parameters

Parameter	Description
cvalue	The data for the cell in the table.
calign	The value for the ALIGN attribute.
cdp	The value for the DP attribute.
cnowrap	If the value of this parameter is not NULL, the NOWRAP attribute is added to the tag.
crispen	The value for the ROWSPAN attribute.
ccolspan	The value for the COLSPAN attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TH ALIGN="calign" DP="cdp" ROWSPAN="crowspan" COLSPAN="ccolspan" NOWRAP
cattributes>cvalue</TH>
```

243.5.112 TABLEOPEN Procedure

This procedure generates the <TABLE> tag which marks the beginning of an HTML table.

To define the end of an HTML table, use the TABLECLOSE Procedure.

Syntax

```
HTP.TABLEOPEN(
   cborder         IN       VARCHAR2   DEFAULT NULL
   calign          IN       VARCHAR2   DEFAULT NULL,
   cnowrap         IN       VARCHAR2   DEFAULT NULL,
   cclear          IN       VARCHAR2   DEFAULT NULL
   cattributes     IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-85 TABLEOPEN Procedure Parameters

Parameter	Description
border	The value for the BORDER attribute.
calign	The value for the ALIGN attribute.
cnowrap	If the value of this parameter is not NULL, the NOWRAP attribute is added to the tag.
cclear	The value for the CLEAR attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TABLE "cborder" NOWRAP ALIGN="calign" CLEAR="cclear" cattributes>
```

243.5.113 TABLEROWCLOSE Procedure

This procedure generates the </TR> tag which marks the end of a new row in an HTML table.

To mark the beginning of a new row, use the TABLEROWOPEN Procedure.

Syntax

```
HTP.TABLEROWCLOSE;
```

Examples

This procedure generates

```
</TABLE>
```

243.5.114 TABLEROWOPEN Procedure

This procedure generates the <TR> tag which marks the beginning of a new row in an HTML table.

To mark the end of a new row, use the TABLEROWCLOSE Procedure.

Syntax

```
HTP.TABLEROWOPEN(
   calign         IN       VARCHAR2   DEFAULT NULL,
   cvalign        IN       VARCHAR2   DEFAULT NULL,
   cdp            IN       VARCHAR2   DEFAULT NULL,
   cnowrap        IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-86 TABLEROWOPEN Procedure Parameters

Parameter	Description
calign	The value for the ALIGN attribute.
cvalign	The value for the VALIGN attribute.
cdp	The value for the DP attribute.
cnowrap	If the value of this parameter is not NULL, the NOWRAP attribute is added to the tag.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<<TR ALIGN="calign" VALIGN="cvalign" DP="cdp" NOWRAP cattributes>
```

243.5.115 TELETYPE Procedure

This procedure generates the <TT> and </TT> tags which direct the browser to render the text they surround in a fixed width typewriter font, for example, the courier font.

Syntax

```
HTP.TELETYPE(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-87 TELETYPE Procedure Parameters

Parameter	Description
ctext	The text to render in a fixed width typewriter font.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<TT cattributes>ctext</TT>
```

243.5.116 TITLE Procedure

This procedure generates the <TITLE> and </TITLE> tags which specify the text to display in the titlebar of the browser window.

Syntax

```
HTP.TITLE(
   ctitle          IN       VARCHAR2);
```

Parameters

Table 243-88 TITLE Procedure Parameters

Parameter	Description
ctitle	The text to display in the titlebar of the browser window.

Examples

This procedure generates

```
<TITLE>ctitle</TITLE>
```

243.5.117 ULISTCLOSE Procedure

This procedure generates the </UL> tag which marks the end of an unordered list. An unordered list presents items with bullets.

To mark the beginning of an unordered list, use the ULISTOPEN Procedure. Add list items with LISTITEM Procedure.

Syntax

```
HTP.ULISTCLOSE;
```

Examples

This procedure generates

```
</TABLE>
```

243.5.118 ULISTOPEN Procedure

This procedure generates the <UL> tag which marks the beginning of an unordered list. An unordered list presents items with bullets.

To mark the end of an unordered list, use the ULISTCLOSE Procedure. Add list items with LISTITEM Procedure.

Syntax

```
HTP.ULISTOPEN(
   cclear         IN       VARCHAR2   DEFAULT NULL,
   cwrap          IN       VARCHAR2   DEFAULT NULL,
   cdingbat       IN       VARCHAR2   DEFAULT NULL,
   csrc           IN       VARCHAR2   DEFAULT NULL,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-89 ULISTOPEN Procedure Parameters

Parameter	Description
cclear	The value for the CLEAR attribute.
cwrap	The value for the WRAP attribute.

Table 243-89 (Cont.) ULISTOPEN Procedure Parameters

Parameter	Description
cdingbat	The value for the DINGBAT attribute.
csrc	The value for the SRC attribute.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<UL CLEAR="cclear" WRAP="cwrap" DINGBAT="cdingbat" SRC="csrc" cattributes>
```

243.5.119 UNDERLINE Procedure

This procedure generates the <U> and </U> tags, which direct the browser to render the text they surround with an underline.

Syntax

```
HTP.UNDERLINE(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-90 UNDERLINE Procedure Parameters

Parameter	Description
ctext	The text to render with an underline.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<U cattributes>ctext</U>
```

243.5.120 VARIABLE Procedure

This procedure generates the <VAR> and </VAR> tags which direct the browser to render the text they surround in italics or however "variable" is defined stylistically.

Syntax

```
HTP.VARIABLE(
   ctext          IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 243-91 VARIABLE Procedure Parameters

Parameter	Description
ctext	The text to render in italics.
cattributes	The other attributes to be included as-is in the tag.

Examples

This procedure generates

```
<VAR cattributes>ctext</VAR>
```

243.5.121 WBR Procedure

This procedure generates the <WBR> tag, which inserts a soft line break within a section of NOBR text.

Syntax

```
HTP.WBR;
```

Examples

This procedure generates

```
<WBR>
```

244
OWA_CACHE

The `OWA_CACHE` package provides an interface that enables the `PL/SQL Gateway` cache to improve the performance of PL/SQL Web applications.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Constants
- Summary of OWA_CACHE Subprograms

244.1 OWA_CACHE Constants

`OWA_CACHE` defines several constants to use when specifying parameter values.

- `system_level CONSTANT VARCHAR(6) := 'SYSTEM';`
- `user_level CONSTANT VARCHAR(4) := 'USER';`

244.2 Summary of OWA_CACHE Subprograms

This table lists the OWA_CACHE subprograms and briefly describes them.

Table 244-1 OWA_CACHE Package Subprograms

Subprogram	Description
DISABLE Procedure	Disables the cache for this particular request
GET_ETAG Function	Returns the tag associated with the cached content (used in the Validation technique model only)
GET_LEVEL Function	Returns the caching level (used in the Validation technique model only)
SET_CACHE Procedure	Sets up the cache headers for validation model cache type
SET_EXPIRES Procedure	Sets up the cache headers for expires model cache type
SET_NOT_MODIFIED Procedure	Sets up the headers for a not modified cache hit (used in the Validation technique model only)
SET_SURROGATE_CONTROL Procedure	Sets up the headers for a surrogate-control header for Web cache

244.2.1 DISABLE Procedure

This procedure disables the cache for this particular request.

Syntax

```
OWA_CACHE.DISABLE;
```

244.2.2 GET_ETAG Function

This function returns the tag associated with the cached content. It is used in the Validation technique only.

Syntax

```
OWA_CACHE.GET_ETAG
  RETURN VARCHAR2;
```

Return Values

The tag for cache hit, otherwise `NULL`.

244.2.3 GET_LEVEL Function

This returns the caching level. It is used in the Validation technique model only.

Syntax

```
OWA_CACHE.GET_LEVEL
  RETURN VARCHAR2;
```

Return Values

The caching level string ('`USER`' or '`SYSTEM`') for cache hit, otherwise `NULL`.

244.2.4 SET_CACHE Procedure

This procedure sets up the cache headers for validation model cache type.

Syntax

```
OWA_CACHE.SET_CACHE(
   p_etag         IN       VARCHAR2,
   p_level        IN       VARCHAR2);
```

Parameters

Table 244-2 *SET_CACHE Procedure Parameters*

Parameter	Description
`p_etag`	The `etag` associated with this content
`p_level`	The caching level ('`USER`' or '`SYSTEM`').

Exceptions

VALUE_ERROR is thrown if

- p_etag is greater than 55
- p_level is not 'USER' or 'SYSTEM'

244.2.5 SET_EXPIRES Procedure

This procedure sets up the cache headers for expires model cache type.

Syntax

```
OWA_CACHE.SET_EXPIRES(
   p_expires      IN       NUMBER,
   p_level        IN       VARCHAR2);
```

Parameters

Table 244-3 SET_EXPIRES Procedure Parameters

Parameter	Description
p_expires	The number of minutes this content is valid.
p_level	The caching level ('USER' or 'SYSTEM').

Exceptions

VALUE_ERROR is thrown if

- p_expires is negative or zero
- p_level is not 'USER' or 'SYSTEM'
- p_expires is > 525600 (1 year)

244.2.6 SET_NOT_MODIFIED Procedure

This procedure sets up the headers for a not-modified cache hit. It is used in the Validation technique only.

Syntax

```
OWA_CACHE.SET_NOT_MODIFIED;
```

Exceptions

VALUE_ERROR is thrown if If the etag was not passed in

244.2.7 SET_SURROGATE_CONTROL Procedure

This procedure sets the headers for a surrogate-control header for Web cache

Syntax

```
OWA_CACHE.SET_SURROGATE_CONTROL(
   p_value         IN       VARCHAR2);
```

Parameters

Table 244-4 SET_SURROGATE_CONTROL Procedure Parameters

Parameter	Description
p_value	The value to be passed as the Surrogate-Control header.

Exceptions

VALUE_ERROR is thrown if If p_value is greater than 55 in length.

245
OWA_COOKIE

The OWA_COOKIE package provides an interface for sending and retrieving HTTP cookies from the client's browser.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Overview
- Types
- Rules and Limits
- Summary of OWA_COOKIE Subprograms

245.1 OWA_CUSTOM Overview

Cookies are opaque strings sent to the browser to maintain state between HTTP calls. State can be maintained throughout the client's sessions, or longer if an expiration date is included. The system date is calculated with reference to the information specified in the OWA_CUSTOM package.

245.2 OWA_COOKIE Types

This datatype contains cookie name-value pairs.

Since the HTTP standard allows cookie names to be overloaded (that is, multiple values can be associated with the same cookie name), there is a PL/SQL `RECORD` holding all values associated with a given cookie name.

```
TYPE vc_arr IS TABLE OF VARCHAR2(4000) INDEX BY BINARY_INTEGER.

TYPE COOKIE IS RECORD (
   name           VARCHAR2(4000),
   vals           vc_arr,
   num_vals       INTEGER);
```

245.3 OWA_COOKIE Rules and Limits

All HTTP headers must be in English and the ASCII character set. If the headers are generated from the database, verify they are created in the English language.

245.4 Summary of OWA_COOKIE Subprograms

This table lists the `OWA_COOKIE` subprograms and briefly describes them.

Table 245-1 *OWA_COOKIE Package Subprograms*

Subprogram	Description
GET Function	Gets the value of the specified cookie
GET_ALL Procedure	Gets all cookie name-value pairs
REMOVE Procedure	Removes the specified cookie
SEND procedure	Generates a "Set-Cookie" line in the HTTP header

245.4.1 GET Function

This function returns the values associated with the specified cookie. The values are returned in a `OWA_COOKIE.COOKIE DATA TYPE` .

Syntax

```
OWA_COOKIE.GET(
    name            IN       VARCHAR2)
  RETURN COOKIE;
```

Parameters

Table 245-2 GET Function Parameters

Parameter	Description
name	The name of the cookie.

Return Values

`OWA_COOKIE.COOKIE DATA TYPE.`

245.4.2 GET_ALL Procedure

This procedure returns all cookie names and their values from the client's browser. The values appear in the order in which they were sent from the browser.

Syntax

```
OWA_COOKIE.GET_ALL(
   names           OUT      vc_arr,
   vals            OUT      vc_arr,
   num_vals        OUT      INTEGER);
```

Parameters

Table 245-3 GET_ALL Procedure Parameters

Parameter	Description
names	The names of the cookies.
vals	The values of the cookies.
num_vals	The number of cookie-value pairs.

245.4.3 REMOVE Procedure

This procedure forces a cookie to expire immediately by setting the "expires" field of a Set-Cookie line in the HTTP header to "01-Jan-1990".

This procedure must be called within the context of an HTTP header.

Syntax

```
OWA_COOKIE.REMOVE(
   name           IN       VARCHAR2,
   val            IN       VARCHAR2,
   path           IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 245-4 REMOVE Procedure Parameters

Parameter	Description
name	The name of the cookie to expire.
val	The value of the cookie.
path	[Currently unused]

245.4.4 SEND procedure

This procedure generates a Set-Cookie line, which transmits a cookie to the client.

This procedure must occur in the context of an HTTP header.

Syntax

```
OWA_COOKIE.SEND(
   name           in       varchar2,
   value          in       varchar2,
   expires        in       date       DEFAULT NULL,
   path           in       varchar2   DEFAULT NULL,
   domain         in       varchar2   DEFAULT NULL,
   secure         in       varchar2   DEFAULT NULL);
```

Parameters

Table 245-5 SEND Procedure Parameters

Parameter	Description
`name`	The name of the cookie.
`value`	The value of the cookie.
`expires`	The date at which the cookie will expire
`path`	The value for the path field.
`domain`	The value for the domain field.
`secure`	If the value of this parameter is not `NULL`, the "secure" field is added to the line.

246
OWA_CUSTOM

The OWA_CUSTOM package provides a Global PLSQL Agent Authorization callback function.

It is used when PLSQL Agent's authorization scheme is set to GLOBAL or CUSTOM when there is no overriding OWA_CUSTOM package.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Constants
- Summary of OWA_CUSTOM Subprograms

246.1 OWA_CUSTOM Constants

OWA_CUSTOM defines several constants to use when specifying parameter values.

- `dbms_server_timezone CONSTANT VARCHAR2(3) := 'PST';`
- `dbms_server_gmtdiff CONSTANT NUMBER := NULL;`

246.2 Summary of OWA_CUSTOM Subprograms

The AUTHORIZE function is the only subprogram in the OWA_CUSTOM package.

Table 246-1 *OWA_CUSTOM Package Subprograms*

Subprogram	Description
AUTHORIZE Function	Provides a Global PLSQL Agent Authorization callback function

246.2.1 AUTHORIZE Function

This function is used when PLSQL Agent's authorization scheme is set to GLOBAL or CUSTOM when there is no overriding OWA_CUSTOM package.

Syntax

```
OWA_CUSTOM.AUTHORIZE
  RETURN BOOLEAN;
```

247
OWA_IMAGE

The `OWA_IMAGE` package provides an interface to access the coordinates where a user clicked on an image.

The chapter contains the following topics:

- Overview
- Types
- Variables
- Examples
- Summary of OWA_IMAGE Subprograms

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

247.1 OWA_IMAGE Overview

Use this package when you have any image map whose destination links invoke the PL/SQL Gateway.

247.2 OWA_IMAGE Types

This datatype (`point`) contain the `X` and `Y` values of a coordinate, and so provides the coordinates of a user's click on an imagemap.

It is defined as:

```
TYPE POINT IS TABLE OF VARCHAR2(32767) INDEX BY BINARY_INTEGER
```

247.3 OWA_IMAGE Variables

This package variable (`null_point`) of `TYPE POINT` is used to default point parameters. Both the X and the Y fields of this variable are `NULL`.

247.4 OWA_IMAGE Examples

This example shows use of `OWA_IMAGE`.

```
CREATE OR REPLACE PROCEDURE process_image
    (my_img in OWA_IMAGE.POINT)
    AS
    x integer := OWA_IMAGE.GET_X(my_img);
    y integer := OWA_IMAGE.GET_Y(my_img);
BEGIN
    /* process the coordinate */
END
```

247.5 Summary of OWA_IMAGE Subprograms

This table lists the OWA_IMAGE subprograms and briefly describes them.

Table 247-1 *OWA_IMAGE Package Subprograms*

Subprogram	Description
GET_X Function	Gets the X value of a point type
GET_Y Function	Gets the Y value of a point type

247.5.1 GET_X Function

This function returns the `X` coordinate of the point where the user clicked on an image map.

Syntax

```
OWA_IMAGE.GET_X(
   p        IN        point)
 RETURN INTEGER;

OWA_IMAGE.GET_X(
   p        IN        point)
 RETURN INTEGER PARALLEL_ENABLE DETERMINISTIC;
```

Parameters

Table 247-2 GET_X Function Parameters

Parameter	Description
p	The point where the user clicked.

Return Values

The X coordinate as an integer.

247.5.2 GET_Y Function

This function returns the Y coordinate of the point where the user clicked on an image map.

Syntax

```
OWA_IMAGE.GET_Y(
   p        IN        point)
 RETURN INTEGER;

OWA_IMAGE.GET_Y(
   p        IN        point)
 RETURN INTEGER PARALLEL_ENABLE DETERMINISTIC;
```

Parameters

Table 247-3 GET_Y Function Parameters

Parameter	Description
p	The point where the user clicked.

Return Values

The Y coordinate as an integer.

248
OWA_OPT_LOCK

The `OWA_OPT_LOCK` package contains subprograms that impose optimistic locking strategies so as to prevent lost updates.

This chapter contains the following topics:

- Overview
- Types
- Summary of OWA_OPT_LOCK Subprograms

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

248.1 OWA_OPT_LOCK Overview

The OWA_OPT_LOCK package contains subprograms that impose optimistic locking strategies, so as to prevent lost updates.

It checks if the row that the user is interested in updating has been changed by someone else in the meantime.

The PL/SQL Gateway cannot use conventional database locking schemes because HTTP is a stateless protocol. The `OWA_OPT_LOCK` package gives you two ways of dealing with the lost update problem:

- The hidden fields method stores the previous values in hidden fields in the HTML page. When the user requests an update, the PL/SQL Gateway checks these values against the current state of the database. The update operation is performed only if the values match. To use this method, call the owa_opt_lock.store_values procedure.
- The checksum method stores a checksum rather than the values themselves. To use this method, call the owa_opt_lock.checksum function.

These methods are optimistic. They do not prevent other users from performing updates, but they do reject the current update if an intervening update has occurred.

248.2 OWA_TEXT Types

This datatype is a PL/SQL table intended to hold ROWIDs.

```
TYPE VCARRAY IS TABLE OF VARCHAR2(2000) INDEX BY BINARY_INTEGER
```

Note that this is different from the `OWA_TEXT.VC_ARR DATA TYPE`.

248.3 Summary of OWA_OPT_LOCK Subprograms

This table lists the `OWA_OPT_LOCK` subprograms and briefly describes them.

Table 248-1 OWA_OPT_LOCK Package Subprograms

Subprogram	Description
CHECKSUM Functions	Returns the checksum value
GET_ROWID Function	Returns the ROWID value
STORE_VALUES Procedure	Stores unmodified values in hidden fields for later verification
VERIFY_VALUES Function	Verifies the stored values against modified values

248.3.1 CHECKSUM Functions

This function returns a `checksum` value for a specified string, or for a row in a table. For a row in a table, the function calculates the `checksum` value based on the values of the columns in the row. This function comes in two versions.

The first version returns a `checksum` based on the specified string. This is a "pure" 32-bit `checksum` executed by the database and based on the Internet 1 protocol.

The second version returns a `checksum` based on the values of a row in a table. This is a "impure" 32-bit checksum based on the Internet 1 protocol.

Syntax

```
OWA_OPT_LOCK.CHECKSUM(
   p_buff        IN       VARCHAR2)
 RETURN NUMBER;

OWA_OPT_LOCK.CHECKSUM(
   p_buff        IN       VARCHAR2)
 RETURN NUMBER PARALLEL_ENABLE DETERMINISTIC;

OWA_OPT_LOCK.CHECKSUM(
   p_owner       IN      VARCHAR2,
   p_tname       IN       VARCHAR2,
   p_rowid       IN       ROWID)
  RETURN NUMBER;
```

Parameters

Table 248-2 CHECKSUM Function Parameters

Parameter	Description
`p_buff`	The nstring where you want to calculate the `checksum`.
`p_owner`	The owner of the table.
`p_tname`	The table name.
`p_rowid`	The row in `p_tname` where you want to calculate the `checksum` value. Use the GET_ROWID Function to convert `VCARRAY` values to proper rowids.

248.3.2 GET_ROWID Function

This function returns the `ROWID` datatype from the specified `OWA_OPT_LOCK.VCARRAY DATA TYPE`.

Syntax

```
OWA_OPT_LOCK.GET_ROWID(
   p_old_values      IN      vcarray)
 RETURN ROWID;

OWA_OPT_LOCK.GET_ROWID(
   p_old_values      IN      vcarray)
 RETURN ROWID PARALLEL_ENABLE DETERMINISTIC;
```

Parameters

Table 248-3 GET_ROWID Function Parameters

Parameter	Description
`p_old_values`	This parameter is usually passed in from an HTML form.

248.3.3 STORE_VALUES Procedure

This procedure stores the column values of the row that you want to update later. The values are stored in hidden HTML form elements.

Syntax

```
OWA_OPT_LOCK.STORE_VALUES(
   p_owner        IN       VARCHAR2,
   p_tname        IN       VARCHAR2,
   p_rowid        IN       ROWID);
```

Parameters

Table 248-4 STORE_VALUES Procedure Parameters

Parameter	Description
`p_owner`	The owner of the table.
`p_tname`	The name of the table.
`p_rowid`	The row where you want to store values.

Usage Notes

Before updating the row, compare these values with the current row values to ensure that the values in the row have not been changed. If the values have changed, you can warn the users and let them decide if the update should take place.

The procedure generates series of hidden form elements:

- One hidden form element is created for the table owner. The name of the element is "*old*_p_tname", where `p_tname` is the name of the table. The value of the element is the owner name.

- One hidden form element is created for the table name. The name of the element is "*old*_p_tname", where p_tname is the name of the table. The value of the element is the table name.
- One element is created for each column in the row. The name of the element is "*old*_p_tname", where p_tname is the name of the table. The value of the element is the column value.

See also the VERIFY_VALUES Function.

248.3.4 VERIFY_VALUES Function

This function verifies whether values in the specified row have been updated since the last query.

Use this function with the STORE_VALUES Procedure.

Syntax

```
OWA_OPT_LOCK.VERIFY_VALUES(
   p_old_values IN vcarray)
 RETURN BOOLEAN;
```

Parameters

Table 248-5 *VERIFY_VALUES Function Parameters*

Parameter	Description
p_old_values	A PL/SQL table containing the following information: • p_old_values(1) specifies the owner of the table. • p_old_values(2) specifies the table. • p_old_values(3) specifies the rowid of the row to verify. The remaining indexes contain values for the columns in the table. Typically, this parameter is passed in from the HTML form, where you have previously called the STORE_VALUES Procedure to store the row values on hidden form elements.

Return Values

TRUE if no other update has been performed, otherwise FALSE.

249
OWA_PATTERN

The OWA_PATTERN package provides an interface to locate text patterns within strings and replace the matched string with another string.

> **See Also:**
>
> For more information about implementation of this package, see the following:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Types
- Operational Notes
- Summary of OWA_PATTERN Subprograms

249.1 OWA_PATTERN Types

You can use a pattern as both an input and output parameter. Because of this, you can pass the same regular expression to `OWA_PATTERN` function calls, and it only has to be parsed once.

- `OWA_PATTERN.PATTERN`

249.2 OWA_PATTERN Operational Notes

The OWA_PATTERN subprograms are overloaded. Specifically, there are six versions of `MATCH`, and four each of `AMATCH` and `CHANGE`.

The subprograms use the following parameters:

- `line` - This is the target to be examined for a match. It can be more than one line of text or a owa_text.multi_line datatype.
- `pat` - This is the pattern that the subprograms attempt to locate in line. The pattern can contain regular expressions. In the owa_pattern.change function and procedure, this parameter is called `from_str`.
- `flags` - This specifies whether the search is case-sensitive or if substitutions are done globally.

Use regular expressions with the subprograms in this package. You Specify a regular expression by creating the string you want to match interspersed with various wildcard tokens and quantifiers.

- Wildcards
- Quantifiers

- Flags

249.2.1 OWA_PATTERN Wildcards

Wildcard tokens match something other than themselves.

Table 249-1 ***Wildcard tokens recognized by OWA_PATTERN package***

Token	Description
^	Matches newline or the beginning of the target
$	Matches newline or the end of the target
\n	Matches newline
.	Matches any character except newline
\t	Matches tab
\d	Matches digits [0-9]
\D	Matches non-digits [not 0-9]
\w	Matches word characters (0-9, a-z, A-Z, or _)
\W	Matches non-word characters (not 0-9, a-z, A-Z, or _)
\s	Matches whitespace characters (blank, tab, or newline).
\S	Matches non-whitespace characters (not blank, tab, or newline)
\b	Matches "word" boundaries (between \w and \W)
\x<HEX>	Matches the value in the current character set of the two hexadecimal digits
\<OCT>	Matches the value in the current character set of the two or three octal digits
\	Followed by any character not covered by another case matches that character
&	Applies only to CHANGE. This causes the string that matched the regular expression to be included in the string that replaces it. This differs from the other tokens in that it specifies how a target is changed rather than how it is matched. This is explained further under CHANGE Functions and Procedures .

249.2.2 OWA_PATTERN Quantifiers

Any tokens except & can have their meaning extended by any of the following quantifiers. You can also apply these quantifiers to literals.

Table 249-2 Quantifiers

Quantifier	Description
?	0 or 1 occurrence(s)
*	0 or more occurrences
+	1 or more occurrence(s)
{n}	Exactly *n* occurrences
(n,}	At least *n* occurrences
{n,m}	At least *n*, but not more than *m*, occurrences

249.2.3 OWA_PATTERN Flags

In addition to targets and regular expressions, the `OWA_PATTERN` functions and procedures use flags to affect how they are interpreted.

Table 249-3 Flags

Flag	Description
`i`	This indicates a case-insensitive search.
`g`	This applies only to CHANGE. It indicates a global replace. That is, all portions of the target that match the regular expression are replaced.

249.3 Summary of OWA_PATTERN Subprograms

This table lists the OWA_PATTERN subprograms and briefly describes them.

Table 249-4 OWA_PATTERN Package Subprograms

Subprogram	Description
AMATCH Function	Determines if a string contains the specified pattern. It lets you specify where in the string the match has to occur
CHANGE Functions and Procedures	Replaces a pattern within a string. If you call it as a function it returns the number of times the regular expression was found and replaced
GETPAT Procedure	Generates a pattern datatype from a `VARCHAR2` type
MATCH Function	Determines if a string contains the specified pattern

249.3.1 AMATCH Function

This function specifies if a pattern occurs in a particular location in a string.

There are four versions to this function:

- The first and second versions of the function do not save the matched tokens (these are saved in the `backrefs` parameters in the third and fourth versions). The difference between the first and second versions is the `pat` parameter, which can be a `VARCHAR2` or a pattern datatype.
- The third and fourth versions of the function save the matched tokens in the `backrefs` parameter. The difference between the third and fourth versions is the `pat` parameter, which can be a VARCHAR2 or a pattern datatype.

> **Note:**
>
> If multiple overlapping strings match the regular expression, this function takes the longest match.

Syntax

```
OWA_PATTERN.AMATCH(
   line           IN       VARCHAR2,
   from_loc       IN       INTEGER,
   pat            IN       VARCHAR2,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;

OWA_PATTERN.AMATCH(
   line           IN       VARCHAR2,
   from_loc       IN       INTEGER,
   pat            IN OUT   PATTERN,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;

OWA_PATTERN.AMATCH(
   line           IN       VARCHAR2
   from_loc       IN       INTEGER
   pat            in       varchar2
   backrefs       OUT      owa_text.vc_arr
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;

OWA_PATTERN.AMATCH(
   line           IN       VARCHAR2
   from_loc       IN       INTEGER
   pat            IN OUT   PATTERN
   backrefs       OUT      owa_text.vc_arr
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;
```

Parameters

Table 249-5 AMATCH Function Parameters

Parameter	Description
`line`	The text to search in.
from_loc	The location (in number of characters) in `line` where the search is to begin.
`pat`	The string to match. It can contain regular expressions. This can be either a `VARCHAR2` or a pattern. If it is a pattern, the output value of this parameter is the pattern matched.
backrefs	The text that is matched. Each token that is matched is placed in a cell in the `OWA_TEXT.VC_ARR DATA TYPE` PL/SQL table.
flags	Whether or not the search is case-sensitive. If the value of this parameter is "i", the search is case-insensitive. Otherwise the search is case-sensitive.

Return Values

The index of the character after the end of the match, counting from the beginning of `line`. If there was no match, the function returns `0`.

249.3.2 CHANGE Functions and Procedures

This function or procedure searches and replaces a string or `multi_line` datatype. If multiple overlapping strings match the regular expression, this subprogram takes the longest match.

Syntax

```
OWA_PATTERN.CHANGE(
   line           IN OUT    VARCHAR2,
   from_str       IN        VARCHAR2,
   to_str         IN        VARCHAR2,
   flags          IN        VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;

OWA_PATTERN.CHANGE(
   line           IN OUT   VARCHAR2,
   from_str       IN       VARCHAR2,
   to_str         IN       VARCHAR2,
   flags          IN       VARCHAR2   DEFAULT NULL);

owa_pattern.change(
   mline          IN OUT   owa_text.multi_line,
   from_str       IN       VARCHAR2,
   to_str         IN       VARCHAR2,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;

OWA_PATTERN.CHANGE(
   mline          IN OUT   owa_text.multi_line,
   from_str       IN       VARCHAR2,
   to_str         IN       VARCHAR2,
   flags          IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 249-6 CHANGE Procedure Parameters

Parameter	Description
`line`	The text to search in. The output value of this parameter is the altered string.
mline	The text to search in. This is a `owa_text.multi_line datatype`. The output value of this parameter is the altered string.
`from_str`	The regular expression to replace.
to_str	The substitution pattern.
flags	Whether or not the search is case-sensitive, and whether or not changes are to be made globally. If "`i`" is specified, the search is case-insensitive. If "`g`" is specified, changes are made to all matches. Otherwise, the function stops after the first substitution is made.

Return Values

As a function, it returns the number of substitutions made. If the flag "g" is not used, this number can only be 0 or 1 and only the first match is replaced. The flag "g" specifies to replace all matches with the regular expression.

Examples

```
OWA_PATTERN.CHANGE('Cats in pajamas', 'C.+in', '& red ')
```

The regular expression matches the substring "`Cats in`". It then replaces this string with "& red". The ampersand character "`&`" indicates "`Cats in`" because that is what matched the regular expression. Thus, this procedure replaces the string "`Cats in pajamas`" with "`Cats in red`" If you call this as a function instead of a procedure, the value returned is 1, indicating that a single substitution has been made.

Example 2:

```
CREATE OR REPLACE PROCEDURE test_pattern as theline VARCHAR2(256);
num_found     INTEGER;
   BEGIN
       theline := 'what is the goal?';
       num_found := OWA_PATTERN.CHANGE(theline, 'goal', 'idea', 'g');
       HTP.PRINT(num_found); -- num_found is 1
       HTP.PRINT(theline); -- theline is 'what is the idea?'
   END;
/
SHOW ERRORS
```

249.3.3 GETPAT Procedure

This procedure converts a VARCHAR2 string into an `OWA_PATTERN.PATTERN DATA TYPE`.

Syntax

```
OWA_PATTERN.GETPAT(
   arg      IN          VARCHAR2,
   pat      IN OUT      pattern);
```

Parameters

Table 249-7 GETPAT Procedure Parameters

Parameter	Description
`arg`	The string to convert.
pat	the `OWA_PATTERN.PATTERN DATA TYPE` initialized with `arg`.

249.3.4 MATCH Function

This function determines if a string contains the specified pattern. The pattern can contain regular expressions. If multiple overlapping strings can match the regular expression, this function takes the longest match.

Syntax

```
owa_pattern.match(
   line           IN       VARCHAR2,
   pat            IN       VARCHAR2,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN BOOLEAN;

owa_pattern.match(
```

```
   line           IN       VARCHAR2,
   pat            IN OUT   PATTERN,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN BOOLEAN;

owa_pattern.match(
   line           IN       VARCHAR2,
   pat            IN       VARCHAR2,
   backrefs       OUT      owa_text.vc_arr,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN BOOLEAN;

OWA_PATTERN.MATCH(
   line           IN       VARCHAR2,
   pat            IN OUT   PATTERN,
   backrefs       OUT      owa_text.vc_arr,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN BOOLEAN;

owa_pattern.match(
   mline          IN       owa_text.multi_line,
   pat            IN       VARCHAR2,
   rlist          OUT      owa_text.row_list,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN BOOLEAN;

OWA_PATTERN.MATCH(
   mline          IN       owa_text.multi_line,
   pat            IN OUT   pattern,
   rlist          OUT      owa_text.row_list,
   flags          IN       VARCHAR2   DEFAULT NULL)
 RETURN BOOLEAN;
```

Parameters

Table 249-8 MATCH Function Parameters

Parameter	Description
line	The line to search in.
mline	The text to search in. This is a owa_text.multi_line datatype..
pat	The pattern to match. This is either a VARCHAR2 or a OWA_PATTERN.PATTERN DATA TYPE. It is a pattern, the output value of this parameter is the pattern matched.
backrefs	The text that is matched. Each token that is matched is placed in a cell in the OWA_TEXT.VC_ARR DATA TYPE PL/SQL table. This parameter is a row_list that holds each string in the target that was matched by a sequence of tokens in the regular expression.
rlist	An output parameter containing a list of matches.
flags	Whether or not the search is case-sensitive. If the value of this parameter is "i", the search is case-insensitive. Otherwise the search is case-sensitive.

Return Values

TRUE if a match was found, FALSE otherwise.

Examples

KAZOO is the target where it is searching for the zoo.* regular expression. The period indicates any character other than newline, and the asterisk matches 0 or more of the preceding characters. In this case, it matches any character other than the newline.

Therefore, this regular expression specifies that a matching target consists of zoo, followed by any set of characters neither ending in nor including a newline (which does not match the period). The i flag indicates to ignore case in the search. In this case, the function returns TRUE, which indicates that a match had been found.

```
boolean foundMatch;
foundMatch := owa_pattern.match('KAZOO', 'zoo.*', 'i');
```

The following example searches for the string "goal" followed by any number of characters in sometext. If found,

```
sometext  VARCHAR2(256);
pat       VARCHAR2(256);

sometext    := 'what is the goal?'
pat         := 'goal.*';
IF OWA_PATTERN.MATCH(sometext, pat)
   THEN
      HTP.PRINT('Match found');
   ELSE
     HTP.PRINT('Match not found');
END IF;
```

Operational Notes

- The regular expression in this function can be either a VARCHAR2 or an OWA_PATTERN.PATTERN DATA TYPE. Create AN OWA_PATTERN.PATTERN DATA TYPE from a string using the OWA_PATTERN.GETPAT procedure.
- Create a MULTI_LINE DATA TYPE from a long string using the OWA_TEXT.STREAM2MULTI procedure. If a multi_line is used, the rlist parameter specifies a list of chunks where matches were found.
- If the line is a string and not a multi_line, you can add an optional output parameter called backrefs. This parameter is a row_list that holds each string in the target that was matched by a sequence of tokens in the regular expression.

250
OWA_SEC

The `OWA_SEC` package provides an interface for custom authentication.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Operational Notes
- Summary of OWA_SEC Subprograms

250.1 OWA_SEC Operational Notes

Parameters that have default values are optional.

250.2 Summary of OWA_SEC Subprograms

This table lists the OWA_SEC subprograms in alphabetical order and briefly describes them.

Table 250-1 *OWA_SEC Package Subprograms*

Subprogram	Description
GET_CLIENT_HOSTNAME Function	Returns the client's hostname
GET_CLIENT_IP Function	Returns the client's IP address
GET_PASSWORD Function	Returns the password that the user entered
GET_USER_ID Function	Returns the username that the user entered
SET_AUTHORIZATION Procedure	Enables the PL/SQL application to use custom authentication
SET_PROTECTION_REALM Procedure	Defines the realm that the page is in

250.2.1 GET_CLIENT_HOSTNAME Function

This function returns the hostname of the client.

Syntax

```
OWA_SEC.GET_CLIENT_HOSTNAME
 RETURN VARCHAR2;
```

Return Values

The hostname.

250.2.2 GET_CLIENT_IP Function

This function returns the IP address of the client.

Syntax

```
OWA_SEC.GET_CLIENT_IP
  RETURN OWA_UTIL.IP_ADDRESS;
```

Return Values

The IP address. The `owa_util.ip_address datatype` is a PL/SQL table where the first four elements contain the four numbers of the IP address. For example, if the IP address is `123.45.67.89` and the variable `ipaddr` is of the owa_util.ip_address datatype, the variable would contain the following values:

```
ipaddr(1) = 123
ipaddr(2) = 45
ipaddr(3) = 67
ipaddr(4) = 89
```

250.2.3 GET_PASSWORD Function

This function returns the password that the user used to log in.

Syntax

```
OWA_SEC.GET_PASSWORD
  RETURN VARCHAR2;
```

Return Values

The password.

Usage Notes

For security reasons, this function returns a true value only when custom authentication is used. If you call this function when you are not using custom authentication, the function returns an undefined value. Thus, the database passwords are not exposed.

250.2.4 GET_USER_ID Function

This function returns the username that the user used to log in.

Syntax

```
OWA_SEC.GET_USER_ID
  RETURN VARCHAR2;
```

Return Values

The username.

250.2.5 SET_AUTHORIZATION Procedure

This procedure, called in the initialization portion of the OWA_CUSTOM package, sets the authorization scheme for the PL/SQL Gateway.

This implements your `authorize` function, which authorizes the user before his requested procedure is run. The placement of the `authorize` function depends on the scheme you select.

Syntax

```
OWA_SEC.SET_AUTHORIZATION(
    scheme          IN       INTEGER);
```

Parameters

Table 250-2 SET_AUTHORIZATION Procedure Parameters

Parameter	Description

`scheme`	The authorization scheme. It is one of the following schemes for `SET_AUTHORIZATION`:
	• `OWA_SEC.NO_CHECK` - Specifies that the PL/SQL application is not to do any custom authentication. This is the default.
	• `OWA_SEC.GLOBAL` - Defines an authorize function that is called for all users and all procedures. This is the `OWA_CUSTOM`.AUTHORIZE Function in the "`sys`" schema.
	• `OWA_SEC.PER_PACKAGE` - Define an `authorize` function that is called when procedures in a package or anonymous procedures are called. If the procedures are in a package, the package.`AUTHORIZE` function in the user's schema is called to authorize the user. If the procedures are not in a package, then the anonymous `authorize` function in the user's schema is called.
	• `OWA_SEC.CUSTOM` - Implements different authorize functions for each user. The function `OWA_CUSTOM`.AUTHORIZE Function in the user's schema is called to authorize the user. If the user's schema does not contain an `OWA_CUSTOM`.AUTHORIZE Function, the PL/SQL Gateway looks for it in the "sys" schema.
	The custom `authorize` function has the following signature:
	`FUNCTION AUTHORIZE`
	`  RETURN BOOLEAN;`
	If the function returns `TRUE`, authentication succeeded. If it returns `FALSE`, authentication failed. If the `authorize` function is not defined, the Gateway returns an error and fails.

Related Topics

- OWA_CUSTOM
  The `OWA_CUSTOM` package provides a Global PLSQL Agent Authorization callback function.

250.2.6 SET_PROTECTION_REALM Procedure

This procedure sets the realm of the page that is returned to the user. The user enters a username and login that already exist in the realm.

Syntax

```
OWA_SEC.SET_PROTECTION_REALM(
   realm      IN      VARCHAR2);
```

Parameters

Table 250-3 SET_PROTECTION_REALM Procedure Parameters

Parameter	Description
realm	The realm where the page belongs. This string is displayed to the user.

251
OWA_TEXT

The `OWA_TEXT` package contains subprograms used by `OWA_PATTERN` for manipulating strings. They are externalized so you can use them directly.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Types
- Summary of OWA_TEXT Subprograms

251.1 OWA_OPT_LOCK Types

There are three OWA_OPT_LOCK lock types.

- MULTI_LINE DATA TYPE
- ROW_LIST DATA TYPE
- VC_ARR DATA TYPE

251.1.1 MULTI_LINE DATA TYPE

This datatype is a PL/SQL record that holds large amounts of text. The rows field, of type `OWA_TEXT.VC_ARR DATA TYPE`, contains the text data in the record.

```
TYPE multi_line IS RECORD (
   rows           vc_arr,
   num_rows       INTEGER,
   partial_row    BOOLEAN);
```

251.1.2 ROW_LIST DATA TYPE

This is the datatype for holding data to be processed.

```
TYPE row_list IS RECORD (
   rows           int_arr,
   num_rows       INTEGER);

int_arr IS DEFINED AS:

TYPE int_arr IS TABLE OF INTEGER INDEX BY BINARY_INTEGER;
```

251.1.3 VC_ARR DATA TYPE

This is a component of the MULTI_LINE DATA TYPE and is used for holding large amounts of text.

```
TYPE vc_arr IS TABLE OF VARCHAR2(32767) INDEX BY BINARY_INTEGER;
```

Related Topics

- MULTI_LINE DATA TYPE
  This datatype is a PL/SQL record that holds large amounts of text. The rows field, of type `OWA_TEXT.VC_ARR DATA TYPE`, contains the text data in the record.

251.2 Summary of OWA_TEXT Subprograms

This table lists the `OWA_TEXT` subprograms and briefly describes them.

Table 251-1 OWA_TEXT Package Subprograms

Subprogram	Description
ADD2MULTI Procedure	Adds text to an existing `multi_line` type
NEW_ROW_LIST Function and Procedure	Creates a new `row_list`
PRINT_MULTI Procedure	Prints out the contents of a `multi_list`
PRINT_ROW_LIST Procedure	Prints out the contents of a `row_list`
STREAM2MULTI Procedure	Converts a varchar2 to a `multi_line` type

251.2.1 ADD2MULTI Procedure

This procedure adds content to an existing Multi_Line Data Type

Syntax

```
OWA_TEXT.ADD2MULTI(
   stream          IN       VARCHAR2,
   mline           IN OUT   multi_line,
```

continue IN BOOLEAN DEFAULT TRUE);

Parameters

Table 251-2 ADD2MULTI Procedure Parameters

Parameter	Description
`stream`	The text to add.
`mline`	The `OWA_TEXT.MULTI_LINE DATA TYPE`. The output of this parameter contains *stream*.
`continue`	If `TRUE`, the procedure appends *stream* within the previous final row (assuming it is less than 32K). If `FALSE`, the procedure places *stream* in a new row.

251.2.2 NEW_ROW_LIST Function and Procedure

This function or procedure creates a new `OWA_TEX.ROW_LIST DATA TYPE`.

The function version uses no parameters and returns a new empty `row_list`. The procedure version creates the `row_list` datatype as an output parameter.

Syntax

```
OWA_TEXT.NEW_ROW_LIST
  RETURN ROW_LIST;

OWA_TEXT.NEW_ROW_LIST(
  rlist    OUT     row_list);
```

Parameters

Table 251-3 *NEW_ROW_LIST Procedure Parameters*

Parameter	Description
`rlist`	This is an output parameter containing the new `row_list` datatype

Return Values

The function version returns the new `row_list` datatype.

251.2.3 PRINT_MULTI Procedure

This procedure uses the PRINT Procedures or the PRN Procedures to print the "rows" field of the `OWA_TEXT.MULTI_LINE DATA TYPE`.

Syntax

```
OWA_TEXT.PRINT_MULTI(
   mline       IN       multi_line);
```

Parameters

Table 251-4 PRINT_MULTI Procedure Parameters

Parameter	Description
`mline`	The multi_line datatype to print.

Return Values

The contents of the `multi_line`.

Related Topics

- PRINT Procedures
  These procedures generate the specified parameter as a string terminated with the `\n` newline character.
- PRN Procedures
  These procedures generate the specified parameter as a string.

251.2.4 PRINT_ROW_LIST Procedure

This procedure uses the PRINT Procedures or the PRN Procedures to print the "rows" field of the `OWA_TEXT.ROW_LIST DATA TYPE`.

Syntax

```
OWA_TEXT.PRINT_ROW_LIST(
   rlist       IN       multi_line);
```

Parameters

Table 251-5 PRINT_ROW_LIST Procedure Parameters

Parameter	Description
`rlist`	The row_list datatype to print.

Return Values

The contents of the `row_list`.

Related Topics

- PRINT Procedures
  These procedures generate the specified parameter as a string terminated with the `\n` newline character.
- PRN Procedures
  These procedures generate the specified parameter as a string.

251.2.5 STREAM2MULTI Procedure

This procedure converts a string to a `multi_line` datatype.

Syntax

```
OWA_TEXT.STREAM2MULTI(
   stream          IN       VARCHAR2
   mline           OUT      multi_line);
```

Parameters

Table 251-6 STREAM2MULTI Procedure Parameters

Parameter	Description
`stream`	The string to convert.
`mline`	The stream in `OWA_TEXT.MULTI_LINE DATA TYPE` format

252
OWA_UTIL

The `OWA_UTIL` package contains utility subprograms for performing operations such as getting the value of CGI environment variables, printing the data that is returned to the client, and printing the results of a query in an HTML table.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

This chapter contains the following topics:

- Overview
- Types
- Summary of OWA_UTIL Subprograms

252.1 OWA_UTIL Overview

The `OWA_UTIL` package contains three types of utility subprograms.

- Dynamic SQL Utilities enable you to produce pages with dynamically generated SQL code.
- HTML utilities enable you to retrieve the values of CGI environment variables and perform URL redirects.
- Date utilities enable correct date-handling. Date values are simple strings in HTML, but are treated as a datatype by the Oracle database.

252.2 OWA_UTIL Types

`OWA_UTIL` uses Types to specify creating information.

- DATETYPE Datatype
- IDENT_ARR Datatype
- IP_ADDRESS Datatype

252.2.1 DATETYPE Datatype

The TODATE Function converts an item of this type to the type `DATE`, which is understood and properly handled as data by the database. The procedure CHOOSE_DATE Procedure enables the user to select the desired date.

```
TYPE dateType IS TABLE OF VARCHAR2(10) INDEX BY BINARY_INTEGER;
```

Related Topics

- TODATE Function
  This function converts the DATETYPE Datatype to the standard Oracle `DATE` type.
- CHOOSE_DATE Procedure
  This procedure generates three HTML form elements that allow the user to select the day, the month, and the year.

252.2.2 IDENT_ARR Datatype

This datatype is used for an array.

```
TYPE ident_arr IS TABLE OF VARCHAR2(30) INDEX BY BINARY_INTEGER;
```

252.2.3 IP_ADDRESS Datatype

This datatype is used by the GET_CLIENT_IP Function in the OWA_SEC package.

```
TYPE ip_address IS TABLE OF INTEGER INDEX BY BINARY_INTEGER;
```

Related Topics

- GET_CLIENT_IP Function
  This function returns the IP address of the client.
- OWA_SEC
  The `OWA_SEC` package provides an interface for custom authentication.

252.3 Summary of OWA_UTIL Subprograms

This table lists the `OWA_UTIL` subprograms and briefly describes them.

Table 252-1 ***OWA_UTIL Package Subprograms***

Subprogram	Description
BIND_VARIABLES Function	Prepares a SQL query and binds variables to it
CALENDARPRINT Procedures	Prints a calendar
CELLSPRINT Procedures	Prints the contents of a query in an HTML table
CHOOSE_DATE Procedure	Generates HTML form elements that allow the user to select a date
GET_CGI_ENV Function	Returns the value of the specified CGI environment variable
GET_OWA_SERVICE_PATH Function	Returns the full virtual path for the PL/SQL Gateway
GET_PROCEDURE Function	Returns the name of the procedure that is invoked by the PL/SQL Gateway
HTTP_HEADER_CLOSE Procedure	Closes the HTTP header
LISTPRINT Procedure	Generates a HTML form element that contains data from a query
MIME_HEADER Procedure	Generates the Content-type line in the HTTP header
PRINT_CGI_ENV Procedure	Generates a list of all CGI environment variables and their values
REDIRECT_URL Procedure	Generates the Location line in the HTTP header
SHOWPAGE Procedure	Prints a page generated by the HTP and HTF packages in SQL*Plus

Table 252-1 (Cont.) *OWA_UTIL Package Subprograms*

Subprogram	Description
SHOWSOURCE Procedure	Prints the source for the specified subprogram
SIGNATURE procedure	Prints a line that says that the page is generated by the PL/SQL Agent
STATUS_LINE Procedure	Generates the Status line in the HTTP header
TABLEPRINT Function	Prints the data from a table in the database as an HTML table
TODATE Function	Converts dateType data to the standard PL/SQL date type
WHO_CALLED_ME Procedure	Returns information on the caller of the procedure.

252.3.1 BIND_VARIABLES Function

This function prepares a SQL query by binding variables to it, and stores the output in an opened cursor. Use this function as a parameter to a procedure sending a dynamically generated query. Specify up to 25 bind variables.

Syntax

```
OWA_UTIL.BIND_VARIABLES(
   theQuery       IN       VARCHAR2   DEFAULT NULL,
   bv1Name        IN       VARCHAR2   DEFAULT NULL,
   bv1Value       IN       VARCHAR2   DEFAULT NULL,
   bv2Name        IN       VARCHAR2   DEFAULT NULL,
   bv2Value       IN       VARCHAR2   DEFAULT NULL,
   bv3Name        IN       VARCHAR2   DEFAULT NULL,
   bv3Value       IN       VARCHAR2   DEFAULT NULL,
        ...
   bv25Name       IN       VARCHAR2   DEFAULT NULL,
   bv25Value      IN       VARCHAR2   DEFAULT NULL)
 RETURN INTEGER;
```

Parameters

Table 252-2 BIND_VARIABLES Function Parameters

Parameter	Description
theQuery	The SQL query statement which must be a SELECT statement
bv1Name	The name of the variable
bv1Value	The value of the variable

Return Values

An integer identifying the opened cursor.

252.3.2 CALENDARPRINT Procedures

These procedures creates a calendar in HTML with a visible border. Each date in the calendar can contain any number of hypertext links.

This procedure has 2 versions.

- Version 1 uses a hard-coded query stored in a varchar2 string.
- Version 2 uses a dynamic query prepared with the BIND_VARIABLES Function.

Syntax

```
OWA_UTIL.CALENDARPRINT(
   p_query        IN       VARCHAR2,
   p_mf_only      IN       VARCHAR2   DEFAULT 'N');

OWA_UTIL.CALENDARPRINT(
   p_cursor       IN       INTEGER,
   p_mf_only      IN       VARCHAR2   DEFAULT 'N');
```

Parameters

Table 252-3 CALENDARPRINT Procedure Parameters

Parameter	Description
p_query	A PL/SQL query.
p_cursor	A PL/SQL cursor containing the same format as p_query.
p_mf_only	If "N" (the default), the generated calendar includes Sunday through Saturday. Otherwise, it includes Monday through Friday only.

Usage Notes

Design your query as follows:

- The first column is a DATE. This correlates the information produced by the query with the calendar output generated by the procedure.
- The query output must be sorted on this column using ORDER BY.
- The second column contains the text, if any, that you want printed for that date.
- The third column contains the destination for generated links. Each item in the second column becomes a hypertext link to the destination given in this column. If this column is omitted, the items in the second column are simple text, not links.

252.3.3 CELLSPRINT Procedures

This procedure generates an HTML table from the output of a SQL query. SQL atomic data items are mapped to HTML cells and SQL rows to HTML rows. You must write the code to begin and end the HTML table.

There are nine versions of this procedure:

- The first version passes the results of a query into an index table. Perform the query and CELLSPRINT does the formatting. To have more control in generating an HTML table from the output of an SQL query, use the FORMAT_CELL Function in the " HTF" package.
- The second and third versions display rows (up to the specified maximum) returned by the query or cursor.
- The fourth and fifth versions exclude a specified number of rows from the HTML table. Use the fourth and fifth versions to scroll through result sets by saving the last row seen in a hidden form element.
- The sixth through ninth versions are the same as the first four versions, except that they return a row count output parameter.

Syntax

```
OWA_UTIL.CELLSPRINT(
   p_colCnt          IN    INTEGER,
   p_resultTbl       IN    vc_arr,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL);

OWA_UTIL.CELLSPRINT(
   p_theQuery        IN    VARCHAR2,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL);

OWA_UTIL.CELLSPRINT(
   p_theCursor       IN    INTEGER,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  iN    VARCHAR2   DEFAULT NULL);

OWA_UTIL.CELLSPRINT(
   p_theQuery        IN    VARCHAR2,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL,
   p_skip_rec        IN    NUMBER      DEFAULT 0,
   p_more_data       OUT   BOOLEAN);

OWA_UTIL.CELLSPRINT(
   p_theCursor       IN    INTEGER,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL,
   p_skip_rec        IN    NUMBER     DEFAULT 0,
   p_more_data       OUT   BOOLEAN);

OWA_UTIL.CELLSPRINT(
   p_theQuery        IN    VARCHAR2,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL,
   p_reccnt          OUT   NUMBER);

OWA_UTIL.CELLSPRINT(
   p_theCursor       IN    INTEGER,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL,
   p_reccnt          OUT   NUMBER);

OWA_UTIL.CELLSPRINT(
   p_theQuery        IN    VARCHAR2,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL,
   p_skip_rec        IN    NUMBER     DEFAULT 0,
   p_more_data       OUT   BOOLEAN
   p_reccnt          OUT   NUMBER);

OWA_UTIL.CELLSPRINT(
   p_theCursor       IN    INTEGER,
   p_max_rows        IN    NUMBER     DEFAULT 100,
   p_format_numbers  IN    VARCHAR2   DEFAULT NULL,
   p_skip_rec        IN    NUMBER     DEFAULT 0,
   p_more_data       OUT   BOOLEAN,
   p_reccnt          OUT   NUMBER);
```

Parameters

Table 252-4 CELLSPRINT Procedure Parameters

Parameter	Description
p_query	A PL/SQL query.
p_colCnt	The number of columns in the table.
p_theQuery	A SQL SELECT statement.
p_theCursor	A cursor ID. This can be the return value from the BIND_VARIABLES Function.
p_max_rows	The maximum number of rows to print.
p_format_numbers	If the value of this parameter is not NULL, number fields are right justified and rounded to two decimal places.
p_skip_rec	The number of rows to exclude from the HTML table.
p_more_data	TRUE if there are more rows in the query or cursor, FALSE otherwise.
p_reccnt	The number of rows that have been returned by the query. This value does not include skipped rows (if any).
p_resultTbl	The index table which will contain the result of the query. Each entry in the query will correspond to one column value.

Examples

This procedure generates

```
<tr><td>QueryResultItem</td><td>QueryResultItem</td></tr>...
```

252.3.4 CHOOSE_DATE Procedure

This procedure generates three HTML form elements that allow the user to select the day, the month, and the year.

Syntax

```
OWA_UTIL.CHOOSE_DATE(
   p_name         IN       VARCHAR2,
   p_date         IN       DATE       DEFAULT SYSDATE);
```

Parameters

Table 252-5 CHOOSE_DATE Procedure Parameters

Parameter	Description
p_name	The name of the form elements.
p_date	The initial date that is selected when the HTML page is displayed.

Usage Notes

- The parameter in the procedure that receives the data from these elements must be a GET_CGI_ENV Function.

- Use the TODATE Function to convert the GET_CGI_ENV Function value to the standard Oracle DATE datatype.

Examples

```
<SELECT NAME="p_name" SIZE="1">
<OPTION value="01">1
     ...
<OPTION value="31">31
</SELECT>
-
<SELECT NAME="p_name" SIZE="1">
<OPTION value="01">JAN
     ...
<OPTION value="12">DEC
</SELECT>
-
<SELECT NAME="p_name" SIZE="1">
<OPTION value="1992">1992
     ...
<OPTION value="2002">2002
</SELECT>
```

252.3.5 GET_CGI_ENV Function

This function returns the value of the specified CGI environment variable.

Syntax

```
OWA_UTIL.GET_CGI_ENV(
   param_name       IN      VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 252-6 GET_CGI_ENV Function Parameters

Parameter	Description
param_name	The name of the CGI environment variable. It is case-insensitive.

Return Values

The value of the specified CGI environment variable. If the variable is not defined, the function returns NULL.

252.3.6 GET_OWA_SERVICE_PATH Function

This function returns the full virtual path of the PL/SQL Gateway that is handling the request.

Syntax

```
OWA_UTIL.GET_OWA_SERVICE_PATH
 RETURN VARCHAR2;
```

Return Values

A virtual path of the PL/SQL Gateway that is handling the request.

252.3.7 GET_PROCEDURE Function

This function returns the name of the procedure that is being invoked by the PL/SQL Gateway.

Syntax

```
OWA_UTIL.GET_PROCEDURE
 RETURN VARCHAR2;
```

Return Values

The name of a procedure, including the package name if the procedure is defined in a package.

252.3.8 HTTP_HEADER_CLOSE Procedure

This procedure generates a newline character to close the HTTP header.

Syntax

```
OWA_UTIL.HTTP_HEADER_CLOSE;
```

Return Values

A newline character, which closes the HTTP header.

Usage Notes

- Use this procedure if you have not closed the header by using the `bclose_header` parameter in calls such as MIME_HEADER Procedure, REDIRECT_URL Procedure, or STATUS_LINE Procedure
- The HTTP header must be closed before any `HTP.PRINT` or `HTP.PRN` calls.

252.3.9 LISTPRINT Procedure

This procedure generates an HTML selection list form element from the output of a SQL query.

There are two versions of this procedure.

- The first version contains a hard-coded SQL query.
- The second version uses a dynamic query prepared with the BIND_VARIABLES Function.

Syntax

```
OWA_UTIL.LISTPRINT(
   p_theQuery     IN       VARCHAR2,
   p_cname        IN       VARCHAR2,
   p_nsize        IN       NUMBER,
   p_multiple     IN       BOOLEAN   DEFAULT FALSE);

OWA_UTIL.LISTPRINT(
   p_theCursor    IN    INTEGER,
   p_cname        IN       VARCHAR2,
   p_nsize        IN       NUMBER,
   p_multiple     IN       BOOLEAN   DEFAULT FALSE);
```

Parameters

Table 252-7 LISTPRINT Procedure Parameters

Parameter	Description
p_theQuery	The SQL query.
p_theCursor	The cursor ID. This can be the return value from the BIND_VARIABLES Function.
p_cname	The name of the HTML form element.
p_nsize	The size of the form element (this controls how many items the user can see without scrolling).
p_multiple	Whether multiple selection is permitted.

Usage Notes

The columns in the output of the query are handled in the following manner:

- The first column specifies the values that are sent back. These values are for the VALUE attribute of the OPTION tag.
- The second column specifies the values that the user sees.
- The third column specifies whether or not the row is marked as SELECTED in the OPTION tag. If the value is not NULL, the row is selected.

Examples

```
<SELECT NAME="p_cname" SIZE="p_nsize">
<OPTION SELECTED value='value_from_the_first_column'>value_from_the_second_column
<OPTION SELECTED value='value_from_the_first_column'>value_from_the_second_column
     ...
</SELECT>
```

252.3.10 MIME_HEADER Procedure

This procedure changes the default MIME header that the script returns. This procedure must come before any HTP.PRINT or HTP.PRN calls to direct the script not to use the default MIME header.

Syntax

```
OWA_UTIL.MIME_HEADER(
   ccontent_type    IN       VARCHAR2   DEFAULT 'text/html',
   bclose_header    IN       BOOLEAN    DEFAULT TRUE,
   ccharset         IN       VARCHAR2   DEFAULT NULL);
```

Parameters

Table 252-8 MIME_HEADER Procedure Parameters

Parameter	Description
ccontent_type	The MIME type to generate

Table 252-8 (Cont.) MIME_HEADER Procedure Parameters

Parameter	Description
bclose_header	Whether or not to close the HTTP header. If `TRUE`, two newlines are sent, which closes the HTTP header. Otherwise, one newline is sent, and the HTTP header remains open.
ccharset	The character set to use.The character set only makes sense if the MIME type is of type 'text'. Therefore, the character set is only tagged on to the Content-Type header only if the MIME type passed in is of type 'text'. Any other MIME type, such as 'image', will not have any character set tagged on.

Examples

```
Content-type: <ccontent_type>; charset=<ccharset>
```

so that

```
owa_util.mime_header('text/plain', false, 'ISO-8859-4')
```

generates

```
Content-type: text/plain; charset=ISO-8859-4\n
```

252.3.11 PRINT_CGI_ENV Procedure

This procedure generates all the CGI environment variables and their values made available by the PL/SQL Gateway to the stored procedure.

Syntax

```
OWA_UTIL.PRINT_CGI_ENV;
```

Examples

This procedure generates a list in the following format:

```
cgi_env_var_name = value\n
```

252.3.12 REDIRECT_URL Procedure

This procedure specifies that the application server is to visit the specified URL. The URL may specify either a Web page to return or a program to execute.

Syntax

```
OWA_UTIL.REDIRECT_URL(
   curl           IN       VARCHAR2
   bclose_header  IN       BOOLEAN    DEFAULT TRUE);
```

Parameters

Table 252-9 REDIRECT_URL Procedure Parameters

Parameter	Description
curl	The URL to visit.
bclose_header	Whether or not to close the HTTP header. If TRUE, two newlines are sent, which closes the HTTP header. Otherwise, one newline is sent, and the HTTP header remains open.

Usage Notes

This procedure must come before any HTP procedure or HTF function call.

Examples

This procedure generates

```
Location: <curl>\n\n
```

252.3.13 SHOWPAGE Procedure

This procedure prints out the HTML output of a procedure in SQL*Plus.

The procedure must use the HTP or HTF packages to generate the HTML page, and this procedure must be issued after the HTP or HTF page-generating subprogram has been called and before any other HTP or HTF subprograms are directly or indirectly called.

Syntax

```
OWA_UTIL.SHOWPAGE;
```

Usage Notes

- This method is useful for generating pages filled with static data.
- This procedure uses the DBMS_OUTPUT package and is limited to 32767 characters for each line and an overall buffer size of 1,000,000 bytes.

Examples

The output of htp procedure is displayed in SQL*Plus, SQL*DBA, or Oracle Server Manager. For example:

```
SQL> set serveroutput on
SQL> spool gretzky.html
SQL> execute hockey.pass("Gretzky")
SQL> execute owa_util.showpage
SQL> exit
```

This would generate an HTML page that could be accessed from Web browsers.

252.3.14 SHOWSOURCE Procedure

This procedure prints the source of the specified procedure, function, or package. If a procedure or function which belongs to a package is specified, then the entire package is displayed.

Syntax

```
OWA_UTIL.SHOWSOURCE (
   cname      IN     VARCHAR2);
```

Parameters

Table 252-10 SHOWSOURCE Procedure Parameters

Parameter	Description
cname	The function or procedure whose source you want to show.

252.3.15 SIGNATURE procedure

This procedure generates an HTML line followed by a signature line on the HTML document.

If a parameter is specified, the procedure also generates a hypertext link to view the PL/SQL source for that procedure. The link calls the SHOWSOURCE Procedure.

Syntax

```
OWA_UTIL.SIGNATURE;

OWA_UTIL.SIGNATURE (
   cname        IN        VARCHAR2);
```

Parameters

Table 252-11 SIGNATURE Procedure Parameters

Parameter	Description
cname	The function or procedure whose source you want to show.

Examples

Without a parameter, the procedure generates a line that looks like the following:

```
This page was produced by the PL/SQL Agent on August 9, 2001 09:30.
```

With a parameter, the procedure generates a signature line in the HTML document that looks like the following:

```
This page was produced by the PL/SQL Agent on 8/09/01 09:30
View PL/SQL Source
```

252.3.16 STATUS_LINE Procedure

This procedure sends a standard HTTP status code to the client.

This procedure must come before any `htp.print` or `htp.prn` calls so that the status code is returned as part of the header, rather than as "content data".

Syntax

```
OWA_UTIL.STATUS_LINE(
   nstatus        IN       INTEGER,
   creason        IN       VARCHAR2   DEFAULT NULL,
   bclose_header  IN       BOOLEAN    DEFAULT TRUE);
```

Parameters

Table 252-12 STATUS_LINE Procedure Parameters

Parameter	Description
`nstatus`	The status code.
`creason`	The string for the status code.
`bclose_header`	Whether or not to close the HTTP header. If `TRUE`, two newlines are sent, which closes the HTTP header. Otherwise, one newline is sent, and the HTTP header remains open.

Examples

This procedure generates

```
Status: <nstatus> <creason>\n\n
```

252.3.17 TABLEPRINT Function

This function generates either preformatted tables or HTML tables (depending on the capabilities of the user's browser) from database tables.

Syntax

```
OWA_UTIL.TABLEPRINT(
   ctable         IN       VARCHAR2,
   cattributes    IN       VARCHAR2   DEFAULT NULL,
   ntable_type    IN       INTEGER    DEFAULT HTML_TABLE,
   ccolumns       IN       VARCHAR2   DEFAULT '*',
   cclauses       IN    VARCHAR2   DEFAULT NULL,
   ccol_aliases   IN       VARCHAR2   DEFAULT NULL,
   nrow_min       IN       NUMBER     DEFAULT 0,
   nrow_max       IN       NUMBER     DEFAULT NULL)
 RETURN BOOLEAN;
```

Parameters

Table 252-13 TABLEPRINT Function Parameters

Parameter	Description
ctable	The database table.
cattributes	Other attributes to be included as-is in the tag.
ntable_type	How to generate the table. Specify HTML_TABLE to generate the table using <TABLE> tags or PRE_TABLE to generate the table using the <PRE> tags. These are constants: • HTML_TABLE CONSTANT INTEGER := 1; • PRE_TABLE CONSTANT INTEGER := 2;
ccolumns	A comma-delimited list of columns from ctable to include in the generated table.
cclauses	WHERE or ORDER BY clauses, which specify which rows to retrieve from the database table, and how to order them.
ccol_aliases	A comma-delimited list of headings for the generated table.
nrow_min	The first row, of those retrieved, to display.
nrow_max	The last row, of those retrieved, to display.

Return Values

Returns TRUE if there are more rows beyond the nrow_max requested, FALSE otherwise.

Usage Notes

- RAW columns are supported, but LONG RAW columns are not. References to LONG RAW columns will print the result 'Not Printable'.
- Note that in this function, cattributes is the second rather than the last parameter.

Examples

For browsers that do not support HTML tables, create the following procedure:

```
CREATE OR REPLACE PROCEDURE showemps IS
  ignore_more BOOLEAN;
BEGIN
  ignore_more := OWA_UTIL.TABLEPRINT('emp', 'BORDER', OWA_UTIL.PRE_TABLE);
END;
```

Requesting a URL such as

```
http://myhost:7777/pls/hr/showemps
```

returns to the following to the client:

```
<PRE>
---------------------------------------------------------------
| EMPNO |ENAME |JOB      |MGR   |HIREDATE   | SAL | COMM | DEPTNO |
---------------------------------------------------------------
| 7369| SMITH  | CLERK    | 7902 | 17-DEC-80 | 800 |      | 20 |
| 7499| ALLEN  | SALESMAN| 7698 | 20-FEB-81 | 1600 | 300 | 30 |
| 7521| WARD   | SALESMAN| 7698 | 22-FEB-81 | 1250 | 500 | 30 |
| 7566| JONES  | MANAGER | 7839 | 02-APR-81 | 2975 |     | 20 |
```

```
| 7654| MARTIN | SALESMAN| 7698 | 28-SEP-81 | 1250 | 1400| 30 |
| 7698| BLAKE  | MANAGER | 7839 | 01-MAY-81 | 2850 |     | 30 |
| 7782| CLARK  | MANAGER | 7839 | 09-JUN-81 | 2450 |     | 10 |
| 7788| SCOTT  | ANALYST | 7566 | 09-DEC-82 | 3000 |     | 20 |
| 7839| KING   | PRESIDENT |    | 17-NOV-81 | 5000 |     | 10 |
| 7844| TURNER | SALESMAN| 7698 | 08-SEP-81 | 1500 | 0   | 30 |
| 7876| ADAMS  | CLERK   | 7788 | 12-JAN-83 | 1100 |     | 20 |
| 7900| JAMES  | CLERK   | 7698 | 03-DEC-81 | 950  |     | 30 |
| 7902| FORD   | ANALYST | 7566 | 03-DEC-81 | 3000 |     | 20 |
| 7934| MILLER | CLERK   | 7782 | 23-JAN-82 | 1300 |     | 10 |
-----------------------------------------------------------------
</PRE>
```

To view the employees in department 10, and only their employee ids, names, and salaries, create the following procedure:

```
CREATE OR REPLACE PROCEDURE showemps_10 IS
 ignore_more BOOLEAN;
begin
 ignore_more := OWA_UTIL.TABLEPRINT
    ('EMP', 'BORDER', OWA_UTIL.PRE_TABLE,
     'empno, ename, sal', 'WHERE deptno=10 ORDER BY empno',
     'Employee Number, Name, Salary');
END;
```

A request for a URL like

```
http://myhost:7777/pls/hr/showemps_10
```

would return the following to the client:

```
<PRE>
-------------------------------------
| Employee Number |Name    | Salary |
-------------------------------------
| 7782            | CLARK  | 2450   |
| 7839            | KING   | 5000   |
| 7934            | MILLER | 1300   |
-------------------------------------
</PRE>
```

For browsers that support HTML tables, to view the department table in an HTML table, create the following procedure:

```
CREATE OR REPLACE PROCEDURE showdept IS
 ignore_more BOOLEAN;
BEGIN
 ignore_more := oWA_UTIL.TABLEPRINT('dept', 'BORDER');
END;
```

A request for a URL like

```
http://myhost:7777/pls/hr/showdept
```

would return the following to the client:

```
<TABLE BORDER>
<TR>
<TH>DEPTNO</TH>
<TH>DNAME</TH>
<TH>LOC</TH>
</TR>
<TR>
```

```
<TD ALIGN="LEFT">10</TD>
<TD ALIGN="LEFT">ACCOUNTING</TD>
<TD ALIGN="LEFT">NEW YORK</TD>
</TR>
<TR>
<TD ALIGN="LEFT">20</TD>
<TD ALIGN="LEFT">RESEARCH</TD>
<TD ALIGN="LEFT">DALLAS</TD>
</TR>
<TR>
<TD ALIGN="LEFT">30</TD>
<TD ALIGN="LEFT">SALES</TD>
<TD ALIGN="LEFT">CHICAGO</TD>
</TR>
<TR>
<TD ALIGN="LEFT">40</TD>
<TD ALIGN="LEFT">OPERATIONS</TD>
<TD ALIGN="LEFT">BOSTON</TD>
</TR>
</TABLE>
```

A Web browser would format this to look like the following table:

DEPTNO	DNAME	LOC
10	ACCOUNTING	NEW YORK
20	RESEARCH	DALLAS
30	SALES	CHICAGO

252.3.18 TODATE Function

This function converts the DATETYPE Datatype to the standard Oracle `DATE` type.

Syntax

```
OWA_UTIL.TODATE(
   p_dateArray      IN     dateType)
 RETURN DATE;
```

Parameters

Table 252-14 TODATE Function Parameters

Parameter	Description
p_dateArray	The value to convert.

Related Topics

- DATETYPE Datatype
  The TODATE Function converts an item of this type to the type `DATE`, which is understood and properly handled as data by the database. The procedure CHOOSE_DATE Procedure enables the user to select the desired date.

252.3.19 WHO_CALLED_ME Procedure

This procedure returns information (in the form of output parameters) about the PL/SQL code unit that invoked it.

Syntax

```
OWA_UTIL.WHO_CALLED_ME(
   owner          OUT      VARCHAR2,
   name           OUT      VARCHAR2,
   lineno         OUT      NUMBER,
   caller_t       OUT      VARCHAR2);
```

Parameters

Table 252-15 WHO_CALLED_ME Procedure Parameters

Parameter	Description
`owner`	The owner of the program unit.
`name`	The name of the program unit. This is the name of the package, if the calling program unit is wrapped in a package, or the name of the procedure or function if the calling program unit is a standalone procedure or function. If the calling program unit is part of an anonymous block, this is `NULL`.
`lineno`	The line number within the program unit where the call was made.
`caller_t`	The type of program unit that made the call. The possibilities are: package body, anonymous block, procedure, and function. Procedure and function are only for standalone procedures and functions.

253
SDO_CS

The `SDO_CS` package contains functions and procedures for working with coordinate systems.

You can perform explicit coordinate transformations on a single geometry or an entire layer of geometries (that is, all geometries in a specified column in a table).

For a complete description of this package within the context of Oracle Spatial, see `SDO_CS` in the *Oracle Spatial and Graph Developer's Guide*.

254
SDO_CSW_PROCESS

The `SDO_CSW_PROCESS` package contains subprograms for various processing operations related to support for Catalog Services for the Web (CSW).

For a complete description of this package within the context of Oracle Spatial, see `SDO_CSW_PROCESS` in the *Oracle Spatial and Graph Developer's Guide*..

255
SDO_GCDR

The `SDO_GCDR` package contains the Oracle Spatial geocoding subprograms, which let you geocode unformatted postal addresses.

For a complete description of this package within the context of Oracle Spatial, see SDO_GCDR in the *Oracle Spatial and Graph Developer's Guide*.

256
SDO_GEOM

The SDO_GEOM package contains the geometry functions.

These can be grouped into the following categories (with examples of each):

- Relationship (True/False) between two objects: RELATE, WITHIN_DISTANCE
- Validation: VALIDATE_GEOMETRY_WITH_CONTEXT, VALIDATE_LAYER_WITH_CONTEXT
- Single-object operations: SDO_ARC_DENSIFY, SDO_AREA, SDO_BUFFER, SDO_CENTROID, SDO_CONVEXHULL, SDO_LENGTH, SDO_MBR, SDO_POINTONSURFACE
- Two-object operations: SDO_DISTANCE, SDO_DIFFERENCE, SDO_INTERSECTION, SDO_UNION, SDO_XOR

For a complete description of this package within the context of Oracle Spatial, see SDO_GEOM in the *Oracle Spatial and Graph Developer's Guide*.

257
SDO_GEOR

The `SDO_GEOR` package contains functions and procedures for the Oracle Spatial GeoRaster feature, which lets you store, index, query, analyze, and deliver raster image data and its associated spatial vector geometry data and metadata.

For complete description of this package within the context of Oracle Spatial, see `SDO_GEOR` in the *Oracle Spatial and Graph GeoRaster Developer's Guide*.

258
SDO_GEOR_ADMIN

The `SDO_GEOR_ADMIN` package contains subprograms for administrative operations related to GeoRaster.

For a complete description of this package within the context of Oracle Spatial, see `SDO_GEOR_ADMIN` in the *Oracle Spatial and Graph GeoRaster Developer's Guide*.

259
SDO_GEOR_AGGR

The `SDO_GEOR_AGGR` package provides an interface to the SDO_GEOR_AGGR package for performing aggregate operations on GeoRaster objects.

For a complete description of this package within the context of Oracle Spatial, see `SDO_GEOR_AGGR` in the *Oracle Spatial and Graph GeoRaster Developer's Guide*.

260
SDO_GEOR_RA

The `SDO_GEOR_RA` package provides an interface to the `SDO_GEOR_RA` package for performing raster algebra and analytic operations related to GeoRaster.

For a complete description of this package within the context of Oracle Spatial, see `SDO_GEOR_RA` in the *Oracle Spatial and Graph GeoRaster Developer's Guide*.

261
SDO_GEOR_UTL

The `SDO_GEOR_UTL` package contains utility functions and procedures for the Oracle Spatial GeoRaster feature, including those related to using triggers with GeoRaster data.

For complete description of this package within the context of Oracle Spatial, see `SDO_GEOR_UTL` in the *Oracle Spatial and Graph GeoRaster Developer's Guide*.

262
SDO_LRS

The `SDO_LRS` package contains functions that create, modify, query, and convert linear referencing elements.

For a complete description of this package within the context of Oracle Spatial, see `SDO_LRS` in the *Oracle Spatial and Graph Developer's Guide*.

263
SDO_MIGRATE

The `SDO_MIGRATE` package lets you upgrade geometry tables from previous releases of Oracle Spatial.

For a complete description of this package within the context of Oracle Spatial, see `SDO_MIGRATE` in the *Oracle Spatial and Graph Developer's Guide*.

264
SDO_NET

The `SDO_NET` package contains functions and procedures for working with data modeled as nodes and links in a network.

For a complete description of this package within the context of Oracle Spatial, see `SDO_NET` in the *Oracle Spatial and Graph Topology Data Model and Network Data Model Graph Developer's Guide*.

265
SDO_NFE

The SDO_NFE package contains functions and procedures for performing network feature editing.

For a complete description of this package within the context of Oracle Spatial and Graph, see SDO_NFE in the *Oracle Spatial and Graph Topology Data Model and Network Data Model Graph Developer's Guide*.

266
SDO_OLS

The `SDO_OLS` package contains subprograms for Spatial OpenLS support.

For a complete description of this package within the context of Oracle Spatial, see `SDO_OLS` in the *Oracle Spatial and Graph Developer's Guide*.

267
SDO_PC_PKG

The `SDO_PC_PKG` package contains subprograms to support the use of point clouds in Spatial.

For a complete description of this package within the context of Oracle Spatial, see `SDO_PC_PKG` in the *Oracle Spatial and Graph Developer's Guide*.

268
SDO_SAM

The `SDO_SAM` package contains functions and procedures for spatial analysis and data mining.

For a complete description of this package within the context of Oracle Spatial, see `SDO_SAM` in the *Oracle Spatial and Graph Developer's Guide*.

269
SDO_TIN_PKG

The `SDO_TIN_PKG` package contains subprograms to support the use of triangulated irregular networks (TINs) in Spatial.

For a complete description of this package within the context of Oracle Spatial, see `SDO_TIN_PKG` in the

270
SDO_TOPO

The `SDO_TOPO` package contains subprograms for creating and managing Oracle Spatial topologies.

For a complete description of this package within the context of Oracle Spatial, see `SDO_TOPO` in the *Oracle Spatial and Graph Topology Data Model and Network Data Model Graph Developer's Guide*.

271
SDO_TOPO_MAP

The `SDO_TOPO_MAP` package contains subprograms for editing Oracle Spatial topologies using a cache (TopoMap object).

For a complete description of this package within the context of Oracle Spatial, see `SDO_TOPO_MAP` in the *Oracle Spatial and Graph Topology Data Model and Network Data Model Graph Developer's Guide*.

272
SDO_TUNE

The `SDO_TUNE` package contains Spatial tuning functions and procedures.

For complete description of this package within the context of Oracle Spatial, see `SDO_TUNE` in the *Oracle Spatial and Graph Developer's Guide*.

273
SDO_UTIL

The `SDO_UTIL` package contains the utility functions and procedures for Oracle Spatial.

For complete description of this package within the context of Oracle Spatial, see `SDO_UTIL` in the *Oracle Spatial and Graph Developer's Guide*.

274
SDO_WFS_LOCK

The `SDO_WFS_LOCK` package contains subprograms for WFS support for registering and unregistering feature tables. Registering a feature table enables the table for WFS transaction locking; unregistering a feature table disables the table for WFS transaction locking.

For a complete description of this package within the context of Oracle Spatial, see `SDO_WFS_LOCK` in the *Oracle Spatial and Graph Developer's Guide*.

275
SDO_WFS_PROCESS

The `SDO_WFS_PROCESS` package contains subprograms for various processing operations related to support for Web Feature Services.

For a complete description of this package within the context of Oracle Spatial, see `SDO_WFS_PROCESS` in the *Oracle Spatial and Graph Developer's Guide*.

276
SEM_APIS

The `SEM_APIS` package contains subprograms for working with the Resource Description Framework (RDF) and Web Ontology Language (OWL) in an Oracle database.

For a complete description of this package within the context of Oracle Database semantic technology support, see `SEM_APIS` in the *Oracle Spatial and Graph RDF Semantic Graph Developer's Guide*.

277
SEM_OLS

The `SEM_OLS` package provides an interface to the `SEM_OLS` package for providing triple-level security to RDF data, using Oracle Label Security (OLS).

For a complete description of this package, see `SEM_OLS` in the *Oracle Spatial and Graph RDF Semantic Graph Developer's Guide*.

278
SEM_PERF

The `SEM_PERF` package contains subprograms for examining and enhancing the performance of the Resource Description Framework (RDF) and Web Ontology Language (OWL) support in an Oracle database.

For a complete description of this package within the context of Oracle Database semantic technology support, see `SEM_PERF` in the *Oracle Spatial and Graph RDF Semantic Graph Developer's Guide*.

279
SEM_RDFCTX

The `SEM_RDFCTX` package contains subprograms for managing extractor policies and semantic indexes created for documents.

For a complete description of this package within the context of Oracle Database semantic technology support, see `SEM_RDFCTX` in the *Oracle Spatial and Graph RDF Semantic Graph Developer's Guide*

280
SEM_RDFSA

The `SEM_RDFSA` package contains subprograms for providing fine-grained access control to RDF data, using either a virtual private database (VPD) or Oracle Label Security (OLS).

For a complete description of this package within the context of Oracle Database semantic technology support, see `SEM_RDFSA` in the *Oracle Spatial and Graph RDF Semantic Graph Developer's Guide*.

281
UTL_CALL_STACK

The `UTL_CALL_STACK` package provides an interface to provide information about currently executing subprograms.

Functions return subprogram names, unit names, owner names, edition names, and line numbers for given dynamic depths. Other functions return error stack information.

This chapter contains the following topics:

- Overview
- Security Model
- UTL_CALL_STACK Operational Notes
- UTL_CALL_STACK Exceptions
- Data Structures
- Summary of UTL_CALL_STACK Subprograms

> **See Also:**
>
> - *Oracle Database PL/SQL Language Reference* regarding Conditional Compilation
> - *Oracle Database Development Guide* regarding Using PL/Scope and Using the PL/SQL Hierarchical Profiler

281.1 UTL_CALL_STACK Overview

The `UTL_CALL_STACK` package provides an interface for PL/SQL programmers to obtain information about currently executing programs including the subprogram name from dynamic and lexical stacks and the depths of those stacks.

Individual functions return subprogram names, unit names, owner names, edition names, and line numbers for given dynamic depths. More functions return error stack information. Such information can be used to create more revealing error logs and application execution traces.

Dynamic Depth

The dynamic depth of an executing instance of a PL/SQL subprogram is defined recursively.

- The dynamic depth of the currently executing subprogram instance is one.
- Otherwise, the dynamic depth of the subprogram instance is one more than the dynamic depth of the subprogram it invoked.
- If there is a SQL, Java, or other non-PL/SQL context that invoked or was invoked by an executing subprogram, it occupies a level on the call stack as if it were a subprogram.

In the case of a call stack in which A calls B, which calls C, which calls D, which calls E, which calls F, which calls E, this stack can be written as a line with the dynamic depths underneath:

```
A B C D E F E
7 6 5 4 3 2 1
```

Lexical Depth

The lexical depth of a PL/SQL subprogram is defined recursively.

- The lexical depth of a unit, an anonymous block, trigger, or ADT is one (1).
- The lexical depth of a subprogram defined within another object is one plus the lexical depth of that object.

Blocks do not affect lexical depth.

Error Depth

The error depth is the number of errors on the error stack.

For example, consider the following anonymous block.

```
BEGIN
  BEGIN
    ... (1)
    raise zero_divide;
  EXCEPTION
    when others then
      raise no_data_found;
  END;
EXCEPTION
  WHEN others THEN
    ... (2)
END;
```

The error depth at (1) is zero and at (2) is two.

Backtrace

The backtrace is a trace from where the exception was thrown to where the backtrace was examined.

Consider a call stack in which A calls B which calls C and C raises an exception. If the backtrace was examined in C, the backtrace would have one unit, C, and the backtrace depth would be one. If it was examined in A, it would have three units, A, B and C, and backtrace depth would be three.

The depth of a backtrace is zero in the absence of an exception.

281.2 UTL_CALL_STACK Security Model

EXECUTE on UTL_CALL_STACK is granted to PUBLIC.

The UTL_CALL_STACK package does not show wrapped program units. For example, consider a call stack in which program unit A calls B, which calls C, and in turn calls UTL_CALL_STACK to determine the subprogram list. If program unit B is wrapped, then the subprogram list onlys shows program unit C.

281.3 UTL_CALL_STACK Operational Notes

Certain operational notes apply to `UTL_CALL_STACK`.

- Compiler optimizations can change lexical, dynamic and backtrace depth.
- `UTL_CALL_STACK` is not supported past RPC boundaries. For example, if A calls remote procedure B, B will not be able to obtain information about A using `UTL_CALL_STACK`.
- Lexical unit information is available through the PL/SQL conditional compilation feature and is therefore not exposed through `UTL_CALL_STACK`.

281.4 UTL_CALL_STACK Exceptions

This table lists the exceptions raised by `UTL_CALL_STACK`.

Table 281-1 Exceptions Raised by UTL_CALL_STACK

Exception	Error Code	Description
BAD_DEPTH_INDICATOR	64610	This exception is raised when a provided depth is out of bounds. Dynamic and lexical depth are positive integer values. Error and backtrace depths are non-negative integer values and are zero only in the absence of an exception.

281.5 UTL_CALL_STACK Data Structures

The `UTL_CALL_STACK` package defines a `VARRAY` type, `UNIT_QUALIFIED_NAME`.

VARRAY Type

- UNIT_QUALIFIED_NAME

281.5.1 UNIT_QUALIFIED_NAME

This data structure is a varray whose individual elements are, in order, the unit name, any lexical parents of the subprogram, and the subprogram name.

```
TYPE UNIT_QUALIFIED_NAME IS VARRAY(256) OF VARCHAR2(32767);
```

Example

Consider the following contrived PL/SQL procedure:

```
PROCEDURE topLevel IS
  FUNCTION localFunction(...) RETURNS VARCHAR2 IS
    FUNCTION innerFunction(...) RETURNS VARCHAR2 IS
      BEGIN
        DECLARE
          localVar PLS_INTEGER;
        BEGIN
          ... (1)
        END;
      END;
  BEGIN
```

```
  ...
END;
```

The unit qualified name at (1) would be

```
["topLevel", "localFunction", "innerFunction"]
```

If the unit were an anonymous block, the unit name would be `"__anonymous_block"`

281.6 Summary of UTL_CALL_STACK Subprograms

This table lists the subprograms in the `UTL_CALL_STACK` package.

Table 281-2 UTL_CALL_STACK Package Subprograms

Subprogram	Description
BACKTRACE_DEPTH Function	Returns the number of backtrace items in the backtrace
BACKTRACE_LINE Function	Returns the line number of the unit at the specified backtrace depth
BACKTRACE_UNIT Function	Returns the name of the unit at the specified backtrace depth
CURRENT_EDITION Function	Returns the current edition name of the unit of the subprogram at the specified dynamic depth
CONCATENATE_SUBPROGRAM Function	Returns a concatenated form of a unit-qualified name
DYNAMIC_DEPTH Function	Returns the number of subprograms on the call stack
ERROR_DEPTH Function	Returns the number of errors on the error stack
ERROR_MSG Function	Returns the error message of the error at the specified error depth
ERROR_NUMBER Function	Returns the error number of the error at the specified error depth
LEXICAL_DEPTH Function	Returns the lexical nesting level of the subprogram at the specified dynamic depth
OWNER Function	Returns the owner name of the unit of the subprogram at the specified dynamic depth
UNIT_LINE Function	Returns the line number of the unit of the subprogram at the specified dynamic depth
SUBPROGRAM Function	Returns the unit-qualified name of the subprogram at the specified dynamic depth

281.6.1 BACKTRACE_DEPTH Function

This function returns the number of backtrace items in the backtrace.

Syntax

```
UTL_CALL_STACK.BACKTRACE_DEPTH
 RETURN PLS_INTEGER;
```

Return Values

The number of backtrace items in the backtrace, zero in the absence of an exception.

281.6.2 BACKTRACE_LINE Function

This function returns the line number of the unit at the specified backtrace depth.

Syntax

```
UTL_CALL_STACK.BACKTRACE_LINE (
   backtrace_depth    IN    PLS_INTEGER)
 RETURN PLS_INTEGER;
```

Parameters

Table 281-3 *BACKTRACE_LINE Function Parameters*

Parameter	Description
backtrace_depth	Depth in backtrace

Return Values

The line number of the unit at the specified backtrace depth

281.6.3 BACKTRACE_UNIT Function

This function returns the name of the unit at the specified backtrace depth.

Syntax

```
UTL_CALL_STACK.BACKTRACE_UNIT (
   backtrace_depth    IN    PLS_INTEGER)
 RETURN VARCHAR2;
```

Parameters

Table 281-4 *BACKTRACE_UNIT Function Parameters*

Parameter	Description
backtrace_depth	Depth in backtrace

Return Values

The name of the unit at the specified backtrace depth

281.6.4 CURRENT_EDITION Function

This function returns the current edition name of the unit of the subprogram at the specified dynamic depth.

Syntax

```
UTL_CALL_STACK.CURRENT_EDITION (
   dynamic_depth     IN     PLS_INTEGER)
 RETURN VARCHAR2;
```

Parameters

Table 281-5 *CURRENT_EDITION Function Parameters*

Parameter	Description
dynamic_depth	Depth in the error stack

Return Values

The current edition name of the unit of the subprogram at the specified dynamic depth

281.6.5 CONCATENATE_SUBPROGRAM Function

This function returns a concatenated form of a unit-qualified name.

Syntax

```
UTL_CALL_STACK.CONCATENATE_SUBPROGRAM (
   qualified_name    IN    UNIT_QUALIFIED_NAME)
 RETURN VARCHAR2;
```

Parameters

Table 281-6 *CONCATENATE_SUBPROGRAM Function Parameters*

Parameter	Description
qualified_name	A unit-qualified name

Return Values

A string of the form `UNIT.SUBPROGRAM.LOCAL_SUBPROGRAM`

281.6.6 DYNAMIC_DEPTH Function

This function returns the number of subprograms on the call stack.

Syntax

```
UTL_CALL_STACK.DYNAMIC_DEPTH
 RETURN PLS_INTEGER;
```

Return Values

The number of subprograms on the call stack

281.6.7 ERROR_DEPTH Function

This function returns the number of errors on the error stack.

Syntax

```
UTL_CALL_STACK.ERROR_DEPTH
 RETURN PLS_INTEGER;
```

Return Values

The number of errors on the error stack

281.6.8 ERROR_MSG Function

This function returns the error message of the error at the specified error depth.

Syntax

```
UTL_CALL_STACK.ERROR_MSG (
   error_depth    IN    PLS_INTEGER)
 RETURN VARCHAR2;
```

Parameters

Table 281-7 *ERROR_MSG Function Parameters*

Parameter	Description
error_depth	Depth in the error stack

Return Values

The error message of the error at the specified error depth.

281.6.9 ERROR_NUMBER Function

This function returns the error number of the error at the specified error depth.

Syntax

```
UTL_CALL_STACK.ERROR_NUMBER (
   error_depth    IN    PLS_INTEGER)
 RETURN PLS_INTEGER;
```

Parameters

Table 281-8 *ERROR_NUMBER Function Parameters*

Parameter	Description
error_depth	Depth in the call stack

Return Values

The error number of the error at the specified error depth

281.6.10 LEXICAL_DEPTH Function

This function returns the lexical nesting level of the subprogram at the specified dynamic depth.

Syntax

```
UTL_CALL_STACK.LEXICAL_DEPTH (
   dynamic_depth    IN    PLS_INTEGER)
 RETURN PLS_INTEGER;
```

Parameters

Table 281-9 *LEXICAL_DEPTH Function Parameters*

Parameter	Description
dynamic_depth	Depth in the call stack

Return Values

The lexical nesting level of the subprogram at the specified dynamic depth

281.6.11 OWNER Function

This function returns the owner name of the unit of the subprogram at the specified dynamic depth.

Syntax

```
UTL_CALL_STACK.OWNER (
   dynamic_depth    IN    PLS_INTEGER)
 RETURN VARCHAR2;
```

Parameters

Table 281-10 *OWNER Function Parameters*

Parameter	Description
dynamic_depth	Depth in the call stack

Return Values

The owner name of the unit of the subprogram at the specified dynamic depth

281.6.12 UNIT_LINE Function

This function returns the line number of the unit of the subprogram at the specified dynamic depth.

Syntax

```
UTL_CALL_STACK.UNIT_LINE (
   dynamic_depth    IN    PLS_INTEGER)
 RETURN PLS_INTEGER;
```

Parameters

Table 281-11 *UNIT_LINE Function Parameters*

Parameter	Description
dynamic_depth	Depth in the call stack

Return Values

The line number of the unit of the subprogram at the specified dynamic depth

281.6.13 SUBPROGRAM Function

This function returns the unit-qualified name of the subprogram at the specified dynamic depth.

Syntax

```
UTL_CALL_STACK.SUBPROGRAM (
   dynamic_depth    IN    PLS_INTEGER)
 RETURN UNIT_QUALIFIED_NAME;
```

Parameters

Table 281-12 *SUBPROGRAM Function Parameters*

Parameter	Description
dynamic_depth	Depth in the call stack

Return Values

Returns the unit-qualified name of the subprogram at the specified dynamic depth

282
UTL_COLL

The UTL_COLL package lets PL/SQL programs use collection locators to query and update.

This chapter contains the following topics:

- Summary of UTL_COLL Subprograms

282.1 Summary of UTL_COLL Subprograms

The UTL_COLL package has one subprogram, the IS_LOCATOR function.

Table 282-1 UTL_COLL Package Subprograms

Subprogram	Description
IS_LOCATOR Function	Determines whether a collection item is actually a locator or not

282.1.1 IS_LOCATOR Function

This function determines whether a collection item is actually a locator or not.

Syntax

```
UTL_COLL.IS_LOCATOR (
   coln IN STANDARD)
  RETURNS BOOLEAN;
```

Pragmas

Asserts WNDS, WNPS and RNPS pragmas

Parameters

Table 282-2 *IS_LOCATOR Function Parameters*

Parameter	Description
coln	Nested table or varray item.

Return Values

Table 282-3 IS_LOCATOR Function Return Values

Return Value	Description
1	Collection item is indeed a locator.
0	Collection item is not a locator.

Examples

```
CREATE OR REPLACE TYPE list_t as TABLE OF VARCHAR2(20);
/

CREATE OR REPLACE TYPE phone_book_t AS OBJECT (
  pno  number,
  ph   list_t );
/

CREATE TABLE phone_book OF phone_book_t
      NESTED TABLE ph STORE AS nt_ph;
CREATE TABLE phone_book1 OF phone_book_t
      NESTED TABLE ph STORE AS nt_ph_1 RETURN LOCATOR;

INSERT INTO phone_book VALUES(1, list_t('650-633-5707','650-323-0953'));
INSERT INTO phone_book1 VALUES(1, list_t('415-555-1212'));

CREATE OR REPLACE PROCEDURE chk_coll IS
  plist list_t;
  plist1 list_t;
BEGIN
  SELECT ph INTO plist FROM phone_book WHERE pno=1;

  SELECT ph INTO plist1 FROM phone_book1 WHERE pno=1;

  IF (UTL_COLL.IS_LOCATOR(plist)) THEN
    DBMS_OUTPUT.PUT_LINE('plist is a locator');
  ELSE
    DBMS_OUTPUT.PUT_LINE('plist is not a locator');
  END IF;

  IF (UTL_COLL.IS_LOCATOR(plist1)) THEN
    DBMS_OUTPUT.PUT_LINE('plist1 is a locator');
  ELSE
    DBMS_OUTPUT.PUT_LINE('plist1 is not a locator');
  END IF;

END chk_coll;

SET SERVEROUTPUT ON
EXECUTE chk_coll;
```

283
UTL_COMPRESS

The UTL_COMPRESS package provides a set of data compression utilities.

This chapter contains the following topics:

- Constants
- Exceptions
- Operational Notes
- Summary of UTL_COMPRESS Subprograms

283.1 UTL_COMPRESS Constants

The maximum number of handles for piecewise operations can be defined by a constant.

```
UTLCOMP_MAX_HANDLE   CONSTANT  PLS_INTEGER := 5;
```

283.2 UTL_COMPRESS Exceptions

This table describes exceptions raised by UTL_COMPRESS subprograms.

Table 283-1 UTL_COMPRESS Exceptions

Exception	Description
BUFFER_TOO_SMALL	The compressed representation is too big.
DATA_ERROR	The input or output data stream was found to be an invalid format.
INVALID_ARGUMENT	One of the arguments was an invalid type or value.
INVALID_HANDLE	Invalid handle for piecewise compress or uncompress.
STREAM_ERROR	An error occurred during compression or uncompression of the data stream

283.3 UTL_COMPRESS Operational Notes

Certain operational notes apply to UTL_COMPRESS.

- It is the caller's responsibility to free the temporary LOB returned by the LZ* functions with DBMS_LOB.FREETEMPORARY call.
- A BFILE passed into LZ_COMPRESS* or lZ_UNCOMPRESS* has to be opened by DBMS_LOB.FILEOPEN.
- Under special circumstances (especially if the input has already been compressed) the output produced by one of the UTL_COMPRESS subprograms may be the same size, or even slightly larger than, the input.

- The output of the `UTL_COMPRESS` compressed data is compatible with `gzip`(with `-n` option)/`gunzip` on a single file.

283.4 Summary of UTL_COMPRESS Subprograms

This table lists the `UTL_COMPRESS` subprograms and briefly describes them.

Table 283-2 *UTL_COMPRESS Package Subprograms*

Subprogram	Description
ISOPEN Function	Checks to see if the handle to a piecewise (un)compress context is open or closed
LZ_COMPRESS Functions and Procedures	Compresses data using Lempel-Ziv compression algorithm
LZ_COMPRESS_ADD Procedure	Adds a piece of compressed data
LZ_COMPRESS_CLOSE	Closes and finishes piecewise compress operation
LZ_COMPRESS_OPEN	Initializes a piecewise context that maintains the compress state and data
LZ_UNCOMPRESS Functions and Procedures	Accepts compressed input, verifies it to be a valid and uncompresses it
LZ_UNCOMPRESS_EXTRACT Procedure	Extracts a piece of uncompressed data
LZ_UNCOMPRESS_OPEN Function	Initializes a piecewise context that maintains the uncompress state and data
LZ_UNCOMPRESS_CLOSE Procedure	Closes and finishes the piecewise uncompress

283.4.1 ISOPEN Function

This function checks to see if the handle to a piecewise (un)compress context is open or closed.

Syntax

```
UTL_COMPRESS.ISOPEN(
   handle in binary_integer)
 RETURN BOOLEAN;
```

Parameters

Table 283-3 ISOPEN Function Parameters

Parameter	Description
`handle`	The handle to a piecewise uncompress context.

Return Values

`TRUE` if the given piecewise handle is opened, otherwise `FALSE`.

Examples

```
IF (UTL_COMPRESS.ISOPEN(myhandle) = TRUE) then
   UTL_COMPRESS.LZ_COMPRESS_CLOSE(myhandle, lob_1);
END IF;
```

Alternatively:

```
IF (UTL_COMPRESS.ISOPEN(myhandle) = TRUE) THEN
   UTL_COMPRESS.LZ_UNCOMPRESS_CLOSE(myhandle);
END IF;
```

283.4.2 LZ_COMPRESS Functions and Procedures

These functions and procedures compress data using Lempel-Ziv compression algorithm.

Syntax

This function accept a `RAW` as input, compress it and return the compressed `RAW` result and metadata:

```
UTL_COMPRESS.LZ_COMPRESS (
   src       IN           RAW,
   quality   IN           BINARY_INTEGER DEFAULT 6)
 RETURN RAW;
```

This function accept a `BLOB` as input, compress it and returns a temporary BLOB for the compressed data:

```
UTL_COMPRESS.LZ_COMPRESS (
  src       IN           BLOB,
  quality   IN           BINARY_INTEGER DEFAULT 6)
 RETURN BLOB;
```

This procedure returns the compressed data into the existing BLOB(dst) which is trimmed to the compressed data size:

```
UTL_COMPRESS.LZ_COMPRESS (
  src      IN            BLOB,
  dst      IN OUT NOCOPY BLOB,
  quality  IN            BINARY_INTEGER DEFAULT 6);
```

This function returns a temporary `BLOB` for the compressed data:

```
UTL_COMPRESS.LZ_COMPRESS (
   src     IN            BFILE,
   quality IN            BINARY_INTEGER DEFAULT 6)
 RETURN BLOB;
```

This procedure will return the compressed data into the existing BLOB(dst) which is trimmed to the compressed data size:

```
UTL_COMPRESS.LZ_COMPRESS (
   src     IN            BFILE,
   dst     IN OUT NOCOPY BLOB,
   quality IN            BINARY_INTEGER DEFAULT 6);
```

Parameters

Table 283-4 LZ_COMPRESS Function and Procedures Parameters

Parameter	Description
src	Data (RAW, BLOB or BFILE) to be compressed.
dst	Destination for compressed data
quality	An integer in the range 1 to 9, 1=fast compression, 9=best compression, default=6

Usage Notes

- quality is an optional compression tuning value. It allows the UTL_COMPRESS user to choose between speed and compression quality, meaning the percentage of reduction in size. A faster compression speed will result in less compression of the data. A slower compression speed will result in more compression of the data. Valid values are [1..9], with 1=fastest and 9=slowest. The default 'quality' value is 6.

283.4.3 LZ_COMPRESS_ADD Procedure

This procedure adds a piece of compressed data.

Syntax

```
UTL_COMPRESS.LZ_COMPRESS_ADD (
   handle IN             BINARY_INTEGER,
   dst    IN OUT NOCOPY  BLOB,
   src    IN             RAW);
```

Parameters

Table 283-5 LZ_COMPRESS_ADD Procedure Parameters

Parameter	Description
handle	The handle to a piecewise compress context.
dst	The opened LOB from LZ_COMPRESS_OPEN to store compressed data.
src	The input data to be compressed.

Exceptions

- invalid_handle - out of range invalid or unopened handle.
- invalid_argument - NULL handle, src, dst, or invalid dst.

283.4.4 LZ_COMPRESS_CLOSE

This procedure closes and finishes piecewise compress operation.

Syntax

```
UTL_COMPRESS.LZ_COMPRESS_CLOSE (
   handle IN             BINARY_INTEGER,
   dst    IN OUT NOCOPY  BLOB);
```

Parameters

Table 283-6 LZ_COMPRESS_CLOSE Procedure Parameters

Parameter	Description
handle	The handle to a piecewise compress context.
dst	The opened LOB from LZ_COMPRESS_OPEN to store compressed data.

Exceptions

- invalid_handle - out of range invalid or uninitialized handle.
- invalid_argument - NULL handle, dst, or invalid dst.

283.4.5 LZ_COMPRESS_OPEN

This function initializes a piecewise context that maintains the compress state and data.

Syntax

```
UTL_COMPRESS.LZ_COMPRESS_OPEN (
   dst       IN OUT NOCOPY BLOB,
   quality   IN            BINARY_INTEGER DEFAULT 6)
 RETURN BINARY_INTEGER;
```

Parameters

Table 283-7 LZ_COMPRESS_OPEN Function Parameters

Parameter	Description
dst	User supplied LOB to store compressed data.
quality	Speed versus efficiency of resulting compressed output. • Valid values are the range 1..9, with a default value of 6. • 1=fastest compression, 9=slowest compression and best compressed file size.

Return Values

A handle to an initialized piecewise compress context.

Exceptions

- invalid_handle - invalid handle, too many open handles.
- invalid_argument - NULL dst or invalid quality specified.

Usage Notes

Close the opened handle with LZ_COMPRESS_CLOSE

- once the piecewise compress is completed
- in the event of an exception in the middle of process

because lack of doing so will cause these handles to leak.

283.4.6 LZ_UNCOMPRESS Functions and Procedures

This procedure accepts as input a RAW, BLOB or BFILE compressed string, verifies it to be a valid compressed value, uncompresses it using Lempel-Ziv compression algorithm, and returns the uncompressed RAW or BLOB result.

Syntax

This function returns uncompressed data as RAW:

```
UTL_COMPRESS.LZ_UNCOMPRESS(
   src  IN  RAW)
  RETURN RAW;
```

This function returns uncompressed data as a temporary BLOB:

```
UTL_COMPRESS.LZ_UNCOMPRESS(
   src  IN  BLOB)
  RETURN BLOB;
```

This procedure returns the uncompressed data into the existing BLOB(dst), which will be trimmed to the uncompressed data size:

```
UTL_COMPRESS.LZ_UNCOMPRESS(
   src  IN  BLOB,
   dst  IN OUT NOCOPY BLOB);
```

This function returns a temporary BLOB for the uncompressed data:

```
UTL_COMPRESS.LZ_UNCOMPRESS(
   src  IN BFILE)
 RETURN BLOB;
```

This procedure returns the uncompressed data into the existing BLOB(dst). The original dst data will be overwritten.

```
UTL_COMPRESS.LZ_UNCOMPRESS(
   src  IN  BFILE,
   dst  IN OUT NOCOPY BLOB);
```

Parameters

Table 283-8 LZ_UNCOMPRESS Function and Procedures Parameters

Parameter	Description
src	Compressed data.
dst	Destination for uncompressed data.

283.4.7 LZ_UNCOMPRESS_EXTRACT Procedure

This procedure extracts a piece of uncompressed data.

Syntax

```
UTL_COMPRESS.LZ_UNCOMPRESS_EXTRACT(
   handle  IN          BINARY_INTEGER,
   dst     OUT NOCOPY  RAW);
```

Parameters

Table 283-9 LZ_UNCOMPRESS_EXTRACT Function Parameters

Parameter	Description
handle	The handle to a piecewise uncompress context.
dst	The uncompressed data.

Exceptions

- no_data_found - finished uncompress.
- invalid_handle - out of range invalid or uninitialized handle.
- invalid_argument - NULL handle.

283.4.8 LZ_UNCOMPRESS_OPEN Function

This function initializes a piecewise context that maintains the uncompress state and data.

Syntax

```
UTL_COMPRESS.LZ_UNCOMPRESS_OPEN(
   src  IN  BLOB)
  RETURN BINARY_INTEGER;
```

Parameters

Table 283-10 LZ_UNCOMPRESS_OPEN Function Parameters

Parameter	Description
src	The input data to be uncompressed.

Return Values

A handle to an initialized piecewise compress context.

Exceptions

- invalid_handle - invalid handle, too many open handles.
- invalid_argument - NULL src.

Usage Notes

Close the opened handle with LZ_UNCOMPRESS_CLOSE

- once the piecewise uncompress is completed
- in the event of an exception in the middle of process

because lack of doing so will cause these handles to leak.

283.4.9 LZ_UNCOMPRESS_CLOSE Procedure

This procedure closes and finishes the piecewise uncompress.

Syntax

```
UTL_COMPRESS.LZ_UNCOMPRESS_CLOSE(
   handle  IN   BINARY_INTEGER);
```

Parameters

Table 283-11 LZ_UNCOMPRESS_CLOSE Procedure Parameters

Parameter	Description
`handle`	The handle to a piecewise uncompress context.

Exceptions

- `invalid_handle` - out of range invalid or uninitialized handle.
- `invalid_argument` - `NULL` handle.

284
UTL_ENCODE

The `UTL_ENCODE` package provides functions that encode `RAW` data into a standard encoded format so that the data can be transported between hosts.

You can use `UTL_ENCODE` functions to encode the body of email text. The package also contains the decode counterpart functions of the encode functions. The functions follow published standards for encoding to accommodate non-Oracle utilities on the sending or receiving ends.

This chapter contains the following topic:

- Summary of UTL_ENCODE Subprograms

284.1 Summary of UTL_ENCODE Subprograms

This table lists the `UTL_ENCODE` subprograms and briefly describes them.

Table 284-1 ***UTL_ENCODE Package Subprograms***

Subprogram	Description
BASE64_DECODE Function	Reads the base 64-encoded `RAW` input string and decodes it to its original `RAW` value
BASE64_ENCODE Function	Encodes the binary representation of the `RAW` value into base 64 elements and returns it in the form of a `RAW` string
MIMEHEADER_DECODE Function	Decodes a string from mime header format
MIMEHEADER_ENCODE Function	Encodes a string into mime header format
QUOTED_PRINTABLE_DECODE Function	Reads the `varchar2` quoted printable format input string and decodes it to the corresponding `RAW` string
QUOTED_PRINTABLE_ENCODE Function	Reads the `RAW` input string and encodes it to the corresponding quoted printable format string
TEXT_DECODE Function	Decodes a character set sensitive text string
TEXT_ENCODE Function	Encodes a character set sensitive text string
UUDECODE Function	Reads the `RAW` uuencode format input string and decodes it to the corresponding `RAW` string
UUENCODE Function	Reads the `RAW` input string and encodes it to the corresponding uuencode format string

284.1.1 BASE64_DECODE Function

This function reads the base 64-encoded `RAW` input string and decodes it to its original `RAW` value.

Syntax

```
UTL_ENCODE.BASE64_DECODE (
   r  IN RAW)
RETURN RAW;
```

Pragmas

```
pragma RESTRICT_REFERENCES(base64_decode, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 284-2 BASE64_DECODE Function Parameters

Parameter	Description
`r`	The RAW string containing base 64-encoded data. There are no defaults or optional parameters.

Return Values

Table 284-3 BASE64_DECODE Function Return Values

Return	Description
`RAW`	Contains the decoded string

284.1.2 BASE64_ENCODE Function

This function encodes the binary representation of the `RAW` value into base 64 elements and returns it in the form of a `RAW` string.

Syntax

```
UTL_ENCODE.BASE64_ENCODE (
   r  IN RAW)
RETURN RAW;
```

Pragmas

```
pragma RESTRICT_REFERENCES(base64_encode, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 284-4 BASE64_ENCODE Function Parameters

Parameter	Description
`r`	The RAW value to be encoded. There are no defaults or optional parameters.

Return Values

Table 284-5 BASE64_ENCODE Function Return Values

Return	Description
RAW	Contains the encoded base 64 elements

284.1.3 MIMEHEADER_DECODE Function

This function accepts as input an "encoded word."

It uses the form:

```
=?<charset>?<encoding>?<encoded text>?=
=?ISO-8859-1?Q?Here is some encoded text?=
```

The `<encoded text>` is encapsulated in mime header tags which give the `MIMEHEADER_DECODE` function information about how to decode the string. The mime header metadata tags are stripped from the input string and the `<encoded text>` is converted to the base database character set as follows:

- If this is a UTF16 platform, convert the encoded text from UTF16 to ASCII
- If this is an EBCDIC platform, convert the encoded text from EBCDIC to ASCII
- If this is an ASCII or UTF8 platform, no conversion needed

The string is decoded using either quoted-printable or base64 decoding, as specified by the `<encoding>` metadata tag in the encoded word. The resulting converted and decoded text is returned to the caller as a `VARCHAR2` string.

Syntax

```
UTL_ENCODE.MIMEHEADER_DECODE (
   buf    IN   VARCHAR2 CHARACTER SET ANY_CS)
  RETURN data VARCHAR2 CHARACTER SET buf%CHARSET;
```

Parameters

Table 284-6 MIMEHEADER_DECODE Function Parameters

Parameter	Description
buf	The encoded text data with mime header format tags.

Return Values

Table 284-7 MIMEHEADER_DECODE Function Return Values

Return	Description
data	The encoded text data with mime header format tags

Examples

```
v2:=utl_encode.mimeheader_decode('=?ISO-8859-1?Q?Here is some encoded text?=');
```

284.1.4 MIMEHEADER_ENCODE Function

This function returns as an output an "encoded word".

The output is in the following form:

```
=?<charset>?<encoding>?<encoded text>?=
=?ISO-8859-1?Q?Here is some text?=
```

The `buf` input parameter is the text to be encoded and becomes the `<encoded text>`.

The <encoding> value is either "Q" or "B" for quoted-printable encode or base64 encoding respectively. The ENCODING input parameter accepts as valid values `UTL_ENCODE.QUOTED_PRINTABLE` or `UTL_ENCODE.BASE64` or `NULL`. If `NULL`, quoted-printable encoding is selected as a default value.

The `<charset>` value is specified as the input parameter `encode_charset`. If `NULL`, the database character set is selected as a default value.

The mimeheader encoding process includes conversion of the `buf` input string to the character set specified by the `encode_charset` parameter. The converted string is encoded to either quoted-printable or base64 encoded format. The mime header tags are appended and prepended.

Finally, the string is converted to the base character set of the database:

- If this is a UTF16 platform, convert the encoded text to UTF16
- If this is an EBCDIC platform, convert the encoded text to EBCDIC
- If this is an ASCII or UTF8 platform, no conversion needed.

Syntax

```
UTL_ENCODE.MIMEHEADER_ENCODE (
   buf            IN  VARCHAR2 CHARACTER SET ANY_CS,
   encode_charset IN  VARCHAR2 DEFAULT NULL,
   encoding       IN  PLS_INTEGER DEFAULT NULL)
 RETURN string VARCHAR2 CHARACTER SET buf%CHARSET;
```

Parameters

Table 284-8 MIMEHEADER_ENCODE Function Parameters

Parameter	Description
`buf`	The text data.
`encode_charset`	The target character set.
`encoding`	The encoding format. Valid values are `UTL_ENCODE.BASE64`, `UTL_ENCODE.QUOTED_PRINTABLE` and `NULL`

Return Values

Table 284-9 MIMEHEADER_ENCODE Function Return Values

Return	Description
`string`	A `VARCHAR2` encoded string with mime header format tags.

284.1.5 QUOTED_PRINTABLE_DECODE Function

This function reads the `varchar2` quoted printable format input string and decodes it to the corresponding `RAW` string.

Syntax

```
UTL_ENCODE.QUOTED_PRINTABLE_DECODE (
   r  IN RAW)
RETURN RAW;
```

Pragmas

```
pragma RESTRICT_REFERENCES(quoted_printable_decode, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 284-10 QUOTED_PRINTABLE_DECODE Function Parameters

Parameters	Description
`r`	The RAW string containing a quoted printable data string. There are no defaults or optional parameters.

Return Values

Table 284-11 QUOTED_PRINTABLE_DECODE Function Return Values

Return	Description
`RAW`	The decoded string

284.1.6 QUOTED_PRINTABLE_ENCODE Function

This function reads the `RAW` input string and encodes it to the corresponding quoted printable format string.

Syntax

```
UTL_ENCODE.QUOTED_PRINTABLE_ENCODE (
   r  IN RAW)
RETURN RAW;
```

Pragmas

```
pragma RESTRICT_REFERENCES(quoted_printable_encode, WNDS, RNDS,WNPS, RNPS);
```

Parameters

Table 284-12 QUOTED_PRINTABLE_ENCODE Function Parameters

Parameter	Description
r	The RAW string. There are no defaults or optional parameters.

Return Values

Table 284-13 QUOTED_PRINTABLE_ENCODE Function Return Values

Return	Description
RAW	Contains the quoted printable string

284.1.7 TEXT_DECODE Function

This function converts the input text to the target character set as specified by the `encode_charset` parameter, if not `NULL`.

The encoded text is converted to the base character set of database, as follows:

- If this is a UTF16 platform, convert the encoded text from UTF16 to ASCII
- If this is an EBCDIC platform, convert the encoded text from EBCDIC to ASCII
- If this is an ASCII or UTF8 platform, no conversion needed

You can decode from either quoted-printable or base64 format, with regard to each `encoding` parameter. If `NULL`, quoted-printable is selected as a default decoding format. If `encode_charset` is not `NULL`, you convert the string from the specified character set to the database character set. The resulting decoded and converted text string is returned to the caller.

Syntax

```
UTL_ENCODE.TEXT_DECODE(
   buf            IN  VARCHAR2 CHARACTER SET ANY_CS,
   encode_charset IN  VARCHAR2 DEFAULT NULL,
   encoding       IN  PLS_INTEGER DEFAULT NULL)
 RETURN string VARCHAR2 CHARACTER SET buf%CHARSET;
```

Parameters

Table 284-14 TEXT_DECODE Function Parameters

Parameter	Description
buf	The encoded text data.
encode_charset	The source character set.
encoding	The encoding format. Valid values are `UTL_ENCODE.BASE64`, `UTL_ENCODE.QUOTED_PRINTABLE` and `NULL`.

Return Values

Table 284-15 TEXT_DECODE Function Return Values

Return	Description
string	A VARCHAR2 decoded text string.

Examples

```
v2:=UTL_ENCODE.TEXT_DECODE(
     'Here is some text',
      WE8ISO8859P1,
      UTL_ENCODE.BASE64);
```

284.1.8 TEXT_ENCODE Function

This function converts the input text to the target character set as specified by the encode_charset parameter, if not NULL.

The text is encoded to either base64 or quoted-printable format, as specified by the encoding parameter. Quoted-printable is selected as a default if ENCODING is NULL.

The encoded text is converted to the base character set of the database:

- If this is a UTF16 platform, convert the encoded text to UTF16
- If this is an EBCDIC platform, convert the encoded text to EBCDIC
- If this is an ASCII or UTF8 platform, no conversion needed

The resulting encoded and converted text string is returned to the caller.

Syntax

```
UTL_ENCODE.TEXT_ENCODE (
   buf            IN  VARCHAR2 CHARACTER SET ANY_CS,
   encode_charset IN  VARCHAR2 DEFAULT NULL,
   encoding       IN  PLS_INTEGER DEFAULT NULL)
 RETURN string VARCHAR2 CHARACTER SET buf%CHARSET;
```

Parameters

Table 284-16 TEXT_ENCODE Function Parameters

Parameter	Description
buf	The text data.
encode_charset	The target character set.
encoding	The encoding format. Valid values are UTL_ENCODE.BASE64, UTL_ENCODE.QUOTED_PRINTABLE and NULL

Return Values

Table 284-17 TEXT_ENCODE Function Return Values

Return	Description
string	A VARCHAR2 encoded string with mime header format tags.

Examples

```
v2:=utl_encode.text_encode(
   'Here is some text',
   'WE8ISO8859P1',
    UTL_ENCODE.BASE64);
```

284.1.9 UUDECODE Function

This function reads the RAW uuencode format input string and decodes it to the corresponding RAW string.

See "UUENCODE Function" for discussion of the cumulative nature of UUENCODE and UUDECODE for data streams.

Syntax

```
UTL_ENCODE.UUDECODE (
   r  IN RAW)
RETURN RAW;
```

Pragmas

```
pragma RESTRICT_REFERENCES(uudecode, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 284-18 UUDECODE Function Parameters

Parameter	Description
r	The RAW string containing the uuencoded data string. There are no defaults or optional parameters.

Return Values

Table 284-19 UUDECODE Function Return Values

Return	Description
RAW	The decoded RAW string

284.1.10 UUENCODE Function

This function reads the `RAW` input string and encodes it to the corresponding uuencode format string.

The output of this function is cumulative, in that it can be used to encode large data streams, by splitting the data stream into acceptably sized `RAW` values, encoded, and concatenated into a single encoded string.

Syntax

```
UTL_ENCODE.UUENCODE (
   r          IN RAW,
   type       IN PLS_INTEGER DEFAULT 1,
   filename   IN VARCHAR2 DEFAULT NULL,
   permission IN VARCHAR2 DEFAULT NULL) RETURN RAW;
```

Pragmas

```
pragma RESTRICT_REFERENCES(uuencode, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 284-20 UUENCODE Function Parameters

Parameter	Description
`r`	`RAW` string
`type`	Optional number parameter containing the type of uuencoded output. Options: complete—a defined PL/SQL constant with a value of `1`. (default) header_piece ...middle_piece ...end_piece
`filename`	Optional `varchar2` parameter containing the uuencode filename; the default is `uuencode.txt`
`permission`	Optional `varchar2` parameter containing the permission mode; the default is `0` (a text string zero).

Return Values

Table 284-21 UUENCODE Function Return Values

Return	Description
`RAW`	Contains the uuencode format string

285
UTL_FILE

With the UTL_FILE package, PL/SQL programs can read and write operating system text files. UTL_FILE provides a restricted version of operating system stream file I/O.

This chapter contains the following topics:

- Security Model
- Operational Notes
- Rules and Limits
- Exceptions
- Examples
- Data Structures
- Summary of UTL_FILE Subprograms

285.1 UTL_FILE Security Model

The set of files and directories that are accessible to the user through UTL_FILE is controlled by a number of factors and database parameters. Foremost of these is the set of directory objects that have been granted to the user.

The nature of directory objects is discussed in the *Oracle Database SQL Language Reference*.

Assuming the user has both READ and WRITE access to the directory object USER_DIR, the user can open a file located in the operating system directory described by USER_DIR, but not in subdirectories or parent directories of this directory.

Lastly, the client (text I/O) and server implementations are subject to operating system file permission checking.

UTL_FILE provides file access both on the client side and on the server side. When run on the server, UTL_FILE provides access to all operating system files that are accessible from the server. On the client side, as in the case for Forms applications, UTL_FILE provides access to operating system files that are accessible from the client.

Directory objects offer more flexibility and granular control to the UTL_FILE application administrator, can be maintained dynamically (that is, without shutting down the database), and are consistent with other Oracle tools. CREATE ANY DIRECTORY privilege is granted only to SYS and SYSTEM by default.

> **Note:**
>
> Use the CREATE DIRECTORY feature for directory access verification.

Note that neither hard nor symbolic links are supported.

On UNIX systems, the owner of a file created by the FOPEN function is the owner of the shadow process running the instance. Normally, this owner is ORACLE. Files created using FOPEN are always writable and readable using the UTL_FILE subprograms. However, non-privileged operating system users who need to read these files outside of PL/SQL may need access from a system administrator.

> **WARNING:**
>
> The privileges needed to access files in a directory object are operating system specific. UTL_FILE directory object privileges give you read and write access to all files within the specified directory.

285.2 UTL_FILE Operational Notes

Keep these notes in mind when using UTL_FILE.

The file location and file name parameters are supplied to the FOPEN function as separate strings, so that the file location can be checked against the list of accessible directories as specified by the ALL_DIRECTORIES view of accessible directory objects. Together, the file location and name must represent a legal filename on the system, and the directory must be accessible. A subdirectory of an accessible directory is not necessarily also accessible; it too must be specified using a complete path name matching an ALL_DIRECTORIES object.

UTL_FILE implicitly interprets line terminators on read requests, thereby affecting the number of bytes returned on a GET_LINE call. For example, the len parameter of UTL_FILE.GET_LINE specifies the requested number of bytes of character data. The number of bytes actually returned to the user will be the lesser of:

- The GET_LINE len parameter, or
- The number of bytes until the next line terminator character, or
- The max_linesize parameter specified by UTL_FILE.FOPEN

The FOPEN max_linesize parameter must be a number in the range 1 and 32767. If unspecified, Oracle supplies a default value of 1024. The GET_LINE len parameter must be a number in the range 1 and 32767. If unspecified, Oracle supplies the default value of max_linesize. If max_linesize and len are defined to be different values, then the lesser value takes precedence.

UTL_FILE.GET_RAW ignores line terminators.

UTL_FILE expects that files opened by UTL_FILE.FOPEN in text mode are encoded in the database character set. It expects that files opened by UTL_FILE.FOPEN_NCHAR in text mode are encoded in the UTF8 character set. If an opened file is not encoded in the expected character set, the result of an attempt to read the file is indeterminate. When data encoded in one character set is read and Globalization Support is told (such as by means of NLS_LANG) that it is encoded in another character set, the result is indeterminate. If NLS_LANG is set, it should therefore be the same as the database character set.

285.3 UTL_FILE Rules and Limits

Operating system-specific parameters, such as C-shell environment variables under UNIX, cannot be used in the file location or file name parameters.

UTL_FILE I/O capabilities are similar to standard operating system stream file I/O (OPEN, GET, PUT, CLOSE) capabilities, but with some limitations. For example, you call the FOPEN function to return a file handle, which you use in subsequent calls to GET_LINE or PUT to perform stream I/O to a file. When file I/O is done, you call FCLOSE to complete any output and free resources associated with the file.

> **Note:**
>
> The UTL_FILE package is similar to the client-side TEXT_IO package currently provided by Oracle Procedure Builder. Restrictions for a server implementation require some API differences between UTL_FILE and TEXT_IO. In PL/SQL file I/O, errors are returned using PL/SQL exceptions.

285.4 UTL_FILE Exceptions

This table describes exceptions raised by UTL_FILE subprograms.

Table 285-1 UTL_FILE Package Exceptions

Exception Name	Description
INVALID_PATH	File location is invalid.
INVALID_MODE	The open_mode parameter in FOPEN is invalid.
INVALID_FILEHANDLE	File handle is invalid.
INVALID_OPERATION	File could not be opened or operated on as requested.
READ_ERROR	Destination buffer too small, or operating system error occurred during the read operation
WRITE_ERROR	Operating system error occurred during the write operation.
INTERNAL_ERROR	Unspecified PL/SQL error
CHARSETMISMATCH	A file is opened using FOPEN_NCHAR, but later I/O operations use nonchar functions such as PUTF or GET_LINE.
FILE_OPEN	The requested operation failed because the file is open.
INVALID_MAXLINESIZE	The MAX_LINESIZE value for FOPEN() is invalid; it should be within the range 1 to 32767.
INVALID_FILENAME	The filename parameter is invalid.
ACCESS_DENIED	Permission to access to the file location is denied.
INVALID_OFFSET	Causes of the INVALID_OFFSET exception: • ABSOLUTE_OFFSET = NULL and RELATIVE_OFFSET = NULL, or • ABSOLUTE_OFFSET < 0, or • Either offset caused a seek past the end of the file

Table 285-1 (Cont.) UTL_FILE Package Exceptions

Exception Name	Description
DELETE_FAILED	The requested file delete operation failed.
RENAME_FAILED	The requested file rename operation failed.

Procedures in UTL_FILE can also raise predefined PL/SQL exceptions such as NO_DATA_FOUND or VALUE_ERROR.

285.5 UTL_FILE Examples

These two examples show use of the procedure.

Example 1

> **Note:**
>
> The examples are UNIX-specific.

Given the following:

```
SQL> CREATE DIRECTORY log_dir AS '/appl/gl/log';
SQL> GRANT READ ON DIRECTORY log_dir TO DBA;
SQL> GRANT WRITE ON DIRECTORY log_dir TO DBA;

SQL> CREATE DIRECTORY USER_DIR AS '/appl/gl/user'';
SQL> GRANT READ ON DIRECTORY USER_DIR TO PUBLIC;
SQL> GRANT WRITE ON DIRECTORY USER_DIR TO PUBLIC;
```

The following file locations and filenames are valid and accessible as follows:

File Location	Filename	READ and WRITE
/appl/gl/log	L12345.log	Users with DBA privilege
/appl/gl/user	u12345.tmp	All users

The following file locations and filenames are invalid:

File Location	Filename	Invalid Because
/appl/gl/log/backup	L12345.log	# subdirectories are not accessible
/APPL/gl/log	L12345.log	# directory strings must follow case sensitivity rules as required by the O/S
/appl/gl/log	backup/L1234.log	# filenames may not include portions of directory paths
/user/tmp	L12345.log	# no corresponding CREATE DIRECTORY command has been issued

Example 2

```
DECLARE
  V1 VARCHAR2(32767);
  F1 UTL_FILE.FILE_TYPE;
BEGIN
  -- In this example MAX_LINESIZE is less than GET_LINE's length request
  -- so the number of bytes returned will be 256 or less if a line terminator is seen.
  F1 := UTL_FILE.FOPEN('USER_DIR','u12345.tmp','R',256);
  UTL_FILE.GET_LINE(F1,V1,32767);
  UTL_FILE.FCLOSE(F1);

  -- In this example, FOPEN's MAX_LINESIZE is NULL and defaults to 1024,
  -- so the number of bytes returned will be 1024 or less if a line terminator is seen.
  F1 := UTL_FILE.FOPEN('USER_DIR','u12345.tmp','R');
  UTL_FILE.GET_LINE(F1,V1,32767);
  UTL_FILE.FCLOSE(F1);

  -- In this example, GET_LINE doesn't specify a number of bytes, so it defaults to
  -- the same value as FOPEN's MAX_LINESIZE which is NULL in this case and defaults to
1024.
  -- So the number of bytes returned will be 1024 or less if a line terminator is seen.
  F1 := UTL_FILE.FOPEN('USER_DIR','u12345.tmp','R');
  UTL_FILE.GET_LINE(F1,V1);
  UTL_FILE.FCLOSE(F1);
END;
```

285.6 UTL_FILE Data Structures

The `UTL_FILE` package defines a RECORD type.

Record Types

- FILETYPE Record Type

285.6.1 FILETYPE Record Type

The contents of `FILE_TYPE` are private to the `UTL_FILE` package. You should not reference or change components of this record.

```
TYPE file_type IS RECORD (
   id          BINARY_INTEGER,
   datatype    BINARY_INTEGER,
   byte_mode   BOOLEAN);
```

Fields

Table 285-2 FILE_TYPE Fields

Field	Description
`id`	A numeric value indicating the internal file handle number
`datatype`	Indicates whether the file is a `CHAR` file, `Nchar` file or other (binary)
`byte_mode`	Indicates whether the file was open as a binary file, or as a text file

> **Caution:**
>
> Oracle does not guarantee the persistence of `FILE_TYPE` values between database sessions or within a single session. Attempts to clone file handles or use dummy file handles may have inderterminate outcomes.

285.7 Summary of UTL_FILE Subprograms

This table lists the `UTL_FILE` subprograms and briefly describes them.

Table 285-3 UTL_FILE Subprograms

Subprogram	Description
FCLOSE Procedure	Closes a file
FCLOSE_ALL Procedure	Closes all open file handles
FCOPY Procedure	Copies a contiguous portion of a file to a newly created file
FFLUSH Procedure	Physically writes all pending output to a file
FGETATTR Procedure	Reads and returns the attributes of a disk file
FGETPOS Function	Returns the current relative offset position within a file, in bytes
FOPEN Function	Opens a file for input or output
FOPEN_NCHAR Function	Opens a file in Unicode for input or output
FREMOVE Procedure	Deletes a disk file, assuming that you have sufficient privileges
FRENAME Procedure	Renames an existing file to a new name, similar to the UNIX `mv` function
FSEEK Procedure	Adjusts the file pointer forward or backward within the file by the number of bytes specified
GET_LINE Procedure	Reads text from an open file
GET_LINE_NCHAR Procedure	Reads text in Unicode from an open file
GET_RAW Procedure	Reads a `RAW` string value from a file and adjusts the file pointer ahead by the number of bytes read
IS_OPEN Function	Determines if a file handle refers to an open file
NEW_LINE Procedure	Writes one or more operating system-specific line terminators to a file
PUT Procedure	Writes a string to a file
PUT_LINE Procedure	Writes a line to a file, and so appends an operating system-specific line terminator
PUT_LINE_NCHAR Procedure	Writes a Unicode line to a file
PUT_NCHAR Procedure	Writes a Unicode string to a file
PUTF Procedure	A `PUT` procedure with formatting
PUTF_NCHAR Procedure	A `PUT_NCHAR` procedure with formatting, and writes a Unicode string to a file, with formatting
PUT_RAW Procedure	Accepts as input a `RAW` data value and writes the value to the output buffer

285.7.1 FCLOSE Procedure

This procedure closes an open file identified by a file handle.

Syntax

```
UTL_FILE.FCLOSE (
   file IN OUT FILE_TYPE);
```

Parameters

Table 285-4 FCLOSE Procedure Parameters

Parameter	Description
`file`	Active file handle returned by an `FOPEN` or `FOPEN_NCHAR` call

Usage Notes

If there is buffered data yet to be written when `FCLOSE` runs, then you may receive a `WRITE_ERROR` exception when closing a file.

Exceptions

```
WRITE_ERROR
INVALID_FILEHANDLE
```

285.7.2 FCLOSE_ALL Procedure

This procedure closes all open file handles for the session. This should be used as an emergency cleanup procedure, for example, when a PL/SQL program exits on an exception.

Syntax

```
UTL_FILE.FCLOSE_ALL;
```

Usage Notes

> **Note:**
>
> `FCLOSE_ALL` does not alter the state of the open file handles held by the user. This means that an `IS_OPEN` test on a file handle after an `FCLOSE_ALL` call still returns `TRUE`, even though the file has been closed. No further read or write operations can be performed on a file that was open before an `FCLOSE_ALL`.

Exceptions

```
WRITE_ERROR
```

285.7.3 FCOPY Procedure

This procedure copies a contiguous portion of a file to a newly created file.

By default, the whole file is copied if the `start_line` and `end_line` parameters are omitted. The source file is opened in read mode. The destination file is opened in write mode. A starting and ending line number can optionally be specified to select a portion from the center of the source file for copying.

Syntax

```
UTL_FILE.FCOPY (
   src_location    IN VARCHAR2,
   src_filename    IN VARCHAR2,
   dest_location   IN VARCHAR2,
   dest_filename   IN VARCHAR2,
   start_line      IN BINARY_INTEGER DEFAULT 1,
   end_line        IN BINARY_INTEGER DEFAULT NULL);
```

Parameters

Table 285-5 FCOPY Procedure Parameters

Parameters	Description
`src_location`	Directory location of the source file, a `DIRECTORY_NAME` from the `ALL_DIRECTORIES` view (case sensitive)
`src_filename`	Source file to be copied
`dest_location`	Destination directory where the destination file is created
`dest_filename`	Destination file created from the source file
`start_line`	Line number at which to begin copying. The default is `1` for the first line
`end_line`	Line number at which to stop copying. The default is `NULL`, signifying end of file

Exceptions

`INVALID_FILENAME`

`INVALID_PATH`

`INVALID_OPERATION`

`INVALID_OFFSET`

`READ_ERROR`

`WRITE_ERROR`

285.7.4 FFLUSH Procedure

`FFLUSH` physically writes pending data to the file identified by the file handle. Normally, data being written to a file is buffered. The `FFLUSH` procedure forces the buffered data to be written to the file. The data must be terminated with a newline character.

Flushing is useful when the file must be read while still open. For example, debugging messages can be flushed to the file so that they can be read immediately.

Syntax

```
UTL_FILE.FFLUSH (
   file  IN FILE_TYPE);
```

Parameters

Table 285-6 FFLUSH Procedure Parameters

Parameters	Description
file	Active file handle returned by an FOPEN or FOPEN_NCHAR call

Exceptions

INVALID_FILENAME

INVALID_MAXLINESIZE

INVALID_OPERATION

WRITE_ERROR

285.7.5 FGETATTR Procedure

This procedure reads and returns the attributes of a disk file.

Syntax

```
UTL_FILE.FGETATTR(
   location     IN VARCHAR2,
   filename     IN VARCHAR2,
   fexists      OUT BOOLEAN,
   file_length  OUT NUMBER,
   block_size   OUT BINARY_INTEGER);
```

Parameters

Table 285-7 FGETATTR Procedure Parameters

Parameters	Description
location	Directory location of the source file, a DIRECTORY_NAME from the ALL_DIRECTORIES view (case sensitive)
filename	Name of the file to be examined
fexists	A BOOLEAN for whether or not the file exists
file_length	Length of the file in bytes. NULL if file does not exist.
block_size	File system block size in bytes. NULL if the file does not exist.

Exceptions

INVALID_PATH

INVALID_FILENAME

INVALID_OPERATION

READ_ERROR

ACCESS_DENIED

285.7.6 FGETPOS Function

This function returns the current relative offset position within a file, in bytes.

Syntax

```
UTL_FILE.FGETPOS (
   file IN FILE_TYPE)
 RETURN PLS_INTEGER;
```

Parameters

Table 285-8 FGETPOS Parameters

Parameters	Description
file	Directory location of the source file

Return Values

FGETPOS returns the relative offset position for an open file, in bytes. It raises an exception if the file is not open. It returns 0 for the beginning of the file.

Exceptions

INVALID_FILEHANDLE

INVALID_OPERATION

READ_ERROR

Usage Notes

If file is opened for byte mode operations, then the INVALID OPERATION exception is raised.

285.7.7 FOPEN Function

This function opens a file. You can specify the maximum line size and have a maximum of 50 files open simultaneously.

See also FOPEN_NCHAR Function.

Syntax

```
UTL_FILE.FOPEN (
   location     IN VARCHAR2,
   filename     IN VARCHAR2,
   open_mode    IN VARCHAR2,
   max_linesize IN BINARY_INTEGER DEFAULT 1024)
  RETURN FILE_TYPE;
```

Parameters

Table 285-9 FOPEN Function Parameters

Parameter	Description
location	Directory location of file. This string is a directory object name and must be specified in upper case. Read privileges must be granted on this directory object for the UTL_FILE user to run FOPEN.
filename	File name, including extension (file type), without directory path. If a directory path is given as a part of the filename, it is ignored by FOPEN. On Unix, the filename cannot end with /.
open_mode	Specifies how the file is opened. Modes include: • r -- read text • w -- write text • a -- append text • rb -- read byte mode • wb -- write byte mode • ab -- append byte mode If you try to open a file specifying 'a' or 'ab' for open_mode but the file does not exist, the file is created in write mode.
max_linesize	Maximum number of bytes for each line, including the newline character, for this file (minimum value 1, maximum value 32767). If unspecified, Oracle supplies a default value of 1024.

Return Values

FOPEN returns a file handle, which must be passed to all subsequent procedures that operate on that file. The specific contents of the file handle are private to the UTL_FILE package, and individual components should not be referenced or changed by the UTL_FILE user.

Table 285-10 FOPEN Function Return Values

Return	Description
FILE_TYPE	Handle to open file

Exceptions

INVALID_MAXILINESIZE

INVALID_MODE

INVALID_OPERATION

INVALID_PATH

INVALID_FILENAME

Usage Notes

The file location and file name parameters must be supplied to the FOPEN function as quoted strings so that the file location can be checked against the list of accessible directories as specified by the ALL_DIRECTORIES view of accessible directory objects.

285.7.8 FOPEN_NCHAR Function

This function opens a file in national character set mode for input or output, with the maximum line size specified. With this function, you can read or write a text file in Unicode instead of in the database character set.

You can have a maximum of 50 files open simultaneously.

Even though the contents of an `NVARCHAR2` buffer may be AL16UTF16 or UTF8 (depending on the national character set of the database), the contents of the file are always read and written in UTF8. `UTL_FILE` converts between UTF8 and AL16UTF16 as necessary.

See also FOPEN Function.

Syntax

```
UTL_FILE.FOPEN_NCHAR (
   location     IN VARCHAR2,
   filename     IN VARCHAR2,
   open_mode    IN VARCHAR2,
   max_linesize IN BINARY_INTEGER DEFAULT 1024)
RETURN FILE_TYPE;
```

Parameters

Table 285-11 FOPEN_NCHAR Function Parameters

Parameter	Description
`location`	Directory location of file
`filename`	File name (including extension)
`open_mode`	Open mode (`r`,`w`,`a`,`rb`,`wb`,`ab`)
`max_linesize`	Maximum number of characters for each line, including the newline character, for this file (minimum value 1, maximum value 32767)

Return Values

`FOPEN_NCHAR` returns a file handle, which must be passed to all subsequent procedures that operate on that file. The specific contents of the file handle are private to the `UTL_FILE` package, and individual components should not be referenced or changed by the `UTL_FILE` user.

Table 285-12 FOPEN_NCHAR Function Return Values

Return	Description
`FILE_TYPE`	Handle to open file

Exceptions

`INVALID_MAXILINESIZE`

`INVALID_MODE`

`INVALID_OPERATION`

```
INVALID_PATH
```

285.7.9 FREMOVE Procedure

This procedure deletes a disk file, assuming that you have sufficient privileges.

Syntax

```
UTL_FILE.FREMOVE (
   location IN VARCHAR2,
   filename IN VARCHAR2);
```

Parameters

Table 285-13 FREMOVE Procedure Parameters

Parameters	Description
location	Directory location of the file, a DIRECTORY_NAME from ALL_DIRECTORIES (case sensitive)
filename	Name of the file to be deleted

Exceptions

```
ACCESS_DENIED

DELETE_FAILED

INVALID_FILENAME

INVALID_OPERATION

INVALID_PATH
```

Usage Notes

The FREMOVE procedure does not verify privileges before deleting a file. The O/S verifies file and directory permissions. An exception is returned on failure.

285.7.10 FRENAME Procedure

This procedure renames an existing file to a new name, similar to the UNIX mv function.

Syntax

```
UTL_FILE.FRENAME (
   src_location      IN   VARCHAR2,
   src_filename      IN   VARCHAR2,
   dest_location     IN   VARCHAR2,
   dest_filename     IN   VARCHAR2,
   overwrite         IN   BOOLEAN DEFAULT FALSE);
```

Parameters

Table 285-14 FRENAME Procedure Parameters

Parameters	Description
src_location	Directory location of the source file, a DIRECTORY_NAME from the ALL_DIRECTORIES view (case sensitive)
src_filename	Source file to be renamed
dest_location	Destination directory of the destination file, a DIRECTORY_NAME from the ALL_DIRECTORIES view (case sensitive)
dest_filename	New name of the file
overwrite	Default is FALSE. Permission on both the source and destination directories must be granted. You can use the overwrite parameter to specify whether or not to overwrite a file if one exists in the destination directory. The default is FALSE for no overwrite.

Exceptions

ACCESS_DENIED

INVALID_FILENAME

INVALID_PATH

RENAME_FAILED

285.7.11 FSEEK Procedure

This procedure adjusts the file pointer forward or backward within the file by the number of bytes specified.

Syntax

```
UTL_FILE.FSEEK (
   file             IN OUT  UTL_FILE.FILE_TYPE,
   absolute_offset  IN      PL_INTEGER DEFAULT NULL,
   relative_offset  IN      PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 285-15 FSEEK Procedure Parameters

Parameters	Description
file	File handle
absolute_offset	Absolute location to which to seek; default = NULL
relative_offset	Number of bytes to seek forward or backward; positive = forward, negative integer = backward, zero = current position, default = NULL

Exceptions

INVALID_FILEHANDLE

INVALID_OFFSET

INVALID_OPERATION

READ_ERROR

Usage Notes

- Using FSEEK, you can read previous lines in the file without first closing and reopening the file. You must know the number of bytes by which you want to navigate.
- If relative_offset, the procedure seeks forward. If relative_offset > 0, or backward, if relative_offset < 0, the procedure seeks through the file by the number of relative_offset bytes specified.
- If the beginning of the file is reached before the number of bytes specified, then the file pointer is placed at the beginning of the file. If the end of the file is reached before the number of bytes specified, then an INVALID_OFFSET error is raised.
- If absolute_offset, the procedure seeks to an absolute location specified in bytes.
- If file is opened for byte mode operations, then the INVALID OPERATION exception is raised.

285.7.12 GET_LINE Procedure

This procedure reads text from the open file identified by the file handle and places the text in the output buffer parameter. Text is read up to, but not including, the line terminator, or up to the end of the file, or up to the end of the len parameter. It cannot exceed the max_linesize specified in FOPEN.

Syntax

```
UTL_FILE.GET_LINE (
   file        IN  FILE_TYPE,
   buffer      OUT VARCHAR2,
   len         IN  PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 285-16 GET_LINE Procedure Parameters

Parameters	Description
file	Active file handle returned by an FOPEN call. The file must be open for reading (mode r); otherwise an INVALID_OPERATION exception is raised.
buffer	Data buffer to receive the line read from the file
len	The number of bytes read from the file. Default is NULL. If NULL, Oracle supplies the value of max_linesize.

Exceptions

INVALID_FILEHANDLE

INVALID_OPERATION

NO_DATA_FOUND

READ_ERROR

Usage Notes

If the line does not fit in the buffer, a `READ_ERROR` exception is raised. If no text was read due to end of file, the `NO_DATA_FOUND` exception is raised. If the file is opened for byte mode operations, the `INVALID_OPERATION` exception is raised.

Because the line terminator character is not read into the buffer, reading blank lines returns empty strings.

The maximum size of the `buffer` parameter is 32767 bytes unless you specify a smaller size in `FOPEN`.If unspecified, Oracle supplies a default value of 1024. See also "GET_LINE_NCHAR Procedure".

285.7.13 GET_LINE_NCHAR Procedure

This procedure reads text from the open file identified by the file handle and places the text in the output buffer parameter. With this function, you can read a text file in Unicode instead of in the database character set.

The file must be opened in national character set mode, and must be encoded in the UTF8 character set. The expected buffer datatype is `NVARCHAR2`. If a variable of another datatype, such as `NCHAR`, `NCLOB`, or `VARCHAR2` is specified, PL/SQL will perform standard implicit conversion from `NVARCHAR2` after the text is read.

See also GET_LINE Procedure

Syntax

```
UTL_FILE.GET_LINE_NCHAR (
   file        IN  FILE_TYPE,
   buffer      OUT NVARCHAR2,
   len         IN  PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 285-17 GET_LINE_NCHAR Procedure Parameters

Parameters	Description
`file`	Active file handle returned by an `FOPEN_NCHAR` call. The file must be open for reading (mode `r`). If the file is opened by `FOPEN` instead of `FOPEN_NCHAR`, a `CHARSETMISMATCH` exception is raised.
`buffer`	Data buffer to receive the line read from the file
`len`	The number of bytes read from the file. Default is `NULL`. If `NULL`, Oracle supplies the value of max_linesize.

Exceptions

`INVALID_FILEHANDLE`

`INVALID_OPERATION`

`NO_DATA_FOUND`

`READ_ERROR`

285.7.14 GET_RAW Procedure

This procedure reads a `RAW` string value from a file and adjusts the file pointer ahead by the number of bytes read. `UTL_FILE.GET_RAW` ignores line terminators.

Syntax

```
UTL_FILE.GET_RAW (
   file       IN            UTL_FILE.FILE_TYPE,
   buffer     OUT NOCOPY    RAW,
   len        IN            PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 285-18 GET_RAW Procedure Parameters

Parameters	Description
`file`	File handle
`buffer`	`RAW` data
`len`	The number of bytes read from the file. Default is `NULL`. If `NULL`, `len` is assumed to be the maximum length of `RAW`.

Exceptions

`INVALID_FILEHANDLE`

`INVALID_OPERATION`

`LENGTH_MISMATCH`

`NO_DATA_FOUND`

`READ_ERROR`

Usage Notes

The subprogram will raise `No_Data_Found` when it attempts to read past the end of the file. Your application should allow for this by catching the exception in its processing loop.

```
PROCEDURE Sys.p (n IN VARCHAR2) IS
      h     UTL_FILE.FILE_TYPE := UTL_FILE.FOPEN('D', n, 'r', 32767);
      Buf   RAW(32767);
      Amnt  CONSTANT PLS_INTEGER := 32767;
    BEGIN
      LOOP
        BEGIN
          Utl_File.Get_Raw(h, Buf, Amnt);
          -- Do something with this chunk
        EXCEPTION WHEN No_Data_Found THEN EXIT; END;
      END LOOP;
      UTL_FILE.FCLOSE (h);
    END;
```

285.7.15 IS_OPEN Function

This function tests a file handle to see if it identifies an open file.

`IS_OPEN` reports only whether a file handle represents a file that has been opened, but not yet closed. It does not guarantee that there will be no operating system errors when you attempt to use the file handle.

Syntax

```
UTL_FILE.IS_OPEN (
   file  IN FILE_TYPE)
  RETURN BOOLEAN;
```

Parameters

Table 285-19 IS_OPEN Function Parameters

Parameter	Description
`file`	Active file handle returned by an `FOPEN` or `FOPEN_NCHAR` call

Return Values

`TRUE` or `FALSE`

Exceptions

`INVALID_FILEHANDLE`

285.7.16 NEW_LINE Procedure

This procedure writes one or more line terminators to the file identified by the input file handle.

This procedure is separate from `PUT` because the line terminator is a platform-specific character or sequence of characters.

Syntax

```
UTL_FILE.NEW_LINE (
   file     IN FILE_TYPE,
   lines    IN BINARY_INTEGER := 1);
```

Parameters

Table 285-20 NEW_LINE Procedure Parameters

Parameters	Description
`file`	Active file handle returned by an `FOPEN` or `FOPEN_NCHAR` call
`lines`	Number of line terminators to be written to the file

Exceptions

`INVALID_FILEHANDLE`

INVALID_OPERATION

WRITE_ERROR

285.7.17 PUT Procedure

PUT writes the text string stored in the buffer parameter to the open file identified by the file handle.

The file must be open for write operations. No line terminator is appended by PUT; use NEW_LINE to terminate the line or use PUT_LINE to write a complete line with a line terminator. See also "PUT_NCHAR Procedure".

Syntax

```
UTL_FILE.PUT (
   file      IN FILE_TYPE,
   buffer    IN VARCHAR2);
```

Parameters

Table 285-21 PUT Procedure Parameters

Parameters	Description
file	Active file handle returned by an FOPEN_NCHAR call. The file must be open for writing.
buffer	Buffer that contains the text to be written to the file. User must have opened the file using mode w or mode a; otherwise, an INVALID_OPERATION exception is raised.

Usage Notes

The maximum size of the buffer parameter is 32767 bytes unless you specify a smaller size in FOPEN. If unspecified, Oracle supplies a default value of 1024. The sum of all sequential PUT calls cannot exceed 32767 without intermediate buffer flushes.

Exceptions

INVALID_FILEHANDLE

INVALID_OPERATION

WRITE_ERROR

285.7.18 PUT_LINE Procedure

This procedure writes the text string stored in the buffer parameter to the open file identified by the file handle.

The file must be open for write operations. PUT_LINE terminates the line with the platform-specific line terminator character or characters.

See also "PUT_LINE_NCHAR Procedure".

Syntax

```
UTL_FILE.PUT_LINE (
   file      IN FILE_TYPE,
   buffer    IN VARCHAR2,
   autoflush IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 285-22 PUT_LINE Procedure Parameters

Parameters	Description
`file`	Active file handle returned by an `FOPEN` call
`buffer`	Text buffer that contains the lines to be written to the file
`autoflush`	Flushes the buffer to disk after the `WRITE`

Exceptions

`INVALID_FILEHANDLE`

`INVALID_OPERATION`

`WRITE_ERROR`

Usage Notes

- The maximum size of the `buffer` parameter is 32767 bytes unless you specify a smaller size in `FOPEN`. If unspecified, Oracle supplies a default value of 1024. The sum of all sequential `PUT` calls cannot exceed 32767 without intermediate buffer flushes.
- If file is opened for byte mode operations, then the `INVALID OPERATION` exception is raised.

285.7.19 PUT_LINE_NCHAR Procedure

This procedure writes the text string stored in the buffer parameter to the open file identified by the file handle. With this function, you can write a text file in Unicode instead of in the database character set.

This procedure is equivalent to the PUT_NCHAR Procedure, except that the line separator is appended to the written text. See also PUT_LINE Procedure.

Syntax

```
UTL_FILE.PUT_LINE_NCHAR (
   file    IN FILE_TYPE,
   buffer  IN NVARCHAR2);
```

Parameters

Table 285-23 PUT_LINE_NCHAR Procedure Parameters

Parameters	Description
`file`	Active file handle returned by an `FOPEN_NCHAR` call. The file must be open for writing.

Table 285-23 (Cont.) PUT_LINE_NCHAR Procedure Parameters

Parameters	Description
buffer	Text buffer that contains the lines to be written to the file

Exceptions

INVALID_FILEHANDLE

INVALID_OPERATION

WRITE_ERROR

Usage Notes

- The maximum size of the buffer parameter is 32767 bytes unless you specify a smaller size in FOPEN. If unspecified, Oracle supplies a default value of 1024. The sum of all sequential PUT calls cannot exceed 32767 without intermediate buffer flushes.
- If file is opened for byte mode operations, then the INVALID OPERATION exception is raised.

285.7.20 PUT_NCHAR Procedure

This procedure writes the text string stored in the buffer parameter to the open file identified by the file handle.

With this function, you can write a text file in Unicode instead of in the database character set. The file must be opened in the national character set mode. The text string will be written in the UTF8 character set. The expected buffer datatype is NVARCHAR2. If a variable of another datatype is specified, PL/SQL will perform implicit conversion to NVARCHAR2 before writing the text.

Syntax

```
UTL_FILE.PUT_NCHAR (
   file      IN FILE_TYPE,
   buffer    IN NVARCHAR2);
```

Parameters

Table 285-24 PUT_NCHAR Procedure Parameters

Parameters	Description
file	Active file handle returned by an FOPEN_NCHAR call. If the file is opened by FOPEN instead of FOPEN_NCHAR, a CHARSETMISMATCH exception is raised.
buffer	Buffer that contains the text to be written to the file. User must have opened the file using mode w or mode a; otherwise, an INVALID_OPERATION exception is raised.

Exceptions

INVALID_FILEHANDLE

INVALID_OPERATION

```
WRITE_ERROR
```

Usage Notes

The maximum size of the `buffer` parameter is 32767 bytes unless you specify a smaller size in `FOPEN`. If unspecified, Oracle supplies a default value of 1024. The sum of all sequential `PUT` calls cannot exceed 32767 without intermediate buffer flushes.

Related Topics

- PUT Procedure
  `PUT` writes the text string stored in the buffer parameter to the open file identified by the file handle.

285.7.21 PUTF Procedure

This procedure is a formatted `PUT` procedure.

It works like a limited `printf`().

Syntax

```
UTL_FILE.PUTF (
   file    IN FILE_TYPE,
   format  IN VARCHAR2,
   [arg1   IN VARCHAR2  DEFAULT NULL,
   . . .
   arg5    IN VARCHAR2  DEFAULT NULL]);
```

Parameters

Table 285-25 PUTF Procedure Parameters

Parameters	Description
`file`	Active file handle returned by an `FOPEN` call
`format`	Format string that can contain text as well as the formatting characters `\n` and `%s`
`arg1..arg5`	From one to five operational argument strings. Argument strings are substituted, in order, for the `%s` formatters in the format string. If there are more formatters in the format parameter string than there are arguments, then an empty string is substituted for each `%s` for which there is no argument.

Usage Notes

- If file is opened for byte mode operations, then the `INVALID OPERATION` exception is raised.
- The format string can contain any text, but the character sequences `%s` and `\n` have special meaning.

Character Sequence	Meaning
`%s`	Substitute this sequence with the string value of the next argument in the argument list.
`\n`	Substitute with the appropriate platform-specific line terminator.

Exceptions

`INVALID_FILEHANDLE`

`INVALID_OPERATION`

`WRITE_ERROR`

Examples

The following example writes the lines:

```
Hello, world!
I come from Zork with greetings for all earthlings.

my_world  varchar2(4) := 'Zork';
...
PUTF(my_handle, 'Hello, world!\nI come from %s with %s.\n',
                my_world,
                'greetings for all earthlings');
```

If there are more `%s` formatters in the format parameter than there are arguments, then an empty string is substituted for each `%s` for which there is no matching argument.

Related Topics

- PUTF_NCHAR Procedure
  This procedure is a formatted version of a PUT_NCHAR Procedure.

285.7.22 PUTF_NCHAR Procedure

This procedure is a formatted version of a PUT_NCHAR Procedure.

Using `PUTF_NCHAR`, you can write a text file in Unicode instead of in the database character set. It accepts a format string with formatting elements `\n` and `%s`, and up to five arguments to be substituted for consecutive instances of `%s` in the format string. The expected datatype of the format string and the arguments is `NVARCHAR2`.

If variables of another datatype are specified, PL/SQL will perform implicit conversion to `NVARCHAR2` before formatting the text. Formatted text is written in the UTF8 character set to the file identified by the file handle. The file must be opened in the national character set mode.

Syntax

```
UTL_FILE.PUTF_NCHAR (
   file    IN FILE_TYPE,
   format  IN NVARCHAR2,
   [arg1   IN NVARCHAR2  DEFAULT NULL,
   . . .
   arg5    IN NVARCHAR2  DEFAULT NULL]);
```

Parameters

Table 285-26 PUTF_NCHAR Procedure Parameters

Parameters	Description
file	Active file handle returned by an FOPEN_NCHAR call. The file must be open for reading (mode r). If the file is opened by FOPEN instead of FOPEN_NCHAR, a CHARSETMISMATCH exception is raised.
format	Format string that can contain text as well as the formatting characters \n and %s
arg1..arg5	From one to five operational argument strings. Argument strings are substituted, in order, for the %s formatters in the format string. If there are more formatters in the format parameter string than there are arguments, then an empty string is substituted for each %s for which there is no argument.

Exceptions

INVALID_FILEHANDLE

INVALID_OPERATION

WRITE_ERROR

Usage Notes

- The maximum size of the buffer parameter is 32767 bytes unless you specify a smaller size in FOPEN. If unspecified, Oracle supplies a default value of 1024. The sum of all sequential PUT calls cannot exceed 32767 without intermediate buffer flushes.
- If file is opened for byte mode operations, then the INVALID OPERATION exception is raised.

Related Topics

- PUT_NCHAR Procedure
  This procedure writes the text string stored in the buffer parameter to the open file identified by the file handle.

285.7.23 PUT_RAW Procedure

This procedure accepts as input a RAW data value and writes the value to the output buffer.

Syntax

```
UTL_FILE.PUT_RAW (
   file          IN    UTL_FILE.FILE_TYPE,
   buffer        IN    RAW,
   autoflush     IN    BOOLEAN DEFAULT FALSE);
```

Parameters

Table 285-27 PUT_RAW Procedure Parameters

Parameters	Description
`file`	File handle
`buffer`	The `RAW` data written to the buffer
`autoflush`	If `TRUE`, then performs a flush after writing the value to the output buffer; default is `FALSE`.

Exceptions

`INVALID_FILEHANDLE`

`INVALID_OPERATION`

`WRITE_ERROR`

Usage Notes

You can request an automatic flush of the buffer by setting the third argument to `TRUE`.

The maximum size of the `buffer` parameter is 32767 bytes unless you specify a smaller size in `FOPEN`. If unspecified, Oracle supplies a default value of 1024. The sum of all sequential `PUT` calls cannot exceed 32767 without intermediate buffer flushes.

286
UTL_HTTP

The `UTL_HTTP` package makes Hypertext Transfer Protocol (HTTP) callouts from SQL and PL/SQL. You can use it to access data on the Internet over HTTP.

When the package fetches data from a Web site using HTTPS, it requires Oracle Wallet which can be created by `orapki` utility. Non-HTTPS fetches do not require an Oracle wallet.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Exceptions
- Examples
- Data Structures
- Operations
- Subprogram Groups
- Summary of UTL_HTTP Subprograms

> **See Also:**
>
> - UTL_URL
> - UTL_SMTP
> - *Oracle Database Enterprise User Security Administrator's Guide* for more information on Wallet Manager

286.1 UTL_HTTP Overview

With the `UTL_HTTP` package, you can write PL/SQL programs that communicate with Web (HTTP) servers. `UTL_HTTP` also contains a function that can be used in SQL queries.

The package supports HTTP over the Secured Socket Layer protocol (SSL), also known as HTTPS. It also supports SSL client authentication by sending the client-certificate in a wallet to authenticate with the remote Web server.

Other Internet-related data-access protocols (such as the File Transfer Protocol (FTP) or the Gopher protocol) are also supported using an HTTP proxy server that supports those protocols.

286.2 UTL_HTTP Security Model

This package is an invoker's rights package. The invoking user will need the `connect` privilege granted in the access control list assigned to the remote network host to which he wants to connect, as well as the `use-client-certificates` or the `use-passwords` privilege to authenticate himself with the remote Web server using the credentials stored in an Oracle wallet.

> **Note:**
>
> For more information about managing fine-grained access, see *Oracle Database Security Guide*

286.3 UTL_HTTP Constants

The `UTL_HTTP` package defines several constants to use when specifying parameter values.

These are shown in following tables.

- Table 286-1
- Table 286-2
- Table 286-3

Table 286-1 UTL_HTTP Constants - HTTP Versions

Name	Type	Value	Description
`HTTP_VERSION_1_0`	`VARCHAR2(10)`	`'HTTP/1.0'`	Denotes HTTP version 1.0 that can be used in the function `BEGIN_REQUEST`.
`HTTP_VERSION_1_1`	`VARCHAR2(10)`	`'HTTP/1.1'`	Denotes HTTP version 1.1 that can be used in the function `BEGIN_REQUEST`.

Table 286-2 UTL_HTTP Constants - Default Ports

Name	Type	Value	Description
`DEFAULT_HTTP_PORT`	`PLS_INTEGER`	`80`	The default TCP/IP port (80) at which a Web server or proxy server listens
`DEFAULT_HTTPS_PORT`	`PLS_INTEGER`	`443`	The default TCP/IP port (443) at which an HTTPS Web server listens

Table 286-3 UTL_HTTP Constants - HTTP 1.1 Status Codes

Name	Type	Value	Description
HTTP_CONTINUE	PLS_INTEGER	100	The client should continue with its request. This interim response is used to inform the client that the initial part of the request has been received and has not yet been rejected by the server.
HTTP_SWITCHING_PROTOCOLS	PLS_INTEGER	101	The server understands and is willing to comply with the client's request, through the Upgrade message header field, for a change in the application protocol being used on this connection. The server will switch protocols to those defined by the response's Upgrade header field immediately after the empty line which terminates the 101 response.
HTTP_OK	PLS_INTEGER	200	The request has succeeded. The information returned with the response is dependent on the method used in the request
HTTP_CREATED CONSTANT	PLS_INTEGER	201	The request has been fulfilled and resulted in a new resource being created.
HTTP_ACCEPTED	PLS_INTEGER	202	The request has been accepted for processing, but the processing has not been completed. The request might or might not eventually be acted upon, as it might be disallowed when processing actually takes place.
HTTP_NON_AUTHORITATIVE_INFO	PLS_INTEGER	203	The returned metainformation in the entity-header is not the definitive set as available from the origin server, but is gathered from a local or a third-party copy.
HTTP_NO_CONTENT	PLS_INTEGER	204	The server has fulfilled the request but does not need to return an entity-body, and might want to return updated metainformation.
HTTP_RESET_CONTENT	PLS_INTEGER	205	The server has fulfilled the request and the user agent should reset the document view which caused the request to be sent. The response must not include an entity.
HTTP_PARTIAL_CONTENT	PLS_INTEGER	206	The server has fulfilled the partial GET request for the resource.
HTTP_MULTIPLE_CHOICES	PLS_INTEGER	300	The requested resource corresponds to any one of a set of representations, each with its own specific location, and agent- driven negotiation information is being provided so that the user (or user agent) can select a preferred representation and redirect its request to that location.
HTTP_MOVED_PERMANENTLY	PLS_INTEGER	301	The requested resource has been assigned a new permanent URI and any future references to this resource should use one of the returned URIs.
HTTP_FOUND CONSTANT	PLS_INTEGER	302	The requested resource resides temporarily under a different URI.

Table 286-3 (Cont.) UTL_HTTP Constants - HTTP 1.1 Status Codes

Name	Type	Value	Description
HTTP_SEE_OTHER	PLS_INTEGER	303	The response to the request can be found under a different URI and should be retrieved using a GET method on that resource.
HTTP_NOT_MODIFIED	PLS_INTEGER	304	If the client has performed a conditional GET request and access is allowed, but the document has not been modified, the server responds with this status code.
HTTP_USE_PROXY	PLS_INTEGER	305	The requested resource must be accessed through the proxy given by the Location field. The Location field gives the URI of the proxy.
HTTP_TEMPORARY_REDIRECT	PLS_INTEGER	307	The requested resource resides temporarily under a different URI.
HTTP_BAD_REQUEST	PLS_INTEGER	400	The request could not be understood by the server due to malformed syntax.
HTTP_UNAUTHORIZED	PLS_INTEGER	401	The request requires user authentication. The client may repeat the request with a suitable Authorization header field. If the request already included Authorization credentials, then the 401 response indicates that authorization has been refused for those credentials.
HTTP_PAYMENT_REQUIRED	PLS_INTEGER	402	This code is reserved for future use.
HTTP_FORBIDDEN	PLS_INTEGER	403	The server understood the request, but is refusing to fulfill it.
HTTP_NOT_FOUND	PLS_INTEGER	404	The server has not found anything matching the Request-URI.
HTTP_NOT_ACCEPTABLE	PLS_INTEGER	406	The resource identified by the request is only capable of generating response entities which have content characteristics not acceptable according to the accept headers sent in the request.
HTTP_PROXY_AUTH_REQUIRED	PLS_INTEGER	407	This code is similar to 401 (Unauthorized), but indicates that the client must first authenticate itself with the proxy.
HTTP_REQUEST_TIME_OUT	PLS_INTEGER	408	The client did not produce a request within the time that the server was prepared to wait.
HTTP_CONFLICT	PLS_INTEGER	409	The request could not be completed due to a conflict with the current state of the resource.
HTTP_GONE	PLS_INTEGER	410	The requested resource is no longer available at the server and no forwarding address is known.
HTTP_LENGTH_REQUIRED	PLS_INTEGER	411	The server refuses to accept the request without a defined `Content-Length`.
HTTP_PRECONDITION_FAILED	PLS_INTEGER	412	The precondition given in one or more of the request-header fields evaluated to false when it was tested on the server.

Table 286-3 (Cont.) UTL_HTTP Constants - HTTP 1.1 Status Codes

Name	Type	Value	Description
HTTP_REQUEST_ENTITY_TOO_LARGE CONSTANT	PLS_INTEGER	413	The server is refusing to process a request because the request entity is larger than the server is willing or able to process.
HTTP_REQUEST_URI_TOO_LARGE	PLS_INTEGER	414	The server is refusing to service the request because the Request-URI is longer than the server is willing to interpret.
HTTP_UNSUPPORTED_MEDIA_TYPE	PLS_INTEGER	415	The server is refusing to service the request because the entity of the request is in a format not supported by the requested resource for the requested method.
HTTP_REQ_RANGE_NOT_SATISFIABLE	PLS_INTEGER	416	A server returns a response with this status code if a request included a Range request-header field, and none of the range-specifier values in this field overlap the current extent of the selected resource, and the request did not include an If-Range request-header field.
HTTP_EXPECTATION_FAILED	PLS_INTEGER	417	The expectation given in an Expect request-header field could not be met by this server, or, if the server is a proxy, the server has unambiguous evidence that the request could not be met by the next-hop server.
HTTP_NOT_IMPLEMENTED	PLS_INTEGER	501	The server does not support the functionality required to fulfill the request.
HTTP_BAD_GATEWAY	PLS_INTEGER	502	The server, while acting as a gateway or proxy, received an invalid response from the upstream server it accessed in attempting to fulfill the request
HTTP_SERVICE_UNAVAILABLE	PLS_INTEGER	503	The server is currently unable to handle the request due to a temporary overloading or maintenance of the server.
HTTP_GATEWAY_TIME_OUT	PLS_INTEGER	504	The server, while acting as a gateway or proxy, did not receive a timely response from the upstream server specified by the URI (for example, HTTP, FTP, LDAP) or some other auxiliary server (for example, DNS) it needed to access in attempting to complete the request.
HTTP_VERSION_NOT_SUPPORTED	PLS_INTEGER	505	The server does not support, or refuses to support, the HTTP protocol version that was used in the request message.

286.4 UTL_HTTP Exceptions

Exceptions indicate that the UTL_HTTP package encountered issues.

The following table lists these exceptions . By default, `UTL_HTTP` raises the exception `request_failed` when a request fails to execute. If the package is set to raise a detailed exception by `set_detailed_excp_support`, the rest of the exceptions will be raised directly (except for the exception `end_of_body`, which will be raised by `READ_TEXT`, `READ_LINE`, and `READ_RAW` regardless of the setting).

Table 286-4 UTL_HTTP Exceptions

Exception	Error Code	Reason	Where Raised
BAD_ARGUMENT	29261	The argument passed to the interface is bad	Any HTTP request or response interface when detailed_exception is enabled
BAD_URL	29262	The requested URL is badly formed	BEGIN_REQUEST, when detailed_exception is enabled
END_OF_BODY	29266	The end of HTTP response body is reached	READ_RAW, READ_TEXT, and READ_LINE, when detailed_exception is enabled
HEADER_NOT_FOUND	29265	The header is not found	GET_HEADER, GET_HEADER_BY_NAME, when detailed_exception is enabled
HTTP_CLIENT_ERROR	29268	From GET_RESPONSE, the response status code indicates that a client error has occurred (status code in 4xx range). Or from begin_request, the HTTP proxy returns a status code in the 4xx range when making an HTTPS request through the proxy.	GET_RESPONSE, BEGIN_REQUEST when detailed_exception is enabled
HTTP_SERVER_ERROR	29269	From GET_RESPONSE, the response status code indicates that a client error has occurred (status code in 5xx range). Or from begin_request, the HTTP proxy returns a status code in the 5xx range when making an HTTPS request through the proxy.	GET_RESPONSE, BEGIN_REQUEST when detailed_exception is enabled
NETWORK_ACCESS_DENIED	24247	Access to the remote network host or credentials in an Oracle wallet is denied	BEGIN_REQUEST and SET_AUTHENTICATION_FROM_WALLET when detailed_exception is enabled
ILLEGAL_CALL	29267	The call to UTL_HTTP is illegal at the current state of the HTTP request	SET_HEADER, SET_AUTHENTICATION, and SET_PERSISTENT_CONN_SUPPORT, when detailed_exception is enabled
PARTIAL_MULTIBYTE_EXCEPTION	29275	No complete character is read and a partial multibyte character is found at the end of the response body	READ_TEXT and READ_LINE, when detailed_exception is enabled
PROTOCOL_ERROR	29263	An HTTP protocol error occurs when communicating with the Web server	SET_HEADER, GET_RESPONSE, READ_RAW, READ_TEXT, and READ_LINE, when detailed_exception is enabled
REQUEST_FAILED	29273	The request fails to executes	Any HTTP request or response interface when detailed_exception is disabled
TOO_MANY_REQUESTS	29270	Too many requests or responses are open	BEGIN_REQUEST, when detailed_exception is enabled
TRANSFER_TIMEOUT	29276	No data is read and a read timeout occurred	READ_TEXT and READ_LINE, when detailed_exception is enabled
UNKNOWN_SCHEME	29264	The scheme of the requested URL is unknown	BEGIN_REQUEST and GET_RESPONSE, when detailed_exception is enabled

> **Note:**
>
> The `partial_multibyte_char` and `transfer_timeout` exceptions are duplicates of the same exceptions defined in `UTL_TCP`. They are defined in this package so that the use of this package does not require the knowledge of the `UTL_TCP`. As those exceptions are duplicates, an exception handle that catches the `partial_multibyte_char` and `transfer_timeout` exceptions in this package also catch the exceptions in the `UTL_TCP`.

For `REQUEST` and `REQUEST_PIECES`, the `request_failed` exception is raised when any exception occurs and `detailed_exception` is disabled.

286.5 UTL_HTTP Examples

These five examples demonstrate how to use `UTL_HTTP`.

- General Usage
- Retrieving HTTP Response Headers
- Handling HTTP Authentication
- Retrieving and Restoring Cookies
- Making HTTP Request with Private Wallet and Cookie Table

286.5.1 UTL_HTTP General Usage

This is an general example of `UTL_HTTP` usage.

```
SET SERVEROUTPUT ON SIZE 40000

DECLARE
  req   UTL_HTTP.REQ;
  resp  UTL_HTTP.RESP;
  value VARCHAR2(1024);
BEGIN
  UTL_HTTP.SET_PROXY('proxy.my-company.com', 'corp.my-company.com');
  req := UTL_HTTP.BEGIN_REQUEST('http://www-hr.corp.my-company.com');
  UTL_HTTP.SET_HEADER(req, 'User-Agent', 'Mozilla/4.0');
  resp := UTL_HTTP.GET_RESPONSE(req);
  LOOP
    UTL_HTTP.READ_LINE(resp, value, TRUE);
    DBMS_OUTPUT.PUT_LINE(value);
  END LOOP;
  UTL_HTTP.END_RESPONSE(resp);
EXCEPTION
  WHEN UTL_HTTP.END_OF_BODY THEN
    UTL_HTTP.END_RESPONSE(resp);
END;
```

286.5.2 UTL_HTTP Retrieving HTTP Response Headers

This example shows how `UTL_HTTP` retrieves HTTP response headers.

```
SET SERVEROUTPUT ON SIZE 40000
```

```
DECLARE
  req   UTL_HTTP.REQ;
  resp  UTL_HTTP.RESP;
  name  VARCHAR2(256);
  value VARCHAR2(1024);
BEGIN
  UTL_HTTP.SET_PROXY('proxy.my-company.com', 'corp.my-company.com');
  req := UTL_HTTP.BEGIN_REQUEST('http://www-hr.corp.my-company.com');
  UTL_HTTP.SET_HEADER(req, 'User-Agent', 'Mozilla/4.0');
  resp := UTL_HTTP.GET_RESPONSE(req);
  DBMS_OUTPUT.PUT_LINE('HTTP response status code: ' || resp.status_code);
  DBMS_OUTPUT.PUT_LINE('HTTP response reason phrase: ' || resp.reason_phrase);
  FOR i IN 1..UTL_HTTP.GET_HEADER_COUNT(resp) LOOP
    UTL_HTTP.GET_HEADER(resp, i, name, value);
    DBMS_OUTPUT.PUT_LINE(name || ': ' || value);
  END LOOP;
  UTL_HTTP.END_RESPONSE(resp);
END;
```

286.5.3 UTL_HTTP Handling HTTP Authentication

This code sample indicates how `UTL_HTTP` handles HTTP authentication.

```
SET serveroutput ON SIZE 40000

CREATE OR REPLACE PROCEDURE get_page (url      IN VARCHAR2,
                                      username IN VARCHAR2 DEFAULT NULL,
                                      password IN VARCHAR2 DEFAULT NULL,
                                      realm    IN VARCHAR2 DEFAULT NULL) AS
  req       UTL_HTTP.REQ;
  resp      UTL_HTTP.RESP;
  my_scheme VARCHAR2(256);
  my_realm  VARCHAR2(256);
  name      VARCHAR2(256);
  value     VARCHAR2(256);
BEGIN
  -- Turn off checking of status code. We will check it by ourselves.
  UTL_HTTP.SET_RESPONSE_ERROR_CHECK(FALSE);
  req := UTL_HTTP.BEGIN_REQUEST(url);
  IF (username IS NOT NULL) THEN
    UTL_HTTP.SET_AUTHENTICATION(req, username, password); -- Use HTTP Basic Authen.
Scheme
  END IF;
  resp := UTL_HTTP.GET_RESPONSE(req);
  IF (resp.status_code = UTL_HTTP.HTTP_UNAUTHORIZED) THEN
    UTL_HTTP.GET_AUTHENTICATION(resp, my_scheme, my_realm, FALSE);
    DBMS_OUTPUT.PUT_LINE('Web proxy server is protected.');
    DBMS_OUTPUT.PUT('Please provide the required ' || my_scheme || ' authentication
username/password for realm ' || my_realm || '
     for the proxy server.');
    UTL_HTTP.END_RESPONSE(resp);
    RETURN;
  ELSIF (resp.status_code = UTL_HTTP.HTTP_PROXY_AUTH_REQUIRED) THEN
    UTL_HTTP.GET_AUTHENTICATION(resp, my_scheme, my_realm, TRUE);
    DBMS_OUTPUT.PUT_LINE('Web page ' || url || ' is protected.');
    DBMS_OUTPUT.PUT('Please provide the required ' || my_scheme || ' authentication
username/password for realm ' || my_realm || '
     for the Web page.');
    UTL_HTTP.END_RESPONSE(resp);
    RETURN;
  END IF;
  FOR i IN 1..UTL_HTTP.GET_HEADER_COUNT(resp) LOOP
```

```
    UTL_HTTP.GET_HEADER(resp, i, name, value);
    DBMS_OUTPUT.PUT_LINE(name || ': ' || value);
  END LOOP;
  UTL_HTTP.END_RESPONSE(resp);
END;
```

286.5.4 UTL_HTTP Handling HTTP Digest Authentication

This code sample indicates how `UTL_HTTP` handles HTTP digest authentication.

```
declare
  url varchar2(32767);
  q utl_http.req;
  p utl_http.resp;
  pstatus pls_integer;
   begin
    url := 'http://slc10tzv.us.oracle.com:3000/digest.html';
    q := utl_http.begin_request(url);
    utl_http.set_authentication(q,
                                username => 'utlhttp_user',
                                password => 'welcome',
                                scheme   => 'Digest');
    p := utl_http.get_response(q);
    pstatus := p.status_code;

  -- status code returned from get_response should be 200
  dbms_output.put_line('-- response status: ' || p.status_code);
  utl_http.end_response(p);
  utl_http.end_request(q);
EXCEPTION WHEN OTHERS THEN
  utl_http.end_request(q);
end;
/
-- response status: 200"
```

286.5.5 UTL_HTTP Retrieving and Restoring Cookies

This example show how UTL_HTTP can be used to retrieve and restore cookies.

```
CREATE TABLE my_cookies (
    session_id  INTEGER,
    name        VARCHAR2(256),
    value       VARCHAR2(1024),
    domain      VARCHAR2(256),
    expire      DATE,
    path        VARCHAR2(1024),
    secure      VARCHAR2(1),
    version     INTEGER);

CREATE SEQUENCE session_id;
SET SERVEROUTPUT ON SIZE 40000

REM Retrieve cookies from UTL_HTTP
CREATE OR REPLACE FUNCTION save_cookies RETURN PLS_INTEGER AS
  cookies        UTL_HTTP.COOKIE_TABLE;
  my_session_id  PLS_INTEGER;
  secure         VARCHAR2(1);
BEGIN
  /* assume that some cookies have been set in previous HTTP requests. */
  UTL_HTTP.GET_COOKIES(cookies);
  SELECT session_id.nextval INTO my_session_id FROM DUAL;
```

```
  FOR i in 1..cookies.count LOOP
    IF (cookies(i).secure) THEN
      secure := 'Y';
    ELSE
      secure := 'N';
    END IF;
    INSERT INTO my_cookies
    VALUES (my_session_id, cookies(i).name, cookies(i).value,
            cookies(i).domain,
            cookies(i).expire, cookies(i).path, secure, cookies(i).version);
  END LOOP;
  RETURN my_session_id;
END;
/

REM Retrieve cookies from UTL_HTTP
CREATE OR REPLACE PROCEDURE restore_cookies (this_session_id IN PLS_INTEGER)
AS
  cookies        UTL_HTTP.COOKIE_TABLE;
  cookie         UTL_HTTP.COOKIE;
  i              PLS_INTEGER := 0;
  CURSOR c (c_session_id PLS_INTEGER) IS
    SELECT * FROM my_cookies WHERE session_id = c_session_id;
BEGIN
  FOR r IN c(this_session_id) LOOP
    i := i + 1;
    cookie.name     := r.name;
    cookie.value    := r.value;
    cookie.domain   := r.domain;
    cookie.expire   := r.expire;
    cookie.path     := r.path;
    IF (r.secure = 'Y') THEN
      cookie.secure := TRUE;
    ELSE
      cookie.secure := FALSE;
    END IF;
    cookie.version := r.version;
    cookies(i) := cookie;
  END LOOP;
  UTL_HTTP.CLEAR_COOKIES;
  UTL_HTTP.ADD_COOKIES(cookies);
END;
/
```

286.5.6 UTL_HTTP Making HTTP Request with Private Wallet and Cookie Table

This example shows how `UTL_HTTP` creates a request context with a wallet and cookie table, then makes an HTTP Request using that wallet and cookie table.

```
SET SERVEROUTPUT ON SIZE 40000

CREATE OR REPLACE PROCEDURE DISPLAY_PAGE(url IN VARCHAR2) AS
  request_context UTL_HTTP.REQUEST_CONTEXT_KEY;
  req             UTL_HTTP.REQ;
  resp            UTL_HTTP.RESP;
  data            VARCHAR2(1024);

BEGIN
```

```
  -- Create a request context with its wallet and cookie table
  request_context := UTL_HTTP.CREATE_REQUEST_CONTEXT(
                       wallet_path          => 'file:/oracle/wallets/test/wallet',
                       wallet_password      => '******',
                       enable_cookies       => TRUE,
                       max_cookies          => 300,
                       max_cookies_per_site => 20);

  -- Make a HTTP request using the private wallet and cookie
  -- table in the request context
  req := UTL_HTTP.BEGIN_REQUEST(
           url             => url,
           request_context => request_context);
  resp := UTL_HTTP.GET_RESPONSE(req);

  BEGIN
    LOOP
       UTL_HTTP.READ_TEXT(resp, data);
       DBMS_OUTPUT.PUT(data);
    END LOOP;
  EXCEPTION
      WHEN UTL_HTTP.END_OF_BODY THEN
         UTL_HTTP.END_RESPONSE(resp);
  END;

  -- Destroy the request context
  UTL_HTTP.DESTROY_REQUEST_CONTEXT(request_context);

END;

BEGIN
  DISPLAY_PAGE('https://www.example.com/');
END;
/
```

286.5.7 UTL_HTTP Using a Proxy Server

This example shows using a proxy server for an Oracle Database connection.

The `UTL_HTTP.REQUEST` procedure uses the proxy server host name and the port number to access the HTTPS URL from within the Oracle Database.

```
SELECT UTL_HTTP.REQUEST('<URL>', '<proxy_hostname>:<proxy_port_number>',
'<wallet_directory>', '<wallet_password>') FROM DUAL;
```

> **Note:**
>
> By default, the `UTL_HTTP.REQUEST` procedure enables you to omit or set the "proxy" argument as `NULL`.

286.6 UTL_HTTP Data Structures

Data structures are used to represent requests, responses, cookies, connections, and request context.

- REQ Type
- RESP Type
- COOKIE and COOKIE_TABLE Types
- CONNECTION Type
- REQUEST_CONTEXT_KEY Type

286.6.1 REQ Type

Use this PL/SQL record type to represent an HTTP request.

Syntax

```
TYPE req IS RECORD (
   url           VARCHAR2(32767),
   method        VARCHAR2(64),
   http_version  VARCHAR2(64));
```

Parameters

Table 286-5 REQ Type Parameters

Parameter	Description
`url`	The URL of the HTTP request. It is set after the request is created by `BEGIN_REQUEST`.
`method`	The method to be performed on the resource identified by the URL. It is set after the request is created by `BEGIN_REQUEST`.
`http_version`	The HTTP protocol version used to send the request. It is set after the request is created by `BEGIN_REQUEST`.

Usage Notes

The information returned in `REQ` from the interface `begin_request` is for read-only. Changing the field values in the record has no effect on the request.

There are other fields in `REQ` record type whose names begin with the prefix `private_`. The fields are private and are intended for use by implementation of the `UTL_HTTP` package. You should not modify the fields.

286.6.2 REQUEST_CONTEXT_KEY Type

This type is used to represent the key to a request context.

A request context is a context that holds a private wallet and cookie table to make a HTTP request. This private wallet and cookie table, unlike the session-wide ones maintained in the package, will not be shared with other HTTP requests within the database session.

Syntax

```
SUBTYPE request_context_key IS PLS_INTEGER;
```

Usage Notes

To provide enhanced security, UTL_HTTP allows PL/SQL programs to create request contexts. A request context is a private context that holds a wallet and a cookie table that will not be shared with other programs in the same database session when making HTTP requests and receiving HTTP responses. PL/SQL programs should use request contexts when they need to use wallets or cookies that contain sensitive information such as authentication credentials.

286.6.3 RESP Type

This PL/SQL record type is used to represent an HTTP response.

Syntax

```
TYPE resp IS RECORD (
   status_code    PLS_INTEGER,
   reason_phrase  VARCHAR2(256),
   http_version   VARCHAR2(64));
```

Parameters

Table 286-6 RESP Type Parameters

Parameter	Description
status_code	The status code returned by the Web server. It is a 3-digit integer that indicates the results of the HTTP request as handled by the Web server. It is set after the response is processed by GET_RESPONSE.
reason_phrase	The short textual message returned by the Web server that describe the status code. It gives a brief description of the results of the HTTP request as handled by the Web server. It is set after the response is processed by GET_RESPONSE.
http_version	The HTTP protocol version used in the HTTP response. It is set after the response is processed by GET_RESPONSE.

Usage Notes

The information returned in RESP from the interface GET_RESPONSE is read-only. There are other fields in the RESP record type whose names begin with the prefix private_. The fields are private and are intended for use by implementation of the UTL_HTTP package. You should not modify the fields.

286.6.4 COOKIE and COOKIE_TABLE Types

The COOKIE type is the PL/SQL record type that represents an HTTP cookie. The COOKIE_TABLE type is a PL/SQL index-by-table type that represents a collection of HTTP cookies.

Syntax

```
TYPE cookie IS RECORD (
   name  VARCHAR2(256),
```

```
  value  VARCHAR2(1024),
  domain  VARCHAR2(256),
  expire  TIMESTAMP WITH TIME ZONE,
  path  VARCHAR2(1024),
  secure  BOOLEAN,
  version  PLS_INTEGER,
  comment  VARCHAR2(1024));

TYPE cookie_table IS TABLE OF cookie INDEX BY binary_integer;
```

Fields of COOKIE Record Type

Table 286-7 shows the fields for the `COOKIE` and `COOKIE_TABLE` record types.

Table 286-7 Fields of COOKIE and COOKIE_TABLE Type

Field	Description
`name`	The name of the HTTP cookie
`value`	The value of the cookie
`domain`	The domain for which the cookie is valid
`expire`	The time by which the cookie will expire
`path`	The subset of URLs to which the cookie applies
`secure`	Should the cookie be returned to the Web server using secured means only.
`version`	The version of the HTTP cookie specification the cookie conforms. This field is `NULL` for Netscape cookies.
`comment`	The comment that describes the intended use of the cookie. This field is `NULL` for Netscape cookies.

Usage Notes

PL/SQL programs do not usually examine or change the cookie information stored in the `UTL_HTTP` package. The cookies are maintained by the package transparently. They are maintained inside the `UTL_HTTP` package, and they last for the duration of the database session only. PL/SQL applications that require cookies to be maintained beyond the lifetime of a database session can read the cookies using `GET_COOKIES`, store them persistently in a database table, and re-store the cookies back in the package using `ADD_COOKIES` in the next database session. All the fields in the `cookie` record, except for the comment field, must be stored. Do not alter the cookie information, which can result in an application error in the Web server or compromise the security of the PL/SQL and the Web server applications. See "Retrieving and Restoring Cookies".

286.6.5 CONNECTION Type

Use the PL/SQL record type to represent the remote hosts and TCP/IP ports of a network connection that is kept persistent after an HTTP request is completed, according to the HTTP 1.1 protocol specification. The persistent network connection may be reused by a subsequent HTTP request to the same host and port. The subsequent HTTP request may be completed faster because the network connection latency is avoided. `connection_table` is a PL/SQL table of `connection`.

For a direct HTTP persistent connection to a Web server, the `host` and `port` fields contain the host name and TCP/IP port number of the Web server. The `proxy_host` and `proxy_port` fields

are not set. For an HTTP persistent connection that was previously used to connect to a Web server using a proxy, the `proxy_host` and `proxy_port` fields contain the host name and TCP/IP port number of the proxy server. The host and port fields are not set, which indicates that the persistent connection, while connected to a proxy server, is not bound to any particular target Web server. An HTTP persistent connection to a proxy server can be used to access any target Web server that is using a proxy.

The `SSL` field indicates if Secured Socket Layer (SSL) is being used in an HTTP persistent connection. An HTTPS request is an HTTP request made over SSL. For an HTTPS (SSL) persistent connection connected using a proxy, the host and port fields contain the host name and TCP/IP port number of the target HTTPS Web server and the fields will always be set. An HTTPS persistent connection to an HTTPS Web server using a proxy server can only be reused to make another request to the same target Web server.

Syntax

```
TYPE connection IS RECORD (
   host  VARCHAR2(256),
   port  PLS_INTEGER,
   proxy_host  VARCHAR2(256),
   proxy_port  PLS_INTEGER,
   ssl  BOOLEAN);

TYPE connection_table IS TABLE OF connection INDEX BY BINARY_INTEGER;
```

286.7 UTL_HTTP Operations

These topics provide information about how `UTL_HTTP` makes HTTP requests from SQL and PL/SQL.

- Operational Flow
- Simple HTTP Fetches
- HTTP Requests
- HTTP Responses
- HTTP Persistent Connections
- Error Conditions
- Session Settings
- Request Context
- External Password Store

286.7.1 UTL_HTTP Operational Flow

The `UTL_HTTP` package provides access to the HTTP protocol.

The interfaces must be called in the order shown in the following illustration, or an exception will be raised.

Figure 286-1 Flow of the Core UTL_HTTP Package

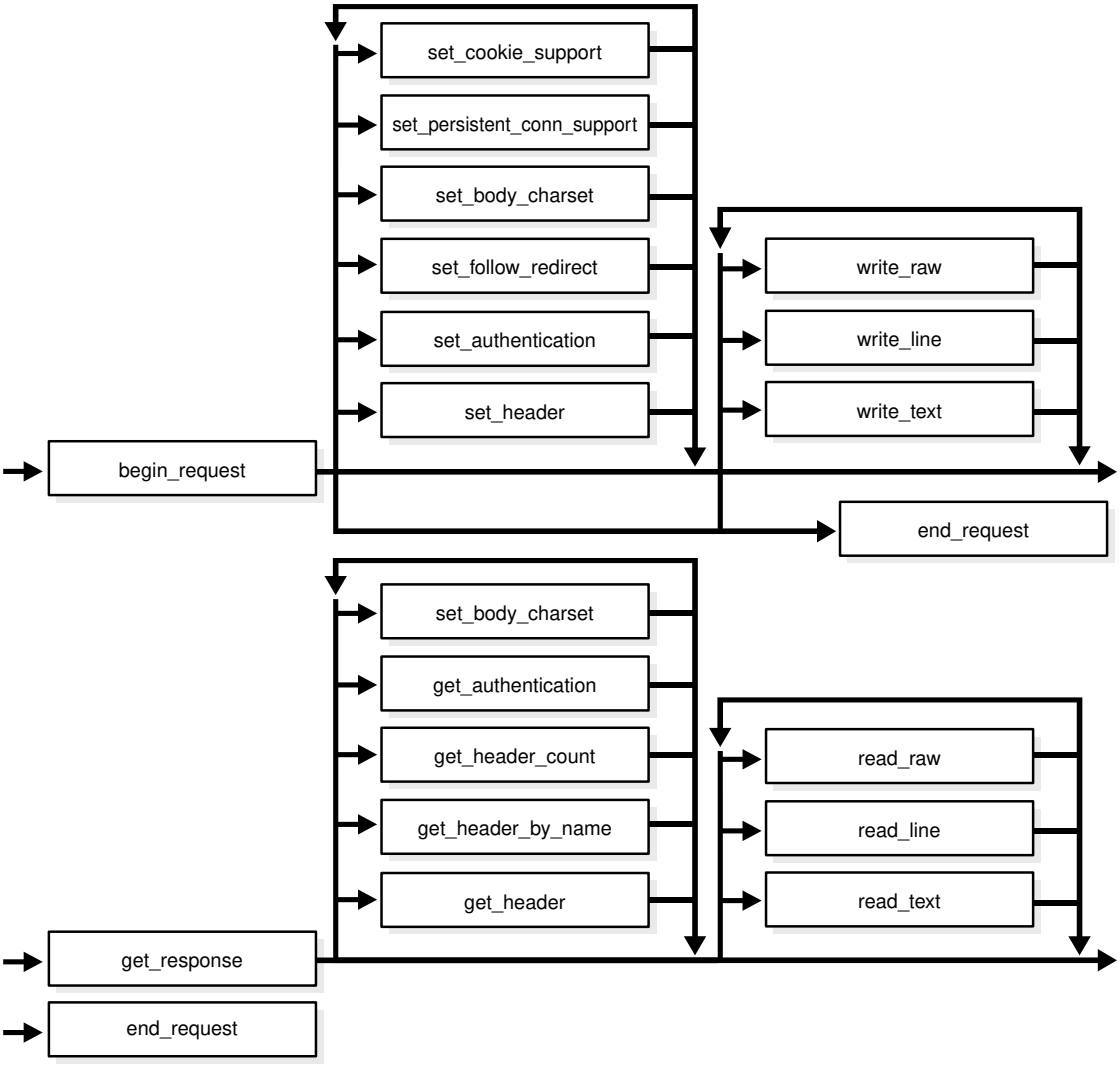

illustration: arpls006
release: 9
caption: The flow of the core UTL_HTTP Package API
date: 9/5/01
platform: pc

The following can be called at any time:

- Non-protocol interfaces that manipulate cookies
  - GET_COOKIE_COUNT
  - GET_COOKIES
  - ADD_COOKIES
  - CLEAR_COOKIES
- Persistent connections
  - GET_PERSISTENT_CONN_COUNT
  - GET_PERSISTENT_CONNS
  - CLOSE_PERSISTENT_CONN
  - CLOSE_PERSISTENT_CONNS
- Interfaces that manipulate attributes and configurations of the UTL_HTTP package in the current session
  - SET_PROXY
  - GET_PROXY
  - SET_COOKIE_SUPPORT
  - GET_COOKIE_SUPPORT
  - SET_FOLLOW_REDIRECT
  - GET_FOLLOW_REDIRECT
  - SET_BODY_CHARSET
  - GET_BODY_CHARSET
  - SET_PERSISTENT_CONN_SUPPORT
  - GET_PERSISTENT_CONN_SUPPORT
  - SET_DETAILED_EXCP_SUPPORT
  - GET_DETAILED_EXCP_SUPPORT
  - SET_WALLET
  - SET_TRANSFER_TIMEOUT
  - GET_TRANSFER_TIMEOUT
- Interfaces that retrieve the last detailed exception code and message UTL_HTTP package in the current session
  - GET_DETAILED_SQLCODE
  - GET_DETAILED_SQLERRM

> **Note:**
>
> Some of the request and response interfaces bear the same name as the interface that manipulates the attributes and configurations of the package in the current session. They are overloaded versions of the interface that manipulate a request or a response.

286.7.2 UTL_HTTP Simple HTTP Fetches

`REQUEST` and `REQUEST_PIECES` take a string uniform resource locator (URL), contact that site, and return the data (typically HTML) obtained from that site.

You should not expect `REQUEST` or `REQUEST_PIECES` to succeed in contacting a URL unless you can contact that URL by using a browser on the same machine (and with the same privileges, environment variables, and so on.)

If `REQUEST` or `REQUEST_PIECES` fails (for example, if it raises an exception, or if it returns an HTML-formatted error message, but you believe that the URL argument is correct), then try contacting that same URL with a browser to verify network availability from your machine. You may have a proxy server set in your browser that needs to be set with each `REQUEST` or `REQUEST_PIECES` call using the optional `proxy` parameter.

> **Note:**
>
> `UTL_HTTP` can also use environment variables to specify its proxy behavior. For example, on UNIX, setting the environment variable `http_proxy` to a URL uses that service as the proxy server for HTTP requests. Setting the environment variable `no_proxy` to a domain name does not use the HTTP proxy server for URLs in that domain. When the `UTL_HTTP` package is executed in the Oracle database server, the environment variables are the ones that are set when the database instance is started.

> **See Also:**
>
> Simple HTTP Fetches in a Single Call Subprograms

286.7.3 UTL_HTTP HTTP Requests

The HTTP Requests group of subprograms begin an HTTP request, manipulate attributes, and send the request information to the Web server. When a request is created, it inherits the default settings of the HTTP cookie support, follow-redirect, body character set, persistent-

connection support, and transfer timeout of the current session. The settings can be changed by calling the request interface.

> **See Also:**
>
> HTTP Requests Subprograms

286.7.4 UTL_HTTP HTTP Responses

The HTTP Responses group of subprograms manipulate an HTTP response obtained from GET_RESPONSE and receive response information from the Web server.

When a response is created for a request, it inherits settings of the HTTP cookie support, follow-redirect, body character set, persistent-connection support, and transfer timeout from the request. Only the body character set can be changed by calling the response interface.

> **See Also:**
>
> HTTP Responses Subprograms

286.7.5 UTL_HTTP HTTP Cookies

The `UTL_HTTP` package provides subprograms to manipulate HTTP cookies.

> **See Also:**
>
> HTTP Cookies Subprograms

286.7.6 UTL_HTTP HTTP Persistent Connections

The `UTL_HTTP` package provides subprograms to manipulate persistent connections.

> **See Also:**
>
> HTTP Persistent Connections Subprograms

286.7.7 UTL_HTTP Error Conditions

The `UTL_HTTP` package provides subprograms to retrieve error information.

> **See Also:**
>
> Error Conditions Subprograms

286.7.8 UTL_HTTP Session Settings

Session settings manipulate the configuration and default behavior of `UTL_HTTP` when HTTP requests are executed within a database user session.

When a request is created, it inherits the default settings of the HTTP cookie support, follow-redirect, body character set, persistent-connection support, and transfer timeout of the current session. Those settings can be changed later by calling the request interface. When a response is created for a request, it inherits those settings from the request. Only the body character set can be changed later by calling the response interface.

> **See Also:**
>
> Session Settings Subprograms

286.7.9 UTL_HTTP Request Context

The `UTL_HTTP` package maintains a common wallet and cookie table within the database session that all HTTP requests and responses share. This makes it easy for users to share the wallet or to maintain application state in the cookies within the session. However, if an application stores private information in the wallet or in the cookies that it does not want to share with other applications in the same database session, it may define a request context to hold its own wallet and cookie table and use this request context to make HTTP requests.

> **See Also:**
>
> HTTP Requests Subprograms

286.7.10 UTL_HTTP External Password Store

The `UTL_HTTP` package allows HTTP password credentials to be stored in an Oracle wallet's external password store. The external password store provides an easy but secure storage for passwords and frees the application developers from the need to maintain their own storage.

> **See Also:**
>
> SET_AUTHENTICATION_FROM_WALLET Procedure

286.8 UTL_HTTP Subprogram Groups

This section describes the `UTL_HTTP` subprograms. They are grouped by function.

- Simple HTTP Fetches in a Single Call Subprograms
- Session Settings Subprograms
- HTTP Requests Subprograms
- HTTP Request Contexts Subprograms
- HTTP Responses Subprograms
- HTTP Cookies Subprograms
- HTTP Persistent Connections Subprograms
- Error Conditions Subprograms

286.8.1 UTL_HTTP Simple HTTP Fetches in a Single Call Subprograms

`REQUEST` and `REQUEST_PIECES` take a string uniform resource locator (URL), contact that site, and return the data (typically HTML) obtained from that site.

Table 286-8 UTL_HTTP Subprograms—Simple HTTP Fetches in a Single Call

Subprogram	Description
REQUEST Function	Returns up to the first 2000 bytes of the data retrieved from the given URL. This function can be used directly in SQL queries.
REQUEST_PIECES Function	Returns a PL/SQL table of 2000-byte pieces of the data retrieved from the given URL

286.8.2 UTL_HTTP Session Settings Subprograms

This table lists and briefly describes the `UTL_HTTP` Session Settings Subprograms.

Table 286-9 UTL_HTTP Subprograms—Session Settings

Subprogram	Description
GET_BODY_CHARSET Procedure	Retrieves the default character set of the body of all future HTTP requests

Table 286-9 (Cont.) UTL_HTTP Subprograms—Session Settings

Subprogram	Description
GET_COOKIE_SUPPORT Procedure	Retrieves the current cookie support settings
GET_DETAILED_EXCP_SUPPORT Procedure	Checks if the `UTL_HTTP` package will raise a detailed exception or not
GET_FOLLOW_REDIRECT Procedure	Retrieves the follow-redirect setting in the current session
GET_PERSISTENT_CONN_SUPPORT Procedure	Checks if the persistent connection support is enabled and gets the maximum number of persistent connections in the current session
GET_PROXY Procedure	Retrieves the current proxy settings
GET_RESPONSE_ERROR_CHECK Procedure	Checks if the response error check is set or not
GET_TRANSFER_TIMEOUT Procedure	Retrieves the current network transfer timeout value
SET_TRANSFER_TIMEOUT Procedure	Sets the default character set of the body of all future HTTP requests when the media type is `text` and the character set is not specified in the `Content-Type` header
SET_COOKIE_SUPPORT Procedures	Sets whether or not future HTTP requests will support HTTP cookies; sets the maximum number of cookies maintained in the current database user session
SET_DETAILED_EXCP_SUPPORT Procedure	Sets the `UTL_HTTP` package to raise a detailed exception
SET_FOLLOW_REDIRECT Procedures	Sets the maximum number of times `UTL_HTTP` follows the HTTP redirect instruction in the HTTP responses to future requests in the `GET_RESPONSE` function
SET_PERSISTENT_CONN_SUPPORT Procedure	Sets whether or not future HTTP requests will support the HTTP 1.1 persistent connection; sets the maximum number of persistent connections maintained in the current database user session
SET_PROXY Procedure	Sets the proxy to be used for requests of HTTP or other protocols
SET_RESPONSE_ERROR_CHECK Procedure	Sets whether or not `GET_RESPONSE` raises an exception when the Web server returns a status code that indicates an error—a status code in the 4xx or 5xx ranges
SET_TRANSFER_TIMEOUT Procedure	Sets the timeout value for `UTL_HTTP` to read the HTTP response from the Web server or proxy server
SET_WALLET Procedure	Sets the Oracle Wallet used for all HTTP requests over Secured Socket Layer (SSL), that is, HTTPS

286.8.3 UTL_HTTP HTTP Requests Subprograms

This table lists and briefly describes the `UTL_HTTP` HTTP Requests.

Table 286-10 UTL_HTTP Subprograms—HTTP Requests

Subprogram	Description
BEGIN_REQUEST Function	Begins a new HTTP request.`UTL_HTTP` establishes the network connection to the target Web server or the proxy server and sends the HTTP request line.
SET_HEADER Procedure	Sets an HTTP request header. The request header is sent to the Web server as soon as it is set.

Table 286-10 (Cont.) UTL_HTTP Subprograms—HTTP Requests

Subprogram	Description
SET_AUTHENTICATION Procedure	Sets HTTP authentication information in the HTTP request header. The Web server needs this information to authorize the request.
SET_AUTHENTICATION_FROM_WALLET Procedure	Sets the HTTP authentication information in the HTTP request header needed for the request to be authorized by the Web server using the username and password credential stored in the Oracle wallet.
SET_BODY_CHARSET Procedures	Sets the character set of the request body when the media type is `text` but the character set is not specified in the `Content-Type` header
SET_COOKIE_SUPPORT Procedures	Enables or disables support for the HTTP cookies in the request
SET_FOLLOW_REDIRECT Procedures	Sets the maximum number of times `UTL_HTTP` follows the HTTP redirect instruction in the HTTP response to this request in the GET_RESPONSE Function
SET_PERSISTENT_CONN_SUPPORT Procedure	Enables or disables support for the HTTP 1.1 persistent-connection in the request
SET_PROXY Procedure	Writes a text line in the HTTP request body and ends the line with new-line characters (CRLF as defined in `UTL_TCP`)
WRITE_RAW Procedure	Writes some binary data in the HTTP request body
WRITE_TEXT Procedure	Writes some text data in the HTTP request body

286.8.4 UTL_HTTP HTTP Request Contexts Subprograms

`UTL_HTTP` HTTP Request Contexts subprograms create or destroy a request context.

The following table lists and briefly describes the `UTL_HTTP` HTTP Request Contexts.

Table 286-11 UTL_HTTP Subprograms—HTTP Request Contexts

Subprogram	Description
CREATE_REQUEST_CONTEXT Function	Creates a request context in `UTL_HTTP` for a wallet and a cookie table
DESTROY_REQUEST_CONTEXT Procedure	Destroys a request context in `UTL_HTTP`

286.8.5 UTL_HTTP HTTP Responses Subprograms

This table lists and briefly describes the HTTP Responses Subprograms of `UTL_HTTP`.

Table 286-12 UTL_HTTP Subprograms—HTTP Responses

Subprogram	Description
END_RESPONSE Procedure	Ends the HTTP response. It completes the HTTP request and response.
GET_AUTHENTICATION Procedure	Retrieves the HTTP authentication information needed for the request to be accepted by the Web server as indicated in the HTTP response header
GET_HEADER Procedure	Returns the n^{th} HTTP response header name and value returned in the response
GET_HEADER_BY_NAME Procedure	Returns the HTTP response header value returned in the response given the name of the header

Table 286-12 (Cont.) UTL_HTTP Subprograms—HTTP Responses

Subprogram	Description
GET_HEADER_COUNT Function	Returns the number of HTTP response headers returned in the response
GET_RESPONSE Function	Reads the HTTP response. When the function returns, the status line and the HTTP response headers have been read and processed.
READ_LINE Procedure	Reads the HTTP response body in text form until the end of line is reached and returns the output in the caller-supplied buffer
READ_RAW Procedure	Reads the HTTP response body in binary form and returns the output in the caller-supplied buffer
READ_TEXT Procedure	Reads the HTTP response body in text form and returns the output in the caller-supplied buffer
SET_BODY_CHARSET Procedures	Sets the character set of the response body when the media type is "text" but the character set is not specified in the `Content-Type` header

286.8.6 UTL_HTTP HTTP Cookies Subprograms

The HTTP cookies subprograms manages cookies in the `UTL_HTTP` package.

The following table lists and briefly describes the HTTP cookies subprograms of `UTL_HTTP`.

Table 286-13 UTL_HTTP Subprograms—HTTP Cookies

Subprogram	Description
ADD_COOKIES Procedure	Add the cookies either to a request context or to the `UTL_HTTP` package's session state
CLEAR_COOKIES Procedure	Clears all the cookies maintained either in a request context or in the `UTL_HTTP` package's session state
GET_COOKIE_COUNT Function	Returns the number of cookies maintained either in a request context or in the `UTL_HTTP` package's session states
GET_COOKIES Function	Returns all the cookies maintained either in a request context or in the `UTL_HTTP` package's session state.

286.8.7 UTL_HTTP HTTP Persistent Connections Subprograms

This table lists and briefly describes the `UTL_HTTP` HTTP Persistent Connections subprograms.

Table 286-14 UTL_HTTP Subprograms—HTTP Persistent Connections

Subprogram	Description
CLOSE_PERSISTENT_CONN Procedure	Closes an HTTP persistent connection maintained by the `UTL_HTTP` package in the current database session
CLOSE_PERSISTENT_CONNS Procedure	Closes a group of HTTP persistent connections maintained by the `UTL_HTTP` package in the current database session
GET_PERSISTENT_CONN_COUNT Function	Returns the number of network connections currently kept persistent by the `UTL_HTTP` package to the Web servers
GET_PERSISTENT_CONNS Procedure	Returns all the network connections currently kept persistent by the `UTL_HTTP` package to the Web servers

286.8.8 UTL_HTTP Error Conditions Subprograms

This table lists and briefly describes error conditions subprograms of `UTL_HTTP` .

Table 286-15 UTL_HTTP Subprograms—Error Conditions

Subprogram	Description
GET_DETAILED_SQLCODE Function	Retrieves the detailed SQLCODE of the last exception raised
GET_DETAILED_SQLERRM Function	Retrieves the detailed SQLERRM of the last exception raised

286.9 Summary of UTL_HTTP Subprograms

This table lists the `UTL_HTTP` subprograms and briefly describes them.

Table 286-16 UTL_HTTP Package Subprograms

Subprogram	Description	Group
ADD_COOKIES Procedure	Add the cookies either to a request context or to the `UTL_HTTP` package's session state	HTTP Cookies Subprograms
BEGIN_REQUEST Function	Begins a new HTTP request. `UTL_HTTP` establishes the network connection to the target Web server or the proxy server and sends the HTTP request line	HTTP Requests Subprograms
CLEAR_COOKIES Procedure	Clears all the cookies maintained either in a request context or in the `UTL_HTTP` package's session state	HTTP Cookies Subprograms
CLOSE_PERSISTENT_CONN Procedure	Closes an HTTP persistent connection maintained by the `UTL_HTTP` package in the current database session	HTTP Persistent Connections Subprograms
CLOSE_PERSISTENT_CONNS Procedure	Closes a group of HTTP persistent connections maintained by the `UTL_HTTP` package in the current database session	HTTP Persistent Connections Subprograms
CREATE_REQUEST_CONTEXT Function	Creates a request context in `UTL_HTTP` for a wallet and a cookie table	HTTP Requests Subprograms
DESTROY_REQUEST_CONTEXT Procedure	Destroys a request context in `UTL_HTTP` for a wallet and a cookie table	HTTP Requests Subprograms
END_REQUEST Procedure	Ends the HTTP request	HTTP Requests Subprograms
END_RESPONSE Procedure	Ends the HTTP response. It completes the HTTP request and response	HTTP Responses Subprograms

Table 286-16 (Cont.) UTL_HTTP Package Subprograms

Subprogram	Description	Group
GET_AUTHENTICATION Procedure	Retrieves the HTTP authentication information needed for the request to be accepted by the Web server as indicated in the HTTP response header	HTTP Responses Subprograms
GET_BODY_CHARSET Procedure	Retrieves the default character set of the body of all future HTTP requests	Session Settings Subprograms
GET_COOKIE_COUNT Function	Returns the number of cookies currently maintained by the `UTL_HTTP` package set by all Web servers	HTTP Cookies Subprograms
GET_COOKIE_SUPPORT Procedure	Retrieves the current cookie support settings	Session Settings Subprograms
GET_COOKIES Function	Returns all the cookies currently maintained by the `UTL_HTTP` package set by all Web servers	HTTP Cookies Subprograms
GET_DETAILED_EXCP_SUPPORT Procedure	Checks if the `UTL_HTTP` package will raise a detailed exception or not	Session Settings Subprograms
GET_DETAILED_SQLCODE Function	Retrieves the detailed `SQLCODE` of the last exception raised	Error Conditions Subprograms
GET_DETAILED_SQLERRM Function	Retrieves the detailed `SQLERRM` of the last exception raised	Error Conditions Subprograms
GET_FOLLOW_REDIRECT Procedure	Retrieves the follow-redirect setting in the current session	Session Settings Subprograms
GET_HEADER Procedure	Returns the n^{th} HTTP response header name and value returned in the response	HTTP Responses Subprograms
GET_HEADER_BY_NAME Procedure	Returns the HTTP response header value returned in the response given the name of the header	HTTP Responses Subprograms
GET_HEADER_COUNT Function	Returns the number of HTTP response headers returned in the response	HTTP Responses and HTTP Responses Subprograms
GET_PERSISTENT_CONN_COUNT Function	Returns the number of network connections currently kept persistent by the `UTL_HTTP` package to the Web servers	HTTP Persistent Connections Subprograms
GET_HEADER_COUNT Function	Sees whether or not future HTTP requests will support the HTTP 1.1 persistent connection; sets the maximum number of persistent connections maintained in the current database user session	Session Settings Subprograms

Table 286-16 (Cont.) UTL_HTTP Package Subprograms

Subprogram	Description	Group
GET_PERSISTENT_CONN_SUPPORT Procedure	Checks if the persistent connection support is enabled and gets the maximum number of persistent connections in the current session (see Session Settings Subprograms)	HTTP Persistent Connections Subprograms
GET_PERSISTENT_CONNS Procedure	Returns all the network connections currently kept persistent by the `UTL_HTTP` package to the Web servers	HTTP Persistent Connections Subprograms
GET_PROXY Procedure	Retrieves the current proxy settings	Session Settings Subprograms
GET_RESPONSE Function	Reads the HTTP response. When the function returns, the status line and the HTTP response headers have been read and processed	HTTP Responses Subprograms
GET_RESPONSE_ERROR_CHECK Procedure	Checks if the response error check is set or no	Session Settings Subprograms
GET_TRANSFER_TIMEOUT Procedure	Retrieves the current network transfer timeout value	Session Settings Subprograms
READ_LINE Procedure	Reads the HTTP response body in text form until the end of line is reached and returns the output in the caller-supplied buffer	HTTP Responses Subprograms
READ_RAW Procedure	Reads the HTTP response body in binary form and returns the output in the caller-supplied buffer	HTTP Responses Subprograms
READ_TEXT Procedure	Reads the HTTP response body in text form and returns the output in the caller-supplied buffer	HTTP Responses Subprograms
REQUEST Function	Returns up to the first 2000 bytes of the data retrieved from the given URL. This function can be used directly in SQL queries.	Simple HTTP Fetches in a Single Call Subprograms
REQUEST_PIECES Function	Returns a PL/SQL table of 2000-byte pieces of the data retrieved from the given URL	Simple HTTP Fetches in a Single Call Subprograms
SET_AUTHENTICATION Procedure	Sets HTTP authentication information in the HTTP request header. The Web server needs this information to authorize the request.	HTTP Requests Subprograms
SET_AUTHENTICATION_FROM_WALLET Procedure	Sets the HTTP authentication information in the HTTP request header needed for the request to be authorized by the Web server using the username and password credential stored in the Oracle wallet.	HTTP Requests Subprograms

Table 286-16 (Cont.) UTL_HTTP Package Subprograms

Subprogram	Description	Group
SET_BODY_CHARSET Procedures	Sets the default character set of the body of all future HTTP requests when the media type is `text` and the character set is not specified in the `Content-Type` header	Session Settings Subprograms
SET_BODY_CHARSET Procedures	Sets the character set of the request body when the media type is `text` but the character set is not specified in the `Content-Type` header	HTTP Requests Subprograms
SET_BODY_CHARSET Procedures	Sets the character set of the response body when the media type is "text" but the character set is not specified in the `Content-Type` header	HTTP Responses Subprograms and Session Settings Subprograms
SET_COOKIE_SUPPORT Procedures	Enables or disables support for the HTTP cookies in the request	HTTP Requests Subprograms
SET_DETAILED_EXCP_SUPPORT Procedure	Sets whether or not future HTTP requests will support HTTP cookies; sets the maximum number of cookies maintained in the current database user session	Session Settings Subprograms
SET_DETAILED_EXCP_SUPPORT Procedure	Sets the `UTL_HTTP` package to raise a detailed exception	Session Settings Subprograms
SET_FOLLOW_REDIRECT Procedures	Sets the maximum number of times `UTL_HTTP` follows the HTTP redirect instruction in the HTTP response to this request in the `GET_RESPONSE` function	HTTP Requests Subprograms
SET_HEADER Procedure	Sets the maximum number of times `UTL_HTTP` follows the HTTP redirect instruction in the HTTP responses to future requests in the `GET_RESPONSE` function	Session Settings Subprograms
SET_HEADER Procedure	Sets an HTTP request header. The request header is sent to the Web server as soon as it is set.	HTTP Requests Subprograms
SET_PERSISTENT_CONN_SUPPORT Procedure	Enables or disables support for the HTTP 1.1 persistent-connection in the request	HTTP Requests Subprograms
SET_PROXY Procedure	Sets the proxy to be used for requests of HTTP or other protocols	Session Settings and Session Settings Subprograms
SET_RESPONSE_ERROR_CHECK Procedure	Sets whether or not `GET_RESPONSE` raises an exception when the Web server returns a status code that indicates an error —a status code in the 4xx or 5xx ranges	Session Settings Subprograms

Table 286-16 (Cont.) UTL_HTTP Package Subprograms

Subprogram	Description	Group
SET_TRANSFER_TIMEOUT Procedure	Sets the timeout value for UTL_HTTP to read the HTTP response from the Web server or proxy server	Session Settings and Session Settings Subprograms
SET_WALLET Procedure	Sets the Oracle Wallet used for all HTTP requests over Secured Socket Layer (SSL), that is, HTTPS	Session Settings Subprograms
WRITE_LINE Procedure	Writes a text line in the HTTP request body and ends the line with new-line characters (CRLF as defined in UTL_TCP	HTTP Requests Subprograms
WRITE_RAW Procedure	Writes some binary data in the HTTP request body	HTTP Requests Subprograms
WRITE_TEXT Procedure	Writes some text data in the HTTP request body	HTTP Requests Subprograms

286.9.1 ADD_COOKIES Procedure

This procedure adds the cookies either to a request context or to the UTL_HTTP package's session state.

> **See Also:**
>
> HTTP Cookies and HTTP Cookies Subprograms

Syntax

```
UTL_HTTP.ADD_COOKIES (
   cookies          IN  cookie_table,
   request_context  IN  request_context_key DEFAULT NULL);
```

Parameters

Table 286-17 ADD_COOKIES Procedure Parameters

Parameter	Description
cookies	The cookies to be added
request_context	Request context to add the cookies. If NULL, the cookies will be added to the UTL_HTTP package's session state instead.

Usage Notes

The cookies that the package currently maintains are not cleared before new cookies are added.

286.9.2 BEGIN_REQUEST Function

This function begins a new HTTP request. `UTL_HTTP` establishes the network connection to the target Web server or the proxy server and sends the HTTP request line. The PL/SQL program continues the request by calling some other interface to complete the request.

The URL may contain the username and password needed to authenticate the request to the server. The format is:

```
scheme://[user[:password]@]host[:port]/[...]
```

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.BEGIN_REQUEST (
   url               IN  VARCHAR2,
   method            IN  VARCHAR2 DEFAULT 'GET',
   http_version      IN  VARCHAR2 DEFAULT NULL,
   request_context   IN  request_context_key DEFAULT NULL,
   https_host        IN  VARCHAR2 DEFAULT NULL)
RETURN req;
```

Parameters

Table 286-18 BEGIN_REQUEST Function Parameters

Parameter	Description
`url`	The URL of the HTTP request
`method`	The method performed on the resource identified by the URL
`http_version`	The HTTP protocol version that sends the request. The format of the protocol version is `HTTP/major-version.minor-version`, where `major-version` and `minor-version` are positive numbers. If this parameter is set to `NULL`, `UTL_HTTP` uses the latest HTTP protocol version that it supports to send the request. The latest version that the package supports is 1.1 and it can be upgraded to a later version. The default is `NULL`.
`request_context`	Request context that holds the private wallet and the cookie table to use in this `HTTP` request. If this parameter is `NULL`, the wallet and cookie table shared in the current database session will be used instead.
`https_host`	A string representing the host name. If the string does not begin with a wildcard, the string will be used as the host name for server name indication (SNI). If the string begins with a wildcard, the string will be used to match against the common name (CN) of the remote server's certificate for an HTTPS request. If NULL, the host name in the given URL will be used for SNI.

Usage Notes

- The URL passed as an argument to this function is not examined for illegal characters, such as spaces, according to URL specification RFC 2396. You should escape those characters with the `UTL_URL` package to return illegal and reserved characters. URLs should consist of US-ASCII characters only. See UTL_URL for a list of legal characters in URLs. Note that URLs should consist of US-ASCII characters only. The use of non-US-ASCII characters in a URL is generally unsafe.
- `BEGIN_REQUEST` can send a URL whose length is up to 32767 bytes. However, different Web servers impose different limits on the length of the URL they can accept. This limit is often about 4000 bytes. If this limit is exceeded, the outcome will depend on the Web server. For example, a Web server might simply drop the HTTP connection without returning a response of any kind. If this happens, a subsequent invocation of the GET_RESPONSE Function will raise the `PROTOCOL_ERROR` exception.

  A URL will be long when its `QUERY_STRING` (that is, the information that follows the question mark (?)) is long. In general, it is better to send this parameterization in the body of the request using the `POST` method.

```
req := UTL_HTTP.BEGIN_REQUEST (url=>the_url, method=>'POST');
UTL_HTTP.SET_HEADER (r      =>  req,
                     name   =>  'Content-Type',
                     value  =>  'application/x-www-form-urlencoded');
UTL_HTTP.SET_HEADER (r      =>   req,
                     name   =>   'Content-Length',
                     value  =>'  <length of data posted in bytes>');
UTL_HTTP.WRITE_TEXT (r      =>   req,
                     data   =>   'p1 = value1&p2=value2...');
resp := UTL_HTTP.GET_RESPONSE
                     (r     =>   req);
...
```

  The programmer must determine whether a particular Web server may, or may not, accept data provided in this way.
- An Oracle wallet must be set before accessing Web servers over HTTPS. See the SET_WALLET Procedure procedure on how to set up an Oracle wallet. To use SSL client authentication, the client certificate should be stored in the wallet and the caller must have the `use-client-certificates` privilege on the wallet. See "Managing Fine-grained Access to External Network Services" in the *Oracle Database Security Guide* to grant the privilege.
- To connect to the remote Web server directly, or indirectly through a HTTP proxy, the `UTL_HTTP` must have the `connect` ACL privilege to the remote Web server host or the proxy host respectively.

286.9.3 CLEAR_COOKIES Procedure

This procedure clears all the cookies maintained either in a request context or in the `UTL_HTTP` package's session state.

> **See Also:**
>
> HTTP Cookies and HTTP Cookies Subprograms

Syntax

```
UTL_HTTP.CLEAR_COOKIES (
   request_context  IN  request_context_key DEFAULT NULL);
```

Parameters

Table 286-19 CLEAR_COOKIES Procedure Parameters

Parameter	Description
request_context	Request context to clear the cookies. If NULL, the cookies maintained in the UTL_HTTP package's session state will be cleared instead.

286.9.4 CLOSE_PERSISTENT_CONN Procedure

This procedure closes an HTTP persistent connection maintained by the UTL_HTTP package in the current database session.

> **See Also:**
>
> HTTP Persistent Connections and HTTP Persistent Connections Subprograms

Syntax

```
UTL_HTTP.CLOSE_PERSISTENT_CONN (
   conn  IN connection);
```

Parameters

Table 286-20 CLOSE_PERSISTENT_CONN Procedure Parameters

Parameter	Description
conn	The HTTP persistent connection to close

286.9.5 CLOSE_PERSISTENT_CONNS Procedure

This procedure closes a group of HTTP persistent connections maintained by the UTL_HTTP package in the current database session. This procedure uses a pattern-match approach to decide which persistent connections to close.

To close a group of HTTP persistent connection that share a common property (for example, all connections to a particular host, or all SSL connections), set the particular parameters and leave the rest of the parameters NULL. If a particular parameter is set to NULL when this procedure is called, that parameter will not be used to decide which connections to close.

For example, the following call to the procedure closes all persistent connections to foobar:

```
UTL_HTTP.CLOSE_PERSISTENT_CONNS(host => 'foobar');
```

And the following call to the procedure closes all persistent connections through the foobar at TCP/IP port 80:

```
UTL_HTTP.CLOSE_PERSISTENT_CONNS(proxy_host => 'foobar',
                                proxy_port => 80);
```

And the following call to the procedure closes all persistent connections:

```
UTL_HTTP.CLOSE_PERSISTENT_CONNS;
```

> **See Also:**
>
> HTTP Persistent Connections and HTTP Persistent Connections Subprograms

Syntax

```
UTL_HTTP.CLOSE_PERSISTENT_CONNS (
   host        IN VARCHAR2 DEFAULT NULL,
   port        IN PLS_INTEGER DEFAULT NULL,
   proxy_host  IN VARCHAR2 DEFAULT NULL,
   proxy_port  IN PLS_INTEGER DEFAULT NULL,
   ssl         IN BOOLEAN DEFAULT NULL);
```

Parameters

Table 286-21 CLOSE_PERSISTENT_CONNS Procedure Parameters

Parameter	Description
`host`	The host for which persistent connections are to be closed
`port`	The port number for which persistent connections are to be closed
`proxy_host`	The proxy host for which persistent connections are to be closed
`proxy_port`	The proxy port for which persistent connections are to be closed
`ssl`	Close persistent SSL connection

Usage Notes

Connections to the same Web server at different TCP/IP ports are counted individually. The host names of the Web servers are identified as specified in the URL of the original HTTP requests. Therefore, fully qualified host names with domain names will be counted differently from the host names without domain names.

Note that the use of a `NULL` value in a parameter when this procedure is called means that the caller does not care about its value when the package decides which persistent connection to close. If you want a `NULL` value in a parameter to match only a `NULL` value of the parameter of a persistent connection (which is when you want to close a specific persistent connection), you should use the `CLOSE_PERSISTENT_CONN` procedure that closes a specific persistent connection.

286.9.6 CREATE_REQUEST_CONTEXT Function

This function creates a request context. A request context is a context that holds a wallet and a cookie for private use in making a HTTP request. This allows the HTTP request to use a wallet

and a cookie table that will not be shared with other applications making HTTP requests in the same database session.

> **See Also:**
>
> Request Context and HTTP Request Contexts Subprograms

Syntax

```
UTL_HTTP.CREATE_REQUEST_CONTEXT (
 wallet_path          IN VARCHAR2 DEFAULT NULL,
 wallet_password      IN VARCHAR2 DEFAULT NULL,
 enable_cookies       IN BOOLEAN  DEFAULT TRUE,
 max_cookies          IN PLS_INTEGER DEFAULT 300,
 max_cookies_per_site IN PLS_INTEGER DEFAULT 20)
RETURN request_context_key;
```

Parameters

Table 286-22 CREATE_REQUEST_CONTEXT Function Parameters

Parameter	Description
`wallet_path`	Directory path that contains the Oracle wallet. The format is `file:`*`directory-path`*
`wallet_password`	The password needed to open the wallet. If the wallet is auto-login enabled, the password may be omitted and should be set to `NULL`. See the *Oracle Database Enterprise User Security Administrator's Guide* for detailed information about wallets.
`enable_cookies`	Sets whether HTTP requests using this request context should support HTTP cookies or not: `TRUE` to enable the support, `FALSE` to disable it.
`max_cookies`	Sets the maximum total number of cookies that will be maintained in this request context
`max_cookies_per_site`	Sets the maximum number of cookies per each Web site that will be maintained in this request context

Return Values

The request context created.

Examples

```
DECLARE
  request_context  UTL_HTTP.REQUEST_CONTEXT_KEY;
  req              utl_http.req;
BEGIN
  request_context := UTL_HTTP.CREATE_REQUEST_CONTEXT(
       wallet_path          => 'file:/oracle/wallets/test_wallets',
       wallet_password      => NULL,
       enable_cookies       => TRUE,
       max_cookies          => 300,
       max_cookies_per_site => 20);
  req := UTL_HTTP.BEGIN_REQUEST(
       url                  => 'http://www.example.com/',
       request_context      => request_context);
END;
```

286.9.7 DESTROY_REQUEST_CONTEXT Procedure

This procedure destroys a request context in `UTL_HTTP`. A request context cannot be destroyed when it is in use by a HTTP request or response.

> **See Also:**
>
> Request Context and HTTP Request Contexts Subprograms

Syntax

```
UTL_HTTP.DESTROY_REQUEST_CONTEXT (
   request_context    request_context_key);
```

Parameters

Table 286-23 DESTROY_REQUEST_CONTEXT Procedure Parameters

Parameter	Description
request_context	Request context to destroy

Examples

```
DECLARE
  request_context  UTL_HTTP.REQUEST_CONTEXT_KEY;
BEGIN
  request_context := UTL_HTTP.CREATE_REQUEST_CONTEXT(…);
  …
  UTL_HTTP.DESTROY_REQUEST_CONTEXT(request_context);
END;
```

286.9.8 END_REQUEST Procedure

This procedure ends the HTTP request. To terminate the HTTP request without completing the request and waiting for the response, the program can call this procedure. Otherwise, the program should go through the normal sequence of beginning a request, getting the response, and closing the response. The network connection will always be closed and will not be reused.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.END_REQUEST (
   r  IN OUT NOCOPY req);
```

Parameters

Table 286-24 END_REQUEST Procedure Parameters

Parameter	Description
r	The HTTP request

286.9.9 END_RESPONSE Procedure

This procedure ends the HTTP response. It completes the HTTP request and response. Unless HTTP 1.1 persistent connection is used in this request, the network connection is also closed.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.END_RESPONSE (
   r  IN OUT NOCOPY resp);
```

Parameters

Table 286-25 END_RESPONSE Procedure Parameters

Parameter	Description
r	The HTTP response

286.9.10 GET_AUTHENTICATION Procedure

This procedure retrieves the HTTP authentication information needed for the request to be accepted by the Web server as indicated in the HTTP response header.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.GET_AUTHENTICATION(
   r          IN OUT NOCOPY resp,
   scheme     OUT VARCHAR2,
   realm      OUT VARCHAR2,
   for_proxy  IN BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 286-26 GET_AUTHENTICATION Procedure Parameters

Parameter	Description
r	The HTTP response
scheme	The scheme for the required HTTP authentication
realm	The realm for the required HTTP authentication
for_proxy	Returns the HTTP authentication information required for the access to the HTTP proxy server instead of the Web server? Default is FALSE

Usage Notes

When a Web client is unaware that a document is protected, at least two HTTP requests are required for the document to be retrieved. In the first HTTP request, the Web client makes the request without supplying required authentication information; so the request is denied. The Web client can determine the authentication information required for the request to be authorized by calling GET_AUTHENTICATION. The Web client makes the second request and supplies the required authentication information with SET_AUTHORIZATION. If the authentication information can be verified by the Web server, the request will succeed and the requested document is returned. Before making the request, if the Web client knows that authentication information is required, it can supply the required authentication information in the first request, thus saving an extra request.

286.9.11 GET_BODY_CHARSET Procedure

This procedure retrieves the default character set of the body of all future HTTP requests.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_BODY_CHARSET (
   charset  OUT NOCOPY VARCHAR2);
```

Parameters

Table 286-27 GET_BODY_CHARSET Procedure Parameters

Parameter	Description
charset	The default character set of the body of all future HTTP requests

286.9.12 GET_COOKIE_COUNT Function

This function returns the number of cookies maintained either in a request context or in the `UTL_HTTP` package's session state.

> **See Also:**
>
> HTTP Cookies and HTTP Cookies Subprograms

Syntax

```
UTL_HTTP.GET_COOKIE_COUNT (
   request_context  IN  request_context_key DEFAULT NULL)
 RETURN PLS_INTEGER;
```

Parameters

Table 286-28 GET_COOKIE_COUNT Function Parameters

Parameter	Description
`request_context`	Request context to return the cookie count for. If `NULL`, the cookie count maintained in the `UTL_HTTP` package's session state will be returned instead.

286.9.13 GET_COOKIE_SUPPORT Procedure

This procedure retrieves the current cookie support settings.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_COOKIE_SUPPORT (
   enable                OUT BOOLEAN,
   max_cookies           OUT PLS_INTEGER,
   max_cookies_per_site  OUT PLS_INTEGER);
```

Parameters

Table 286-29 GET_COOKIE_SUPPORT Procedure Parameters

Parameter	Description
`enable`	Indicates whether future HTTP requests should support HTTP cookies (`TRUE`) or not (`FALSE`)
`max_cookies`	Indicates the maximum total number of cookies maintained in the current session

Table 286-29 (Cont.) GET_COOKIE_SUPPORT Procedure Parameters

Parameter	Description
max_cookies_per_site	Indicates the maximum number of cookies maintained in the current session for each Web site

286.9.14 GET_COOKIES Function

This function returns all the cookies maintained either in a request context or in the UTL_HTTP package's session state.

> **See Also:**
>
> HTTP Cookies and HTTP Cookies Subprograms

Syntax

```
UTL_HTTP.GET_COOKIES (
   cookies          IN  OUT NOCOPY cookie_table,
   request_context  IN             request_context_key DEFAULT NULL);
```

Parameters

Table 286-30 GET_COOKIES Function Parameters

Parameter	Description
cookies	The cookies returned
request_context	Request context to return the cookies for. If NULL, the cookies maintained in the UTL_HTTP package's session state will be returned instead.

286.9.15 GET_DETAILED_EXCP_SUPPORT Procedure

This procedure checks if the UTL_HTTP package will raise a detailed exception or not.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_DETAILED_EXCP_SUPPORT (
   enable  OUT BOOLEAN);
```

Parameters

Table 286-31 GET_DETAILED_EXCP_SUPPORT Procedure Parameters

Parameter	Description
enable	TRUE if UTL_HTTP raises a detailed exception; otherwise FALSE

286.9.16 GET_DETAILED_SQLCODE Function

This function retrieves the detailed SQLCODE of the last exception raised.

> **See Also:**
>
> Error Conditions andError Conditions Subprograms

Syntax

```
UTL_HTTP.GET_DETAILED_SQLCODE
RETURN PLS_INTEGER;
```

286.9.17 GET_DETAILED_SQLERRM Function

This function retrieves the detailed SQLERRM of the last exception raised.

> **See Also:**
>
> Error Conditions and Error Conditions Subprograms

Syntax

```
UTL_HTTP.GET_DETAILED_SQLERRM
RETURN VARCHAR2;
```

286.9.18 GET_FOLLOW_REDIRECT Procedure

This procedure retrieves the follow-redirect setting in the current session

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_FOLLOW_REDIRECT (
   max_redirects  OUT PLS_INTEGER);
```

Parameters

Table 286-32 GET_FOLLOW_REDIRECT Procedure Parameters

Parameter	Description
max_redirects	The maximum number of redirections for all future HTTP requests

286.9.19 GET_HEADER Procedure

This procedure returns the nth HTTP response header name and value returned in the response.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.GET_HEADER (
   r      IN OUT NOCOPY resp,
   n      IN PLS_INTEGER,
   name   OUT NOCOPY VARCHAR2,
   value  OUT NOCOPY VARCHAR2);
```

Parameters

Table 286-33 GET_HEADER Procedure Parameters

Parameter	Description
r	The HTTP response
n	The nth header to return
name	The name of the HTTP response header
value	The value of the HTTP response header

Usage Notes

If the response body returned by the remote Web server is encoded in chunked transfer encoding format, the trailer headers that are returned at the end of the response body will be added to the response, and the response header count will be updated. You can retrieve the additional headers after the end of the response body is reached and before you end the response.

286.9.20 GET_HEADER_BY_NAME Procedure

This procedure returns the HTTP response header value returned in the response given the name of the header.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.GET_HEADER_BY_NAME(
   r      IN OUT NOCOPY resp,
   name   IN VARCHAR2,
   value  OUT NOCOPY VARCHAR2,
   n      IN PLS_INTEGER DEFAULT 1);
```

Parameters

Table 286-34 GET_HEADER_BY_NAME Procedure Parameters

Parameter	Description
r	The HTTP response
name	The name of the HTTP response header for which the value is to return
value	The value of the HTTP response header
n	The n^{th} occurrence of an HTTP response header by the specified name to return. The default is 1.

Usage Notes

If the response body returned by the remote Web server is encoded in chunked transfer encoding format, the trailer headers that are returned at the end of the response body will be added to the response, and the response header count will be updated. You can retrieve the additional headers after the end of the response body is reached and before you end the response.

286.9.21 GET_HEADER_COUNT Function

This function returns the number of HTTP response headers returned in the response.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.GET_HEADER_COUNT (
   r  IN OUT NOCOPY resp)
RETURN PLS_INTEGER;
```

Parameters

Table 286-35 GET_HEADER_COUNT Function Parameters

Parameter	Description
r	The HTTP response

Usage Notes

If the response body returned by the remote Web server is encoded in chunked transfer encoding format, the trailer headers that are returned at the end of the response body will be added to the response, and the response header count will be updated. You can retrieve the additional headers after the end of the response body is reached and before you end the response.

286.9.22 GET_PERSISTENT_CONN_COUNT Function

This function returns the number of network connections currently kept persistent by the UTL_HTTP package to the Web servers.

> **See Also:**
>
> HTTP Persistent Connections and HTTP Persistent Connections Subprograms

Syntax

```
UTL_HTTP.GET_PERSISTENT_CONN_COUNT
RETURN PLS_INTEGER;
```

Usage Notes

Connections to the same Web server at different TCP/IP ports are counted individually. The host names of the Web servers are identified as specified in the URL of the original HTTP requests. Therefore, fully qualified host names with domain names will be counted differently from the host names without domain names.

286.9.23 GET_PERSISTENT_CONN_SUPPORT Procedure

This procedure checks if the persistent connection support is enabled, and gets the maximum number of persistent connections in the current session.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_PERSISTENT_CONN_SUPPORT (
   enable     OUT BOOLEAN,
   max_conns  OUT PLS_INTEGER);
```

Parameters

Table 286-36 GET_PERSISTENT_CONN_SUPPORT Procedure Parameters

Parameter	Description
`enable`	`TRUE` if persistent connection support is enabled; otherwise `FALSE`
`max_conns`	the maximum number of persistent connections maintained in the current session

286.9.24 GET_PERSISTENT_CONNS Procedure

This procedure returns all the network connections currently kept persistent by the `UTL_HTTP` package to the Web servers.

> **See Also:**
>
> HTTP Persistent Connections and HTTP Persistent Connections Subprograms

Syntax

```
UTL_HTTP.get_persistent_conns (
   connections  IN OUT NOCOPY connection_table);
```

Parameters

Table 286-37 GET_PERSISTENT_CONNS Procedure Parameters

Parameter	Description
`connections`	The network connections kept persistent

Usage Notes

Connections to the same Web server at different TCP/IP ports are counted individually. The host names of the Web servers are identified as specified in the URL of the original HTTP requests. Therefore, fully qualified host names with domain names will be counted differently from the host names without domain names.

286.9.25 GET_PROXY Procedure

This procedure retrieves the current proxy settings.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_PROXY (
   proxy             OUT NOCOPY VARCHAR2,
   no_proxy_domains  OUT NOCOPY VARCHAR2);
```

Parameters

Table 286-38 GET_PROXY Procedure Parameters

Parameter	Description
`proxy`	The proxy (host and an optional port number) currently used by the `UTL_HTTP` package
`no_proxy_domains`	The list of hosts and domains for which no proxy is used for all requests

286.9.26 GET_RESPONSE Function

This function reads the HTTP response.

When the function returns, the status line and the HTTP response headers have been read and processed. The status code, reason phrase, and the HTTP protocol version are stored in the response record. This function completes the HTTP headers section.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.GET_RESPONSE (
   r                      IN OUT NOCOPY req,
   return_info_response   IN BOOLEAN DEFAULT FALSE)
RETURN resp;
```

Parameters

Table 286-39 GET_RESPONSE Function Parameters

Parameter	Description
`r`	The HTTP response

Table 286-39 (Cont.) GET_RESPONSE Function Parameters

Parameter	Description
return_info_response	Return 100 informational response or not. • TRUE means get_response should return 100 informational response when it is received from the HTTP server. The request will not be ended if a 100 response is returned. • FALSE means the API should ignore any 100 informational response received from the HTTP server and should return the following non-100 response instead. The default is FALSE.

Exceptions

- When detailed-exception is disabled:

  ORA-29273 REQUEST_FAILED - the request fails to execute. Use the GET_DETAILED_EXCP_SUPPORT Procedure and the GET_DETAILED_SQLERRM Function to get the detailed error message.

- When detailed-exception is enabled:

  ORA-29261 BAD_ARGUMENT - some arguments passed are not valid

- When response error check is enabled:

  ORA-29268 HTTP_CLIENT_ERROR - the response code is in 400 range

  ORA-29269 HTTP_SERVER_ERROR - the response code is in 500 range

Usage Notes

- The request will be ended when this functions returns regardless of whether an exception is raised or not. There is no need to invoke the END_REQUEST Procedure.
- If URL redirection occurs, the URL and method fields in the req record will be updated to the last redirected URL and the method used to access the URL.

Examples

In certain situations (initiated by the HTTP client or not), the HTTP server may return a 1xx informational response. The user who does not expect such a response may indicate to GET_RESPONSE to ignore the response and proceed to receive the regular response. In the case when the user expects such a response, the user can indicate to GET_RESPONSE to return the response.

For example, when a user is issuing a HTTP POST request with a large request body, the user may want to check with the HTTP server to ensure that the server will accept the request before sending the data. To do so, the user will send the additional EXPECT: 100-CONTINUE request header, and check for 100 CONTINUE response from the server before proceeding to send the request body. Then, the user will get the regular HTTP response.

The following code example illustrates this:

```
DECLARE
 data  VARCHAR2(1024) := '...';
 req   utl_http.req;
 resp  utl_http.resp;
BEGIN
```

```
 req := utl_http.begin_request('http://www.acme.com/receiver', 'POST');
 utl_http.set_header(req, 'Content-Length', length(data));
 -- Ask HTTP server to return "100 Continue" response
 utl_http.set_header(req, 'Expect', '100-continue');
 resp := utl_http.get_response(req, TRUE);

 -- Check for and dispose "100 Continue" response
 IF (resp.status_code <> 100) THEN
   utl_http.end_response(resp);
   raise_application_error(20000, 'Request rejected');
 END IF;
 utl_http.end_response(resp);

 -- Now, send the request body
 utl_http.write_text(req, data);

 -- Get the regular response
 resp := utl_http.get_response(req);
 utl_http.read_text(resp, data);

 utl_http.end_response(resp);

END;
```

286.9.27 GET_RESPONSE_ERROR_CHECK Procedure

This procedure checks if the response error check is set or not.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_RESPONSE_ERROR_CHECK (
   enable  OUT BOOLEAN);
```

Parameters

Table 286-40 GET_RESPONSE_ERROR_CHECK Procedure Parameters

Parameter	Description
enable	TRUE if the response error check is set; otherwise FALSE

286.9.28 GET_TRANSFER_TIMEOUT Procedure

This procedure retrieves the default timeout value for all future HTTP requests.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.GET_TRANSFER_TIMEOUT (
   timeout  OUT PLS_INTEGER);
```

Parameters

Table 286-41 GET_TRANSFER_TIMEOUT Procedure Parameters

Parameter	Description
timeout	The network transfer timeout value in seconds

286.9.29 READ_LINE Procedure

This procedure reads the HTTP response body in text form until the end of line is reached and returns the output in the caller-supplied buffer.

The end of line is as defined in the function `read_line` of UTL_TCP. The `end_of_body` exception will be raised if the end of the HTTP response body is reached. Text data is automatically converted from the response body character set to the database character set.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.READ_LINE(
   r            IN OUT NOCOPY resp,
   data         OUT NOCOPY  VARCHAR2 CHARACTER SET ANY_CS,
   remove_crlf  IN  BOOLEAN DEFAULT FALSE);
```

Parameters

Table 286-42 READ_LINE Procedure Parameters

Parameter	Description
r	The HTTP response
data	The HTTP response body in text form
remove_crlf	Removes the newline characters if set to `TRUE`

Usage Notes

The `UTL_HTTP` package supports HTTP 1.1 chunked transfer-encoding. When the response body is returned in chunked transfer-encoding format as indicated in the response header, the package automatically decodes the chunks and returns the response body in de-chunked format.

If transfer timeout is set in the request of this response, `read_line` waits for each data packet to be ready to read until timeout occurs. If it occurs, this procedure stops reading and returns

all the data read successfully. If no data is read successfully, the `transfer_timeout` exception is raised. The exception can be handled and the read operation can be retried later.

If a partial multibyte character is found at the end of the response body, `read_line` stops reading and returns all the complete multibyte characters read successfully. If no complete character is read successfully, the `partial_multibyte_char` exception is raised. The exception can be handled and the bytes of that partial multibyte character can be read as binary by the `read_raw` procedure. If a partial multibyte character is seen in the middle of the response body because the remaining bytes of the character have not arrived and read timeout occurs, the `transfer_timeout` exception is raised instead. The exception can be handled and the read operation can be retried later.

When the `Content-Type` response header specifies the character set of the response body and the character set is unknown or unsupported by Oracle, the "ORA-01482: unsupported character set" exception is raised if you try to read the response body as text. You can either read the response body as binary using the `READ_RAW` procedure, or set the character set of the response body explicitly using the `SET_BODY_CHARSET` procedure and read the response body as text again.

286.9.30 READ_RAW Procedure

This procedure reads the HTTP response body in binary form and returns the output in the caller-supplied buffer.

The `end_of_body` exception is raised if the end of the HTTP response body is reached.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.READ_RAW(
   r     IN OUT NOCOPY resp,
   data  OUT NOCOPY RAW,
   len   IN PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 286-43 READ_RAW Procedure Parameters

Parameter	Description
`r`	The HTTP response
`data`	The HTTP response body in binary form
`len`	The number of bytes of data to read. If `len` is `NULL`, this procedure will read as much input as possible to fill the buffer allocated in `data`. The actual amount of data returned may be less than that specified if not much data is available before the end of the HTTP response body is reached or the `transfer_timeout` amount of time has elapsed. The default is `NULL`

Usage Notes

The `UTL_HTTP` package supports HTTP 1.1 chunked transfer-encoding. When the response body is returned in chunked transfer-encoding format as indicated in the response header, the package automatically decodes the chunks and returns the response body in de-chunked format.

If transfer timeout is set in the request of this response, `read_raw` waits for each data packet to be ready to read until timeout occurs. If it occurs, `read_raw` stops reading and returns all the data read successfully. If no data is read successfully, the `transfer_timeout` exception is raised. The exception can be handled and the read operation can be retried later.

286.9.31 READ_TEXT Procedure

This procedure reads the HTTP response body in text form and returns the output in the caller-supplied buffer.

The `end_of_body` exception is raised if the end of the HTTP response body is reached. Text data is automatically converted from the response body character set to the database character set.

> **See Also:**
>
> HTTP Responses and HTTP Responses Subprograms

Syntax

```
UTL_HTTP.READ_TEXT(
   r     IN OUT NOCOPY resp,
   data  OUT NOCOPY VARCHAR2 CHARACTER SET ANY_CS,
   len   IN PLS_INTEGER DEFAULT NULL);
```

Parameters

Table 286-44 READ_TEXT Procedure Parameters

Parameter	Description
`r`	The HTTP response
`data`	The HTTP response body in text form
`len`	The maximum number of characters of data to read. If `len` is NULL, this procedure will read as much input as possible to fill the buffer allocated in `data`. The actual amount of data returned may be less than that specified if little data is available before the end of the HTTP response body is reached or the `transfer_timeout` amount of time has elapsed. The default is NULL.

Usage Notes

The `UTL_HTTP` package supports HTTP 1.1 chunked transfer-encoding. When the response body is returned in chunked transfer-encoding format as indicated in the response header, the package automatically decodes the chunks and returns the response body in de-chunked format.

If transfer timeout is set in the request of this response, `read_text` waits for each data packet to be ready to read until timeout occurs. If it occurs, this procedure stops reading and returns all the data read successfully. If no data is read successfully, the `transfer_timeout` exception is raised. The exception can be handled and the read operation can be retried later.

If a partial multibyte character is found at the end of the response body, `read_text` stops reading and returns all the complete multibyte characters read successfully. If no complete character is read successfully, the `partial_multibyte_char` exception is raised. The exception can be handled and the bytes of that partial multibyte character can be read as binary by the `read_raw` procedure. If a partial multibyte character is seen in the middle of the response body because the remaining bytes of the character have not arrived and read timeout occurs, the `transfer_timeout` exception is raised instead. The exception can be handled and the read operation can be retried later.

When the `Content-Type` response header specifies the character set of the response body and the character set is unknown or unsupported by Oracle, the "`ORA-01482: unsupported character set`" exception is raised if you try to read the response body as text. You can either read the response body as binary using the `READ_RAW` procedure, or set the character set of the response body explicitly using the `SET_BODY_CHARSET` procedure and read the response body as text again.

286.9.32 REQUEST Function

This function returns up to the first 2000 bytes of data retrieved from the given URL.

This function can be used directly in SQL queries. The URL may contain the username and password needed to authenticate the request to the server. The format is

```
scheme://[user[:password]@]host[:port]/[...]
```

You can define a username/password for the proxy to be specified in the proxy string. The format is

```
[http://][user[:password]@]host[:port][/]
```

> **See Also:**
>
> Simple HTTP Fetches and Simple HTTP Fetches in a Single Call Subprograms

Syntax

```
UTL_HTTP.REQUEST (
   url              IN VARCHAR2,
   proxy            IN VARCHAR2 DEFAULT NULL,
   wallet_path      IN VARCHAR2 DEFAULT NULL,
   wallet_password  IN VARCHAR2 DEFAULT NULL,
   https_host       IN VARCHAR2 DEFAULT NULL)
RETURN VARCHAR2;
```

Pragmas

```
pragma restrict_references (request, wnds, rnds, wnps, rnps);
```

Parameters

Table 286-45 REQUEST Function Parameters

Parameter	Description
`url`	Uniform resource locator
`proxy`	(Optional) Specifies a proxy server to use when making the HTTP request. See `SET_PROXY` for the full format of the proxy setting.
`wallet_path`	(Optional) Specifies a client-side wallet. The client-side wallet contains the list of trusted certificate authorities required for HTTPS request. The format of `wallet_path` on a PC is, for example, file:`c:\WINNT\Profiles\`*`username`*`\WALLETS`, and in Unix is, for example, `file:/home/`*`username`*`/wallets` When the `UTL_HTTP` package is executed in the Oracle database server, the wallet is accessed from the database server. Therefore, the wallet path must be accessible from the database server. See `SET_WALLET` for a description on how to set up an Oracle wallet. Non-HTTPS requests do not require an Oracle wallet.
`wallet_password`	(Optional) Specifies the password required to open the wallet
`https_host`	A string representing the host name. If the string does not begin with a wildcard, the string will be used as the host name for server name indication (SNI). If the string begins with a wildcard, the string will be used to match against the common name (CN) of the remote server's certificate for an HTTPS request. If NULL, the host name in the given URL will be used for SNI.

Return Values

The return type is a string of length 2000 or less, which contains up to the first 2000 bytes of the HTML result returned from the HTTP request to the argument URL.

Exceptions

```
INIT_FAILED
REQUEST_FAILED
```

Usage Notes

The URL passed as an argument to this function is not examined for illegal characters, for example, spaces, according to URL specification RFC 2396. The caller should escape those characters with the UTL_URL package. See the comments of the package for the list of legal characters in URLs. Note that URLs should consist of US-ASCII characters only. The use of non-US-ASCII characters in a URL is generally unsafe.

Please see the documentation of the function `SET_WALLET` on the use of an Oracle wallet, which is required for accessing HTTPS Web servers.

Unless response error check is turned on, this function does not raise an exception when a 4xx or 5xx response is received from the Web server. Instead, it returns the formatted error message from the Web server:

```
<HTML>
<HEAD>
<TITLE>Error Message</TITLE>
```

```
</HEAD>
<BODY>
<H1>Fatal Error 500</H1>
Can't Access Document:  http://home.nothing.comm.
<P>
<B>Reason:</B> Can't locate remote host:  home.nothing.comm.
<P>
<P><HR>
<ADDRESS><A HREF="http://www.w3.org">
CERN-HTTPD3.0A</A></ADDRESS>
</BODY>
</HTML>
```

Examples

```
SQL> SELECT UTL_HTTP.REQUEST('http://www.my-company.com/') FROM DUAL;
UTL_HTTP.REQUEST('HTTP://WWW.MY-COMPANY.COM/')
<html>
<head><title>My Company Home Page</title>
<!--changed Jan. 16, 19
1 row selected.
```

If you are behind a firewall, include the `proxy` parameter. For example, from within the Oracle firewall, where there might be a proxy server named `www-proxy.my-company.com`:

```
SQLPLUS> SELECT
UTL_HTTP.REQUEST('http://www.my-company.com', 'www-proxy.us.my-company.com') FROM DUAL;
```

286.9.33 REQUEST_PIECES Function

This function returns a PL/SQL table of 2000-byte pieces of the data retrieved from the given URL.

You can define a username/password for the proxy to be specified in the proxy string. The format is

```
[http://][user[:password]@]host[:port][/]
```

> **See Also:**
>
> Simple HTTP Fetches and Simple HTTP Fetches in a Single Call Subprograms

Syntax

```
TYPE html_pieces IS TABLE OF VARCHAR2(2000) INDEX BY BINARY_INTEGER;

UTL_HTTP.REQUEST_PIECES (
   url             IN VARCHAR2,
   max_pieces      IN NATURAL DEFAULT 32767,
   proxy           IN VARCHAR2 DEFAULT NULL,
   wallet_path     IN VARCHAR2 DEFAULT NULL,
   wallet_password IN VARCHAR2 DEFAULT NULL,
   https_host      IN VARCHAR2 DEFAULT NULL)
RETURN html_pieces;
```

Pragmas

```
PRAGMA RESTRICT_REFERENCES (request_pieces, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 286-46 REQUEST_PIECES Function Parameters

Parameter	Description
`url`	Uniform resource locator
`max_pieces`	(Optional) The maximum number of pieces (each 2000 characters in length, except for the last, which may be shorter), that `REQUEST_PIECES` should return. If provided, then that argument should be a positive integer.
`proxy`	(Optional) Specifies a proxy server to use when making the HTTP request. See `SET_PROXY` for the full format of the proxy setting.
`wallet_path`	(Optional) Specifies a client-side wallet. The client-side wallet contains the list of trusted certificate authorities required for HTTPS request. The format of wallet_path on a PC is, for example, `file:c:\WINNT\Profiles\`*`username`*`\WALLETS`, and in Unix is, for example, `file:/home/`*`username`*`/wallets`. When the `UTL_HTTP` package is executed in the Oracle database server, the wallet is accessed from the database server. Therefore, the wallet path must be accessible from the database server. See `SET_WALLET` for the description on how to set up an Oracle wallet. Non-HTTPS requests do not require an Oracle wallet.
`wallet_password`	(Optional) Specifies the password required to open the wallet
`https_host`	A string representing the host name. If the string does not begin with a wildcard, the string will be used as the host name for server name indication (SNI). If the string begins with a wildcard, the string will be used to match against the common name (CN) of the remote server's certificate for an HTTPS request. If NULL, the host name in the given URL will be used for SNI.

Return Values

`REQUEST_PIECES` returns a PL/SQL table of type `UTL_HTTP.HTML_PIECES`. Each element of that PL/SQL table is a string of maximum length 2000. The elements of the PL/SQL table returned by `REQUEST_PIECES` are successive pieces of the data obtained from the HTTP request to that URL.

Exceptions

```
INIT_FAILED
REQUEST_FAILED
```

Usage Notes

The URL passed as an argument to this function will not be examined for illegal characters, for example, spaces, according to URL specification RFC 2396. The caller should escape those characters with the `UTL_URL` package. See the comments of the package for the list of legal characters in URLs. Note that URLs should consist of US-ASCII characters only. The use of non-US-ASCII characters in a URL is generally unsafe.

Each entry of the PL/SQL table (the "pieces") returned by this function may not be filled to their fullest capacity. The function may start filling the data in the next piece before the previous "piece" is totally full.

Please see the documentation of the function `SET_WALLET` on the use of an Oracle wallet, which is required for accessing HTTPS Web servers.

Unless response error check is turned on, this function does not raise an exception when a 4xx or 5xx response is received from the Web server. Instead, it returns the formatted error message from the Web server:

```
<HTML>
<HEAD>
<TITLE>Error Message</TITLE>
</HEAD>
<BODY>
<H1>Fatal Error 500</H1>
Can't Access Document:  http://home.nothing.comm.
<P>
<B>Reason:</B> Can't locate remote host:  home.nothing.comm.
<P>
<P><HR>
<ADDRESS><A HREF="http://www.w3.org">
CERN-HTTPD3.0A</A></ADDRESS>
</BODY>
</HTML>
```

Examples

```
SET SERVEROUTPUT ON

DECLARE
   x   UTL_HTTP.HTML_PIECES;
   len PLS_INTEGER;
BEGIN
   x := UTL_HTTP.REQUEST_PIECES('http://www.oracle.com/', 100);
   DBMS_OUTPUT.PUT_LINE(x.count || ' pieces were retrieved.');
   DBMS_OUTPUT.PUT_LINE('with total length ');
   IF x.count < 1 THEN
      DBMS_OUTPUT.PUT_LINE('0');
  ELSE
   len := 0;
   FOR i in 1..x.count LOOP
      len := len + length(x(i));
   END LOOP;
   DBMS_OUTPUT.PUT_LINE(len);
  END IF;
END;
/
-- Output
Statement processed.
4 pieces were retrieved.
with total length
7687
```

286.9.34 SET_AUTHENTICATION Procedure

This procedure sets HTTP authentication information in the HTTP request header. The Web server needs this information to authorize the request.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

UTL_HTTP Authentication Scheme

The following is the challenge and response work flow:

1. Client sends a HTTP request to the server.
2. The server responds to the client with a 401 (Unauthorized) response status. It also provides information on how to authorize using a WWW-Authenticate response header that contains at least one authentication scheme challenge.

   > **Note:**
   >
   > When server responds with multiple challenges, it will be in the order of preference, starting with the most preferred algorithm followed by the less preferred algorithm.

3. Client authenticates itself with the server by issuing the request that includes an Authorization request header along with the credentials build by using the values from challenge headers, such as algorithm, realm, and nonce.

   > **Note:**
   >
   > UTL_HTTP supports digest SHA-256 algorithm.

Digest Authentication

An authentication scheme is set in the HTTP request header that is authorized by the Web server using the UTL_HTTP.SET_AUTHENTICATION call. Digest is one of the authentication schemes that UTL_HTTP supports. The other authentication schemes are basic, AWS,AWS4, BMC, and AZURE.

The following is the UTL_HTTP request and response flow using the digest authentication scheme:

1. UTL_HTTP.BEGIN_REQUEST- This function sends an HTTP request to the Web server.
2. UTL_HTTP.SET_AUTHENTICATION- This function sets authentication information in the HTTP request header
3. ULT_HTTP.GET_RESPONSE- This function gets the response from the Web server
   a. Server replies “401 Unauthorized” along with the challenge headers.

b. Build digests credentials with values from the challenge headers using the SHA 256 algorithm.

c. Resend request includes an authorization request header with the digest credentials.

d. Server processes request and replies (typically 200-OK).

Syntax

```
UTL_HTTP.SET_AUTHENTICATION(
   r         IN OUT NOCOPY req,
   username  IN VARCHAR2,
   password  IN VARCHAR2,
   scheme    IN VARCHAR2 DEFAULT 'Basic',
   for_proxy IN BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 286-47 SET_AUTHENTICATION Procedure Parameters

Parameter	Description
`r`	HTTP request
`username`	Username for the HTTP authentication
`password`	Password for the HTTP authentication
`scheme`	HTTP authentication scheme. Either `Basic` for the HTTP basic or `AWS` for Amazon S3 authentication scheme. Default is basic.
`for_proxy`	Identifies if the HTTP authentication information is for access to the HTTP proxy server instead of the Web server. Default is `FALSE`.

Usage Notes

The supported authentication schemes are HTTP basic and Amazon S3 authentication.

286.9.35 SET_AUTHENTICATION_FROM_WALLET Procedure

This procedure sets the HTTP authentication information in the HTTP request header needed for the request to be authorized by the Web server using the username and password credential stored in the Oracle wallet.

> **See Also:**
>
> External Password Store on , and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.SET_AUTHENTICATION_FROM_WALLET(
   r         IN OUT NOCOPY req,
   alias     IN VARCHAR2,
   scheme    IN VARCHAR2 DEFAULT 'Basic',
   for_proxy IN BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 286-48 SET_AUTHENTICATION_FROM_WALLET Procedure Parameters

Parameter	Description
r	The HTTP request
alias	Alias to identify and retrieve the username and password credential stored in the Oracle wallet
scheme	HTTP authentication scheme. Either Basic for the HTTP basic or AWS for Amazon S3 authentication scheme. Default is basic.
for_proxy	Identifies if the HTTP authentication information is for access to the HTTP proxy server instead of the Web server. Default is FALSE.

Usage Notes

- To use the password credentials in a wallet, the UTL_HTTP user must have the use-passwords privilege on the wallet.
- The supported authentication schemes are HTTP basic and Amazon S3 authentication schemes.

Examples

Creating a wallet and entering username and password in the wallet

```
> mkstore -wrl /oracle/wallets/test_wallet -create
Enter password: ******
Enter password again: ******
> mkstore -wrl /oracle/wallets/test_wallet -createCredential hr-access jsmith
Your secret/Password is missing in the command line
Enter your secret/Password: ****
Re-enter your secret/Password: ****
Enter wallet password: ******
```

Granting the use-passwords privilege on the wallet to a user by the database administrator

```
BEGIN
  DBMS_NETWORK_ACL_ADMIN.CREATE_ACL(
    acl         => 'wallet-acl.xml',
    description => 'Wallet ACL',
    principal   => 'SCOTT',
    is_grant    => TRUE,
    privilege   => 'use-passwords');
  DBMS_NETWORK_ACL_ADMIN.ASSIGN_WALLET_acl(
      acl          => 'wallet-acl.xml',
      wallet_path  => 'file: /oracle/wallets/test_wallet');
END;
```

Using username and password from the wallet

```
DECLARE
  req  UTL_HTTP.req;
BEGIN
  UTL_HTTP.SET_WALLET(path => 'file:/oracle/wallets/test_wallet');
  req := UTL_HTTP.BEGIN_REQUEST(…);
  UTL_HTTP.SET_AUTHENTICATION_FROM_WALLET(req, 'hr-access');
  …
END;
```

286.9.36 SET_BODY_CHARSET Procedures

This procedure is overloaded. The description of different functionality is located alongside the syntax declarations.

> **See Also:**
>
> - HTTP Responses and HTTP Responses Subprograms
> - Session Settings and Session Settings Subprograms

Syntax

Sets the default character set of the body of all future HTTP requests when the media type is `text` and the character set is not specified in the `Content-Type` header. Following the HTTP protocol standard specification, if the media type of a request or a response is `text`, but the character set information is missing in the `Content-Type` header, the character set of the request or response body should default to `ISO-8859-1`. A response created for a request inherits the default body character set of the request instead of the body character set of the current session. The default body character set is ISO-8859-1 in a database user session. The default body character set setting affects only future requests and has no effect on existing requests. After a request is created, the body character set can be changed by using the other `SET_BODY_CHARSET` procedure that operates on a request:

```
UTL_HTTP.SET_BODY_CHARSET (
   charset  IN VARCHAR2 DEFAULT NULL);
```

Sets the character set of the request body when the media type is `text` but the character set is not specified in the `Content-Type` header. According to the HTTP protocol standard specification, if the media type of a request or a response is "text" but the character set information is missing in the `Content-Type` header, the character set of the request or response body should default to "`ISO-8859-1`". Use this procedure to change the default body character set a request inherits from the session default setting:

```
UTL_HTTP.SET_BODY_CHARSET(
   r        IN OUT NOCOPY req,
   charset  IN VARCHAR2 DEFAULT NULL);
```

Sets the character set of the response body when the media type is "text" but the character set is not specified in the `Content-Type` header. For each the HTTP protocol standard specification, if the media type of a request or a response is "text" but the character set information is missing in the `Content-Type` header, the character set of the request or response body should default to "ISO-8859-1". Use this procedure to change the default body character set a response inherits from the request:

```
UTL_HTTP.SET_BODY_CHARSET(
   r        IN OUT NOCOPY resp,
   charset  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 286-49 SET_BODY_CHARSET Procedure Parameters

Parameter	Description
r	The HTTP response.
charset	The default character set of the response body. The character set can be in Oracle or Internet Assigned Numbers Authority (IANA) naming convention. If charset is NULL, the database character set is assumed.

286.9.37 SET_COOKIE_SUPPORT Procedures

This overloaded procedure handles cookie support. The description of different functionality is located alongside the syntax declarations.

> **See Also:**
>
> - HTTP Requests andHTTP Requests Subprograms
> - Session Settings and Session Settings Subprograms

Syntax

Enables or disables support for the HTTP cookies in the request. Use this procedure to change the cookie support setting a request inherits from the session default setting:

```
UTL_HTTP.SET_COOKIE_SUPPORT(
   r       IN OUT NOCOPY REQ,
   enable  IN BOOLEAN DEFAULT TRUE);
```

Sets whether or not future HTTP requests will support HTTP cookies, and the maximum number of cookies maintained in the current database user session:

```
UTL_HTTP.SET_COOKIE_SUPPORT (
   enable       IN BOOLEAN,
   max_cookies  IN PLS_INTEGER DEFAULT 300,
   max_cookies_per_site  IN PLS_INTEGER DEFAULT 20);
```

Parameters

Table 286-50 SET_COOKIE_SUPPORT Procedure Parameters

Parameter	Description
r	The HTTP request
enable	Set enable to TRUE to enable HTTP cookie support; FALSE to disable
max_cookies	Sets the maximum total number of cookies maintained in the current session
max_cookies_per_site	Sets the maximum number of cookies maintained in the current session for each Web site

Usage Notes

If cookie support is enabled for an HTTP request, all cookies saved in the current session and applicable to the request are returned to the Web server in the request in accordance with HTTP cookie specification standards. Cookies set in the response to the request are saved in the current session for return to the Web server in the subsequent requests if cookie support is enabled for those requests. If the cookie support is disabled for an HTTP request, no cookies are returned to the Web server in the request and the cookies set in the response to the request are not saved in the current session, although the `Set-Cookie` HTTP headers can still be retrieved from the response.

Cookie support is enabled by default for all HTTP requests in a database user session. The default setting of the cookie support (enabled versus disabled) affects only the future requests and has no effect on the existing ones. After your request is created, the cookie support setting may be changed by using the other `SET_COOKIE_SUPPORT` procedure that operates on a request.

The default maximum number of cookies saved in the current session is 20 for each site and 300 total.

If you lower the maximum total number of cookies or the maximum number of cookies for each Web site, the oldest cookies will be purged first to reduce the number of cookies to the lowered maximum. HTTP cookies saved in the current session last for the duration of the database session only; there is no persistent storage for the cookies. Cookies saved in the current session are not cleared if you disable cookie support.

See "Examples" for how to use `GET_COOKIES` and `ADD_COOKIES` to retrieve, save, and restore cookies.

286.9.38 SET_DETAILED_EXCP_SUPPORT Procedure

This procedure sets the `UTL_HTTP` package to raise a detailed exception.

By default, `UTL_HTTP` raises the `request_failed` exception when an HTTP request fails. Use `GET_DETAILED_SQLCODE` and `GET_DETAILED_SQLEERM` for more detailed information about the error.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.SET_DETAILED_EXCP_SUPPORT (
   enable  IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 286-51 SET_DETAILED_EXCP_SUPPORT Procedure Parameters

Parameter	Description
`enable`	Asks `UTL_HTTP` to raise a detailed exception directly if set to `TRUE`; otherwise `FALSE`

286.9.39 SET_FOLLOW_REDIRECT Procedures

This procedure sets the maximum number of times `UTL_HTTP` follows the HTTP redirect instruction in the HTTP response to this request, or future requests, in the `GET_RESPONSE` function.

> **See Also:**
>
> - HTTP Requests and HTTP Requests Subprograms
> - Session Settings and Session Settings Subprograms

Syntax

Use this procedure to set the maximum number of redirections:

```
UTL_HTTP.SET_FOLLOW_REDIRECT (
   max_redirects  IN PLS_INTEGER DEFAULT 3);
```

Use this procedure to change the maximum number of redirections a request inherits from the session default setting:

```
UTL_HTTP.SET_FOLLOW_REDIRECT(
   r              IN OUT NOCOPY req,
   max_redirects  IN PLS_INTEGER DEFAULT 3);
```

Parameters

Table 286-52 SET_FOLLOW_REDIRECT Procedure Parameters

Parameter	Description
`r`	The HTTP request
`max_redirects`	The maximum number of redirects. Set to zero to disable redirects.

Usage Notes

If `max_redirects` is set to a positive number, the GET_RESPONSE Function will automatically follow the redirected URL for the HTTP response status code 301, 302, and 307 for the HTTP HEAD and GET methods, and 303 for all HTTP methods, and retry the HTTP request (the request method will be changed to HTTP GET for the status code 303) at the new location. It follows the redirection until the final, non-redirect location is reached, or an error occurs, or the maximum number of redirections has been reached (to prevent an infinite loop). The URL and method fields in the `REQ` record will be updated to the last redirected URL and the method used to access the URL. Set the maximum number of redirects to zero to disable automatic redirection.

While it is set not to follow redirect automatically in the current session, it is possible to specify individual HTTP requests to follow redirect instructions the function `FOLLOW_REDIRECT` and vice versa.

The default maximum number of redirections in a database user session is 3. The default value affects only future requests and has no effect on existing requests.

The SET_FOLLOW_REDIRECT procedure must be called before GET_RESPONSE for any redirection to take effect.

286.9.40 SET_HEADER Procedure

This procedure sets an HTTP request header. The request header is sent to the Web server as soon as it is set.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.SET_HEADER (
   r      IN OUT NOCOPY req,
   name   IN VARCHAR2,
   value  IN VARCHAR2);
```

Parameters

Table 286-53 SET_HEADER Procedure Parameters

Parameter	Description
r	The HTTP request
name	The name of the HTTP request header
value	The value of the HTTP request header

Usage Notes

Multiple HTTP headers with the same name are allowed in the HTTP protocol standard. Therefore, setting a header does not replace a prior header with the same name.

If the request is made using HTTP 1.1, UTL_HTTP sets the Host header automatically for you.

When you set the Content-Type header with this procedure, UTL_HTTP looks for the character set information in the header value. If the character set information is present, it is set as the character set of the request body. It can be overridden later by using the SET_BODY_CHARSET procedure.

When you set the Transfer-Encoding header with the value chunked, UTL_HTTP automatically encodes the request body written by the WRITE_TEXT, WRITE_LINE and WRITE_RAW procedures. Note that some HTTP-1.1-based Web servers or CGI programs do not support or accept the request body encoding in the HTTP 1.1 chunked transfer-encoding format.

286.9.41 SET_PERSISTENT_CONN_SUPPORT Procedure

This overloaded procedure provides persistent connection support. Descriptions of the different functionality are given in the syntax declarations.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

Sets whether future HTTP requests should support the HTTP 1.1 persistent connection or not, and the maximum numbers of persistent connections to be maintained in the current database user session.

```
UTL_HTTP.SET_PERSISTENT_CONN_SUPPORT(
   enable      IN BOOLEAN DEFAULT FALSE,
   max_conns   IN PLS_INTEGER DEFAULT 0);
```

Enables or disables support for the HTTP 1.1 persistent-connection in the request.

```
UTL_HTTP.SET_PERSISTENT_CONN_SUPPORT(
   r           IN OUT NOCOPY req,
   enable      IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 286-54 SET_PERSISTENT_CONN_SUPPORT Procedure Parameters

Parameter	Description
enable	TRUE to keep the network connection persistent. FALSE otherwise.
maximum_conns	Maximum number of connections
r	The HTTP request

Usage Notes

If the persistent-connection support is enabled for an HTTP request, the package will keep the network connections to a Web server or the proxy server open in the package after the request is completed properly for a subsequent request to the same server to reuse for each HTTP 1.1 protocol specification. With the persistent connection support, subsequent HTTP requests may be completed faster because the network connection latency is avoided. If the persistent-connection support is disabled for a request, the package will always send the HTTP header "Connection: close" automatically in the HTTP request and close the network connection when the request is completed. This setting has no effect on HTTP requests that follows HTTP 1.0 protocol, for which the network connections will always be closed after the requests are completed.

When a request is being made, the package attempts to reuse an existing persistent connection to the target Web server (or proxy server) if one is available. If none is available, a new network connection will be initiated. The persistent-connection support setting for a request affects only whether the network connection should be closed after a request completes.

Use this procedure to change the persistent-connection support setting a request inherits from the session default setting.

Users should note that while the use of persistent connections in `UTL_HTTP` may reduce the time it takes to fetch multiple Web pages from the same server, it consumes precious system resources (network connections) in the database server. Also, excessive use of persistent connections may reduce the scalability of the database server when too many network connections are kept open in the database server. Network connections should be kept open only if they will be used immediately by subsequent requests and should be closed immediately when they are no longer needed. Set the default persistent connection support as disabled in the session, and enable persistent connection in individual HTTP requests as shown in "Examples".

The default value of the maximum number of persistent connections in a database session is zero. To truly enable persistent connections, you must also set the maximum number of persistent connections to a positive value or no connections will be kept persistent.

Note that if you want to use persistent connections, you must call the overload that takes the maximum_conns parameter prior to calling the BEGIN_REQUEST Function, otherwise persistent connections will not be enabled for the current request even if the other form of `SET_PERSISTENT_CONN_SUPPORT` is called.

Examples

Using SET_PERSISTENT_CONN_SUPPORT in HTTP requests at the session level, showing the active persistent connection after each request

```
DECLARE
  pieces utl_http.html_pieces;
  conns  utl_http.connection_table;
BEGIN

  -- Turns on persistent connection support for the request_pieces call.
  utl_http.set_persistent_conn_support(true, 1);

  FOR i IN 1..10 LOOP

    pieces := utl_http.request_pieces('http://www.example.com/');

    -- Shows the active persistent connection
    utl_http.get_persistent_conns(conns);
    FOR j IN 1..conns.count LOOP
       dbms_output.put_line('Persistent connection '||j||': '||conns(j).host||':'||
conns(j).port);
    END LOOP;

  END LOOP;


  -- Turns off persistent connection support. Set active max persistent connection to 0
to close all active connections.
  utl_http.set_persistent_conn_support(false, 0);

END;
/
```

Using SET_PERSISTENT_CONN_SUPPORT in HTTP requests showing how to use persistent connection individually in each request to fetch multiple URLs at the same host

```
DECLARE
-- Table to store the URLs
```

```
TYPE vc2_table IS TABLE OF VARCHAR2(256) INDEX BY BINARY_INTEGER;
paths VC2_TABLE;

PROCEDURE fetch_pages(paths IN vc2_table) AS
  req  UTL_HTTP.REQ;
  resp UTL_HTTP.RESP;
  data VARCHAR2(1024);

BEGIN

  -- Set the proxy server
  UTL_HTTP.SET_PROXY('www-proxy.example.com:80', '');

  FOR i IN 1..paths.count LOOP

    req := UTL_HTTP.BEGIN_REQUEST(paths(i));

    -- Use persistent connections except for the last request
    IF (i < paths.count) THEN
      -- Use a persistent connection for the current request
      UTL_HTTP.SET_PERSISTENT_CONN_SUPPORT(req, TRUE);
    END IF;

    resp := UTL_HTTP.GET_RESPONSE(req);

    -- Display the results of the response
    DBMS_OUTPUT.PUT_LINE('-');
    DBMS_OUTPUT.PUT_LINE('URL: ' || paths(i));
    DBMS_OUTPUT.PUT_LINE('HTTP Response Status Code:   ' || resp.status_code);
    DBMS_OUTPUT.PUT_LINE('HTTP Response Reason Phrase: ' || resp.reason_phrase);
    DBMS_OUTPUT.PUT_LINE('HTTP Response Version:       ' || resp.http_version);

    BEGIN
      LOOP
        UTL_HTTP.READ_TEXT(resp, data);
        -- do something with the data
      END LOOP;
    EXCEPTION
      WHEN UTL_HTTP.END_OF_BODY THEN
        NULL;
    END;
    UTL_HTTP.END_RESPONSE(resp);
  END LOOP;
END;

BEGIN
-- Set a maximum of 1 persistent connection, but start with persistent connections
-- off
   UTL_HTTP.SET_PERSISTENT_CONN_SUPPORT(FALSE, 1);

   -- Create a list of URLs
   paths(1) := 'http://www.example.com/technetwork/index.html';
   paths(2) := 'http://www.example.com/us/products/index.html';

   fetch_pages(paths);
 END;
/
```

286.9.42 SET_PROXY Procedure

This procedure sets the proxy to be used for requests of the HTTP or other protocols, excluding those for hosts that belong to the domain specified in `no_proxy_domains`.

`no_proxy_domains` is a comma-, semi-colon-, or space-separated list of domains or hosts for which HTTP requests should be sent directly to the destination HTTP server instead of going through a proxy server.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.SET_PROXY (
   proxy             IN VARCHAR2,
   no_proxy_domains  IN VARCHAR2);
```

Parameters

Table 286-55 SET_PROXY Procedure Parameters

Parameter	Description
`proxy`	The proxy (host and an optional port number) to be used by the `UTL_HTTP` package
`no_proxy_domains`	The list of hosts and domains for which no proxy should be used for all requests

Usage Notes

The proxy may include an optional TCP/IP port number at which the proxy server listens. The syntax is `[http://]host[:port][/]`, for example, `www-proxy.my-company.com:80`. If the port is not specified for the proxy, port 80 is assumed.

Optionally, a port number can be specified for each domain or host. If the port number is specified, the no-proxy restriction is only applied to the request at the port of the particular domain or host, for example, `corp.my-company.com, eng.my-company.com:80`. When `no_proxy_domains` is `NULL` and the proxy is set, all requests go through the proxy. When the proxy is not set, `UTL_HTTP` sends requests to the target Web servers directly.

You can define a username/password for the proxy to be specified in the proxy string. The format is

```
[http://][user[:password]@]host[:port][/]
```

If proxy settings are set when the database server instance is started, the proxy settings in the environment variables `http_proxy` and `no_proxy` are assumed. Proxy settings set by this procedure override the initial settings.

286.9.43 SET_RESPONSE_ERROR_CHECK Procedure

This procedure sets whether or not `GET_RESPONSE` raises an exception when the Web server returns a status code that indicates an error—a status code in the 4xx or 5xx ranges.

For example, when the requested URL is not found in the destination Web server, a 404 (document not found) response status code is returned.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.SET_RESPONSE_ERROR_CHECK (
   enable  IN BOOLEAN DEFAULT FALSE);
```

Parameters

Table 286-56 SET_RESPONSE_ERROR_CHECK Procedure Parameters

Parameter	Description
`enable`	`TRUE` to check for response errors; otherwise `FALSE`

Usage Notes

If the status code indicates an error—a 4xx or 5xx code—and this procedure is enabled, `GET_RESPONSE` will raise the `HTTP_CLIENT_ERROR` or `HTTP_SERVER_ERROR` exception. If `SET_RESPONSE_ERROR_CHECK` is set to `FALSE`, `GET_RESPONSE` will not raise an exception when the status code indicates an error.

Response error check is turned off by default.

The `GET_RESPONSE` function can raise other exceptions when `SET_RESPONSE_ERROR_CHECK` is set to `FALSE`.

286.9.44 SET_TRANSFER_TIMEOUT Procedure

This procedure sets the default time out value for all future HTTP requests that the `UTL_HTTP` package should attempt while reading the HTTP response from the Web server or proxy server.

This time out value may be used to avoid the PL/SQL programs from being blocked by busy Web servers or heavy network traffic while retrieving Web pages from the Web servers.

> **See Also:**
>
> Session Settings and Session Settings Subprograms

Syntax

```
UTL_HTTP.SET_TRANSFER_TIMEOUT (
   timeout  IN PLS_INTEGER DEFAULT 60);
```

Parameters

Table 286-57 SET_TRANSFER_TIMEOUT Procedure Parameters

Parameter	Description
`timeout`	The network transfer timeout value in seconds.

Usage Notes

The default value of the time out is 60 seconds.

286.9.45 SET_WALLET Procedure

This procedure sets the Oracle wallet used for all HTTP requests over Secured Socket Layer (SSL), namely HTTPS.

When the `UTL_HTTP` package communicates with an HTTP server over SSL, the HTTP server presents its digital certificate, which is signed by a certificate authority, to the `UTL_HTTP` package for identification purpose. The Oracle wallet contains the list of certificate authorities that are trusted by the user of the `UTL_HTTP` package. An Oracle wallet is required to make an HTTPS request.

> **See Also:**
>
> - Session Settings and Session Settings Subprograms
> - *Oracle Database Security Guide* managing fine-grained access

Syntax

```
UTL_HTTP.SET_WALLET (
   path      IN VARCHAR2,
   password  IN VARCHAR2 DEFAULT NULL);
```

Parameters

Table 286-58 SET_WALLET Procedure Parameters

Parameter	Description
`path`	The directory path that contains the Oracle wallet. The format is `file:`*`directory-path`*. The format of wallet_path on a PC is, for example, `file:c:\WINNT\Profiles\`*`username`*`\WALLETS`, and in Unix is, for example, `file:/home/`*`username`*`/wallets`. When the `UTL_HTTP` package is executed in the Oracle database server, the wallet is accessed from the database server. Therefore, the wallet path must be accessible from the database server. If you want to use the operating system certificate store to act in place of the Oracle wallet, then set the `path` parameter to `system:` (include the colon). Doing so greatly improves performance in the database.
`password`	The password needed to open the wallet. If the wallet is auto-login enabled, the password may be omitted and should be set to `NULL`. See *Oracle Database Security Guide* for information about using the `orapki` utility to create an auto-login wallet. If you set `path` to `system:`, then omit the password by setting it to `NULL`.

Usage Notes

To set up an Oracle wallet, use the `ORAPKI` utility to create a wallet. In order for the HTTPS request to succeed, the certificate authority that signs the certificate of the remote HTTPS Web server must be a trust point set in the wallet.

When a wallet is created, it is populated with a set of well-known certificate authorities as trust points. If the certificate authority that signs the certificate of the remote HTTPS Web server is not among the trust points, or the certificate authority has new root certificates, you should obtain the root certificate of that certificate authority and install it as a trust point in the wallet.

> **See Also:**
>
> *Oracle Database Transparent Data Encryption Guide* for more information on Wallet Manager

286.9.46 WRITE_LINE Procedure

This procedure writes a text line in the HTTP request body and ends the line with new-line characters (CRLF as defined in `UTL_TCP`).

As soon as some data is sent as the HTTP request body, the HTTP request headers section is completed. Text data is automatically converted from the database character set to the request body character set.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.WRITE_LINE(
   r     IN OUT NOCOPY req,
   data  IN VARCHAR2 CHARACTER SET ANY_CS);
```

Parameters

Table 286-59 WRITE_LINE Procedure Parameters

Parameter	Description
r	The HTTP request
data	The text line to send in the HTTP request body

Usage Notes

An HTTP client must always let the remote Web server know the length of the request body it is sending. If the amount of data is known beforehand, you can set the `Content-Length` header in the request, where the length of the content is measured in bytes instead of characters. If the length of the request body is not known beforehand, you can send the request body using the HTTP 1.1 chunked transfer-encoding format. The request body is sent in chunks, where the length of each chunk is sent before the chunk is sent. The `UTL_HTTP` package performs chunked transfer-encoding on the request body transparently when the Transfer-Encoding: chunked header is set. Note that some HTTP-1.1-based Web servers or CGI programs do not support or accept the request body encoding in the HTTP 1.1 chunked transfer-encoding format. See the `SET_HEADER` procedure for details.

If you send the `Content-Length` header, you should note that the length specified in the header should be the byte-length of the textual request body after it is converted from the database character set to the request body character set. When either one of the two character sets is a multibyte character set, the precise byte-length of the request body in the request body character set cannot be known beforehand. In this case, you can perform the character set conversion explicitly, determine the byte-length of the results, send the `Content-Length` header, and the results using the `WRITE_RAW` procedure to avoid the automatic character set conversion. Or, if the remove Web server or CGI programs allow, you can send the request body using the HTTP 1.1 chunked transfer-encoding format, where `UTL_HTTP` handles the length of the chunks transparently.

286.9.47 WRITE_RAW Procedure

This procedure writes some binary data in the HTTP request body. As soon as some data is sent as the HTTP request body, the HTTP request headers section is completed.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.WRITE_RAW(
   r     IN OUT NOCOPY REQ,
   data  IN            RAW);
```

Parameters

Table 286-60 WRITE_RAW Procedure Parameters

Parameter	Description
r	The HTTP request
data	The binary data to send in the HTTP request body

Usage Notes

An HTTP client must always let the remote Web server know the length of the request body it is sending. If the amount of data is known beforehand, you can set the `Content-Length` header in the request, where the length of the content is measured in bytes instead of characters. If the length of the request body is not known beforehand, you can send the request body using the HTTP 1.1 chunked transfer-encoding format. The request body is sent in chunks, where the length of each chunk is sent before the chunk is sent. `UTL_HTTP` performs chunked transfer-encoding on the request body transparently when the `Transfer-Encoding:chunked` header is set. Note that some HTTP-1.1-based Web servers or CGI programs do not support or accept the request body encoding in the HTTP 1.1 chunked transfer-encoding format. See the `SET_HEADER` procedure for details.

286.9.48 WRITE_TEXT Procedure

This procedure writes some text data in the HTTP request body.

As soon as some data is sent as the HTTP request body, the HTTP request headers section is completed. Text data is automatically converted from the database character set to the request body character set.

> **See Also:**
>
> HTTP Requests and HTTP Requests Subprograms

Syntax

```
UTL_HTTP.WRITE_TEXT(
   r     IN OUT NOCOPY REQ,
   data  IN            VARCHAR2 CHARACTER SET ANY_CS);
```

Parameters

Table 286-61 WRITE_TEXT Procedure Parameters

Parameter	Description
r	The HTTP request

Table 286-61 (Cont.) WRITE_TEXT Procedure Parameters

Parameter	Description
`data`	The text data to send in the HTTP request body

Usage Notes

An HTTP client must always let the remote Web server know the length of the request body it is sending. If the amount of data is known beforehand, you can set the `Content-Length` header in the request, where the length of the content is measured in bytes instead of characters. If the length of the request body is not known beforehand, you can send the request body using the HTTP 1.1 chunked transfer-encoding format. The request body is sent in chunks, where the length of each chunk is sent before the chunk is sent. `UTL_HTTP` performs chunked transfer-encoding on the request body transparently when the `Transfer-Encoding: chunked` header is set. Note that some HTTP-1.1-based Web servers or CGI programs do not support or accept the request body encoding in the HTTP 1.1 chunked transfer-encoding format. See the `SET_HEADER` procedure for details.

If you send the `Content-Length` header, you should note that the length specified in the header should be the byte-length of the textual request body after it is converted from the database character set to the request body character set. When either one of the two character sets is a multibyte character set, the precise byte-length of the request body in the request body character set cannot be known beforehand. In this case, you can perform the character set conversion explicitly, determine the byte-length of the results, send the `Content-Length` header, and the results using the `WRITE_RAW` procedure to avoid the automatic character set conversion. Or, if the remove Web server or CGI programs allow, you can send the request body using the HTTP 1.1 chunked transfer-encoding format, where `UTL_HTTP` handles the length of the chunks transparently.

287
UTL_I18N

`UTL_I18N` is a set of services that provides additional globalization functionality for applications written in PL/SQL.

> **See Also:**
>
> *Oracle Database Globalization Support Guide*

The chapter contains the following topics:

- Overview
- Security Model
- Constants
- Data Structures
- Summary of UTL_I18N Subprograms

287.1 UTL_I18N Overview

`UTL_I18N` services provide additional globalization functionality for applications written in PL/SQL.

The `UTL_I18N` PL/SQL package consists of the following categories of services:

- String conversion functions for various datatypes.
- Functions that convert a text string to character references and vice versa.
- Functions that map between Oracle, Java, and ISO languages and territories.
- Functions that map between Oracle, Internet Assigned Numbers Authority (IANA), and e-mail safe character sets.
- A function that returns the Oracle character set name from an Oracle language name.
- A function that returns the maximum number of bytes for a character of an Oracle character set.
- A function that performs script transliteration.
- Functions that return the ISO currency code, local time zones, and local languages supported for a given territory.
- Functions that return the most appropriate linguistic sort, a listing of all the applicable linguistic sorts, and the local territories supported for a given language.
- Functions that map between the Oracle full and short language names.
- A function that returns the language translation of a given language and territory name.
- A function that returns a listing of the most commonly used time zones.

- Procedures that detect the most likely Oracle character sets and Oracle languages for a given data sample.

287.2 UTL_I18N Security Model

The functions and procedures of the `UTL_I18N` package neither read database contents nor modify them. They operate on their arguments only and/or they retrieve static internationalization information from NLS Data files. The execution privilege for the package is granted to `PUBLIC` by default

287.3 UTL_I18N Constants

`UTL_I18N` defines constants to use when specifying parameter values.

These constants are shown in the following table.

Table 287-1 UTL_I18N Constants

Constant	Type	Value	Description
GENERIC_CONTEXT	PLS_INTEGER	0	Returns the default character set for general cases.
MAIL_GENERIC	PLS_INTEGER	0	Map from an Oracle character set name to an email safe character set name on a non-Windows platform.
ORACLE_TO_IANA	PLS_INTEGER	0	Map from an Oracle character set name to an IANA character set name.
SHIFT_IN	PLS_INTEGER	0	Used with `shift_status`. Must be set the first time it is called in piecewise conversion.
IANA_TO_ORACLE	PLS_INTEGER	1	Map from an IANA character set name to an Oracle character set name.
MAIL_CONTEXT	PLS_INTEGER	1	The mapping is between an Oracle character set name and an email safe character set name.
MAIL_WINDOWS	PLS_INTEGER	1	Map from an Oracle character set name to an email safe character set name on a Windows platform.
SHIFT_OUT	PLS_INTEGER	1	
FWKATAKANA_HIRAGANA	VARCHAR2(30)	'fwkatakana_hiragana'	Converts only fullwidth Katakana characters to fullwidth Hiragana characters.
FWKATAKANA_HWKATAKANA	VARCHAR2(30)	'fwkatakana_hwkatakana'	Converts only fullwidth Katakana characters to halfwidth Katakana characters.
HIRAGANA_FWKATAKANA	VARCHAR2(30)	'hiragana_fwkatakana'	Converts only fullwidth Hiragana characters to fullwidth Katakana characters.
HIRAGANA_HWKATAKANA	VARCHAR2(30)	'hiragana_hwkatakana'	Converts only fullwidth Hiragana characters to halfwidth Katakana characters.

Table 287-1 (Cont.) UTL_I18N Constants

Constant	Type	Value	Description
HWKATAKANA_FWKATAKANA	VARCHAR2(30)	'hwkatakana_fwkatakana'	Converts only halfwidth Katakana characters to fullwidth Katakana characters.
HWKATAKANA_HIRAGANA	VARCHAR2(30)	'hwkatakana_hiragana'	Converts only halfwidth Katakana characters to fullwidth Hiragana characters.
KANA_FWKATAKANA	VARCHAR2(30)	'kana_fwkatakana'	Converts any type of Kana character to a fullwidth Katakana character.
KANA_HIRAGANA	VARCHAR2(30)	'kana_hiragana'	Converts any type of Kana character to a fullwidth Hiragana character.
KANA_HWKATAKANA	VARCHAR2(30)	'kana_hwkatakana'	Converts any type of Kana character to a halfwidth Katakana character.
LATIN_ASCII_DIN91379	VARCHAR2(30)	'latin_ascii_din91379'	Converts Latin characters to ASCII characters.
CYRILLIC_LATIN_ISO9	VARCHAR2(30)	'cyrillic_latin_iso9'	Converts Cyrillic characters to Latin characters.
GREEK_LATIN_ISO843	VARCHAR2(30)	'greek_latin_iso843'	Converts Greek characters to Latin characters.
ARABIC_LATIN_ISO233	VARCHAR2(30)	'arabic_latin_iso233'	Converts Arabic characters to Latin characters.
HEBREW_LATIN_ISO259	VARCHAR2(30)	'hebrew_latin_iso259'	Converts Hebrew characters to Latin characters.
MODERN_HEBREW_LATIN_ISO259_2	VARCHAR2(30)	'modern_hebrew_latin_iso259_2'	Converts Modern Hebrew characters to Latin characters.
CYR_ASCII_ICAO9303	VARCHAR2(30)	'cyr_ascii_ICAO9303'	Converts Cyrillic characters to ASCII characters.
CYR_BY_ASCII_ICAO9303	VARCHAR2(30)	'cyr_by_ascii_ICAO9303'	Converts Cyrillic Belarus characters to ASCII characters.
CYR_UA_ASCII_ICAO9303	VARCHAR2(30)	'cyr_ua_ascii_ICAO9303'	Converts Cyrillic Ukraine characters to ASCII characters.
CYR_RS_ASCII_ICAO9303	VARCHAR2(30)	'cyr_rs_ascii_ICAO9303'	Converts Cyrillic Serbia characters to ASCII characters.
CYR_BG_ASCII_ICAO9303	VARCHAR2(30)	'cyr_bg_ascii_ICAO9303'	Converts Cyrillic Bulgaria characters to ASCII characters.
CYR_MK_ASCII_ICAO9303	VARCHAR2(30)	'cyr_mk_ascii_ICAO9303'	Converts Cyrillic North Macedonia characters to ASCII characters.

287.4 UTL_I18N Data Structures

The `UTL_I18N` package defines `RECORD` types and `TABLE` types.

RECORD Types

- CHARSET_RESULT_RECORD Record Type
- LANGUAGE_CHARSET_RESULT_RECORD Record Type
- LANGUAGE_RESULT_RECORD Record Type

TABLE Types

- CHARSET_RESULT_TABLE Table Type
- LANGUAGE_CHARSET_RESULT_TABLE Table Type
- LANGUAGE_RESULT_TABLE Table Type

287.4.1 CHARSET_RESULT_RECORD Record Type

This record type contains information returned by the `DETECT_CHARSET` procedure in the `UTL_I18N` package.

Syntax

```
TYPE CHARSET_RESULT_RECORD IS RECORD (
  charset      VARCHAR2(28),
  score        NUMBER);
```

Fields

Table 287-2 CHARSET_RESULT_RECORD Fields

Field	Description
`charset`	Oracle character set
`score`	A numeric value representing the likelihood that `charset` is the Oracle character set for the given data sample. The value is between 0 (least likely) and 1 (most likely).

287.4.2 LANGUAGE_CHARSET_RESULT_RECORD Record Type

This record type contains information returned by the `DETECT_LANGUAGE_CHARSET` procedure in the `UTL_I18N` package.

Syntax

```
TYPE LANGUAGE_CHARSET_RESULT_RECORD IS RECORD (
  language      VARCHAR2(28),
  charset       VARCHAR2(28),
  score         NUMBER);
```

Fields

Table 287-3 LANGUAGE_CHARSET_RESULT_RECORD Fields

Field	Description
`language`	Oracle language
`charset`	Oracle character set
`score`	A numeric value representing the likelihood that `language` is the Oracle language and `charset` is the Oracle character set for the given data sample. The value is between 0 (least likely) and 1 (most likely).

287.4.3 LANGUAGE_RESULT_RECORD Record Type

This record type contains information returned by the `DETECT_LANGUAGE` procedure in the `UTL_I18N` package.

Syntax

```
TYPE LANGUAGE_RESULT_RECORD IS RECORD (
  language      VARCHAR2(28),
  score         NUMBER);
```

Fields

Table 287-4 LANGUAGE_RESULT_RECORD Fields

Field	Description
`language`	Oracle language
`score`	A numeric value representing the likelihood that `language` is the Oracle language for the given data sample. The value is between 0 (least likely) and 1 (most likely).

287.4.4 CHARSET_RESULT_TABLE Table Type

This is a table of `CHARSET_RESULT_RECORD` Record Type.

Syntax

```
TYPE charset_result_table IS TABLE OF charset_result_record;
```

Related Topics

- CHARSET_RESULT_RECORD Record Type
  This record type contains information returned by the `DETECT_CHARSET` procedure in the `UTL_I18N` package.

287.4.5 LANGUAGE_CHARSET_RESULT_TABLE Table Type

This is a table of `LANGUAGE_CHARSET_RESULT_RECORD` Record Type.

Syntax

```
TYPE language_charset_result_table IS TABLE OF language_charset_result_record;
```

Related Topics

- LANGUAGE_CHARSET_RESULT_RECORD Record Type
  This record type contains information returned by the `DETECT_LANGUAGE_CHARSET` procedure in the `UTL_I18N` package.

287.4.6 LANGUAGE_RESULT_TABLE Table Type

This is a table of `LANGUAGE_RESULT_RECORD` Record Type.

Syntax

```
TYPE language_result_table IS TABLE OF language_result_record;
```

Related Topics

- LANGUAGE_RESULT_RECORD Record Type
  This record type contains information returned by the `DETECT_LANGUAGE` procedure in the `UTL_I18N` package.

287.5 Summary of UTL_I18N Subprograms

This table lists the `UTL_I18N` subprograms and briefly describes them.

Table 287-5 UTL_I18N Package Subprograms

Procedure	Description
DETECT_CHARSET Procedure	Detects the most likely Oracle character sets for a given data sample.
DETECT_LANGUAGE Procedure	Detects the most likely Oracle languages for a given data sample.
DETECT_LANGUAGE_CHARSET Procedure	Detects the most likely Oracle languages and Oracle character sets for a given data sample.
ESCAPE_REFERENCE Function	Converts a given text string to its character reference counterparts, for characters that fall outside the document character set.
GET_COMMON_TIME_ZONES Function	Returns the list of common time zone IDs that are independent of the locales.
GET_DEFAULT_CHARSET Function	Returns the default Oracle character set name or the default e-mail safe character set name from an Oracle language name.
GET_DEFAULT_ISO_CURRENCY Function	Returns the default ISO 4217 currency code for the specified territory.
GET_DEFAULT_LINGUISTIC_SORT Function	Returns the default linguistic sort name for the specified language.
GET_LOCAL_LANGUAGES Function	Returns the local language names for the specified territory.
GET_LOCAL_LINGUISTIC_SORTS Function	Returns the local linguistic sort names for the specified language.
GET_LOCAL_TERRITORIES Function	Returns the local territory names for the specified language.
GET_LOCAL_TIME_ZONES Function	Returns the local time zone IDs for the specified territory.
GET_MAX_CHARACTER_SIZE Function	Returns the maximum character size of a given character set.
GET_TRANSLATION Function	Returns the translation of the language and territory name in the specified translation language.

Table 287-5 (Cont.) UTL_I18N Package Subprograms

Procedure	Description
MAP_CHARSET Function	• Maps an Oracle character set name to an IANA character set name. • Maps an IANA character set name to an Oracle character set name. • Maps an Oracle character set name to an e-mail safe character set name.
MAP_FROM_SHORT_LANGUAGE Function	Maps an Oracle short language name to an Oracle language name.
MAP_LANGUAGE_FROM_ISO Function	Returns an Oracle language name from an ISO locale name.
MAP_LOCALE_TO_ISO Function	Returns an ISO locale name from the Oracle language and territory name.
MAP_TERRITORY_FROM_ISO Function	Returns an Oracle territory name from an ISO locale name.
MAP_TO_SHORT_LANGUAGE Function	Maps an Oracle language name to an Oracle short language name.
RAW_TO_CHAR Functions	Converts `RAW` data that is not encoded in the database character set into a `VARCHAR2` string
RAW_TO_NCHAR Functions	Converts `RAW` data that is not encoded in the national character set into an `NVARCHAR2` string
STRING_TO_RAW Function	Converts a `VARCHAR2` or `NVARCHAR2` string to another character set. The result is returned as a `RAW` datatype.
TRANSLITERATE Function	Performs script transliteration.
UNESCAPE_REFERENCE Function	Converts an input string that contains character references to a text string.
VALIDATE_CHARACTER_ENCODING Function	Validates the character encoding of VARCHAR2, NVARCHAR2, CLOB, and NCLOB data.

287.5.1 DETECT_CHARSET Procedure

This procedure detects the most likely Oracle character sets for a given data sample.

The procedure is overloaded to accept data of type BFILE, BLOB, CLOB, VARCHAR2, or NVARCHAR2. The procedure analyzes the data and returns its possible Oracle character sets, ranked by likelihood. It stores the results in a table of type `CHARSET_RESULT_TABLE`.

Syntax

This procedure analyzes BFILE data:

```
UTL_I18N.DETECT_CHARSET (
    result         OUT charset_result_table,
    src            IN BFILE,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes BLOB data:

```
UTL_I18N.DETECT_CHARSET (
    result         OUT charset_result_table,
```

```
    src             IN BLOB,
    profile         IN BINARY_INTEGER DEFAULT 1,
    num_results     IN BINARY_INTEGER DEFAULT NULL,
    sample_size     IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes CLOB data:

```
UTL_I18N.DETECT_CHARSET (
    result          OUT charset_result_table,
    src             IN CLOB,
    profile         IN BINARY_INTEGER DEFAULT 1,
    num_results     IN BINARY_INTEGER DEFAULT NULL,
    sample_size     IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes VARCHAR2 or NVARCHAR2 data:

```
UTL_I18N.DETECT_CHARSET (
    result          OUT charset_result_table,
    src             IN VARCHAR2,
    profile         IN BINARY_INTEGER DEFAULT 1,
    num_results     IN BINARY_INTEGER DEFAULT NULL,
    sample_size     IN BINARY_INTEGER DEFAULT NULL
);
```

Parameters

Table 287-6 DETECT_CHARSET Procedure Parameters

Parameter	Description
`result`	Result table of type `CHARSET_RESULT_TABLE`. Each result row consists of an Oracle character set name and its score. The rows are ordered by score in descending order, that is, the character set with the highest score (most likely) is listed first, and the character set with the lowest score (least likely) is listed last.
`src`	Data sample whose character set is to be detected.
`profile`	Language and character set detection profile identifier: • `1` - Supports the languages and character sets supported by LCSSCAN and GDK. When this profile is used and the given sample data is binary data, the procedure does not return a reasonable result. • `2` - Supports the languages and character sets supported by LCSSCAN and GDK, as well as the pseudo BINARY language and the pseudo BINARY character set. This profile does not support the AL16UTF16 character set. The default profile identifier is `1`. **See Also:** *Oracle Database Globalization Support Guide* for the list of languages and character sets supported by LCSSCAN and GDK
`num_results`	Number of result rows to store in the result table. The default is NULL, meaning the procedure stores all result rows in the result table.
`sample_size`	Sampling size of the data (in bytes). If this value is greater than the length of the data, then the data length is used. The default is NULL, meaning the full length of the data is used.

Usage Notes

If the given data sample is NULL, then the procedure returns a result table with no rows.

Example

Create a table that contains a data sample of type VARCHAR2.

```
CREATE TABLE src_data (sample VARCHAR2(40));
INSERT INTO src_data VALUES ('Hello! How are you?');
```

Use the `DETECT_CHARSET` procedure to analyze the data sample.

```
VAR l_char VARCHAR2(40)
EXEC SELECT sample INTO :l_char FROM src_data;


VAR l_cur REFCURSOR

DECLARE
  l_result     UTL_I18N.CHARSET_RESULT_TABLE;
BEGIN
  UTL_I18N.DETECT_CHARSET(l_result, :l_char, 1, 5);
  OPEN :l_cur FOR SELECT * FROM TABLE(l_result);
END;
/

PRINT l_cur
```

The result table displays the top five likeliest Oracle character sets for the given data sample.

```
CHARSET                          SCORE
-------------------------- ----------
US7ASCII                       .853497
EE8ISO8859P2                  .0933974
NEE8ISO8859P4                 .0224543
AL16UTF16                     .0100793
WE8ISO8859P9                 .00778038
```

287.5.2 DETECT_LANGUAGE Procedure

This procedure detects the most likely Oracle languages for a given data sample.

The procedure is overloaded to accept data of type BFILE, BLOB, CLOB, VARCHAR2, or NVARCHAR2. The procedure analyzes the data and returns its possible Oracle languages, ranked by likelihood. It stores the results in a table of type `LANGUAGE_RESULT_TABLE`.

Syntax

This procedure analyzes BFILE data:

```
UTL_I18N.DETECT_LANGUAGE (
    result         OUT language_result_table,
    src            IN BFILE,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes BLOB data:

```
UTL_I18N.DETECT_LANGUAGE (
    result         OUT language_result_table,
    src            IN BLOB,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes CLOB data:

```
UTL_I18N.DETECT_LANGUAGE (
    result         OUT language_result_table,
    src            IN CLOB,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes VARCHAR2 or NVARCHAR2 data:

```
UTL_I18N.DETECT_LANGUAGE (
    result         OUT language_result_table,
    src            IN VARCHAR2,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

Parameters

Table 287-7 DETECT_LANGUAGE Procedure Parameters

Parameter	Description
`result`	Result table of type `LANGUAGE_RESULT_TABLE`. Each result row consists of an Oracle language name and its score. The rows are ordered by score in descending order, that is, the language with the highest score (most likely) is listed first, and the language with the lowest score (least likely) is listed last.
`src`	Data sample whose language is to be detected.
`profile`	Language and character set detection profile identifier: • `1` - Supports the languages and character sets supported by LCSSCAN and GDK. When this profile is used and the given sample data is binary data, the procedure does not return a reasonable result. • `2` - Supports the languages and character sets supported by LCSSCAN and GDK, as well as the pseudo BINARY language and the pseudo BINARY character set. This profile does not support the AL16UTF16 character set. The default profile identifier is `1`. **See Also:** *Oracle Database Globalization Support Guide* for the list of languages and character sets supported by LCSSCAN and GDK
`num_results`	Number of result rows to store in the result table. The default is NULL, meaning the procedure stores all result rows in the result table.
`sample_size`	Sampling size of the data (in bytes). If this value is greater than the length of the data, then the data length is used. The default is NULL, meaning the full length of the data is used.

Usage Notes

If the given data sample is NULL, then the procedure returns a result table with no rows.

Examples

Create a table that contains a data sample of type VARCHAR2.

```
CREATE TABLE src_data (sample VARCHAR2(40));
INSERT INTO src_data VALUES ('Hello! How are you?');
```

Use the `DETECT_LANGUAGE` procedure to analyze the data sample.

```
VAR l_char VARCHAR2(40)
EXEC SELECT sample INTO :l_char FROM src_data;


VAR l_cur REFCURSOR

DECLARE
  l_result     UTL_I18N.LANGUAGE_RESULT_TABLE;
BEGIN
  UTL_I18N.DETECT_LANGUAGE(l_result, :l_char, 1, 5);
  OPEN :l_cur FOR SELECT * FROM TABLE(l_result);
END;
/

PRINT l_cur
```

The result table displays the top five likeliest Oracle languages for the given data sample.

```
LANGUAGE                          SCORE
------------------------------ ----------
ENGLISH                         .245638
ITALIAN                        .0876835
POLISH                         .0758762
FRENCH                         .0744972
INDONESIAN                     .0538222
```

287.5.3 DETECT_LANGUAGE_CHARSET Procedure

This procedure detects the most likely Oracle languages and Oracle character sets for a given data sample.

The procedure is overloaded to accept data of type BFILE, BLOB, CLOB, VARCHAR2, or NVARCHAR2. The procedure analyzes the data and returns its possible Oracle languages and Oracle character sets, ranked by likelihood. It stores the results in a table of type `LANGUAGE_CHARSET_RESULT_TABLE`.

Syntax

This procedure analyzes BFILE data:

```
UTL_I18N.DETECT_LANGUAGE_CHARSET (
    result         OUT language_charset_result_table,
    src            IN BFILE,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes BLOB data:

```
UTL_I18N.DETECT_LANGUAGE_CHARSET (
    result         OUT language_charset_result_table,
    src            IN BLOB,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes CLOB data:

```
UTL_I18N.DETECT_LANGUAGE_CHARSET (
    result         OUT language_charset_result_table,
    src            IN CLOB,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

This procedure analyzes VARCHAR2 or NVARCHAR2 data:

```
UTL_I18N.DETECT_LANGUAGE_CHARSET (
    result         OUT language_charset_result_table,
    src            IN VARCHAR2,
    profile        IN BINARY_INTEGER DEFAULT 1,
    num_results    IN BINARY_INTEGER DEFAULT NULL,
    sample_size    IN BINARY_INTEGER DEFAULT NULL
);
```

Parameters

Table 287-8 DETECT_LANGUAGE_CHARSET Procedure Parameters

Parameter	Description
`result`	Result table of type `LANGUAGE_CHARSET_RESULT_TABLE`. Each result row consists of an Oracle language, an Oracle character set name, and their score. The rows are ordered by score in descending order, that is, the language and character set with the highest score (most likely) is listed first, and the language and character set with the lowest score (least likely) is listed last.
`src`	Data sample whose language and character set are to be detected.

Table 287-8 (Cont.) DETECT_LANGUAGE_CHARSET Procedure Parameters

Parameter	Description
profile	Language and character set detection profile identifier: • 1 - Supports the languages and character sets supported by LCSSCAN and GDK. When this profile is used and the given sample data is binary data, the procedure does not return a reasonable result. • 2 - Supports the languages and character sets supported by LCSSCAN and GDK, as well as the pseudo BINARY language and the pseudo BINARY character set. This profile does not support the AL16UTF16 character set. The default profile identifier is 1. **See Also:** *Oracle Database Globalization Support Guide* for the list of languages and character sets supported by LCSSCAN and GDK
num_results	Number of result rows to store in the result table. The default is NULL, meaning the procedure stores all result rows in the result table.
sample_size	Sampling size of the data (in bytes). If this value is greater than the length of the data, then the data length is used. The default is NULL, meaning the full length of the data is used.

Usage Notes

If the given data sample is NULL, then the procedure returns a result table with no rows.

Examples

Create a table that contains a data sample of type VARCHAR2.

```
CREATE TABLE src_data (sample VARCHAR2(40));
INSERT INTO src_data VALUES ('Hello! How are you?');
```

Use the DETECT_LANGUAGE_CHARSET procedure to analyze the data sample.

```
VAR l_char VARCHAR2(40)
EXEC SELECT sample INTO :l_char FROM src_data;


VAR l_cur REFCURSOR

DECLARE
  l_result     UTL_I18N.LANGUAGE_CHARSET_RESULT_TABLE;
BEGIN
  UTL_I18N.DETECT_LANGUAGE_CHARSET(l_result, :l_char, 1, 5);
  OPEN :l_cur FOR SELECT * FROM TABLE(l_result);
END;
/

PRINT l_cur
```

The result table displays the top five likeliest Oracle languages and Oracle character sets for the given data sample.

```
LANGUAGE                       CHARSET                          SCORE
------------------------------ ------------------------------ ----------
ENGLISH                        US7ASCII                        .245023
ITALIAN                        US7ASCII                       .0873574
FRENCH                         US7ASCII                       .0741715
INDONESIAN                     US7ASCII                       .0534795
POLISH                         EE8ISO8859P2                   .0380973
```

287.5.4 ESCAPE_REFERENCE Function

This function converts a text string to its character reference counterparts for characters that fall outside the character set used by the current document.

Character references are mainly used in HTML and XML documents to represent characters independently of the encoding of the document. Character references may appear in two forms, numeric character references and character entity references. Numeric character references specify the Unicode code point value of a character, while character entity references use symbolic names to refer to the same character. For example, `å` is the numeric character reference for the small letter "a" with a ring above, whereas `å` is the character entity reference for the same character. Character entity references are also used to escape special characters, as an example, `<` represents the < (less than) sign. This is to avoid possible confusion with the beginning of a tag in Markup languages.

Syntax

```
UTL_I18N.ESCAPE_REFERENCE(
    str            IN VARCHAR2 CHARACTER SET ANY_CS,
    page_cs_name   IN VARCHAR2 DEFAULT NULL)
 RETURN VARCHAR2 CHARACTER SET str%CHARSET;
```

Parameters

Table 287-9 ESCAPE_REFERENCE Function Parameters

Parameter	Description
`str`	Specifies the input string
`page_cs_name`	Specifies the character set of the document. If `page_cs_name` is `NULL`, then the database character set is used for `CHAR` data and the national character set is used for `NCHAR` data.

Usage Notes

If the user specifies an invalid character set or a `NULL` string, then the function returns a `NULL` string.

Examples

```
UTL_I18N.ESCAPE_REFERENCE('hello < '||chr(229),'us7ascii')
```

This returns `'hello < å'`.

287.5.5 GET_COMMON_TIME_ZONES Function

This function returns a listing of the most coemmonly used time zones. This list contains a subset of the time zones that are supported in the database.

Syntax

```
UTL_I18N.GET_COMMON_TIME_ZONES
 RETURN STRING_ARRAY;
```

Examples

Returns the list of the most commonly used time zones.

```
DECLARE
   retval UTL_I18N.STRING_ARRAY;
BEGIN
   retval := UTL_I18N.GET_COMMON_TIME_ZONES;
END;
/
```

287.5.6 GET_DEFAULT_CHARSET Function

This function returns the default Oracle character set name or the default e-mail safe character set name from an Oracle language name.

> **See Also:**
>
> "MAP_CHARSET Function" for an explanation of an e-mail safe character set

Syntax

```
UTL_I18N.GET_DEFAULT_CHARSET(
   language  IN VARCHAR2,
   context   IN PLS_INTEGER DEFAULT GENERIC_CONTEXT,
   iswindows IN BOOLEAN DEFAULT FALSE)
 RETURN VARCHAR2;
```

Parameters

Table 287-10 GET_DEFAULT_CHARSET Function Parameters

Parameter	Description
`language`	Specifies a valid Oracle language
`context`	`GENERIC_CONTEXT | MAIL_CONTEXT` `GENERIC_CONTEXT`: Returns the default character set for general cases `MAIL_CONTEXT`: Returns the default e-mail safe character set name
`iswindows`	If `context` is set as `MAIL_CONTEXT`, then `iswindows` should be set to `TRUE` if the platform is Windows and `FALSE` if the platform is not Windows. The default is `FALSE`. `iswindows` has no effect if `context` is set as `GENERIC_CONTEXT`.

Usage Notes

If the user specifies an invalid language name or an invalid flag, then the function returns a `NULL` string.

Examples

GENERIC_CONTEXT, iswindows=FALSE

```
UTL_I18N.GET_DEFAULT_CHARSET('French', UTL_I18N.GENERIC_CONTEXT, FALSE)
```

This returns `'WE8ISO8859P1'`.

MAIL_CONTEXT, iswindows=TRUE

```
UTL_I18N.GET_DEFAULT_CHARSET('French', UTL_I18N.MAIL_CONTEXT, TRUE)
```

This returns 'WE8MSWIN1252'.

MAIL_CONTEXT, iswindows=FALSE

```
UTL_I18N.GET_DEFAULT_CHARSET('French', UTL_I18N.MAIL_CONTEXT, FALSE)
```

This returns 'WE8ISO8859P1'.

287.5.7 GET_DEFAULT_ISO_CURRENCY Function

This function returns the default ISO 4217 currency code for the specified territory.

Syntax

```
UTL_I18N.GET_DEFAULT_ISO_CURRENCY (
   territory    IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN VARCHAR2;
```

Parameters

Table 287-11 GET_DEFAULT_ISO_CURRENCY Function Parameters

Parameter	Description
`territory`	Specifies a valid Oracle territory. It is case-insensitive.

Usage Notes

If the user specifies an invalid territory name, then the function returns a `NULL` string.

Examples

Displays the default ISO currency code for China.

```
DECLARE
   retval VARCHAR2(50);
BEGIN
   retval := UTL_I18N.GET_DEFAULT_ISO_CURRENCY('CHINA');
   DBMS_OUTPUT.PUT_LINE(retval);
END;
/
```

287.5.8 GET_DEFAULT_LINGUISTIC_SORT Function

This function returns the most commonly used Oracle linguistic sort for the specified language.

Syntax

```
UTL_I18N.GET_DEFAULT_LINGUISTIC_SORT (
   language  IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN VARCHAR2;
```

Parameters

Table 287-12 GET_DEFAULT_LINGUISTIC_SORT Function Parameters

Parameter	Description
`language`	Specifies a valid Oracle language. It is case-insensitive.

Usage Notes

If the user specifies an invalid language name, then the function returns a `NULL` string.

Examples

Displays the name of the most appropriate linguistic sort name for the language used in the current SQL session.

```
DECLARE
   retval VARCHAR2(50);
BEGIN
  SELECT value INTO retval FROM nls_database_parameters
   WHERE parameter = 'NLS_LANGUAGE';
   retval := UTL_I18N.GET_DEFAULT_LINGUISTIC_SORT(retval);
   DBMS_OUTPUT.PUT_LINE(retval);
END;
/
```

287.5.9 GET_LOCAL_LANGUAGES Function

This function returns the local language names for the specified territory.

Syntax

```
UTL_I18N.GET_LOCAL_LANGUAGES (
   territory    IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN STRING_ARRAY;
```

Parameters

Table 287-13 GET_LOCAL_LANGUAGES Function Parameters

Parameter	Description
`territory`	Specifies a valid Oracle territory. It is case-insensitive.

Usage Notes

If the user specifies an invalid territory name, then the function returns a `NULL` string.

Examples

Returns the list of local languages used in Belgium.

```
DECLARE
   retval UTL_I18N.STRING_ARRAY;
   cnt    INTEGER;
BEGIN
   retval   :=  UTL_I18N.GET_LOCAL_LANGUAGES('BELGIUM');
   DBMS_OUTPUT.PUT('Count = ');
   DBMS_OUTPUT.PUT_LINE(retval.LAST);
   cnt := retval.FIRST;
   WHILE cnt IS NOT NULL LOOP
      DBMS_OUTPUT.PUT_LINE(retval(cnt));
      cnt := retval.NEXT(cnt);
   END LOOP;
END;
/
...
Count = 2
DUTCH
FRENCH
```

287.5.10 GET_LOCAL_LINGUISTIC_SORTS Function

This function returns a list of the Oracle linguistic sort names that are appropriate for the specified language. A `BINARY` sort is included for all languages.

Syntax

```
UTL_I18N.GET_LOCAL_LINGUISTIC_SORTS (
   language  IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN STRING_ARRAY;
```

Parameters

Table 287-14 GET_LOCAL_LINGUISTIC_SORTS Function Parameters

Parameter	Description
`language`	Specifies a valid Oracle language. It is case-insensitive.

Usage Notes

If the user specifies an invalid language name, then the function returns a `NULL` string.

Examples

Displays the local linguistic sort names for `JAPANESE`.

```
DECLARE
   retval UTL_I18N.STRING_ARRAY;
   cnt INTEGER;
BEGIN
  retval := UTL_I18N.GET_LOCAL_LINGUISTIC_SORTS('Japanese');
```

```
  DBMS_OUTPUT.PUT('Count = ');
  DBMS_OUTPUT.PUT_LINE(retval.COUNT);
  cnt := retval.FIRST;
   WHILE cnt IS NOT NULL LOOP
      DBMS_OUTPUT.PUT_LINE(retval(cnt));
      cnt := retval.NEXT(cnt);
    END LOOP;
END;
/

...
Count = 2
JAPANESE_M
BINARY
```

287.5.11 GET_LOCAL_TERRITORIES Function

This function returns the local territory names for the specified language.

Syntax

```
UTL_I18N.GET_LOCAL_TERRITORIES (
   language  IN VARCHAR2 CHARACTER SET ANY_CS)
 RETURN STRING_ARRAY;
```

Parameters

Table 287-15 GET_LOCAL_TERRITORIES Function Parameters

Parameter	Description
`language`	Specifies a valid Oracle language. It is case-insensitive.

Usage Notes

If the user specifies an invalid language name, then the function returns a `NULL` string.

Examples

Returns the list of Oracle territories that use German as one of their local languages.

```
DECLARE
   retval  UTL_I18N.STRING_ARRAY;
   cnt     INTEGER;
BEGIN
   retval  := UTL_I18N.GET_LCOAL_TERRITORIIES('GERMAN');
   DBMS_OUTPUT.PUT('Count = ');
   DBMS_OUTPUT.PUT_LINE(retval.LAST);
  cnt := retval.FIRST;
   WHILE cnt IS NOT NULL LOOP
      DBMS_OUTPUT.PUT_LINE(retval(cnt));
      cnt := retval.NEXT(cnt));
   END LOOP;
END;
/
...
Count = 4
GERMANY
AUSTRIA
```

```
LUXEMBOURG
SWITZERLAND
```

287.5.12 GET_LOCAL_TIME_ZONES Function

This function returns the local time zone IDs for the specified territory.

Syntax

```
UTL_I18N.GET_LOCAL_TIME_ZONES (
   territory      IN VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL)
RETURN STRING_ARRAY;
```

Parameters

Table 287-16 GET_LOCAL_TIME_ZONES Function Parameters

Parameter	Description
territory	Specifies a valid Oracle territory. It is case-insensitive.

Usage Notes

If the user specifies an invalid territory name, then the function returns a NULL string.

Examples

Creates a function that returns the list of time zones locally used in the territory AZERBAIJAN followed by the general common time zones. This is useful for when the user's territory is known and the application still allows the user to choose other time zones as a user's preference.

```
CREATE OR REPLACE FUNCTION get_time_zones
(territory IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN utl_i18n.string_array
IS
  retval  utl_i18n.string_array;
  retval2 utl_i18n.string_array;
  stpos   INTEGER;
BEGIN
  retval  := utl_i18n.get_local_time_zones(
    territory);
  retval2 := utl_i18n.get_common_time_zones;
  stpos := retval.LAST + 1;
  retval(stpos) := '-----'; -- a separator
  FOR i IN retval2.FIRST..retval2.LAST LOOP
    stpos := stpos + 1;
    retval(stpos) := retval2(i);
  END LOOP;
  RETURN retval;
END;
/
```

Returns the list of local time zones for AZERBAIJAN followed by the common time zones with a separator string of five dashes (-----).

```
DECLARE
   retval UTL_I18N.STRING_ARRAY;
   cnt INTEGER;
BEGIN
```

```
   DBMS_OUTPUT.ENABLE(100000);
   retval UTL_I18N.GET_TIME_ZONES('AZERBAIJAN');
   cnt := retval.FIRST;
  WHILE cnt IS NOT NULL LOOP
    DBMS_OUTPUT.PUT_LINE(retval(cnt));
    cnt := retval.NEXT(cnt);
  END LOOP;
END;
/

Asia/Baku
-----
Pacific/Pago_Pago
Pacific/Honolulu
America/Anchorage
America/Vancouver
America/Los_Angeles
America/Tijuana
America/Edmonton
America/Denver
America/Phoenix
America/Mazatlan
America/Winnipeg
America/Regina
America/Chicago
America/Mexico_City
America/Guatemala
America/El_Salvador
America/Managua
America/Costa_Rica
America/Montreal
...
```

287.5.13 GET_MAX_CHARACTER_SIZE Function

This function returns the maximum character size of a given character set.

Syntax

```
UTL_I18N.GET_MAX_CHARACTER_SIZE(
    charset_name       IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN PLS_INTEGER;
```

Parameters

Table 287-17 GET_MAX_CHARACTER_SIZE Function Parameters

Parameter	Description
charset_name	Specifies a valid character set name. It is case-insensitive.

Usage Notes

For shift-sensitive character sets, the returned maximum character size will include the possible extra shift characters.

Examples

```
UTL_I18N.GET_MAX_CHARACTER_SIZE('AL32UTF8');
```

This returns 4.

287.5.14 GET_TRANSLATION Function

This function returns the translation of the language and territory name in the specified translation language.

Syntax

```
UTL_I18N.GET_TRANSLATION (
   parameter        IN VARCHAR2 CHARACTER SET ANY_CS,
   trans_language   IN VARCHAR2 'AMERICAN',
   flag             IN PLS_INTEGER DEFAULT LANGUAGE_TRANS)
RETURN VARCHAR2 CHARACTER SET parameter%CHARSET;
```

Parameters

Table 287-18 GET_TRANSLATION Function Parameters

Parameter	Description
`parameter`	Specifies a valid language name, territory name, or a combined string in the form of *`language_territory`*. It is case-insensitive.
`trans_language`	Specifies a translation language name. For example, `ITALIAN` is for the Italian language. The default is `AMERICAN`, which indicates American English.
`flag`	Specifies the translation type: • `LANGUAGE_TRANS`: The function returns the language translation. • `TERRITORY_TRANS`: The function returns the territory translation. • `LANGUAGE_TERRITORY_TRANS`: The function returns the language and territory translation. The default translation type is `LANGUAGE_TRANS`.

Usage Notes

If `VARCHAR2` is used as a parameter type, the returned translation text can be corrupted due to the conversion to the database character set. Using `NVARCHAR2` as the parameter type will preserve the translation text because Unicode can encode all translated languages.

If the specified translation language is not available or an invalid name is provided, the default "American English" translations are returned. For example, Oracle does not provide `GUJARATI` translations, so the returned translation would be in American English.

Examples

The following returns the names of all the Oracle-supported languages in Italian.

```
DECLARE
   CURSOR c1 IS
     SELECT value FROM V$NLS_VALID_VALUES
     WHERE parameter = 'LANGUAGE'
  ORDER BY value;
  retval NVARCHAR2(100);
BEGIN
   FOR item IN c1 LOOP
     retval := UTL_I18N.GET_TRANSLATION (TO_NCHAR(item.value), 'italian');
   END LOOP;
END;
```

287.5.15 MAP_CHARSET Function

This function maps a character set to another character set.

It maps the following:

- An Oracle character set name to an IANA character set name.
- An IANA character set name to an Oracle character set name.
- An Oracle character set to an e-mail safe character set name.

Syntax

```
UTL_I18N.MAP_CHARSET(
   charset   IN VARCHAR2,
   context   IN PLS_INTEGER DEFAULT GENERIC_CONTEXT,
   flag      IN PLS_INTEGER DEFAULT ORACLE_TO_IANA)
RETURN VARCHAR2;
```

Parameters

Table 287-19 MAP_CHARSET Function Parameters

Parameter	Description
charset	Specifies the character set name to be mapped. The mapping is case-insensitive.
context	GENERIC_CONTEXT | MAIL_CONTEXT GENERIC_CONTEXT: The mapping is between an Oracle character set name and an IANA character set name. This is the default value. MAIL_CONTEXT: The mapping is between an Oracle character set name and an email safe character set name.
flag	• ORACLE_TO_IANA | IANA_TO_ORACLE if GENERIC_CONTEXT is set ORACLE_TO_IANA: Map from an Oracle character set name to an IANA character set name. This is the default. IANA_TO_ORACLE: Map from an IANA character set name to an Oracle character set name. • MAIL_GENERIC | MAIL_WINDOWS if MAIL_CONTEXT is set MAIL_GENERIC: Map from an Oracle character set name to an email safe character set name on a non-Windows platform. MAIL_WINDOWS: Map from an Oracle character set name to an email safe character set name on a Windows platform.

Usage Notes

An e-mail safe character set is an Oracle character set that is commonly used by applications when they submit e-mail messages. The character set is usually used to convert contents in the database character set to e-mail safe contents. To specify the character set name in the mail header, you should use the corresponding IANA character set name obtained by calling the MAP_CHARSET function with the ORACLE_TO_IANA option, providing the e-mail safe character set name as input.

For example, no e-mail client recognizes message contents in the WE8DEC character set, whose corresponding IANA name is DEC-MCS. If WE8DEC is passed to the MAP_CHARSET function with the MAIL_CONTEXT option, then the function returns WE8ISO8859P1. Its corresponding IANA name, ISO-8859-1, is recognized by most e-mail clients.

The steps in this example are as follows:

1. Call the `MAP_CHARSET` function with the `MAIL_CONTEXT | MAIL_GENERIC` option with the database character set name, `WE8DEC`. The result is `WE8ISO8859P1`.
2. Convert the contents stored in the database to `WE8ISO8859P1`.
3. Call the `MAP_CHARSET` function with the `ORACLE_TO_IANA | GENERIC_CONTEXT` option with the e-mail safe character set, `WE8ISO8859P1`. The result is `ISO-8859-1`.
4. Specify `ISO-8859-1` in the mail header when the e-mail message is submitted.

The function returns a character set name if a match is found. If no match is found or if the flag is invalid, then it returns `NULL`.

> **Note:**
>
> Many Oracle character sets can map to one e-mail safe character set. There is no function that maps an e-mail safe character set to an Oracle character set name.

Examples

Generic Context

```
UTL_I18N.MAP_CHARSET('iso-8859-1',UTL_I18N.GENERIC_CONTEXT,UTL_I18N.IANA_TO_ORACLE)
```

This returns `'WE8ISO8859P1'`.

Context

```
UTL_I18N.MAP_CHARSET('WE8DEC', utl_i18n.mail_context,  utl_i18n.mail_generic)
```

This returns `'WE8ISO8859P1'`.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for a list of valid Oracle character sets

287.5.16 MAP_FROM_SHORT_LANGUAGE Function

This function maps an Oracle short language name to an Oracle language name.

Syntax

```
UTL_I18N.MAP_FROM_SHORT_LANGUAGE (
   language           IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN VARCHAR2;
```

Parameters

Table 287-20 MAP_FROM_SHORT_LANGUAGE Function Parameters

Parameter	Description
language	Specifies a valid short language name. It is case-insensitive.

Usage Notes

If the user specifies an invalid language name, then the function returns a NULL string.

Examples

Returns the default linguistic sort name for the customer with the ID of 9000. Note that the table customers is from the oe user in the Common Schema. Because the customer's language preference is stored using a short language name, you need to convert to a full language name by calling the GET_DEFAULT_LINGUISTIC_SORT procedure.

```
DECLARE
   short_n VARCHAR2(10);
   ling_n VARCHAR2(50);
BEGIN
   SELECT nls_language INTO short
   FROM customers WHERE customer_id = 9000;
   ling_n := UTL_I18N.GET_DEFAULT_LINGUISTIC_SORT (
   UTL_I18N.MAP_FROM_SHORT_LANGUAGE(short_n));
   DBMS_OUTPUT.PUT_LINE(ling_n);
END;
/
```

287.5.17 MAP_LANGUAGE_FROM_ISO Function

This function returns an Oracle language name from an ISO locale name.

Syntax

```
UTL_I18N.MAP_LANGUAGE_FROM_ISO(
  isolocale IN VARCHAR2)
RETURN VARCHAR2;
```

Parameters

Table 287-21 MAP_LANGUAGE_FROM_ISO Function Parameters

Parameter	Description
isolocale	Specifies the ISO locale. The mapping is case-insensitive.

Usage Notes

If the user specifies an invalid locale string, then the function returns a NULL string.

If the user specifies a locale string that includes only the language (for example, en_ instead of en_US), then the function returns the default language name for the specified language (for example, American).

Examples

```
UTL_I18N.MAP_LANGUAGE_FROM_ISO('en_US')
```

This returns `'American'`.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for a list of valid Oracle languages

287.5.18 MAP_LOCALE_TO_ISO Function

This function returns an ISO locale name from an Oracle language name and an Oracle territory name.

A valid string must include at least one of the following: a valid Oracle language name or a valid Oracle territory name.

Syntax

```
UTL_I18N.MAP_LOCALE_TO_ISO (
   ora_language   IN VARCHAR2,
   ora_territory  IN VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 287-22 MAP_LOCALE_TO_ISO Function Parameters

Parameter	Description
`ora_language`	Specifies an Oracle language name. It is case-insensitive.
ora_territory	Specifies an Oracle territory name. It is case-insensitive.

Usage Notes

If the user specifies an invalid string, then the function returns a `NULL` string.

Examples

```
UTL_I18N.MAP_LOCALE_TO_ISO('American','America')
```

This returns `'en_US'`.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for a list of valid Oracle languages and territories

287.5.19 MAP_TERRITORY_FROM_ISO Function

This function returns an Oracle territory name from an ISO locale.

Syntax

```
UTL_I18N.MAP_TERRITORY_FROM_ISO (
  isolocale IN VARCHAR2)
 RETURN VARCHAR2;
```

Parameters

Table 287-23 MAP_TERRITORY_FROM_ISO Function Parameters

Parameter	Description
`isolocale`	Specifies the ISO locale. The mapping is case-insensitive.

Usage Notes

If the user specifies an invalid locale string, then the function returns a `NULL` string.

If the user specifies a locale string that includes only the territory (for example, `_fr` instead of `fr_fr`), then the function returns the default territory name for the specified territory (for example, `France`).

Examples

```
UTL_I18N.MAP_TERRITORY_FROM_ISO('en_US')
```

This returns '`America`'.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for a list of valid Oracle territories

287.5.20 MAP_TO_SHORT_LANGUAGE Function

This function maps an Oracle language name to an Oracle short language name.

Syntax

```
UTL_I18N.MAP_TO_SHORT_LANGUAGE (
   language    IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN VARCHAR2;
```

Parameters

Table 287-24 MAP_TO_SHORT_LANGUAGE Function Parameters

Parameter	Description
`language`	Specifies a valid full language name. It is case-insensitive.

Usage Notes

If the user specifies an invalid language name, then the function returns a `NULL` string.

Examples

Returns the short language name for the language.

```
DECLARE  retval VARCHAR2(100);BEGIN  retval :=
UTL_I18N.MAP_TO_SHORT_LANGUAGE('american');  DBMS_OUTPUT.PUT_LINE(retval);END;/US
```

287.5.21 RAW_TO_CHAR Functions

This function converts `RAW` data from a valid Oracle character set to a `VARCHAR2` string in the database character set.

The function is overloaded. The different forms of functionality are described along with the syntax declarations.

Syntax

Buffer Conversion:

```
UTL_I18N.RAW_TO_CHAR(
   data          IN RAW,
   src_charset   IN VARCHAR2 DEFAULT NULL)
 RETURN VARCHAR2;
```

Piecewise conversion converts raw data into character data piece by piece:

```
UTL_I18N.RAW_TO_CHAR (
   data            IN RAW,
   src_charset     IN VARCHAR2 DEFAULT NULL,
   scanned_length  OUT PLS_INTEGER,
   shift_status    IN OUT PLS_INTEGER)
RETURN VARCHAR2;
```

Parameters

Table 287-25 RAW_TO_CHAR Function Parameters

Parameter	Description
`data`	Specifies the `RAW` data to be converted to a `VARCHAR2` string
`src_charset`	Specifies the character set that the `RAW` data was derived from. If src_charset is `NULL`, then the database character set is used.
`scanned_length`	Specifies the number of bytes of source data scanned
`shift_status`	Specifies the shift status at the end of the scan. The user must set it to `SHIFT_IN` the first time it is called in piecewise conversion. **Note:** ISO 2022 character sets use escape sequences instead of shift characters to indicate the encoding method. `shift_status` cannot hold the encoding method information that is provided by the escape sequences for the next function call. As a result, this function cannot be used to reconstruct ISO 2022 character from raw data in a piecewise way unless each unit of input can be guaranteed to be a closed string. A closed string begins and ends in a 7-bit escape state.

Usage Notes

If the user specifies an invalid character set, `NULL` data, or data whose length is 0, then the function returns a `NULL` string.

Examples

Buffer Conversion

```
UTL_I18N.RAW_TO_CHAR(hextoraw('616263646566C2AA'), 'utf8')
```

This returns the following string in the database character set:

```
'abcde'||chr(170)
```

Piecewise Conversion

```
UTL_I18N.RAW_TO_CHAR(hextoraw('616263646566C2AA'),'utf8',shf,slen)
```

This expression returns the following string in the database character set:

```
'abcde'||chr(170)
```

It also sets `shf` to `SHIFT_IN` and `slen` to `8`.

The following example converts data from the Internet piece by piece to the database character set.

```
rvalue RAW(1050);
  nvalue VARCHAR2(1024);
  conversion_state  PLS_INTEGER = 0;
  converted_len   PLS_INTEGER;
  rtemp  RAW(10) = '';
  conn   utl_tcp.connection;
  tlen PLS_INTEGER;

  ...
  conn := utl_tcp.open_connection ( remote_host => 'localhost',
                                    remote_port => 2000);
  LOOP
      tlen := utl_tcp.read_raw(conn, rvalue, 1024);
      rvalue := utl_raw.concat(rtemp, rvalue);
      nvalue := utl_i18n.raw_to_char(rvalue, 'JA16SJIS', converted_len,
conversion_stat);
      if (converted_len < utl_raw.length(rvalue) )
      then
        rtemp := utl_raw.substr(rvalue, converted_len+1);
      else
        rtemp := '';
      end if;
      /* do anything you want with nvalue */
      /* e.g htp.prn(nvalue); */
    END LOOP;
    utl_tcp.close_connection(conn);
  EXCEPTION
    WHEN utl_tcp.end_of_input THEN
      utl_tcp.close_connection(conn);
END;
```

287.5.22 RAW_TO_NCHAR Functions

This function converts `RAW` data from a valid Oracle character set to an `NVARCHAR2` string in the national character set.

The function is overloaded. The different forms of functionality are described along with the syntax declarations.

Syntax

Buffer Conversion:

```
UTL_I18N.RAW_TO_NCHAR (
   data         IN RAW,
   src_charset  IN VARCHAR2 DEFAULT NULL)
 RETURN NVARCHAR2;
```

Piecewise conversion converts raw data into character data piece by piece:

```
UTL_I18N.RAW_TO_NCHAR (
   data            IN RAW,
   src_charset      IN VARCHAR2 DEFAULT NULL,
   scanned_length   OUT PLS_INTEGER,
   shift_status     IN OUT PLS_INTEGER)
 RETURN NVARCHAR2;
```

Parameters

Table 287-26 RAW_TO_NCHAR Function Parameters

Parameter	Description
`data`	Specifies the `RAW` data to be converted to an `NVARCHAR2` string
`src_charset`	Specifies the character set that the `RAW` data was derived from. If `src_charset` is `NULL`, then the database character set is used.
`scanned_length`	Specifies the number of bytes of source data scanned
`shift_status`	Specifies the shift status at the end of the scan. The user must set it to `SHIFT_IN` the first time it is called in piecewise conversion. **Note:** ISO 2022 character sets use escape sequences instead of shift characters to indicate the encoding method. `shift_status` cannot hold the encoding method information that is provided by the escape sequences for the next function call. As a result, this function cannot be used to reconstruct ISO 2022 character from raw data in a piecewise way unless each unit of input can be guaranteed to be a closed string. A closed string begins and ends in a 7-bit escape state.

Usage Notes

If the user specifies an invalid character set, `NULL` data, or data whose length is 0, then the function returns a `NULL` string.

Examples

Buffer Conversion

```
UTL_I18N.RAW_TO_NCHAR(hextoraw('616263646566C2AA'),'utf8')
```

This returns the following string in the national character set:

```
'abcde'||chr(170)
```

Piecewise Conversion

```
UTL_I18N.RAW_TO_NCHAR(hextoraw('616263646566C2AA'),'utf8', shf, slen)
```

This expression returns the following string in the national character set:

```
'abcde'||chr(170)
```

It also sets `shf` to `SHIFT_IN` and `slen` to `8`.

The following example converts data from the Internet piece by piece to the national character set.

```
rvalue RAW(1050);
  nvalue NVARCHAR2(1024);
  converstion_state  PLS_INTEGER = 0;
  converted_len   PLS_INTEGER;
  rtemp  RAW(10) = '';
  conn   utl_tcp.connection;
  tlen PLS_INTEGER;

  ...
  conn := utl_tcp.open_connection ( remote_host => 'localhost',
                                    remote_port => 2000);
  LOOP
      tlen := utl_tcp.read_raw(conn, rvalue, 1024);
      rvalue := utl_raw.concat(rtemp, rvalue);
      nvalue := utl_i18n.raw_to_nchar(rvalue, 'JA16SJIS', converted_len,
conversion_stat);
      if (converted_len < utl_raw.length(rvalue) )
      then
        rtemp := utl_raw.substr(rvalue, converted_len+1);
      else
        rtemp := '';
      end if;
      /* do anything you want with nvalue */
      /* e.g htp.prn(nvalue); */
    END LOOP;
    utl_tcp.close_connection(conn);
  EXCEPTION
    WHEN utl_tcp.end_of_input THEN
      utl_tcp.close_connection(conn);
  END;
```

287.5.23 STRING_TO_RAW Function

This function converts a `VARCHAR2` or `NVARCHAR2` string to another valid Oracle character set and returns the result as `RAW` data.

Syntax

```
UTL_I18N.STRING_TO_RAW(
   data          IN VARCHAR2 CHARACTER SET ANY_CS,
   dst_charset   IN VARCHAR2 DEFAULT NULL)
RETURN RAW;
```

Parameters

Table 287-27 STRING_TO_RAW Function Parameters

Parameter	Description
data	Specifies the VARCHAR2 or NVARCHAR2 string to convert.
dst_charset	Specifies the destination character set. If dst_charset is NULL, then the database character set is used for CHAR data and the national character set is used for NCHAR data.

Usage Notes

If the user specifies an invalid character set, a NULL string, or a string whose length is 0, then the function returns a NULL string.

Examples

```
DECLARE
    r raw(50);
    s varchar2(20);
  BEGIN
    s:='abcdef'||chr(170);
    r:=utl_i18n.string_to_raw(s,'utf8');
    dbms_output.put_line(rawtohex(r));
  end;
/
```

This returns a hex value of '616263646566C2AA'.

287.5.24 TRANSLITERATE Function

This function performs script transliteration.

Syntax

```
UTL_I18N.TRANSLITERATE (
  data IN VARCHAR2 CHARACTER SET ANY_CS,
  name IN VARCHAR2)
RETURN VARCHAR2 CHARACTER SET data%CHARSET;
```

Parameters

Table 287-28 TRANSLITERATE Function Parameters

Parameter	Description
data	Specifies the data to be converted. Either CHAR or NCHAR datatype can be specified.
name	Specifies the transliteration name string. For a list of valid names, see Table 287-29.

Constants

Table 287-29 TRANSLITERATE Function Constants

Constant Name	Value	Description
ARABIC_LATIN_ISO233	'arabic_latin_iso233'	Converts Arabic characters to Latin characters.
CYR_ASCII_ICAO9303	'cyr_ascii_ICAO9303'	Converts Cyrillic characters to ASCII characters.
CYR_BG_ASCII_ICAO9303	'cyr_bg_ascii_ICAO9303'	Converts Cyrillic Bulgaria characters to ASCII characters.
CYR_BY_ASCII_ICAO9303	'cyr_by_ascii_ICAO9303'	Converts Cyrillic Belarus characters to ASCII characters.
CYR_MK_ASCII_ICAO9303	'cyr_mk_ascii_ICAO9303'	Converts Cyrillic North Macedonia characters to ASCII characters.
CYR_RS_ASCII_ICAO9303	'cyr_rs_ascii_ICAO9303'	Converts Cyrillic Serbia characters to ASCII characters.
CYR_UA_ASCII_ICAO9303	'cyr_ua_ascii_ICAO9303'	Converts Cyrillic Ukraine characters to ASCII characters.
CYRILLIC_LATIN_ISO9	'cyrillic_latin_iso9'	Converts Cyrillic characters to Latin characters.
FWKATAKANA_HIRAGANA	'fwkatakana_hiragana'	Converts only fullwidth Katakana characters to fullwidth Hiragana characters.
FWKATAKANA_HWKATAKANA	'fwkatakana_hwkatakana'	Converts only fullwidth Katakana characters to halfwidth Katakana characters.
GREEK_LATIN_ISO843	'greek_latin_iso843'	Converts Greek characters to Latin characters.
HEBREW_LATIN_ISO259	'hebrew_latin_iso259'	Converts Hebrew characters to Latin characters.
HIRAGANA_FWKATAKANA	'hiragana_fwkatakana'	Converts only fullwidth Hiragana characters to fullwidth Katakana characters.
HIRAGANA_HWKATAKANA	'hiragana_hwkatakana'	Converts only fullwidth Hiragana characters to halfwidth Katakana characters.
HWKATAKANA_FWKATAKANA	'hwkatakana_fwkatakana'	Converts only halfwidth Katakana characters to fullwidth Katakana characters.
HWKATAKANA_HIRAGANA	'hwkatakana_hiragana'	Converts only halfwidth Katakana characters to fullwidth Hiragana characters.
KANA_FWKATAKANA	'kana_fwkatakana'	Converts any type of Kana character to a fullwidth Katakana character.
KANA_HIRAGANA	'kana_hiragana'	Converts any type of Kana character to a fullwidth Hiragana character.
KANA_HWKATAKANA	'kana_hwkatakana'	Converts any type of Kana character to a halfwidth Katakana character.
LATIN_ASCII_DIN91379	'latin_ascii_din91379'	Converts Latin characters to ASCII characters.
MODERN_HEBREW_LATIN_ISO259_2	'modern_hebrew_latin_iso259_2'	Converts Modern Hebrew characters to Latin Characters.

Usage Notes

- The user must provide the input data for transliteration in NFC form.

- The function returns the converted string.

Examples

Given a table `japanese_emp`, containing an `NVARCHAR2` column `ename`, the following statement can be used to normalize all the kana names in `ename` to hiragana:

```
UPDATE japanese_emp
   SET ename = UTL_I18N.TRANSLITERATE (ename, 'kana_hiragana');
```

The following figure shows how this output might look.

Figure 287-1 Loading Locale-Specific Data to the Database

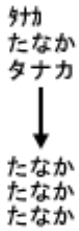

The following statement normalizes one kana name to hiragana:

```
DECLARE
   Name  japanese_emp.ename%TYPE;
   Eno   CONSTANT  NUMBER(4) := 1;
BEGIN
   SELECT ename INTO name FROM japanese_emp WHERE enumber = eno;
   name := UTL_I18N.TRANSLITERATE(name, UTL_I18N.KANA_HIRAGANA);
   UPDATE japanese_emp SET ename = name WHERE enumber = eno;
EXCEPTION
   WHEN  UTL_I18N.UNSUPPORTED_TRANSLITERATION THEN
   DBMS_OUTPUT.PUT_LINE('transliteration not supported');
END;
/
```

287.5.25 UNESCAPE_REFERENCE Function

This function returns a string from an input string that contains character references. It decodes each character reference to the corresponding character value.

> **See Also:**
>
> "ESCAPE_REFERENCE Function" for more information about escape sequences

Syntax

```
UTL_I18N.UNESCAPE_REFERENCE (
   str IN VARCHAR2 CHARACTER SET ANY_CS)
 RETURN VARCHAR2 CHARACTER SET str%CHARSET;
```

Parameters

Table 287-30 UNESCAPE_REFERENCE Function Parameters

Parameter	Description
`str`	Specifies the input string

Usage Notes

If the user specifies a `NULL` string or a string whose length is 0, then the function returns a `NULL` string. If the function fails, then it returns the original string.

Examples

```
UTL_I18N.UNESCAPE_REFERENCE('hello &lt; &#xe5;')
```

This returns `'hello <'||chr(229).`

287.5.26 VALIDATE_CHARACTER_ENCODING Functions

This function validates the character encoding of VARCHAR2, NVARCHAR2, CLOB, and NCLOB data. The validation is based on the database character set for VARCHAR2 and CLOB data and national character set for NVARCHAR2 and NCLOB data.

For Unicode character sets, such as AL32UTF8, AL16UTF16, AL16UTF16LE, UTF8, and UTFE, any byte sequences mapped to the following Unicode code points are considered invalid:

- Unpaired surrogate code point
- Non-character code point

In addition, any irregular or illegal UTF-8 byte sequence is considered invalid for AL32UTF8 and UTF8 character sets.

The `VALIDATE_CHARACTER_ENCODING` function is overloaded. One function is for validating VARCHAR2 and NVARCHAR2 data, while the other function is for validating CLOB and NCLOB data.

- **Validating VARCHAR2 and NVARCHAR2 data**

  A VARCHAR2 or NVARCHAR2 byte or its byte sequence is considered invalid for a character set, if it does not map to any of the characters defined in the character set.

- **Validating CLOB and NCLOB data**

  A LOB character is considered invalid for a character set if a byte (in case of a single-byte database character set) or a byte pair (in case of UTF-16 encoding used with a multibyte database character set) corresponding to the encoding of the LOB character does not map to any of the characters defined in the character set.

Syntax

This function validates VARCHAR2 and NVARCHAR2 data:

```
UTL_I18N.VALIDATE_CHARACTER_ENCODING (
   data IN VARCHAR2 CHARACTER SET ANY_CS)
RETURN PLS_INTEGER;
```

This function validates CLOB and NCLOB data:

```
UTL_I18N.VALIDATE_CHARACTER_ENCODING (
   lob_loc IN CLOB CHARACTER SET ANY_CS)
RETURN PLS_INTEGER;
```

Parameters

Table 287-31 VALIDATE_CHARACTER_ENCODING Function Parameters

Parameter	Description
`data`	VARCHAR2 or NVARCHAR2 data to validate.
`lob_loc`	CLOB or NCLOB data to validate.

Usage Notes

This function returns the offset of the first invalid byte for the VARCHAR2 or NVARCHAR2 data. It returns the offset of the first invalid character for the CLOB or NCLOB data. It returns 0, if all the bytes in the character data are valid. It returns `NULL`, if the value of the parameter `data` or `lob_loc` is `NULL`.

Examples

This example validates the character encoding of NVARCHAR2 and CLOB data where the database character set is AL32UTF8 while the national character set is AL16UTF16.

```
CREATE TABLE temp(col1 NVARCHAR2(20), col2 CLOB);
INSERT INTO temp VALUES(UNISTR('foo\D800bar'), UNISTR('foo\D800bar'));
COMMIT;
SELECT UTL_I18N.VALIDATE_CHARACTER_ENCODING(col1) invalid_offset_column1,
       UTL_I18N.VALIDATE_CHARACTER_ENCODING(col2) invalid_offset_column2
FROM temp;
```

The query returns:

```
INVALID_OFFSET_COLUMN1 INVALID_OFFSET_COLUMN2
---------------------- ----------------------
                     7                      4
```

Here, the surrogate code point `U+D800` is invalid. The number 7 is returned as `INVALID_OFFSET_COLUMN1`, because for `col1`, `'foo'` is encoded in 6 bytes in NVARCHAR2 and the invalid code point `U+D800` starts at offset 7. The number 4 is returned as `INVALID_OFFSET_COLUMN2`, because for `col2`, `'foo'` is encoded in 3 UTF-16 code points in CLOB and the invalid code point `U+D800` starts at offset 4.

288
UTL_IDENT

The `UTL_IDENT` package specifies which Database or client PL/SQL is running.

This chapter contains the following topics:

- Overview
- Security Model
- Constants

288.1 UTL_IDENT Overview

The `UTL_IDENT` package is intended for use for conditional compilation of PL/SQL packages that are supported by Oracle, TimesTen Database, and clients such as Oracle Forms.

288.2 UTL_IDENT Security Model

The `UTL_IDENT` package runs as the package owner `SYS`. The public synonym `UTL_IDENT`, and `EXECUTE` permission on this package is granted to `PUBLIC`.

288.3 UTL_IDENT Constants

The `UTL_IDENT` package defines several constants to use when specifying parameter values.

These constants are shown in the following table.

Table 288-1 UTL_IDENT Constants

Constant	Type	Value	Description
IS_ORACLE_SERVER	BOOLEAN	TRUE/FALSE	Stipulates if Oracle Server or not
IS_ORACLE_CLIENT	BOOLEAN	TRUE/FALSE	Stipulates if Oracle Client or not
IS_TIMESTEN	BOOLEAN	TRUE/FALSE	Stipulates if TimesTen or not
IS_ORACLE_FORMS	BOOLEAN	TRUE/FALSE	Stipulates if Oracle Forms or not

289
UTL_INADDR

The `UTL_INADDR` package provides a PL/SQL procedures to support internet addressing. It provides an API to retrieve host names and IP addresses of local and remote hosts.

This chapter contains the following topics:

- Security Model
- Exceptions
- Examples
- Summary of UTL_INADDR Subprograms

289.1 UTL_INADDR Security Model

This package is an invoker's rights package, which means that the invoking user must be granted the `connect` privilege in the access control list assigned to the remote network host to which he or she wishes to connect.

> **Note:**
>
> For more information about managing fine-grained access, see *Oracle Database Security Guide*.

289.2 UTL_INADDR Exceptions

This table describes cxceptions raised by `UTL_INADDR` subprograms.

Table 289-1 ***Exception - Internet Address Package***

Number	Exception	Description
ORA-24247	NETWORK_ACCESS_DENIED	Access to network is denied.
ORA-29257	UNKNOWN_HOST	The host is unknown.

289.3 UTL_INADDR Examples

This UTL_INADDR example retrieves the local host name and IP address.

```
SET serveroutput on
BEGIN
  DBMS_OUTPUT.PUT_LINE(UTL_INADDR.GET_HOST_NAME);  -- get local host name
  DBMS_OUTPUT.PUT_LINE(UTL_INADDR.GET_HOST_ADDRESS);  -- get local IP addr
END;
/
```

289.4 Summary of UTL_INADDR Subprograms

This table lists the `UTL_INADDR` subprograms and briefly describes them.

Table 289-2 UTL_INADDR Package Subprograms

Subprogram	Description
GET_HOST_ADDRESS Function	Retrieves the IP address of the local or remote host given its name
GET_HOST_NAME Function	Retrieves the name of the local or remote host given its IP address

289.4.1 GET_HOST_ADDRESS Function

This function retrieves the IP address of the specified host.

Syntax

```
UTL_INADDR.GET_HOST_ADDRESS (
   host  IN VARCHAR2 DEFAULT NULL)
RETURN host_address VARCHAR2;
```

Parameters

Table 289-3 GET_HOST_ADDRESS Function Parameters

Parameter	Description
`host`	The name of the host to retrieve the IP address.

Return Values

Table 289-4 GET_HOST_ADDRESS Function Return Values

Parameter	Description
`host_address`	The IP address of the specified host, or that of the local host if host is `NULL`.

Exceptions

`UNKNOWN_HOST`: The specified IP address is unknown

Usage Notes

The permission to obtain the host name or IP address of the current host is controlled by the `resolve` privilege on `LOCALHOST`.

289.4.2 GET_HOST_NAME Function

This function retrieves the name of the local or remote host given its IP address.

Syntax

```
UTL_INADDR.GET_HOST_NAME (
   ip  IN VARCHAR2 DEFAULT NULL)
RETURN host_name VARCHAR2;
```

Parameters

Table 289-5 GET_HOST_NAME Function Parameters

Parameter	Description
ip	The IP address of the host used to determine its host name. If ip is not NULL, the official name of the host with its domain name is returned. If this is NULL, the name of the local host is returned and the name does not contain the domain to which the local host belongs.

Return Values

Table 289-6 GET_HOST_NAME Function Return Values

Parameter	Description
host_name	The name of the local or remote host of the specified IP address.

Exceptions

UNKNOWN_HOST: The specified IP address is unknown

Usage Notes

The permission to obtain the host name or IP address of the current host is controlled by the resolve privilege granted through DBMS_NETWORK_ACL_ADMIN on LOCALHOST.

290
UTL_LMS

`UTL_LMS` retrieves and formats error messages in different languages.

This chapter contains the following topics:

- Security Model
- Summary of UTL_LMS Subprograms

> **See Also:**
>
> *Oracle Database Globalization Support Guide*

290.1 UTL_LMS Security Model

This package must be created as the user `SYS`.

290.2 Summary of UTL_LMS Subprograms

This table lists the `UTL_LMS` subprograms and briefly describes them.

Table 290-1 UTL_LMS Package Subprograms

Function	Description
FORMAT_MESSAGE Function	Formats a retrieved error message
GET_MESSAGE Function	Retrieves an error message based on error number, product, facility, language, and message specified

290.2.1 FORMAT_MESSAGE Function

This function formats a message retrieved by the `GET_MESSAGE` function and returns the formatted message. If the function fails, then it returns a `NULL` result.

The following table shows special characters that can be used in the format string.

Special Character	Description
`'%s'`	Substitute the next string argument
`'%d'`	Substitute the next integer argument
`'%%'`	Represents the special character `%`

Syntax

```
UTL_LMS.FORMAT_MESSAGE (
   format IN VARCHAR2 CHARACTER SET ANY_CS,
   args   IN VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL)
 RETURN VARCHAR2 CHARACTER SET format%CHARSET;
```

Parameters

Table 290-2 FORMAT_MESSAGE Procedure Parameters

Parameter	Description
`format`	Specifies the string to format
`args`	Specifies the list of arguments

Examples

```
DECLARE
   s varchar2(200);
   i  pls_integer;
BEGIN
   i:= utl_lms.get_message(26052, 'rdbms', 'ora', 'french', s);
   dbms_output.put_line('before format, message is: '||s);
   dbms_output.put_line('formatted message is: '||utl_lms.format_message(s, 9,
'my_column_name');
END;
/
```

The following is an unformatted message:

```
Type %d non pris en charge pour l'expression SQL sur la colonne %s.
```

The following is the formatted message:

```
Type 9 non pris en charge pour l'expression SQL sur la colonne my_column_name.
```

290.2.2 GET_MESSAGE Function

This function retrieves an Oracle error message. The user can define user-specific error messages with the `lmsgen` utility.

It returns `0` when it is successful. It returns `-1` when it fails.

> **See Also:**
>
> *Oracle Database Globalization Support Guide* for more information about the `lmsgen` utility

Syntax

```
UTL_LMS.GET_MESSAGE (
   errnum    IN PLS_INTEGER,
   product   IN VARCHAR2,
   facility  IN VARCHAR2,
```

```
   language  IN VARCHAR2,
   message   OUT NOCOPY VARCHAR2CHARCTER SET ANY_CS)
RETURN PLS_INTEGER;
```

Parameters

Table 290-3 GET_MESSAGE Function Parameters

Parameter	Description
`errnum`	Specifies the error number. Example: `'972'` (for `ORA-00972`)
`product`	Specifies the product to which the error message applies Example: `'rdbms'`
`facility`	Specifies the error message prefix Example: `'ora'`
`language`	Specifies the language of the message. The parameter is case-insensitive. The default is `NULL`, which causes `GET_MESSAGE` to use the value of the `NLS_LANGUAGE` session parameter.
`message`	Specifies the output buffer for the retrieved message

Usage Notes

If the `language` parameter is set to `NULL`, then the value of the `NLS_LANGUAGE` session parameter is used as the default.

Examples

```
DECLARE
  s varchar2(200);
  i pls_integer;
BEGIN
   i:=utl_lms.get_message(601, 'rdbms', 'oci', 'french', s);
   dbms_output.put_line('OCI--00601 is: '||s);
END
/
```

The following output results:

```
OCI--00601 is: Echec du processus de nettoyage.
```

291
UTL_MAIL

The UTL_MAIL package is a utility for managing email which includes commonly used email features, such as attachments, CC, and BCC.

This chapter contains the following topics:

- Security Model
- Operational Notes
- Rules and Limits
- Summary of UTL_MAIL Subprograms

291.1 UTL_MAIL Security Model

UTL_MAIL is not installed by default because of the SMTP_OUT_SERVER configuration requirement and the security exposure this involves. In installing UTL_MAIL, you should take steps to prevent the port defined by SMTP_OUT_SERVER being swamped by data transmissions.

This package is now an invoker's rights package and the invoking user will need the connect privilege granted in the access control list assigned to the remote network host to which he wants to connect.

> **Note:**
>
> For more information about managing fine-grained access, see *Oracle Database Security Guide*.

291.2 UTL_MAIL Operational Notes

You must both install UTL_MAIL and define the SMTP_OUT_SERVER.

- To install UTL_MAIL:

```
sqlplus sys/<pwd>
SQL> @$ORACLE_HOME/rdbms/admin/utlmail.sql
SQL> @$ORACLE_HOME/rdbms/admin/prvtmail.plb
```

- You define the SMTP_OUT_SERVER parameter in the init.ora rdbms initialization file. However, if SMTP_OUT_SERVER is not defined, this invokes a default of DB_DOMAIN which is guaranteed to be defined to perform appropriately.

291.3 UTL_MAIL Rules and Limits

Use UTL_MAIL only within the context of the ASCII (American Standard Code for Information Interchange) and EBCDIC (Extended Binary-Coded Decimal Interchange Code) codes.

291.4 Summary of UTL_MAIL Subprograms

This table lists the `UTL_MAIL` subprograms and briefly describes them.

Table 291-1 UTL_MAIL Package Subprograms

Subprogram	Description
SEND Procedure	Packages an email message into the appropriate format, locates SMTP information, and delivers the message to the SMTP server for forwarding to the recipients
SEND_ATTACH_RAW Procedure	Represents the `SEND` Procedure overloaded for `RAW` attachments
SEND_ATTACH_VARCHAR2 Procedure	Represents the `SEND` Procedure overloaded for `VARCHAR2` attachments

291.4.1 SEND Procedure

This procedure packages an email message into the appropriate format, locates SMTP information, and delivers the message to the SMTP server for forwarding to the recipients.

It hides the SMTP API and exposes a one-line email facility for ease of use.

Syntax

```
UTL_MAIL.SEND (
   sender      IN    VARCHAR2 CHARACTER SET ANY_CS,
   recipients  IN    VARCHAR2 CHARACTER SET ANY_CS,
   cc          IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   bcc         IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   subject     IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   message     IN    VARCHAR2 CHARACTER SET ANY_CS,
   mime_type   IN    VARCHAR2 DEFAULT 'text/plain; charset=us-ascii',
   priority    IN    PLS_INTEGER DEFAULT 3,
   replyto     IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL);
```

Parameters

Table 291-2 SEND Procedure Parameters

Parameter	Description
`sender`	Email address of the sender
`recipients`	Email addresses of the recipient(s), separated by commas
`cc`	Email addresses of the CC recipient(s), separated by commas, default is `NULL`
`bcc`	Email addresses of the BCC recipient(s), separated by commas, default is `NULL`
`subject`	String to be included as email subject string, default is `NULL`
`message`	Text message body
`mime_type`	Mime type of the message, default is `'text/plain; charset=us-ascii'`

Table 291-2 (Cont.) SEND Procedure Parameters

Parameter	Description
priority	Message priority, which maps to the X-priority field. 1 is the highest priority and 5 the lowest. The default is 3.
replyto	Defines to whom the reply email is to be sent

291.4.2 SEND_ATTACH_RAW Procedure

This procedure is the SEND Procedure overloaded for RAW attachments.

Syntax

```
UTL_MAIL.SEND_ATTACH_RAW (
   sender           IN    VARCHAR2 CHARACTER SET ANY_CS,
   recipients       IN    VARCHAR2 CHARACTER SET ANY_CS,
   cc               IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   bcc              IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   subject          IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   message          IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   mime_type        IN    VARCHAR2 DEFAULT CHARACTER SET ANY_CS
                              DEFAULT 'text/plain; charset=us-ascii',
   priority         IN    PLS_INTEGER DEFAULT 3,
   attachment       IN    RAW,
   att_inline       IN    BOOLEAN DEFAULT TRUE,
   att_mime_type    IN    VARCHAR2 CHARACTER SET ANY_CS
                                       DEFAULT 'text/plain; charset=us-ascii',
   att_filename     IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   replyto          IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL);
```

Parameters

Table 291-3 *SEND_ATTACH_RAW Procedure Parameters*

Parameter	Description
sender	Email address of the sender
recipients	Email addresses of the recipient(s), separated by commas
cc	Email addresses of the CC recipient(s), separated by commas, default is NULL
bcc	Email addresses of the BCC recipient(s), separated by commas, default is NULL
subject	String to be included as email subject string, default is NULL
message	Text message body
mime_type	Mime type of the message, default is 'text/plain; charset=us-ascii'
priority	Message priority, which maps to the X-priority field. 1 is the highest priority and 5 the lowest. The default is 3.
attachment	RAW attachment
att_inline	Specifies whether the attachment is viewable inline with the message body, default is TRUE

Table 291-3 (Cont.) *SEND_ATTACH_RAW Procedure Parameters*

Parameter	Description
att_mime_type	Mime type of the attachment, default is 'application/octet'
att_filename	String specifying a filename containing the attachment, default is NULL
replyto	Defines to whom the reply email is to be sent

291.4.3 SEND_ATTACH_VARCHAR2 Procedure

This procedure is the SEND Procedure overloaded for VARCHAR2 attachments.

Syntax

```
UTL_MAIL.SEND_ATTACH_VARCHAR2 (
   sender            IN    VARCHAR2 CHARACTER SET ANY_CS,
   recipients        IN    VARCHAR2 CHARACTER SET ANY_CS,
   cc                IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   bcc               IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   subject           IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   message           IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL,
   mime_type         IN    VARCHAR2 CHARACTER SET ANY_CS
                                    DEFAULT 'text/plain; charset=us-ascii',
   priority          IN    PLS_INTEGER DEFAULT 3,
   attachment        IN    VARCHAR2 CHARACTER SET ANY_CS, ,
   att_inline        IN    BOOLEAN DEFAULT TRUE,
   att_mime_type     IN    VARCHAR2 CHARACTER SET ANY_CS
                                    DEFAULT 'text/plain; charset=us-ascii,
   att_filename      IN    VARCHAR2CHARACTER SET ANY_CS DEFAULT NULL,
   replyto           IN    VARCHAR2 CHARACTER SET ANY_CS DEFAULT NULL);
```

Parameters

Table 291-4 SEND_ATTACH_VARCHAR2 Procedure Parameters

Parameter	Description
sender	Email address of the sender
recipients	Email addresses of the recipient(s), separated by commas
cc	Email addresses of the CC recipient(s), separated by commas, default is NULL
bcc	Email addresses of the BCC recipient(s), separated by commas, default is NULL
subject	String to be included as email subject string, default is NULL
message	Text message body
mime_type	Mime type of the message, default is 'text/plain; charset=us-ascii
priority	Message priority, which maps to the X-priority field. 1 is the highest priority and 5 the lowest. The default is 3.
attachment	Text attachment
att_inline	Specifies whether the attachment is inline, default TRUE

Table 291-4 (Cont.) SEND_ATTACH_VARCHAR2 Procedure Parameters

Parameter	Description
`att_mime_type`	Mime type of the attachment, default is `'text/plain; charset=us-ascii'`
`att_filename`	String specifying a filename containing the attachment, default is `NULL`
`replyto`	Defines to whom the reply email is to be sent

292
UTL_MATCH

The `UTL_MATCH` package facilitates matching two records. This is typically used to match names, such as two First Names or two Last Names.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of UTL_MATCH Subprograms

292.1 UTL_MATCH Overview

`UTL_MATCH` can use either the Edit Distance algorithm or Jaro-Winkler algorithm when determining matches.

Edit Distance, also known as Levenshtein Distance (named after the Russian scientist Vladimir Levenshtein, who devised the algorithm in 1965), is a measure of similarity between two strings, `s1` and `s2`. The distance is the number of insertions, deletions or substitutions required to transform s1 to s2.

The Edit Distance between strings `shackleford` and `shackelford` = 2.

The "Jaro-Winkler algorithm" is another way of calculating Edit distance between two strings. This method, developed at the U.S. Census, is a String Comparator measure that gives values of partial agreement between two strings. The string comparator accounts for length of strings and partially accounts for typical human errors made in alphanumeric strings.

The following table shows similarity values returned by Jaro-Winkler and Edit Distance

Table 292-1 Comparison between normalized values returned by Jaro-Winkler and Edit Distance algorithms

String 1	String 2	Jaro Winkler	Edit Distance
Dunningham	Cunnigham	89	80
Abroms	Abrams	92	83
Lampley	Campley	90	86
Marhta	Martha	96	67
Jonathon	Jonathan	95	88
Jeraldine	Geraldine	92	89

292.2 UTL_MATCH Security Model

The `UTL_MATCH` package runs with definer's rights. `UTL_MATCH` must be created under `SYS`. Operations provided by this package are performed with `SYS` privileges.

292.3 Summary of UTL_MATCH Subprograms

This table lists the `UTL_MATCH` subprograms and briefly describes them.

Table 292-2 UTL_MATCH Package Subprograms

Subprogram	Description
EDIT_DISTANCE Function	Calculates the number of changes required to transform string-1 into string-2
EDIT_DISTANCE_SIMILARITY Function	Calculates the number of changes required to transform string-1 into string-2, returning a value between 0 (no match) and 100 (perfect match)
JARO_WINKLER Function	Calculates the measure of agreement between string-1 and string-2
JARO_WINKLER_SIMILARITY Function	Calculates the measure of agreement between string-1 and string-2, returning a value between 0 (no match) and 100 (perfect match)

292.3.1 EDIT_DISTANCE Function

This function calculates the number of insertions, deletions or substitutions required to transform string-1 into string-2.

Syntax

```
UTL_MATCH.EDIT_DISTANCE (
   s1  IN  VARCHAR2,
   s2  IN  VARCHAR2)
 RETURN PLS_INTEGER;
```

Parameters

Table 292-3 EDIT_DISTANCE Function Parameters

Parameter	Description
`s1`	The string to be transformed
`s2`	The string into which `s1` is to be transformed

Examples

```
SELECT UTL_MATCH.EDIT_DISTANCE('shackleford', 'shackelford') FROM DUAL;
-------------
returns 2
```

292.3.2 EDIT_DISTANCE_SIMILARITY Function

This function calculates the number of insertions, deletions or substations required to transform string-1 into string-2, and returns the Normalized value of the Edit Distance between two strings.

The value is typically between 0 (no match) and 100 (perfect match).

Syntax

```
UTL_MATCH.EDIT_DISTANCE_SIMILARITY (
   s1  IN  VARCHAR2,
   s2  IN  VARCHAR2)
 RETURN PLS_INTEGER;
```

Parameters

Table 292-4 EDIT_DISTANCE_SIMILARITY Function Parameters

Parameter	Description
s1	The string to be transformed
s2	The string into which s1 is to be transformed

Examples

```
SELECT UTL_MATCH.EDIT_DISTANCE_SIMILARITY('shackleford', 'shackelford') FROM DUAL;
--------------
returns 82
```

292.3.3 JARO_WINKLER Function

This function calculates the measure of agreement between two strings.

Syntax

```
UTL_MATCH.JARO_WINKLER (
   s1  IN  VARCHAR2,
   s2  IN  VARCHAR2)
 RETURN BINARY_DOUBLE;
```

Parameters

Table 292-5 JARO_WINKLER Function Parameters

Parameter	Description
s1	Input
s2	input

Examples

```
SELECT UTL_MATCH.JARO_WINKLER('shackleford', 'shackelford') FROM DUAL;
--------------
returns 9.818E-001
```

292.3.4 JARO_WINKLER_SIMILARITY Function

This function calculates the measure of agreement between two strings, and returns a score between 0 (no match) and 100 (perfect match).

Syntax

```
UTL_MATCH.JARO_WINKLER_SIMILARITY (
   s1  IN  VARCHAR2,
   s2  IN  VARCHAR2)
 RETURN PLS_INTEGER;
```

Parameters

Table 292-6 JARO_WINKLER_SIMILARITY Function Parameters

Parameter	Description
s1	Input
s2	input

Examples

```
SELECT UTL_MATCH.JARO_WINKLER_SIMILARITY('shackleford', 'shackelford') FROM DUAL;
--------------
returns 98
```

293
UTL_NLA

The `UTL_NLA` package exposes a subset of the BLAS and LAPACK (Version 3.0) operations on vectors and matrices represented as `VARRAY`s.

This chapter contains the following topics:

- Overview
- Rules and Limits
- Security Model
- Subprogram Groups
  - BLAS Level 1 (Vector-Vector Operations) Subprograms
  - BLAS Level 2 (Matrix-Vector Operations) Subprograms
  - BLAS Level 3 (Matrix-Matrix Operations) Subprograms
  - LAPACK Driver Routines (Linear Equations) Subprograms
  - LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
- Summary of UTL_NLA Subprograms

293.1 UTL_NLA Overview

The `UTL_NLA` package exposes a subset of the BLAS (Basic Linear Algebra Subprograms) and LAPACK (Linear Algebra PACKage) (Version 3.0) operations on vectors and matrices represented as `VARRAY`s.

Standards

For more information on the BLAS and LAPACK standards see

```
http://www.netlib.org/blas/
http://www.netlib.org/lapack/
```

Required Expertise

Users of this package are expected to have a sound grasp of linear algebra in general and of the BLAS and LAPACK libraries in particular.

Implementation

The mapping between BLAS and LAPACK procedures and their corresponding PL/SQL calls is one-to-one.

- All BLAS functions have the `BLAS_` prefix (for example, the BLAS_ASUM Functions). The subroutines and functions in BLAS are mapped to PL/SQL procedures and functions, respectively.
- All LAPACK functions have the `LAPACK_` prefix (for example, the LAPACK_GBSV Procedures). The subroutines in LAPACK are mapped to PL/SQL procedures. Procedures

that perform the same operation but differ only on the datatype of the arguments have the same overloaded names.

The mapping between BLAS and LAPACK procedure parameters and those of their corresponding PL/SQL subprograms is almost one-to-one.

- Also in the PL/SQL interface for LAPACK, all `/work/` arguments have been removed. The `UTL_NLA` package manages the allocation and de-allocation of all work areas required by the libraries.
- A new optional parameter, `pack`, has been added to the end of each LAPACK procedure that specifies if the matrix has been linearized in the row-major or column-major (default) format.

293.2 UTL_NLA Rules and Limits

Vectors and matrices are stored in `VARRAY`s with a maximum size of one million entries. Given this restriction, `UTL_NLA` vectors can be up to one million entries but matrices need to be of size RxC <= 1,000,000.

293.3 UTL_NLA Security Model

The `UTL_NLA` package is owned by user `SYS` and is installed as part of database installation. Execution privilege on the package is granted to public. The routines in the package are run with invokers' rights (run with the privileges of the current user).

293.4 Subprogram Groups

The `UTL_NLA` package contains subprogram groups for BLAS and LAPACK operations.

- BLAS Level 1 (Vector-Vector Operations) Subprograms
- BLAS Level 2 (Matrix-Vector Operations) Subprograms
- BLAS Level 3 (Matrix-Matrix Operations) Subprograms
- LAPACK Driver Routines (Linear Equations) Subprograms
- LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms

293.4.1 UTL_NLA BLAS Level 1 (Vector-Vector Operations) Subprograms

This table lists and briefly describes the UTL_NLA BLAS Level 1 Vector-Vector Operations subprograms.

Table 293-1 BLAS Level 1 (Vector-Vector Operations) Subprograms

Subprogram	Description
BLAS_ASUM Functions	Computes the sum of the absolute values of the vector components
BLAS_AXPY Procedures	Copies `alpha*X + Y` into vector `Y`
BLAS_COPY Procedures	Copies the contents of vector `X` to vector `Y`
BLAS_DOT Functions	Returns the dot (scalar) product of two vectors `X` and `Y`

Table 293-1 (Cont.) BLAS Level 1 (Vector-Vector Operations) Subprograms

Subprogram	Description
BLAS_IAMAX Functions	Computes the index of the first element of a vector that has the largest absolute value
BLAS_NRM2 Functions	Computes the vector 2-norm (Euclidean norm)
BLAS_ROT Procedures	Returns the plane rotation of points
BLAS_ROTG Procedures	Returns the Givens rotation of points
BLAS_SCAL Procedures	Scales a vector by a constant
BLAS_SWAP Procedures	Swaps the contents of two vectors each of size n

293.4.2 UTL_NLA BLAS Level 2 (Matrix-Vector Operations) Subprograms

This table lists and briefly describes the UTL_NLA BLAS Level 2 Matrix-Vector Operations subprograms.

Table 293-2 BLAS Level 2 (Matrix-Vector Operations) Subprograms

Subprogram	Description
BLAS_GBMV Procedures	Performs the matrix-vector operation `y := alpha*A*x + beta*y` or `y := alpha*A'*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are vectors and `A` is an `m` by `n` band matrix, with `kl` sub-diagonals and `ku` super-diagonals
BLAS_GEMV Procedures	Performs the matrix-vector operations `y := alpha*A*x + beta*y` or `y := alpha*A'*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are vectors and `A` is an `m` by `n` matrix
BLAS_GER Procedures	Performs a rank 1 operation `A := alpha*x*y' + A` where `alpha` is a scalar, `x` is an `m` element vector, `y` is an n element vector and `A` is an `m` by `n` matrix
BLAS_SBMV Procedures	Performs a matrix-vector operation `y := alpha*A*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are n element vectors and `A` is an `n` by `n` symmetric band matrix, with `k` super-diagonals
BLAS_SPMV Procedures	Performs a matrix-vector operation `y := alpha*A*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are `n` element vectors and `A` is an `n` by `n` symmetric matrix, supplied in packed form
BLAS_SPR Procedures	Performs a symmetric rank 1 operation `A := alpha*x*x' + A` where `alpha` is a real scalar, `x` is an `n` element vector, and `A` is an `n` by `n` symmetric matrix, supplied in packed form
BLAS_SPR2 Procedures	Performs a symmetric rank 2 operation `A := alpha*x*y' + alpha*y*x' +A` where `alpha` is a scalar, `x` and `y` are `n` element vectors, and `A` is an `n` by `n` symmetric matrix, supplied in packed form
BLAS_SBMV Procedures	Performs a matrix-vector operation `y := alpha*A*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are n element vectors and `A` is an `n` by `n` symmetric band matrix, with `k` super-diagonals
BLAS_SYMV Procedures	Performs a matrix-vector operation `y := alpha*A*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are `n` element vectors and `A` is an `n` by `n` symmetric matrix

Table 293-2 (Cont.) BLAS Level 2 (Matrix-Vector Operations) Subprograms

Subprogram	Description
BLAS_SYR Procedures	Performs a symmetric rank 1 operation `A := alpha*x*x' + A` where `alpha` is a real scalar, `x` is an `n` element vector, and `A` is an `n` by `n` symmetric matrix
BLAS_SYR2 Procedures	Performs a symmetric rank 2 operation `A := alpha*x*y' + alpha*y*x' + A` where `alpha` is a scalar, `x` and `y` are `n` element vectors, and `A` is an `n` by `n` symmetric matrix
BLAS_TBMV Procedures	Performs a matrix-vector operation `x := A*x` or `A'*x = b` where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular band matrix, with `(k + 1)` diagonals
BLAS_TBSV Procedures	Solves one of the systems of equation `A*x = b` or `A'*x = b` where `b` and `x` are `n` element vectors and `A` is an `n` by n unit, or non-unit, upper or lower triangular band matrix, with `(k + 1)` diagonals
BLAS_TPMV Procedures	Performs a matrix-vector operation `x := A*x` or `x := A'*x` where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix, supplied in packed form
BLAS_TPSV Procedures	Solves one of the systems of equation `A*x = b` or `A'*x = b` where `b` and `x` are `n` element vectors and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix, supplied in packed form
BLAS_TRMV Procedures	Performs a matrix-vector operation `x := A*x` or `x := A'*x` where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix
BLAS_TRSV Procedures	Solves one of the systems of equation `A*x = b` or `A'*x = b` where `b` and `x` are `n` element vectors and `A` is an n by `n` unit, or non-unit, upper or lower triangular matrix

293.4.3 UTL_NLA BLAS Level 3 (Matrix-Matrix Operations) Subprograms

This table lists and briefly describes the UTL_NLA BLAS Level 3 Matrix-Matrix Operations subprograms.

Table 293-3 BLAS Level 3 (Matrix-Matrix Operations) Subprograms

Subprogram	Description
BLAS_GEMM Procedures	Performs one of the matrix-vector operations `C := alpha*op(A)*op(B) + beta*C` where `op(X) is one of op(X) = X or op(X) = X'` where `alpha` and `beta` are scalars, and `A`, `B` and `C` are matrices, with `op(A)` an `m` by `k` matrix, `op(B)` a `k` by `n` matrix and `C` an `m` by `n` matrix
BLAS_SYMM Procedures	Performs one of the matrix-vector operations `C := alpha*A*B + beta*C` or `C := alpha*B*A + beta*C` where `alpha` and `beta` are scalars, `A` is a symmetric matrix, and `B` and `C` are `m` by `n` matrices

Table 293-3 (Cont.) BLAS Level 3 (Matrix-Matrix Operations) Subprograms

Subprogram	Description
BLAS_SYR2K Procedures	Performs one of the symmetric rank2 k operations `C := alpha*A*B' + alpha*B*A' + beta*C` or `C := alpha*A'*B + alpha*B'*A + beta*C` where `alpha` and `beta` are scalars, `C` is an `n` by `n` symmetric matrix and `A` and `B` are `n` by `k` matrices in the first case and `k` by `n` matrices in the second case
BLAS_SYRK Procedures	Performs one of the symmetric rank `k` operations `C := alpha*A*A' + beta*C` or `C := alpha*A'*A + beta*C` where `alpha` and `beta` are scalars, `C` is an `n` by `n` symmetric matrix and `A` is an `n` by `k` matrix in the first case and a `k` by `n` matrix in the second case
BLAS_TRMM Procedures	Performs one of the matrix-vector operations `B := alpha*op(A)*B` or `B := alpha*B*op(A)` where `alpha` is a scalar, `B` is an `m` by `n` matrix, `A` is a unit, or non-unit, upper or lower triangular matrix and `op(A)` is one of two alternatives
BLAS_TRSM Procedures	Performs one of the matrix-vector operations `op(A)*X = alpha*B` or `X*op(A) = alpha*B` where `alpha` is a scalar, `X` and `B` are `m` by `n` matrices, `A` is a unit, or non-unit, upper or lower triangular matrix, `op(A)` is one of two alternatives. The matrix `X` is overwritten on `B`

293.4.4 UTL_NLA LAPACK Driver Routines (Linear Equations) Subprograms

This table lists and briefly describes the LAPACK Driver Routines (Linear Equations) subprograms.

Table 293-4 LAPACK Driver Routines (Linear Equations) Subprograms

Subprogram	Description
LAPACK_GBSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an `n` by `n` matrix and `x` and `b` are `n` by `nrhs` matrices. The `LU` decomposition with partial pivoting and row interchanges is used to factor `A`.
LAPACK_GESV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an `n` by `n` matrix and `x` and `b` are `n` by `nrhs` matrices. The `LU` decomposition with partial pivoting and row interchanges is used to factor `A`.
LAPACK_GTSV Procedures	This procedure solves the equation `a * x = b` where `a` is an `n` by `n` tridiagonal matrix, by Gaussian elimination with partial pivoting.
LAPACK_PBSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite band matrix and `x` and `b` are `n` by `nrhs` matrices. The Cholesky decomposition is used to factor `A`.
LAPACK_POSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite matrix and `x` and `b` are `n` by `nrhs` matrices. The Cholesky decomposition is used to factor `A`.

Table 293-4 (Cont.) LAPACK Driver Routines (Linear Equations) Subprograms

Subprogram	Description
LAPACK_PPSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite matrix stored in packed format and `x` and `b` are `n` by `nrhs` matrices. The Cholesky decomposition is used to factor `A`.
LAPACK_PTSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite tridiagonal matrix, and `x` and `b` are `n` by `nrhs` matrices.
LAPACK_SPSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric matrix stored in packed format, and `x` and `b` are `n` by `nrhs` matrices. The diagonal pivoting method is used to factor `A`.
LAPACK_SYSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric matrix, and `x` and `b` are `n` by `nrhs` matrices. The diagonal pivoting method is used to factor `A`.

293.4.5 UTL_NLA LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms

This table lists and briefly describes the LAPACK Driver Routines (LLS and Eigenvalue) subprograms.

Table 293-5 LAPACK Driver Routines (LLS and Eigenvalue Problems)

Subprogram	Description
LAPACK_GEES Procedures	Computes for an `n` by `n` real nonsymmetric matrix `A`, the eigenvalues, the real Schur form `T`, and, optionally, the matrix of Schur vectors `Z`. This gives the Schur factorization `A = Z*T*(Z**T)`.
LAPACK_GEEV Procedures	Computes for an `n` by `n` real nonsymmetric matrix `A`, the eigenvalues and, optionally, the left and/or right eigenvectors.
LAPACK_GELS Procedures	Solves overdetermined or underdetermined real linear systems involving an `m` by `n` matrix `A`, or its transpose, using a `QR` or `LQ` factorization of `A`. It is assumed that `A` has full rank.
LAPACK_GESDD Procedures	Computes the singular value decomposition (SVD) of a real `m` by `n` matrix `A`, optionally computing the left and right singular vectors. If singular vectors are desired, it uses a divide-and-conquer algorithm that makes mild assumptions about floating point arithmetic.
LAPACK_GESVD Procedures	Computes the singular value decomposition (SVD) of a real `m` by `n` matrix `A`, optionally computing the left and/or right singular vectors. The SVD is written `A = U * SIGMA * transpose(V)`.
LAPACK_SBEV Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric band matrix `A`
LAPACK_SBEVD Procedures	Ccomputes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.

Table 293-5 (Cont.) LAPACK Driver Routines (LLS and Eigenvalue Problems)

Subprogram	Description
LAPACK_SPEV Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A` in packed storage
LAPACK_SPEVD Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A` in packed storage. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.
LAPACK_STEV Procedures	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix `A`
LAPACK_STEVD Procedures	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.
LAPACK_SYEV Procedures	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`
LAPACK_SYEVD Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.

293.5 Summary of UTL_NLA Subprograms

This table lists the `UTL_NLA` subprograms and briefly describes them.

Table 293-6 UTL_NLA Package Subprograms

Subprogram	Description	Group
BLAS_ASUM Functions	Computes the sum of the absolute values of the vector components	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_AXPY Procedures	Copies `alpha*X + Y` into vector `Y`	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_COPY Procedures	Copies the contents of vector `X` to vector `Y`	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_DOT Functions	Returns the dot (scalar) product of two vectors `X` and `Y`	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_GBMV Procedures	Performs the matrix-vector operation `y := alpha*A*x + beta*y` or `y := alpha*A'*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are vectors and `A` is an `m` by `n` band matrix, with `kl` sub-diagonals and `ku` super-diagonals	BLAS Level 2 (Matrix-Vector Operations) Subprograms

Table 293-6 (Cont.) UTL_NLA Package Subprograms

Subprogram	Description	Group
BLAS_GEMM Procedures	Performs one of the matrix-vector operations where `alpha` and `beta` are scalars, and `A`, `B` and `C` are matrices, with `op(A)` an `m` by `k` matrix, `op(B)` a `k` by `n` matrix and `C` an `m` by `n` matrix	BLAS Level 3 (Matrix-Matrix Operations) Subprograms
BLAS_GEMV Procedures	Performs the matrix-vector operations `y := alpha*A*x + beta*y` or `y := alpha*A'*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are vectors and `A` is an `m` by `n` matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_GER Procedures	Performs a rank 1 operation `A := alpha*x*y' + A` where `alpha` is a scalar, `x` is an `m` element vector, `y` is an n element vector and `A` is an `m` by `n` matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_IAMAX Functions	Computes the index of the first element of a vector that has the largest absolute value	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_NRM2 Functions	Computes the vector 2-norm (Euclidean norm)	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_ROT Procedures	Returns the plane rotation of points	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_ROTG Procedures	Returns the Givens rotation of points	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_SBMV Procedures	Performs a matrix-vector operation `y := alpha*A*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are n element vectors and `A` is an `n` by `n` symmetric band matrix, with `k` super-diagonals	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_SCAL Procedures	Scales a vector by a constant	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_SPMV Procedures	Performs a matrix-vector operation `y := alpha*A*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are `n` element vectors and `A` is an `n` by `n` symmetric matrix, supplied in packed form	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_SPR Procedures	Performs a symmetric rank 1 operation `A := alpha*x*x' + A` where `alpha` is a real scalar, `x` is an `n` element vector, and `A` is an `n` by `n` symmetric matrix, supplied in packed form	BLAS Level 2 (Matrix-Vector Operations) Subprograms

Table 293-6 (Cont.) UTL_NLA Package Subprograms

Subprogram	Description	Group
BLAS_SPR2 Procedures	Performs a symmetric rank 2 operation where `alpha` is a scalar, `x` and `y` are `n` element vectors, and `A` is an `n` by `n` symmetric matrix, supplied in packed form	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_SWAP Procedures	Swaps the contents of two vectors each of size n	BLAS Level 1 (Vector-Vector Operations) Subprograms
BLAS_SYMM Procedures	Performs one of the matrix-vector operations where `alpha` and `beta` are scalars, `A` is a symmetric matrix, and `B` and `C` are `m` by `n` matrices	BLAS Level 3 (Matrix-Matrix Operations) Subprograms
BLAS_SYMV Procedures	Performs a matrix-vector operation where `alpha` and `beta` are scalars, `x` and `y` are `n` element vectors and `A` is an `n` by `n` symmetric matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_SYR Procedures	Performs a symmetric rank 1 operation where `alpha` is a real scalar, `x` is an `n` element vector, and `A` is an `n` by `n` symmetric matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_SYR2 Procedures	Performs a symmetric rank 2 operation where `alpha` is a scalar, `x` and `y` are `n` element vectors, and `A` is an `n` by `n` symmetric matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_SYR2K Procedures	Performs one of the symmetric rank2 k operations where `alpha` and `beta` are scalars, `C` is an `n` by `n` symmetric matrix and `A` and `B` are `n` by `k` matrices in the first case and `k` by `n` matrices in the second case	BLAS Level 3 (Matrix-Matrix Operations) Subprograms
BLAS_SYRK Procedures	Performs one of the symmetric rank `k` operations where `alpha` and `beta` are scalars, `C` is an `n` by `n` symmetric matrix and `A` is an `n` by `k` matrix in the first case and a `k` by `n` matrix in the second case	BLAS Level 3 (Matrix-Matrix Operations) Subprograms
BLAS_TBMV Procedures	Performs a matrix-vector operation where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular band matrix, with `(k + 1)` diagonals	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_TBSV Procedures	Solves one of the systems of equation where `b` and `x` are `n` element vectors and `A` is an `n` by n unit, or non-unit, upper or lower triangular band matrix, with `(k + 1)` diagonals	BLAS Level 2 (Matrix-Vector Operations) Subprograms

Table 293-6 (Cont.) UTL_NLA Package Subprograms

Subprogram	Description	Group
BLAS_TPMV Procedures	Performs a matrix-vector operation where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix, supplied in packed form	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_TPSV Procedures	Solves one of the systems of equation where `b` and `x` are `n` element vectors and `A` is an `n` by n unit, or non-unit, upper or lower triangular matrix, supplied in packed form	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_TRMM Procedures	Performs one of the matrix-vector operations where`alpha` is a scalar, `B` is an `m` by `n` matrix, `A` is a unit, or non-unit, upper or lower triangular matrix and `op(A)` is one of two alternatives	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_TRMV Procedures	Performs a matrix-vector operation where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
BLAS_TRSM Procedures	Performs one of the matrix-vector operations `op( A )*X = alpha*B` or `X*op(A) = alpha*B` where `alpha` is a scalar, `X` and `B` are `m` by `n` matrices, `A` is a unit, or non-unit, upper or lower triangular matrix, `op(A)` is one of two alternatives. The matrix `X` is overwritten on `B`	BLAS Level 3 (Matrix-Matrix Operations) Subprograms
BLAS_TRSV Procedures	Solves one of the systems of equation where `b` and `x` are `n` element vectors and `A` is an n by `n` unit, or non-unit, upper or lower triangular matrix	BLAS Level 2 (Matrix-Vector Operations) Subprograms
LAPACK_GBSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an `n` by `n` matrix and `x` and `b` are `n` by `nrhs` matrices. The `LU` decomposition with partial pivoting and row interchanges is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_GEES Procedures	Computes for an `n` by `n` real nonsymmetric matrix `A`, the eigenvalues, the real Schur form `T`, and, optionally, the matrix of Schur vectors `Z`. This gives the Schur factorization `A = Z*T*(Z**T)`.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms

Table 293-6 (Cont.) UTL_NLA Package Subprograms

Subprogram	Description	Group
LAPACK_GEEV Procedures	Computes for an `n` by `n` real nonsymmetric matrix `A`, the eigenvalues and, optionally, the left and/or right eigenvectors.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_GELS Procedures	Solves overdetermined or underdetermined real linear systems involving an `m` by `n` matrix `A`, or its transpose, using a `QR` or `LQ` factorization of `A`. It is assumed that `A` has full rank.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_GESDD Procedures	Computes the singular value decomposition (SVD) of a real `m` by `n` matrix `A`, optionally computing the left and right singular vectors. If singular vectors are desired, it uses a divide-and-conquer algorithm that makes mild assumptions about floating point arithmetic.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_GESV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an `n` by `n` matrix and `x` and `b` are `n` by `nrhs` matrices. The `LU` decomposition with partial pivoting and row interchanges is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_GESVD Procedures	Computes the singular value decomposition (SVD) of a real `m` by `n` matrix `A`, optionally computing the left and/or right singular vectors. The SVD is written `A = U * SIGMA * transpose(V)`.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_GTSV Procedures	This procedure solves the equation `a * x = b` where `a` is an `n` by `n` tridiagonal matrix, by Gaussian elimination with partial pivoting.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_PBSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite band matrix and `x` and `b` are `n` by `nrhs` matrices. The Cholesky decomposition is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_POSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite matrix and `x` and `b` are `n` by `nrhs` matrices. The Cholesky decomposition is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms

Table 293-6 (Cont.) UTL_NLA Package Subprograms

Subprogram	Description	Group
LAPACK_PPSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite matrix stored in packed format and `x` and `b` are `n` by `nrhs` matrices. The Cholesky decomposition is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_PTSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric positive definite tridiagonal matrix, and `x` and `b` are `n` by `nrhs` matrices.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_SBEV Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric band matrix `A`	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_SBEVD Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_SPEV Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A` in packed storage	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_SPEVD Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A` in packed storage. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_SPSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric matrix stored in packed format, and `x` and `b` are `n` by `nrhs` matrices. The diagonal pivoting method is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms
LAPACK_STEV Procedures	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix `A`	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms

Table 293-6 (Cont.) UTL_NLA Package Subprograms

Subprogram	Description	Group
LAPACK_STEVD Procedures	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_SYEVD Procedures	Computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.	LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms
LAPACK_SYSV Procedures	This procedure computes the solution to a real system of linear equations `a * x = b` where `a` is an n by n symmetric matrix, and `x` and `b` are `n` by `nrhs` matrices. The diagonal pivoting method is used to factor `A`.	LAPACK Driver Routines (Linear Equations) Subprograms

293.5.1 BLAS_ASUM Functions

This procedure computes the sum of the absolute values of the vector components.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_ASUM (
   n      IN       POSITIVEN,
   x      IN       UTL_NLA_ARRAY_DBL,
   incx   IN       POSITIVEN)
 RETURN BINARY_DOUBLE;

UTL_NLA.BLAS_ASUM (
   n      IN       POSITIVEN,
   alpha  IN       SCALAR_DOUBLE,
   x      IN       UTL_NLA_ARRAY_FLT)
 RETURN BINARY_FLOAT
```

Parameters

Table 293-7 BLAS_ASUM Function Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.

293.5.2 BLAS_AXPY Procedures

This procedure copies alpha*X + Y into vector Y.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_AXPY (
   n      IN       POSITIVEN,
   alpha  IN       SCALAR_DOUBLE,
   x      IN       UTL_NLA_ARRAY_DBL,
   incx   IN       POSITIVEN,
   y      IN OUT   UTL_NLA_ARRAY_DBL,
   incy   IN       POSITIVEN);

UTL_NLA.BLAS_AXPY (
   n      IN       POSITIVEN,
   alpha  IN       SCALAR_DOUBLE,
   x      IN       UTL_NLA_ARRAY_FLT,
   incx   IN       POSITIVEN,
   y      IN OUT   UTL_NLA_ARRAY_FLT,
   incy   IN       POSITIVEN);
```

Parameters

Table 293-8 BLAS_AXPY Procedure Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
alpha	Specifies the scalar alpha.

Table 293-8 (Cont.) BLAS_AXPY Procedure Parameters

Parameter	Description
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.
y	UTL_NLA_ARRAY_FLT/DBL of DIMENSION at least (1 + (n - 1)*abs(incy))
incy	Specifies the increment for the elements of y. incy must not be zero.

293.5.3 BLAS_COPY Procedures

This procedure copies the contents of vector X to vector Y.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_COPY (
   n     IN     POSITIVEN,
   x     IN     UTL_NLA_ARRAY_DBL,
   incx  IN     POSITIVEN,
   y     IN OUT UTL_NLA_ARRAY_DBL,
   incy  IN     POSITIVEN);

UTL_NLA.BLAS_COPY (
   n     IN     POSITIVEN,
   x     IN     UTL_NLA_ARRAY_FLT,
   incx  IN     POSITIVEN,
   y     IN OUT UTL_NLA_ARRAY_FLT,
   incy  IN     POSITIVEN);
```

Parameters

Table 293-9 BLAS_COPY Procedure Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.

Table 293-9 (Cont.) BLAS_COPY Procedure Parameters

Parameter	Description
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incy))
incy	Specifies the increment for the elements of y. incy must not be zero.

293.5.4 BLAS_DOT Functions

This function returns the dot (scalar) product of two vectors X and Y.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_DOT (
   n     IN   POSITIVEN,
   x     IN   UTL_NLA_ARRAY_DBL,
   incx  IN   POSITIVEN,
   y     IN   UTL_NLA_ARRAY_DBL,
   incy  IN   POSITIVEN)
  RETURN BINARY_DOUBLE;

UTL_NLA.BLAS_DOT (
   n     IN   POSITIVEN,
   x     IN   UTL_NLA_ARRAY_FLT,
   incx  IN   POSITIVEN,
   y     IN   UTL_NLA_ARRAY_FLT,
   incy  IN   POSITIVEN)
  RETURN BINARY_FLOAT;
```

Parameters

Table 293-10 BLAS_DOT Function Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incy))

Table 293-10 (Cont.) BLAS_DOT Function Parameters

Parameter	Description
incy	Specifies the increment for the elements of y. incy must not be zero.

293.5.5 BLAS_GBMV Procedures

This procedure performs one of the matrix-vector operations `y := alpha*A*x + beta*y` or `y := alpha*A'*x + beta*y`, where `alpha` and `beta` are scalars, `x` and `y` are vectors and `A` is an `m` by `n` band matrix, with `kl` sub-diagonals and `ku` super-diagonals.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_GBMV (
   trans  IN      flag,
   m      IN      POSITIVEN,    n      IN      POSITIVEN,
   kl     IN      NATURALN,
   ku     IN      NATURALN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN      UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   y      IN OUT  UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_GBMV (
   trans  IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   kl     IN      NATURALN,
   ku     IN      NATURALN,
   alpha  IN      SCALAR_FLOAT,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   x      IN      UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_FLOAT,
   y      IN OUT  UTL_NLA_ARRAY_FLT,
   incy   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-11 BLAS_GBMV Procedure Parameters

Parameter	Description
trans	Specifies the operation to be performed: • trans = 'N' or 'n'y := alpha*A*x + beta*y • trans = 'T' or 't'y := alpha*A'*x + beta*y • trans = 'C' or 'c'y := alpha*A'*x + beta*y
m	Specifies the number of rows of the matrix A. m must be at least zero.
n	Specifies the number of columns of the matrix A. n must be at least zero.
kl	Specifies the number of sub-diagonals of the matrix A. kl must satisfy 0. le. kl.
ku	Specifies the number of super-diagonals of the matrix A. ku must satisfy 0 .le. ku.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,n). Before entry, the leading (kl + ku + 1) by n part of the array A must contain the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row (ku+1) of the array, the first super-diagonal starting at position 2 in row ku, the first sub-diagonal starting at position 1 in row (ku+2), and so on. Elements in the array A that do not correspond to elements in the band matrix (such as the top left ku by ku triangle) are not referenced.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least (kl+ku+1).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx)) when trans = ''N' or 'n' and at least (1 + (m - 1)*abs(incx)) otherwise. Before entry, the incremented array X must contain the vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta. When beta is supplied as zero then y need not be set on input.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (m - 1)*abs(incy)) when trans = 'N' or 'n' and at least (1+(n-1)*abs(incy)) otherwise. Before entry with beta nonzero, the incremented array Y must contain the vector y. On exit, Y is overwritten by the updated vector y.

Table 293-11 (Cont.) BLAS_GBMV Procedure Parameters

Parameter	Description
incy	Specifies the increment for the elements of y. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.6 BLAS_GEMM Procedures

This procedure performs one of the matrix-matrix operations.

```
C := alpha*op( A )*op( B ) + beta*C
```

where op(X) is one of

```
op(X) = X
```

or

```
op(X) = X'
```

where alpha and beta are scalars, and A, B and C are matrices, with op(A) an m by k matrix, op(B) a k by n matrix and C an m by n matrix.

> **See Also:**
>
> BLAS Level 3 (Matrix-Matrix Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_GEMM (
   transa IN      flag,
   transb IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   k      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   b      IN      UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   c      IN OUT  UTL_NLA_ARRAY_DBL,
   ldc    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_GEMM (
   transa IN      flag,
   transb IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   k      IN      POSITIVEN,
```

```
  alpha  IN      SCALAR_FLOAT,
  a      IN      UTL_NLA_ARRAY_FLT,
  lda    IN      POSITIVEN,
  b      IN      UTL_NLA_ARRAY_FLT,
  ldb    IN      POSITIVEN,
  beta   IN      SCALAR_FLOAT,
  c      IN OUT  UTL_NLA_ARRAY_FLT,
  ldc    IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-12 BLAS_GEMM Procedure Parameters

Parameter	Description
transa	Specifies the form of op(A) to be used in the matrix multiplication as follows: • transa = 'N' or 'n' : op(A) = 'A' • transa = 'T' or 't' : op(A) = 'A' • transa = 'C' or 'c' : op(A) = 'A'
transb	Specifies the form of op(B) to be used in the matrix multiplication as follows: • transb = 'N' or 'n' : op(B) = B • transb ='T' or 't' : op(B) = B' • transb = 'C' or 'c': op(B) = B'
m	Specifies the number of rows of the matrix op(A) and of the matrix C. m must be at least zero.
n	Specifies the number of columns of the matrix op(B) and of the matrix C. n must be at least zero.
k	Specifies the rows of the matrix op(A) and the number of columns of the matrix op(B). k must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, ka) where ka is k when transa = 'N' or 'n', and is m otherwise. Before entry with transa = 'N' or 'n', the leading m by k part of the array A must contain the matrix A, otherwise the leading k by m part of the array A must contain the matrix A.
lda	Specifies the first dimension of a as declared in the calling (sub) program. When transa = 'N' or 'n', lda must be at least max (1,k).
b	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, kb) where kb is n when transb = ''N' or 'n', and is k otherwise. Before entry with transb = 'N' or 'n', the leading k by n part of the array b must contain the matrix B, otherwise the leading n by k part of the arrayb must contain the matrix B.
ldb	Specifies the first dimension of b as declared in the calling (sub) program. When transb = 'N' or 'n', ldb must be at least max (1, n).
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta. When beta is supplied as zero then c need not be set on input.

Table 293-12 (Cont.) BLAS_GEMM Procedure Parameters

Parameter	Description
c	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldc, n). Before entry, the leading m by n part of the array C must contain the matrix C, except when beta is zero, in which case C need not be set on entry. On exit, the arrayC is overwritten by the m by n matrix (alpha*op(A)*op(B) + beta*C).
ldc	Specifies the first dimension of C as declared in the calling (sub) program. ldc must be at least max(1, m).
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.7 BLAS_GEMV Procedures

This procedure performs one of the matrix-vector operations: `y := alpha*A*x + beta*y` or `y := alpha*A'*x + beta*y` where `alpha` and `beta` are scalars, `x` and `y` are vectors and `A` is an `m` by `n` matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_GEMV (
   trans  IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN      UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   y      IN OUT  UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_GEMV (
   trans  IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   x      IN      UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_FLOAT,
   y      IN OUT  UTL_NLA_ARRAY_FLT,
```

```
  incy   IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-13 BLAS_GEMV Procedure Parameters

Parameter	Description
trans	Specifies the operation to be performed: • trans = 'N' or 'n',y := alpha*A*x + beta*y • trans = 'T' or 't'y := alpha*A'*x + beta*y • trans = 'C' or 'c'y := alpha*A'*x + beta*y
m	Specifies the number of rows of the matrix A. m must be at least zero.
n	Specifies the number of columns of the matrix A. n must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n). Before entry, the leading m by n part of the array a must contain the matrix of coefficients.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least max(1, m).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx)) when trans = ''N' or 'n' and at least (1+(m-1)*abs(incx)) otherwise. Before entry, the incremented array X must contain the vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta. When beta is supplied as zero then y need not be set on input.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (m - 1)*abs(incy)) when trans = 'N' or 'n' and at least (1 + (n - 1)*abs(incy)) otherwise. Before entry with beta nonzero, the incremented array Y must contain the vector y. On exit, Y is overwritten by the updated vector y.
incy	Specifies the increment for the elements of y. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.8 BLAS_GER Procedures

This procedure performs the rank 1 operation: `A := alpha*x*y' + A` where `alpha` is a scalar, `x` is an `m` element vector, `y` is an n element vector and `A` is an `m` by `n` matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_GER (
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DBL,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   y      IN      UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_GER (
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLT,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   y      IN      UTL_NLA_ARRAY_FLT,
   incy   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-14 BLAS_GER Procedure Parameters

Parameter	Description
`m`	Specifies the number of rows of the matrix `A`. `m` must be at least zero.
`n`	Specifies the number of columns of the matrix `A`. `n` must be at least zero.
`alpha`	Specifies the scalar alpha.
`x`	`UTL_NLA_ARRAY_FLT/DBL` of dimension at least `( 1 + ( m - 1 )*abs( incx ) )` Before entry, the incremented array `X` must contain the m element vector `x`.
`incx`	Specifies the increment for the elements of `x`. incx must not be zero.

Table 293-14 (Cont.) BLAS_GER Procedure Parameters

Parameter	Description
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least `( 1 + ( n - 1 )*abs( incy ) )` Before entry, the incremented array Y must contain the m element vector y.
incy	Specifies the increment for the elements of y. incx must not be zero.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n). Before entry, the leading m by n part of the array a must contain the matrix of coefficients. On exit, a is overwritten by the updated matrix.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least `max( 1, m )`
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.9 BLAS_IAMAX Functions

This function computes the index of first element of a vector that has the largest absolute value.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_IAMAX (
   n     IN   POSITIVEN,
   x     IN   UTL_NLA_ARRAY_DBL,
   incx  IN   POSITIVEN,
  RETURN POSITIVEN;

UTL_NLA.BLAS_IAMAX (
   n     IN   POSITIVEN,
   x     IN   UTL_NLA_ARRAY_FLT,
   incx  IN   POSITIVEN,
  RETURN POSITIVEN;
```

Parameters

Table 293-15 BLAS_IAMAX Function Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of DIMENSION at least (1 + (n - 1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.

293.5.10 BLAS_NRM2 Functions

This function computes the vector 2-norm (Euclidean norm).

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_NRM2 (
   n     IN   POSITIVEN,
   x     IN   UTL_NLA_ARRAY_DBL,
   incx  IN   POSITIVEN)
  RETURN BINARY_DOUBLE;

UTL_NLA.BLAS_NRM2 (
   n     IN   POSITIVEN,
   x     IN   UTL_NLA_ARRAY_FLT,
   incx  IN   POSITIVEN)
  RETURN BINARY_FLOAT;
```

Parameters

Table 293-16 BLAS_NRM2 Function Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.

293.5.11 BLAS_ROT Procedures

This procedure returns the plane rotation of points.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_ROT (
   n      IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   y      IN OUT  UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   c      IN      SCALAR_DOUBLE,
   s      IN      SCALAR_DOUBLE);

UTL_NLA.BLAS_ROT (
   n      IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   y      IN OUT  UTL_NLA_ARRAY_FLT,
   incy   IN      POSITIVEN,
   c      IN      SCALAR_DOUBLE,
   s      IN      SCALAR_DOUBLE);
```

Parameters

Table 293-17 BLAS_ROT Procedure Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)* abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.
y	UTL_NLA_ARRAY_FLT/DBL of DIMENSION at least (1+(n-1)*abs(incy))
incy	Specifies the increment for the elements of y. incy must not be zero.
c	SCALAR_FLOAT/DOUBLE.Specifies the scalar C.
s	SCALAR_FLOAT/DOUBLE.Specifies the scalar S.

293.5.12 BLAS_ROTG Procedures

This procedure returns the Givens rotation of points.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_ROTG (
   a    IN OUT   SCALAR_DOUBLE,
   b    IN OUT   SCALAR_DOUBLE,
   c    IN OUT   SCALAR_DOUBLE,
   s    IN OUT   SCALAR_DOUBLE);

UTL_NLA.BLAS_ROTG (
   a    IN OUT   SCALAR_FLOAT,
   b    IN OUT   SCALAR_FLOAT,
   c    IN OUT   SCALAR_FLOAT,
   s    IN OUT   SCALAR_FLOAT);
```

Parameters

Table 293-18 BLAS_ROTG Procedure Parameters

Parameter	Description
a	SCALAR_FLOAT/DOUBLE. Specifies the scalar A.
b	SCALAR_FLOAT/DOUBLE. Specifies the scalar B.
c	SCALAR_FLOAT/DOUBLE. Specifies the scalar C.
s	SCALAR_FLOAT/DOUBLE. Specifies the scalar S.

293.5.13 BLAS_SCAL Procedures

This procedure scales a vector by a constant.

> **See Also:**
>
> BLAS Level 1 (Vector-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SCAL (
   n      IN  POSITIVEN,
   alpha  IN  SCALAR_DOUBLE,
   x      IN  OUT UTL_NLA_ARRAY_DBL,
   incx   IN  POSITIVEN);
```

```
UTL_NLA.BLAS_SCAL (
   n      IN  POSITIVEN,
   alpha  IN  SCALAR_FLOAT,
   x      IN  OUT UTL_NLA_ARRAY_FLT,
   incx   IN  POSITIVEN);
```

Parameters

Table 293-19 BLAS_SCAL Procedure Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
alpha	Specifies the scalar alpha.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.

293.5.14 BLAS_SPMV Procedures

This procedure performs the matrix-vector operation `y := alpha*A*x + beta*y`, where `alpha` and `beta` are scalars, `x` and `y` are n element vectors and `A` is an `n` by `n` symmetric matrix, supplied in packed form.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SPMV (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   ap     IN      UTL_NLA_ARRAY_DBL,
   x      IN      UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   y      IN OUT  UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SPMV (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
   ap      IN     UTL_NLA_ARRAY_FLT,
   x      IN      UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_FLOAT,
   y      IN OUT  UTL_NLA_ARRAY_FLT,
```

```
  incy   IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-20 BLAS_SPMV Procedure Parameters

Parameter	Description
uplo	Specifies the upper or lower triangular part of the matrix A is supplied in the packed array AP: • uplo = 'U' or 'u'. The upper triangular part of A is supplied in AP. • uplo = 'L' or 'l'. The lower triangular part of A is supplied in AP.
n	Specifies the order of the matrix A. n must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
ap	UTL_NLA_ARRAY_FLT/DBL of dimension at least ((n*(n+1))/2) Before entry with uplo = 'U' or 'u', the array ap must contain the upper triangular part of the symmetric matrix packed sequentially, column by column, so that ap(1) contains a(1,1), ap(2) and ap(3) contain a(1,2) and a(2,2) respectively, and so on. Before entry with uplo = 'L' or 'l', the array ap must contain the lower triangular part of the symmetric matrix packed sequentially, column by column, so that ap(1) contains, ap(2) and ap(3) contain a(2,1) and a(3,1) respectively, and so on.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx)) Before entry, the incremented array X must contain the n element vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta. When beta is supplied as zero then Y need not be set on input.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at leasT (1+(n-1)*abs(incy)) Before entry, the incremented array Y must contain the n element vector y. On exit, Y is overwritten by the updated vector y.
incy	Specifies the increment for the elements of y. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.15 BLAS_SPR Procedures

This procedure performs the rank 1 operation `A := alpha*x*x' + A`, where `alpha` is a real scalar, `x` is an `n` element vector, and `A` is an `n` by `n` symmetric matrix, supplied in packed form.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SPR (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DBL,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   ap     IN OUT  UTL_NLA_ARRAY_DBL,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SPR (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLT,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   ap     IN OUT  UTL_NLA_ARRAY_FLT,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-21 BLAS_SPR Procedure Parameters

Parameter	Description
`uplo`	Specifies whether the upper or lower triangular part of the matrix `A` is supplied in the packed array `ap`: • `uplo` = `'U'` or `'u'`: The upper triangular part of `A` is supplied in `ap`. • `uplo` = `'L'` or `'l'` : The lower triangular part of `A` is supplied in `ap`.
`n`	Specifies the order of the matrix `A`. `n` must be at least zero.
`alpha`	Specifies the scalar alpha.
`x`	`UTL_NLA_ARRAY_FLT/DBL` of dimension at least `(1+(n-1)*abs(incx))` Before entry, the incremented array `X` must contain the m element vector `x`.
`incx`	Specifies the increment for the elements of `x`. `incx` must not be zero.

Table 293-21 (Cont.) BLAS_SPR Procedure Parameters

Parameter	Description
ap	UTL_NLA_ARRAY_FLT/DBL of dimension at least ((n*(n +1))/2) Before entry with uplo = 'U' or 'u', the array ap must contain the upper triangular part of the symmetric matrix packed sequentially, column by column, so that ap(1) contains a(1,1), ap(2) and ap(3) contain a(1,2) and a(2,2) respectively, and so on. On exit, the array ap is overwritten by the upper triangular part of the updated matrix. Before entry with uplo = 'L' or 'l', the array ap must contain the lower triangular part of the symmetric matrix packed sequentially, column by column, so that ap(1) contains a(1,1), ap(2) and ap(3) contain a(2,1) and a(3,1) respectively, and so on. On exit, the array ap is overwritten by the lower triangular part of the updated matrix
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.16 BLAS_SPR2 Procedures

This procedure performs the rank 2 operation A := alpha*x*y' + alpha*y*x' +A, where alpha is a scalar, x and y are n element vectors, and A is an n by n symmetric matrix, supplied in packed form.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SPR2 (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DBL,
   x      IN      UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   y      IN      UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SPR2 (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLT,
   x      IN      UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   y      IN      UTL_NLA_ARRAY_FLT,
```

```
  incy   IN      POSITIVEN,
  a      IN OUT  UTL_NLA_ARRAY_FLT,
  lda    IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-22 BLAS_SPR2 Procedure Parameters

Parameter	Description
uplo	Specifies whether the upper or lower triangular part of the matrix A is supplied in the packed array ap : • uplo = 'U' or 'u' : The upper triangular part of A is supplied in ap. • uplo = 'L' or 'l' : The lower triangular part of A is supplied in ap.
n	Specifies the order of the matrix A. n must be at least zero.
alpha	Specifies the scalar alpha.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx)) Before entry, the incremented array X must contain the m element vector x.
incx	Specifies the increment for the elements of x. incx must not be zero.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incy)) Before entry, the incremented array X must contain the m element vector y.
incy	Specifies the increment for the elements of y. incy must not be zero.
ap	UTL_NLA_ARRAY_FLT/DBL of dimension at least ((n*(n+1))/2) Before entry with uplo = 'U' or 'u', the array ap must contain the upper triangular part of the symmetric matrix packed sequentially, column by column, so that ap(1) contains ap(1) contains a(1,1), ap(2) and ap(3) contain a(1,2) and a(2,2) respectively, and so on. On exit, the array ap is overwritten by the upper triangular part of the updated matrix. Before entry with uplo = 'L' or 'l', the array ap must contain the lower triangular part of the symmetric matrix packed sequentially, column by column, so that ap(1) contains a(1,1), ap(2) and ap(3) contain a(2,1) and a(3,1) respectively, and so on. On exit, the array ap is overwritten by the lower triangular part of the updated matrix
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least (k + 1).
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.17 BLAS_SBMV Procedures

This procedure performs the matrix-vector operation `y := alpha*A*x + beta*y`, where `alpha` and `beta` are scalars, `x` and `y` are n element vectors and `A` is an `n` by `n` symmetric band matrix, with `k` super-diagonals.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SBMV (
   uplo    IN      flag,
   n       IN      POSITIVEN,
   k       IN      NATURALN,
   alpha   IN      SCALAR_DOUBLE,
   a       IN      UTL_NLA_ARRAY_DBL,
   lda     IN      POSITIVEN,
   x       IN      UTL_NLA_ARRAY_DBL,
   incx    IN      POSITIVEN,
   beta    IN      SCALAR_DOUBLE,
   y       IN OUT  UTL_NLA_ARRAY_DBL,
   incy    IN      POSITIVEN,
   pack    IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SBMV (
   uplo    IN      flag,
   n       IN      POSITIVEN,
   k       IN      NATURALN,
   alpha   IN      SCALAR_FLOAT,
   a       IN      UTL_NLA_ARRAY_FLT,
   lda     IN      POSITIVEN,
   x       IN      UTL_NLA_ARRAY_FLT,
   incx    IN      POSITIVEN,
   beta    IN      SCALAR_FLOAT,
   y       IN OUT  UTL_NLA_ARRAY_FLT,
   incy    IN      POSITIVEN,
   pack    IN      flag DEFAULT 'C');
```

Parameters

Table 293-23 BLAS_SBMV Procedure Parameters

Parameter	Description
`uplo`	Specifies whether the upper or lower triangular part of the band matrix `A` is being supplied: • `uplo` = `'U'` or `'u'`. The upper triangular part of `A` is supplied. • `uplo` = `'L'` or `'l'`. The lower triangular part of `A` is supplied.
`n`	Specifies the order of the matrix `A`. `n` must be at least zero.
`k`	Specifies the number of super-diagonals of the matrix `A`. `k` must satisfy `0 .le. k`.

Table 293-23 (Cont.) BLAS_SBMV Procedure Parameters

Parameter	Description
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,n). Before entry with uplo = 'U' or 'u', the leading (k+1) by n part of the array A must contain the upper triangular band part of the symmetric matrix, supplied column by column, with the leading diagonal of the matrix in row (k+1) of the array, the first super-diagonal starting at position 2 in rowk, and so on. The top left k by k triangle of the array A is not referenced. Before entry with uplo = 'L' or 'l', the leading (k+1) by n part of the array A must contain the lower triangular band part of the symmetric matrix, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right k by k triangle of the array A is not referenced. Unchanged on exit
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least (k + 1).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx)) Before entry, the incremented array X must contain the n element vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incy)) Before entry, the incremented array Y must contain the n element vector y. On exit, Y is overwritten by the updated vector y.
incy	Specifies the increment for the elements of y. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.18 BLAS_SWAP Procedures

This procedure swaps the contents of two vectors each of size n.

Syntax

```
UTL_NLA.BLAS_SWAP (
   n     IN     POSITIVEN,
   x     IN OUT UTL_NLA_ARRAY_DBL,
   incx  IN     POSITIVEN,
   y     IN OUT UTL_NLA_ARRAY_DBL,
   incy  IN     POSITIVEN);
```

```
UTL_NLA.BLAS_SWAP (
   n     IN     POSITIVEN,
   x     IN OUT UTL_NLA_ARRAY_FLT,
   incx  IN     POSITIVEN,
   y     IN OUT UTL_NLA_ARRAY_FLT,
   incy  IN     POSITIVEN);
```

Parameters

Table 293-24 BLAS_SWAP Procedure Parameters

Parameter	Description
n	Specifies the number of elements of the vectors x and y. n must be at least zero.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx))
incx	Specifies the increment for the elements of x. incx must not be zero.
y	UTL_NLA_ARRAY_FLT/DBL of DIMENSION at least (1+(n-1)*abs(incy))
incy	Specifies the increment for the elements of y. incy must not be zero.

293.5.19 BLAS_SYMM Procedures

This procedure performs one of the matrix-matrix operations `C := alpha*A*B + beta*C` or `C := alpha*B*A + beta*C`, where `alpha` and `beta` are scalars, `A` is a symmetric matrix, and `B` and `C` are `m` by `n` matrices.

> **See Also:**
>
> BLAS Level 3 (Matrix-Matrix Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SYMM (
   side   IN     flag,
   uplo   IN     flag,
   m      IN     POSITIVEN,
   n      IN     POSITIVEN,
   alpha  IN     SCALAR_DOUBLE,
   a      IN     UTL_NLA_ARRAY_DBL,
   lda    IN     POSITIVEN,
   b      IN     UTL_NLA_ARRAY_DBL,
   ldb    IN     POSITIVEN,
   beta   IN     SCALAR_DOUBLE,
   c      IN OUT UTL_NLA_ARRAY_DBL,
   ldc    IN     POSITIVEN,
   pack   IN     flag DEFAULT 'C');
```

```
UTL_NLA.BLAS_SYMM (
   side   IN      flag,
   uplo   IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   b      IN      UTL_NLA_ARRAY_FLT,
   ldb    IN      POSITIVEN,
   beta   IN      SCALAR_FLOAT,
   c      IN OUT  UTL_NLA_ARRAY_FLT,
   ldc    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-25 BLAS_SYMM Procedure Parameters

Parameter	Description
side	Specifies whether the symmetric matrix A appears on the left or right in the operation: • side = 'L' or 'l' : C := alpha*A*B + beta*C • side = 'R' or 'r' : C := alpha*B*A + beta*C
uplo	Specifies whether the upper or lower triangular part of the array A is to be referenced: • uplo = 'U' or 'u' : Only the upper triangular part of the symmetric matrix is to be referenced. • uplo = 'L' or 'l' : Only the lower triangular part of the symmetric matrix is to be referenced.
m	Specifies the number of rows of the matrix C. m must be at least zero.
n	Specifies the number of columns of the matrix C. n must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,ka) where ka is m when side = 'L' or 'l', and is n otherwise. Before entry with side = 'L' or 'l', the leading m by m part of the array A must contain the symmetric matrix, such that when uplo = 'U' or 'u', the leading m by m upper triangular part of the array A must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of A is not referenced, and when uplo = 'L' or 'l', the leading m by m lower triangular part of the array A must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of A is not referenced. Before entry with side = 'R' or 'r', the n by n part of the array A must contain the symmetric matrix, such that when uplo = 'U' or 'u', the leading n by n upper triangular part of the array A must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of A is not referenced, and when uplo = 'L' or 'l', the leading n by n lower triangular part of the array A must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of A is not referenced.

Table 293-25 (Cont.) BLAS_SYMM Procedure Parameters

Parameter	Description
lda	Specifies the first dimension of a as declared in the calling (sub) program. When side = 'L' or 'l', lda must be at least max(1,m), otherwise lda must be at least max(1,n).
b	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldb,n). Before entry, the leading m by n part of the array B must contain the matrix B.
ldb	Specifies the first dimension of b as declared in the calling (sub) program. ldb must be at least max(1,m).
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta. When beta is supplied as zero then c need not be set on input.
c	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldc,n). Before entry, the leading m by n part of the array C must contain the matrix C, except when beta is zero, in which case C need not be set on entry. On exit, the array C is overwritten by the m by n updated matrix.
ldc	Specifies the first dimension of C as declared in the calling (sub) program. ldc must be at least max (1,m).
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.20 BLAS_SYMV Procedures

This procedure performs the matrix-vector operation y := alpha*A*x + beta*y , where alpha and beta are scalars, x and y are n element vectors and A is an n by n symmetric matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SYMV (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN      UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   y      IN OUT  UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SYMV (
   uplo   IN      flag,
```

```
  n      IN      POSITIVEN,
  alpha  IN      SCALAR_FLOAT,
  a      IN      UTL_NLA_ARRAY_FLT,
  lda    IN      POSITIVEN,
  x      IN      UTL_NLA_ARRAY_FLT,
  incx   IN      POSITIVEN,
  beta   IN      SCALAR_FLOAT,
  y      IN OUT  UTL_NLA_ARRAY_FLT,
  incy   IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-26 BLAS_SYMV Procedure Parameters

Parameter	Description
uplo	Specifies whether the upper or lower triangular part of the array A is to be referenced: • uplo = 'U'or 'u'. Only the upper triangular part ofA is to be referenced. • uplo = 'L' or 'l'. Only the lower triangular part of A is to be referenced.
n	Specifies the order of the matrix A. n must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,n). Before entry with uplo = 'U' or 'u', the leading n by n upper triangular part of the array A must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of A is not referenced. Before entry with uplo = 'L' or 'l', the leading n by n lower triangular part of the array A must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of A is not referenced.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least max(1,n).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx)) Before entry, the incremented array X must contain the n element vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta. When beta is supplied as zero then y need not be set on input.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incy)) Before entry, the incremented array Y must contain the n element vector y. On exit, Y is overwritten by the updated vector y.
incy	Specifies the increment for the elements of y. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.21 BLAS_SYR Procedures

This procedure performs the rank 1 operation `A := alpha*x*x' + A` , where `alpha` is a real scalar, `x` is an `n` element vector, and `A` is an `n` by `n` symmetric matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SYR (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DBL,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SYR (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLT,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-27 BLAS_SYR Procedure Parameters

Parameter	Description
`uplo`	Specifies whether the upper or lower triangular part of the array `A` is to be referenced: • `uplo` = `'U'` or `'u'` : Only the upper triangular part of `A` is to be referenced. • `uplo` = `'L'` or `'l'` : Only the lower triangular part of A is to be referenced.
`n`	Specifies the order of the matrix `A`. `n` must be at least zero.
`alpha`	Specifies the scalar alpha.
`x`	`UTL_NLA_ARRAY_FLT/DBL` of dimension at least `(1+(n-1)*abs(incx))` Before entry, the incremented array `X` must contain the m element vector `x`.
`incx`	Specifies the increment for the elements of `x`. `incx` must not be zero.

Table 293-27 (Cont.) BLAS_SYR Procedure Parameters

Parameter	Description
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n) Before entry with uplo = 'U' or 'u', the leading n by n upper triangular part of the array A must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of A is not referenced. On exit, the upper triangular part of the array A is overwritten by the upper triangular part of the updated matrix. Before entry with uplo = 'L' or 'l', the leading n by n lower triangular part of the array A must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of A is not referenced. On exit, the lower triangular part of the array A is overwritten by the lower triangular part of the updated matrix.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least max(1, n)
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.22 BLAS_SYR2 Procedures

This procedure performs the rank 2 operation A := alpha*x*y' + alpha*y*x' + A , where alpha is a scalar, x and y are n element vectors, and A is an n by n symmetric matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SYR2 (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DBL,
   x      IN      UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   y      IN      UTL_NLA_ARRAY_DBL,
   incy   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SYR2 (
   uplo   IN      flag,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLT,
   x      IN      UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
```

```
  y      IN      UTL_NLA_ARRAY_FLT,
  incy   IN      POSITIVEN,
  a      IN OUT  UTL_NLA_ARRAY_FLT,
  lda    IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-28 BLAS_SYR2 Procedure Parameters

Parameter	Description
uplo	Specifies whether the upper or lower triangular part of the array A is to be referenced: • uplo = 'U' or 'u' : Only the upper triangular part of A is to be referenced. • uplo = 'L' or 'l' : Only the lower triangular part of A is to be referenced.
n	Specifies the order of the matrix A. n must be at least zero.
alpha	Specifies the scalar alpha.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incx)) Before entry, the incremented array X must contain the m element vector x.
incx	Specifies the increment for the elements of x. incx must not be zero.
y	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1)*abs(incy)) Before entry, the incremented array Y must contain the m element vector y.
incy	Specifies the increment for the elements of y. incy must not be zero.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n) With uplo = 'U' or 'u', the leading n by n upper triangular part of the array A must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of A is not referenced. On exit, the upper triangular part of the array A is overwritten by the upper triangular part of the updated matrix. With uplo = 'L' or 'l', the leading n by n lower triangular part of the array A must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of A is not referenced. On exit, the lower triangular part of the array A is overwritten by the lower triangular part of the updated matrix.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least max(1, n)
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.23 BLAS_SYR2K Procedures

It performs one of the symmetric rank2 k operations `C := alpha*A*B' + alpha*B*A' + beta*C` or `C := alpha*A'*B + alpha*B'*A + beta*C`, where `alpha` and `beta` are scalars, `C` is an `n` by `n` symmetric matrix and `A` and `B` are `n` by `k` matrices in the first case and `k` by `n` matrices in the second case.

> **See Also:**
>
> BLAS Level 3 (Matrix-Matrix Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SYR2K (
   uplo   IN      flag,
   trans  IN      flag,
   n      IN      POSITIVEN,
   k      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   b      IN      UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   c      IN OUT  UTL_NLA_ARRAY_DBL,
   ldc    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SYR2K (
   uplo   IN      flag,
   trans  IN      flag,
   n      IN      POSITIVEN,
   k      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   b      IN OUT  UTL_NLA_ARRAY_FLT,
   ldb    IN      POSITIVEN,
   beta   IN      SCALAR_FLOAT,
   c      IN OUT  UTL_NLA_ARRAY_FLT,
   ldc    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-29 BLAS_SYR2K Procedure Parameters

Parameter	Description
uplo	Specifies whether the upper or lower triangular part of the array C is to be referenced: • uplo = 'U' or 'u' : Only the upper triangular part of C is to be referenced. • uplo = 'L' or 'l' : Only the lower triangular part of C is to be referenced.
trans	Specifies the operations to be performed: • trans = 'N' or 'n'C : C := alpha*A*B' + alpha*B*A' + beta*C • trans = 'T' or 't'C : C := alpha*A'*B + alpha*B'*A + beta*C • trans = 'C' or 'c'C : C := alpha*A'*B + alpha*B'*A + beta*C
n	Specifies the order of matrix C. n must be at least zero.
k	On entry with trans = 'N' or 'n', k specifies the number of columns of the matrices A and B. On entry with trans = 'T' or 't' or trans = 'C' or 'c', k specifies the number of rows of the matrices A and B. k must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,ka) where kb is k when trans = 'N' or 'n', and is n otherwise. Before entry with trans = 'N' or 'n', the leading n byk part of the array A must contain the matrix A, otherwise the leading k by n part of the array A must contain the matrix A.
lda	Specifies the first dimension of a as declared in the calling (sub) program. When trans = 'N' or 'n', lda must be at leastmax(1,n), otherwise lda must be at least max(1,k).
b	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,kb) where kb is k when trans = 'N' or 'n', and is n otherwise. Before entry with trans = 'N' or 'n', the leading n byk part of the array B must contain the matrix B, otherwise the leading k by n part of the array B must contain the matrix B.
ldb	Specifies the first dimension of b as declared in the calling (sub) program. When trans = 'N' or 'n', ldb must be at least max (1,n), otherwise ldb must be at least max (1,k).
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta.

Table 293-29 (Cont.) BLAS_SYR2K Procedure Parameters

Parameter	Description
c	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldc,n). Before entry with uplo = 'U' or 'u', the leading n by n upper triangular part of the array C must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of C is not referenced. On exit, the upper triangular part of the array C is overwritten by the upper triangular part of the updated matrix. Before entry with uplo = 'L' or 'l', the leading n by n lower triangular part of the array C must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of C is not referenced. On exit, the lower triangular part of the array C is overwritten by the lower triangular part of the updated matrix.
ldc	Specifies the first dimension of C as declared in the calling (sub) program. ldc must be at least max(1,n).
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.24 BLAS_SYRK Procedures

This procedure performs one of the symmetric rank k operations C := alpha*A*A' + beta*C or C := alpha*A'*A + beta*C, where alpha and beta are scalars, C is an n by n symmetric matrix andA is an n by k matrix in the first case and a k byn matrix in the second case.

> **See Also:**
>
> BLAS Level 3 (Matrix-Matrix Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_SYRK (
   uplo   IN      flag,
   trans  IN      flag,
   n      IN      POSITIVEN,
   k      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   beta   IN      SCALAR_DOUBLE,
   c      IN OUT  UTL_NLA_ARRAY_DBL,
   ldc    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_SYRK (
   uplo   IN      flag,
   trans  IN      flag,
   n      IN      POSITIVEN,
   k      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
```

```
  a      IN      UTL_NLA_ARRAY_FLT,
  lda    IN      POSITIVEN,
  beta   IN      SCALAR_FLOAT,
  c      IN OUT  UTL_NLA_ARRAY_DBL,
  ldc    IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-30 BLAS_SYRK Procedure Parameters

Parameter	Description
uplo	Specifies whether the upper or lower triangular part of the array C is to be referenced: • uplo = 'U' or 'u' : Only the upper triangular part of C is to be referenced. • uplo = 'L' or 'l' : Only the lower triangular part of C is to be referenced.
trans	Specifies the operations to be performed: • trans = 'N' or 'n' : C := alpha*A*A' + beta*C • trans = 'T' or 't' : C := alpha*A'*A + beta*C • trans ='C' or 'c' : C := alpha*A'*A + beta*C
n	Specifies the order of matrix C. n must be at least zero.
k	On entry with trans = 'N' or 'n', k specifies the number of columns of the matrix A. On entry with trans = 'T' or 't' or trans = 'C' or 'c', k specifies the number of rows of the matrix A. k must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,ka) where ka is k when trans = 'N' or 'n', and is n otherwise. Before entry with trans = 'N' or 'n', the leading n byk part of the array A must contain the matrix A, otherwise the leading k by n part of the array A must contain the matrix A.
lda	Specifies the first dimension of a as declared in the calling (sub) program. When trans = 'N' or 'n', lda must be at least max(1,n), otherwiselda must be at least max(1,k).
beta	SCALAR_FLOAT/DOUBLE. Specifies the scalar beta.
c	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldc,n). Before entry with uplo = 'U' or 'u', the leading n by n upper triangular part of the array C must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of C is not referenced. On exit, the upper triangular part of the array C is overwritten by the upper triangular part of the updated matrix. Before entry with uplo = 'L' or 'l', the leading n by n lower triangular part of the array C must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of C is not referenced. On exit, the lower triangular part of the array C is overwritten by the lower triangular part of the updated matrix.
ldc	Specifies the first dimension of C as declared in the calling (sub) program. ldc must be at least max(1,n).

Table 293-30 (Cont.) BLAS_SYRK Procedure Parameters

Parameter	Description
`pack`	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.25 BLAS_TBMV Procedures

This procedure performs the matrix-vector operations `x := A*x` or `x := A'*x`, where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular band matrix, with `(k+1)` diagonals.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TBMV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   k      IN      NATURALN,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_TBMV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   k      IN      NATURALN,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-31 BLAS_TBMV Procedure Parameters

Parameter	Description
uplo	Specifies whether the matrix is an upper or lower triangular matrix: • uplo = 'U' or 'u'. A is an upper triangular matrix. • uplo = 'L' or 'l'. A is a lower triangular matrix.
trans	Specifies the operation to be performed: • trans = 'N' or 'n'x := A*x • trans = 'T' or 't'x := A'*x • trans = 'C' or 'c'x := A'*x
diag	Specifies whether or not A is unit triangular: • diag = 'U' or 'u'. A is assumed to be unit triangular. • diag = 'N' or 'n'. A is not assumed to be unit triangular.
n	Specifies the order of the matrix A. n must be at least zero.
k	Specifies whether or not A is unit triangular: • with uplo = 'U' or 'u', K specifies the number of super-diagonals of the matrix A. • with uplo = 'L' or 'l', K specifies the number of sub-diagonals of the matrix A. K must satisfy 0 .le. k.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n). Before entry with uplo = 'U' or 'u', the leading (k+1) by n part of the array A must contain the upper triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row (k+1) of the array, the first super-diagonal starting at position 2 in row k, and so on. The top left k by k triangle of the array A is not referenced. Before entry with uplo = 'L' or 'l', the leading (k+1) by n part of the array A must contain the lower triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right k by k triangle of the array A is not referenced. Note that when diag = ''U' or 'u', the elements of the array A corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity.
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least (k+1).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx)). Before entry, the incremented array X must contain the n element vector x. On exit, X is overwritten with the transformed vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.26 BLAS_TBSV Procedures

This procedure solves one of the systems of equations `A*x = b` or `A'*x = b`, where `b` and `x` are `n` element vectors and `A` is an `n` by n unit, or non-unit, upper or lower triangular band matrix, with `(k+1)` diagonals.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TBSV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   k      IN      NATURALN,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_STBSV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   k      IN      NATURALN,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-32 BLAS_TBSV Procedure Parameters

Parameter	Description
`uplo`	Specifies whether the matrix is an upper or lower triangular matrix: • `uplo` = `'U'` or `'u'`. `A` is an upper triangular matrix. • `uplo` = `'L'` or `'l'`. `A` is a lower triangular matrix.
`trans`	Specifies the equations to be solved: • `trans` = `'N'` or `'n : 'A*x = b` • `trans` = `'T'` or `'t': A'*x = b` • `trans` = `'C'` or `'c': A'*x = b`
`diag`	Specifies whether or not `A` is unit triangular: • `diag` = `'U'` or `'u'` : A is assumed to be unit triangular. • `diag` = `'N'` or `'n'` : `A` is not assumed to be unit triangular.

Table 293-32 (Cont.) BLAS_TBSV Procedure Parameters

Parameter	Description
n	Specifies the order of the matrix A. n must be at least zero.
k	Specifies whether or not A is unit triangular: • with uplo = 'U' or 'u', K specifies the number of super-diagonals of the matrix A. • with uplo = 'L' or 'l', K specifies the number of sub-diagonals of the matrix A. K must satisfy 0 .le. k.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,n). Before entry with uplo = 'U' or 'u', the leading (k+1) by n part of the array A must contain the upper triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row (k+1) of the array, the first super-diagonal starting at position 2 in row k, and so on. The top left k by k triangle of the array A is not referenced. Before entry with uplo = 'L' or 'l', the leading (k+1) by n part of the array A must contain the lower triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right k by k triangle of the array A is not referenced. Note that when diag = 'U' or 'u', the elements of the array A corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity.
lda	On entry, lda specifies the first dimension of A as declared in the calling (sub) program. lda must be at least (k+1).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1) *abs(incx)) Before entry, the incremented array X must contain the n element right-hand side vector b. On exit, X is overwritten with the solution vector x.
incx	Specifies the increment for the elements of x. incx must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

Usage Notes

No test for singularity or near-singularity is included in this routine. Such tests must be performed before calling this routine.

293.5.27 BLAS_TPMV Procedures

This procedure performs the matrix-vector operations `x := A*x` or `x := A'*x`, where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix, supplied in packed form.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TPMV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   ap     IN      UTL_NLA_ARRAY_DBL,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_TBMV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   ap     IN      UTL_NLA_ARRAY_FLT,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-33 BLAS_TPMV Procedure Parameters

Parameter	Description
`uplo`	Specifies whether the matrix is an upper or lower triangular matrix: • `uplo` = `'U'` or `'u'`. `A` is an upper triangular matrix. • `uplo` = `'L'` or `'l'`. `A` is a lower triangular matrix.
`trans`	Specifies the operation to e performed: • `trans` = `'N'` or `'n'` `x := A*x` • `trans` = `'T'` or `'t'` `x := A'*x` • `trans` = `'C'` or `'c'` `x := A'*x`
`diag`	Specifies whether or not `A` is unit triangular: • `diag` = `'U'` or `'u'`. A is assumed to be unit triangular. • `diag` = `'N'` or `'n'`. `A` is not assumed to be unit triangular.
`n`	Specifies the order of the matrix `A`. `n` must be at least zero.

Table 293-33 (Cont.) BLAS_TPMV Procedure Parameters

Parameter	Description
ap	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,n). Before entry with uplo = 'U' or 'u', the leading (k+1) by n part of the array A must contain the upper triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row (k+1) of the array, the first super-diagonal starting at position 2 in row k, and so on. The top left k by k triangle of the array A is not referenced. Before entry with uplo = 'L' or 'l', the leading (k+1) by n part of the array A must contain the lower triangular band part of the matrix of coefficients, supplied column by column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right k by k triangle of the array A is not referenced. Note that when diag = 'U' or 'u', the elements of the array A corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*abs(incx)). Before entry, the incremented array X must contain the n element vector x. On exit, X is overwritten with the transformed vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.28 BLAS_TPSV Procedures

This procedure solves one of the systems of equations A*x = b or A'*x = b, where b and x are n element vectors and A is an n by n unit, or non-unit, upper or lower triangular matrix, supplied in packed form.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TPSV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   ap     IN      UTL_NLA_ARRAY_DBL,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

```
UTL_NLA.BLAS_TPSV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   ap     IN      UTL_NLA_ARRAY_FLT,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-34 BLAS_TPSV Procedure Parameters

Parameter	Description
uplo	Specifies whether the matrix is an upper or lower triangular matrix: • uplo = 'U' or 'u' : A is an upper triangular matrix. • uplo = 'L' or 'l' : A is a lower triangular matrix.
trans	Specifies the operation to be performed: • trans = 'N' or 'n' : A*x = b • trans = 'T' or 't' : A'*x = b • trans = 'C' or 'c' : A'*x = b
diag	Specifies whether or not A is unit triangular: • diag = 'U' or 'u' : A is assumed to be unit triangular. • diag = 'N' or 'n' : 'A is not assumed to be unit triangular.
n	Specifies the order of the matrix A. n must be at least zero.
ap	UTL_NLA_ARRAY_FLT/DBL of dimension at least ((n*(n+1))/2) Before entry with uplo = 'U' or 'u', the array ap must contain the upper triangular matrix packed sequentially, column by column, so that ap(1) contains a(1,1), ap(2) and ap(3) contain a(1,2) and a(2,2) respectively, and so on. Before entry with uplo = 'L' or 'l', the array ap must contain the lower triangular matrix packed sequentially, column by column, so that ap(1) contains a(1,1), ap(2) and ap(3) contain a(2,1) and a(3,1) respectively, and so on. Note that when diag = 'U' or 'u', the diagonal elements of A are not referenced, but are assumed to be unity.
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1) *abs(incx)) Before entry, the incremented array X must contain the n element right-hand side vector b. On exit, X is overwritten with the solution vector x.
incx	Specifies the increment for the elements of x. incx must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

Usage Notes

No test for singularity or near-singularity is included in this routine. Such tests must be performed before calling this routine.

293.5.29 BLAS_TRMM Procedures

This procedure performs a matrix-matrix operation.

It performs one of the following matrix-matrix operations:

```
B := alpha*op( A )*B
```

or

```
B := alpha*B*op( A )
```

where `alpha` is a scalar, `B` is an `m` by `n` matrix, `A` is a unit, or non-unit, upper or lower triangular matrix and `op(A)` is one of

```
op( A ) = A
```

or

```
op( A ) = A'
```

> **See Also:**
>
> BLAS Level 3 (Matrix-Matrix Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TRMM (
   side   IN      flag,
   uplo   IN      flag,
   transa IN      flag,
   diag   IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   b      IN OUT  UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_TRMM (
   side   IN      flag,
   uplo   IN      flag,
   transa IN      flag,
   diag   IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
```

```
  b      IN OUT  UTL_NLA_ARRAY_FLT,
  ldb    IN      POSITIVEN,
  pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-35 BLAS_TRMM Procedure Parameters

Parameter	Description
side	Specifies whether the symmetric matrix A appears on the left or right in the operation: • side = 'L' or 'l' : B := alpha*op(A)*B • side = 'R' or 'r' : B := alpha*B*op(A)
uplo	Specifies whether the upper or lower triangular part of the array A is to be referenced: • uplo = 'U' or 'u' : A is an upper triangular matrix. • uplo = 'L' or 'l'' : A is a lower triangular matrix.
transa	Specifies the form of op(A) to be used in the matrix multiplication as follows: • transa = 'N' or 'n' : op(A) = A • transa = 'T' or 't' : op(A) = A' • transa ='C' or 'c' : op(A) = A'
diag	Specifies whether or not A is unit triangular: • diag = 'U' or 'u'. A is assumed to be unit triangular. • diag = 'N' or 'n'. A is not assumed to be unit triangular.
m	Specifies the number of rows of the B. m must be at least zero.
n	Specifies the number of columns of B. n must be at least zero.
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha. When alpha is zero then A is not referenced and B need not be set before entry.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda,k) where k is m when side = 'L' or 'l', and is n when side = ''R' or 'r'. Before entry with uplo = 'U' or 'u' , the leading k by k upper triangular part of the array A must contain the upper triangular matrix, and the strictly lower triangular part of A is not referenced. Before entry with uplo = 'L' or 'l', the leading k byk lower triangular part of the array A must contain the lower triangular matrix and the strictly upper triangular part of A is not referenced. Note that when diag = ''U' or 'u', the diagonal elements of A are not referenced either, but are assumed to be unity.
lda	Specifies the first dimension of a as declared in the calling (sub) program. When side = 'L' or 'l', lda must be at least max(1,m), otherwise lda must be at least max(1,n).
b	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldb,n). Before entry, the leading m by n part of the array B must contain the matrix B, and on exit is overwritten by the transformed matrix.
ldb	Specifies the first dimension of b as declared in the calling (sub) program. ldb must be at least max(1,m).
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.30 BLAS_TRMV Procedures

This procedure performs the matrix-vector operations `x := A*x` or `x := A'*x`, where `x` is an `n` element vector and `A` is an `n` by `n` unit, or non-unit, upper or lower triangular matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TRMV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_TRMV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-36 BLAS_TRMV Procedure Parameters

Parameter	Description
`uplo`	Specifies whether the matrix is an upper or lower triangular matrix: • `uplo` = `'U'` or `'u'`. `A` is an upper triangular matrix. • `uplo` = `'L'` or `'l'`. `A` is a lower triangular matrix.
`trans`	Specifies the operation to be performed: • `trans` = `'N'` or `'n'` `x := A*x` • `trans`= `'T'` or `'t'` `x := A'*x` • `trans` = `'C'` or `'c'` `x := A'*x`
`diag`	Specifies whether or not `A` is unit triangular: • `diag` = `'U'` or `'u'`. A is assumed to be unit triangular. • `diag` = `'N'` or `'n'`. `A` is not assumed to be unit triangular.
`n`	Specifies the order of the matrix `A`. `n` must be at least zero.

Table 293-36 (Cont.) BLAS_TRMV Procedure Parameters

Parameter	Description
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n). Before entry with uplo = 'U' or 'u', the leading n by n upper triangular part of the array A must contain the upper triangular matrix and the strictly lower triangular part of A is not referenced. Before entry with uplo = 'L' or 'l', the leading n by n lower triangular part of the array A must contain the lower triangular matrix and the strictly upper triangular part of A is not referenced. Note that when diag = 'U' or 'u', the diagonal elements of A are not referenced either, but are assumed to be unity
lda	Specifies the first dimension of a as declared in the calling (sub) program. lda must be at least max(1,n).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1+(n-1)*as(incx)). Before entry, the incremented array X must contain the n element vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.31 BLAS_TRSM Procedures

This procedure performs a matrix-matrix operation.

It performs one of the matrix-matrix operations:

```
op( A )*X = alpha*B
```

or

```
X*op( A ) = alpha*B
```

where alpha is a scalar, X and B are m by n matrices, A is a unit, or non-unit, upper or lower triangular matrix and op(A) is one of

```
op( A ) = A
```

or

```
op( A ) = A'
```

The matrix X is overwritten on B.

> **See Also:**
>
> BLAS Level 3 (Matrix-Matrix Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TRSM (
   side   IN      flag,
   uplo   IN      flag,
   transa IN      flag,
   diag   IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_DOUBLE,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   b      IN OUT  UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.BLAS_TRSM (
   side   IN      flag,
   uplo   IN      flag,
   transa IN      flag,
   diag   IN      flag,
   m      IN      POSITIVEN,
   n      IN      POSITIVEN,
   alpha  IN      SCALAR_FLOAT,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   b      IN OUT  UTL_NLA_ARRAY_FLT,
   ldb    IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-37 BLAS_TRSM Procedure Parameters

Parameter	Description
side	Specifies whether the symmetric matrix A appears on the left or right in the operation: • side = 'L' or 'l' : op(A)*X = alpha*B • side = 'R' or 'r' : X*op(A) = alpha*B
uplo	Specifies whether the upper or lower triangular part of the array A is to be referenced: • uplo = 'U' or 'u' : A is an upper triangular matrix. • uplo = 'L' or 'l' : A is a lower triangular matrix.
transa	Specifies the form of op(A) to be used in the matrix multiplication as follows: • transa = 'N' or 'n' : op(A) = A • transa = 'T' or 't' : op(A) = A' • transa = 'C' or 'c' : op(A) = A'
diag	Specifies whether or not A is unit triangular: • diag = 'U' or 'u'. A is assumed to be unit triangular. • diag = 'N' or 'n'. A is not assumed to be unit triangular.
m	Specifies the number of rows of the B. m must be at least zero.
n	Specifies the number of columns of B. n must be at least zero.

Table 293-37 (Cont.) BLAS_TRSM Procedure Parameters

Parameter	Description
alpha	SCALAR_FLOAT/DOUBLE. Specifies the scalar alpha. When alpha is zero then A is not referenced and B need not be set before entry.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, k) where k is m when side = 'L' or 'l', and is n when side = 'R' or 'r'. Before entry with uplo = 'U' or 'u', the leading k by k upper triangular part of the array A must contain the upper triangular matrix, and the strictly lower triangular part of A is not referenced. Before entry with uplo = 'L' or 'l', the leading k by k lower triangular part of the array A must contain the lower triangular matrix and the strictly upper triangular part of A is not referenced. Note that when diag = 'U' or 'u', the diagonal elements of A are not referenced either, but are assumed to be unity.
lda	Specifies the first dimension of a as declared in the calling (sub) program. When side = 'L' or 'l', lda must be at least max(1, m), otherwise lda must be at least max (1, n).
b	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (ldb, n). Before entry, the leading m by n part of the array B must contain the matrix B, and on exit is overwritten by the solution matrix X.
ldb	Specifies the first dimension of b as declared in the calling (sub) program. ldb must be at least max(1, m).
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.32 BLAS_TRSV Procedures

This procedure solves one of the systems of equations A*x = b or A'*x = b, where b and x are n element vectors and A is an n by n unit, or non-unit, upper or lower triangular matrix.

> **See Also:**
>
> BLAS Level 2 (Matrix-Vector Operations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.BLAS_TRSV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   a      IN      UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_DBL,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

```
UTL_NLA.BLAS_TRSV (
   uplo   IN      flag,
   trans  IN      flag,
   diag   IN      flag,
   n      IN      POSITIVEN,
   a      IN      UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   x      IN OUT  UTL_NLA_ARRAY_FLT,
   incx   IN      POSITIVEN,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-38 BLAS_TRSV Procedure Parameters

Parameter	Description
uplo	Specifies whether the matrix is an upper or lower triangular matrix: • uplo = 'U' or 'u'. A is an upper triangular matrix. • uplo = 'L' or 'l'. A is a lower triangular matrix.
trans	Specifies the operation to be performed: • trans = 'N' or 'n'A*x = b • trans = 'T' or 't'A'*x = b • trans = 'C' or 'c'A'*x = b
diag	Specifies whether or not A is unit triangular: • diag = 'U' or 'u'. A is assumed to be unit triangular. • diag = 'N' or 'n'. A is not assumed to be unit triangular.
n	Specifies the order of the matrix A. n must be at least zero.
a	UTL_NLA_ARRAY_FLT/DBL of DIMENSION (lda, n). Before entry with uplo = 'U' or 'u', the leading n by n upper triangular part of the array A must contain the upper triangular matrix and the strictly lower triangular part of A is not referenced. Before entry with uplo = 'L' or 'l', the leading n by n lower triangular part of the array A must contain the lower triangular matrix and the strictly upper triangular part of A is not referenced. Note that when diag = 'U' or 'u', the diagonal elements of A are not referenced either, but are assumed to be unity.
lda	Specifies the first dimension of A as declared in the calling (sub) program. lda must be at least max(1, n).
x	UTL_NLA_ARRAY_FLT/DBL of dimension at least (1 + (n - 1) * abs (incx)) Before entry, the incremented array X must contain the n element right-hand side vector b. On exit, X is overwritten with the solution vector x.
incx	Specifies the increment for the elements of x. Must not be zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

Usage Notes

No test for singularity or near-singularity is included in this routine. Such tests must be performed before calling this routine.

293.5.33 LAPACK_GBSV Procedures

This procedure computes the solution to a real system of linear equations `a * x = b`, where `a` is a band matrix of order `n` with `kl` sub diagonals and `ku` superdiagonals, and `x` and `b` are `n` by `nrhs` matrices.

The `LU` decomposition with partial pivoting and row interchanges is used to factor `A` as

```
a = L * U
```

where L is a product of permutation and unit lower triangular matrices with kl sub diagonals, and `U` is upper triangular with `kl+ku` superdiagonals. The factored form of `a` is then used to solve the system of equations

```
a * x = b
```

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GBSV (
   n      IN      POSITIVEN,
   kl     IN      NATURALN,
   ku     IN      NATURALN,
   nrhs   IN      POSITIVEN,
   ab     IN OUT  UTL_NLA_ARRAY_DBL,
   ldab   IN      POSITIVEN,
   ipiv   IN OUT  UTL_NLA_ARRAY_INT,
   b      IN OUT  UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
   info   OUT  INTEGER,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_GBSV (
   n      IN      POSITIVEN,
   kl     IN      NATURALN,
   ku     IN      NATURALN,
   nrhs   IN      POSITIVEN,
   ab     IN OUT  UTL_NLA_ARRAY_FLT,
   ldab   IN      POSITIVEN,
   ipiv   IN OUT  UTL_NLA_ARRAY_INT,
   b      IN OUT  UTL_NLA_ARRAY_FLT,
   ldb    IN      POSITIVEN,
   info   OUT  INTEGER,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-39 LAPACK_GBSV Procedure Parameters

Parameter	Description
n	The number of linear equations, equivalent to the order of the matrixa .n >= 0
kl	The number of sub diagonals within the band of a. kl >= 0.
ku	The number of superdiagonals within the band of a . ku >= 0.
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
ab	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldab, n). On entry, the matrix a in band storage, in rows kl+1 to 2*kl+ku+1; rows 1 to kl of the array need not be set. The j-th column of A is stored in the j-th column of the array ab: ab(kl+ku+1+i-j,j) = a(i,j) for max(1,j-ku)<=i<=min(n,j+kl) On exit, details of the factorization: U is stored as an upper triangular band matrix with kl+ku superdiagonals in rows 1 to KL+KU+1, and the multipliers used during the factorization are stored in rows: kl+ku+2 to 2*kl+ku+1
ldab	The leading dimension of the array ab. ldab >= 2*kl+ku+1
ipiv	INTEGER array, DIMENSION (n). The pivot indices that define the permutation matrix P; row iof the matrix was interchanged with row ipiv(i).
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, U(i,i) is exactly zero. The factorization has been completed, but the factor U is exactly singular, and the solution has not been computed
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.34 LAPACK_GEES Procedures

This procedure computes for an n by n real nonsymmetric matrix A, the eigenvalues, the real Schur form T, and, optionally, the matrix of Schur vectors Z.

This gives the Schur factorization A = Z*T*(Z**T).

A matrix is in real Schur form if it is upper quasi-triangular with `1` by `1` and `2` by `2` blocks. `2` by `2` blocks will be standardized in the form

```
[ a b ]

[ c a ]
```

where `b*c < 0`. The eigenvalues of such a block are `a +- sqrt(bc)`.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GEES (
   jobvs    IN       flag,
   n        IN       POSITIVEN,
   a        IN  OUT  UTL_NLA_ARRAY_DBL,
   lda      IN       POSITIVEN,
   wr       IN OUT   UTL_NLA_ARRAY_DBL,
   wi       IN OUT   UTL_NLA_ARRAY_DBL,
   vs       IN OUT   UTL_NLA_ARRAY_DBL,
   ldvs     IN       POSITIVEN,
   info     OUT      INTEGER,
   pack     IN       flag DEFAULT 'C');

UTL_NLA.LAPACK_GEES (
   jobvs    IN       flag,
   n        IN       POSITIVEN,
   a        IN OUT   UTL_NLA_ARRAY_FLT,
   lda      IN       POSITIVEN,
   wr       IN       OUT UTL_NLA_ARRAY_FLT,
   wi       IN       OUT UTL_NLA_ARRAY_FLT,
   vs       IN OUT   UTL_NLA_ARRAY_FLT,
   ldvs     IN       POSITIVEN,
   info     OUT      integer,
   pack     IN       flag DEFAULT 'C');
```

Parameters

Table 293-40 LAPACK_GEES Procedure Parameters

Parameter	Description
`jobz`	• `'N'`: Schur vectors are not computed. • `'V'`: Schur vectors are computed.
`n`	The order of the matrix `a`. `N >= 0`.
a	`UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n)`. • On entry, the `n` by `n` matrix `A`. • On exit, `A` has been overwritten by its real Schur form `T`.
`lda`	The leading dimension of the array `a`. `lda >= max(1,n)`.

Table 293-40 (Cont.) LAPACK_GEES Procedure Parameters

Parameter	Description
wr	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). wr and wi contain the real and imaginary parts respectively of the computed eigenvalues in the same order that they appear on the diagonal of the output Schur form T. Complex conjugate pairs of eigenvalues will appear consecutively with the eigenvalue having the positive imaginary part first.
wi	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz, n). wr and wi contain the real and imaginary parts respectively of the computed eigenvalues in the same order that they appear on the diagonal of the output Schur form T. Complex conjugate pairs of eigenvalues will appear consecutively with the eigenvalue having the positive imaginary part first.
vs	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • If jobvs = 'V', vs contains the orthogonal matrix Z of Schur vectors. • If jobvs = 'N', vs is not referenced.
ldvs	The leading dimension of the array vs. VS. ldvs >= 1. If jobvs = 'V', ldvs >= N
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, and i is <= N: the QR algorithm failed to compute all the eigenvalues. Elements 1:ILO-1 and i+1:N of wr and wi contain those eigenvalues which have converged. If jobvs = 'V', vs contains the matrix which reduces A to its partially converged Schur form.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.35 LAPACK_GELS Procedures

This procedure solves overdetermined or underdetermined real linear systems involving an m by n matrix A, or its transpose, using a QR or LQ factorization of A. It is assumed that A has full rank.

The following options are provided:

1. If TRANS = 'N' and m >= n: find the least squares solution of an overdetermined system, that is, solve the least squares problem.

   ```
   minimize || B - A*X ||
   ```

2. If TRANS = 'N' and m < n: find the minimum norm solution of an underdetermined system A * X = B.
3. If TRANS = 'T' and m >= n: find the minimum norm solution of an undetermined system A**T * X = B.
4. If TRANS = 'T' and m < n: find the least squares solution of an overdetermined system, that is, solve the least squares problem minimize || B - A**T * X ||.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GELS (
   trans    IN       flag,
   m        IN       POSITIVEN,
   n        IN       POSITIVEN,
   nrhs     IN       POSITIVEN,
   a        IN OUT   UTL_NLA_ARRAY_DBL,
   lda      IN       POSITIVEN,
   b        IN OUT   UTL_NLA_ARRAY_DBL,
   ldb      IN       POSITIVEN,
   info     OUT      INTEGER,
   pack     IN       flag DEFAULT 'C');

UTL_NLA.LAPACK_GELS (
  trans     IN       flag,
  m         IN       POSITIVEN,
  n         IN       POSITIVEN,
  nrhs      IN       POSITIVEN,
  a         IN OUT   UTL_NLA_ARRAY_FLT,
  lda       IN       POSITIVEN,
  b         IN OUT   UTL_NLA_ARRAY_FLT,
  ldb       IN       POSITIVEN,
  info      OUT      INTEGER,
  pack      IN       flag DEFAULT 'C');
```

Parameters

Table 293-41 LAPACK_GELS Procedure Parameters

Parameter	Description
trans	• CHARACTER = 'N': The linear system involves A. • CHARACTER = 'T': The linear system involves A**T .
m	The number of rows of the matrix a. M >= 0.
n	The number of columns of the matrix a. N >= 0.
nrhs	The number of right-hand sides, which is the number of columns of the matrix band x.nrhs >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). On entry, the matrix b of right hand side vectors, stored columnwise; b is m by nrhs if TRANS = 'N', or n by nrhsif trans = 'T'. On exit, if m >= n, a is overwritten by details of its QR factorization as returned by SGEQRF. If m < n, A is overwritten by details of its LQ factorization as returned by SGELQF.
lda	The leading dimension of the array A. lda >= max(1,m).

Table 293-41 (Cont.) LAPACK_GELS Procedure Parameters

Parameter	Description
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the matrix b of right hand side vectors, stored columnwise. b is m bynrhs if trans = 'n', or n by nrhs if trans = 'T'. On exit, b is overwritten by the solution vectors, stored columnwise: • If trans = 'n' and m >= n, rows 1 to n of b contain the least squares solution vectors; the residual sum of squares for the solution in each column is given by the sum of squares of elements n+1 to m in that column. • If trans = 'n' and m < n, rows 1 to n of b contain the minimum norm solution vectors. • If trans = 'T' and m >= n, rows 1 to m of b contain the minimum norm solution vectors. • If trans = 'T' and m < n, rows 1 to m of b contain the least squares solution vectors; the residual sum of squares for the solution in each column is given by the sum of squares of elements m+1 to n in that column.
ldb	The leading dimension of the array b. ldb >= max(1,m,n)
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.36 LAPACK_GESDD Procedures

This procedures computes the singular value decomposition (SVD) of a real m by n matrix A, optionally computing the left and right singular vectors. If singular vectors are desired, it uses a divide-and-conquer algorithm that makes mild assumptions about floating point arithmetic.

The SVD is written

```
A = U * SIGMA * transpose(V)
```

where SIGMA is an m by n matrix which is zero except for its min(m,n) diagonal elements, U is anm by m orthogonal matrix, and V is ann by n orthogonal matrix. The diagonal elements of SIGMA are the singular values of A, they are real and non-negative, and are returned in descending order. The first min(m,n) columns of U and V are the left and right singular vectors of A.

Note that the routine returns V**T, not V.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GESDD (
   jobz   IN       flag,
   m      IN       POSITIVEN,
   n      IN       POSITIVEN,
   a      IN OUT   UTL_NLA_ARRAY_DBL,
   lda    IN       POSITIVEN,
   s      IN OUT   UTL_NLA_ARRAY_DBL,
   u      IN OUT   UTL_NLA_ARRAY_DBL,
   ldu    IN       POSITIVEN,
   vt     IN OUT   UTL_NLA_ARRAY_DBL,
   ldvt   IN       POSITIVEN,
   info   OUT      INTEGER,
   pack   IN       flag DEFAULT 'C');

UTL_NLA.LAPACK_GESDD (
   jobz   IN       flag,
   m      IN       POSITIVEN,
   n      IN       POSITIVEN,
   a      IN OUT   UTL_NLA_ARRAY_FLT,
   lda    IN       POSITIVEN,
   s      IN OUT   UTL_NLA_ARRAY_FLT,
   u      IN OUT   UTL_NLA_ARRAY_FLT,
   ldu    IN       POSITIVEN,
   vt     IN OUT   UTL_NLA_ARRAY_FLT,
   ldvt   IN       POSITIVEN,
   info   OUT      INTEGER,
   pack   IN       flag DEFAULT 'C');
```

Parameters

Table 293-42 LAPACK_GESDD Procedure Parameters

Parameter	Description
jobz	Specifies options for computing all or part of the matrix U: • 'A': All m columns of u and all n rows of V**T are returned in arrays u and vt. • 'S': The first min(m,n) columns of u and the first min(m,n) rows of V**T are returned in the arrays u and vt. • 'O': The first min(m,n) columns of u (the left singular vectors) are overwritten on the array a. jobu and jobvt cannot both be 'O • 'N': No columns of u (no left singular vectors) are computed.
m	The order of the matrix a. m >= 0.
n	The order of the matrix a. n >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). On entry, the n by n matrix A. On exit: • If jobz = 'O', a is overwritten with the first min(m,n) columns of u (the left singular vectors, stored columnwise). • If m >= n, a is overwritten with the first m rows of V**T (the right singular vectors, stored rowwise). • If jobz .ne. 'O', the contents of a are destroyed.
lda	The leading dimension of the array a. lda >= max(1,m).

Table 293-42 (Cont.) LAPACK_GESDD Procedure Parameters

Parameter	Description
s	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (min(m,n)). The singular values of a, sorted so that S(i) >= S(i+1).
u	UTL_NLA_ARRAY_FLT/DBL. ucol = m if jobz = 'A' or jobz = 'O' and m < n; ucol = min(m,n) if jobz = 'S'. • If jobz = 'A' or jobz = 'O' and m < n, u contains the m by m orthogonal matrix u. • If jobz = 'S', u contains the first min(m,n) columns of u (the left singular vectors, stored columnwise). • If jobz = 'O' and m >= n, or jobz = 'n', u is not referenced.
ldu	The leading dimension of the array U. ldu >= 1. If jobz = 'S' or 'A',or jobz = 'O' and m < n, ldu >= m.
vt	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldvt, n). • If jobz = 'A' or jobz = 'O' and m >= n, vt contains the n by n orthogonal matrix V**T. • If jobz = 'S', vt contains the first min(m,n) rows of V**T (the right singular vectors, stored rowwise). • If jobz = 'O' and m < n, or jobz = 'N', vt is not referenced.
ldvt	The leading dimension of the array vt. ldvt >= 1. • If jobz = 'A', or jobz = 'O' and m >= n, ldvt >= n. • If jobz = 'S', ldvt >= min(m,n).
info	• = 0 : successful exit • < 0 : If info = -i, the i-th argument had an illegal value • > 0 : SBDSDC did not converge, updating process failed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.37 LAPACK_GESV Procedure

This procedure computes the solution to a real system of linear equations a * x = b, where a is an n by n matrix and x and b are n by nrhs matrices.

The LU decomposition with partial pivoting and row interchanges is used to factor A as

```
a = P * L * U
```

where P is a permutation matrix, L is unit lower triangular, and U is upper triangular. The factored form of a is then used to solve the system of equations

```
a * x = b
```

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GESV (
   n      IN      POSITIVEN,
   nrhs   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_DBL,
   lda    IN      POSITIVEN,
   ipiv   IN OUT  UTL_NLA_ARRAY_INT,
   b      IN OUT  UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
   info      OUT  INTEGER,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_GESV (
   n      IN      POSITIVEN,
   nrhs   IN      POSITIVEN,
   a      IN OUT  UTL_NLA_ARRAY_FLT,
   lda    IN      POSITIVEN,
   ipiv   IN OUT  UTL_NLA_ARRAY_INT,
   b      IN OUT  UTL_NLA_ARRAY_FLT,
   ldb    IN      POSITIVEN,
   info      OUT  INTEGER,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-43 LAPACK_GESV Procedure Parameters

Parameter	Description
n	The number of linear equations, equivalent to the order of the matrix a. n >= 0
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). On entry, the n by n coefficient matrix a. On exit, the factors L and U from the factorization a = P*L*U; the unit diagonal elements of L are not stored.
lda	The leading dimension of the array a. lda >= max(1,n)
ipiv	INTEGER array, DIMENSION (n). The pivot indices that define the permutation matrix P; row iof the matrix was interchanged with row ipiv(i).
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, U(i,i) is exactly zero. The factorization has been completed, but the factor U is exactly singular, so the solution could not be computed.

Table 293-43 (Cont.) LAPACK_GESV Procedure Parameters

Parameter	Description
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.38 LAPACK_GESVD Procedures

This procedures computes the singular value decomposition (SVD) of a real m by n matrix A, optionally computing the left and/or right singular vectors.

The SVD is written

```
A = U * SIGMA * transpose(V)
```

where SIGMA is an m by n matrix which is zero except for its min(m,n) diagonal elements, U is anm by m orthogonal matrix, and V is ann by n orthogonal matrix. The diagonal elements of SIGMA are the singular values of A, they are real and non-negative, and are returned in descending order. The first min(m,n) columns of U and V are the left and right singular vectors of A.

Note that the routine returns V**T, not V.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GESVD (
   jobu   IN       flag,
   jobvt  IN       flag,
   m      IN       POSITIVEN,
   n      IN       POSITIVEN,
   a      IN OUT   UTL_NLA_ARRAY_DBL,
   lda    IN       POSITIVEN,
   s      IN OUT   UTL_NLA_ARRAY_DBL,
   u      IN OUT   UTL_NLA_ARRAY_DBL,
   ldu    IN       POSITIVEN,
   vt     IN OUT   UTL_NLA_ARRAY_DBL,
   ldvt   IN       POSITIVEN,
   info   OUT      INTEGER,
   pack   IN       flag DEFAULT 'C');

UTL_NLA.LAPACK_GESVD (
   jobu   IN       flag,
   jobvt  IN       flag,
   m      IN       POSITIVEN,
   n      IN       POSITIVEN,
   a      IN OUT   UTL_NLA_ARRAY_FLT,
   lda    IN       POSITIVEN,
   s      IN OUT   UTL_NLA_ARRAY_FLT,
```

```
  u      IN OUT   UTL_NLA_ARRAY_FLT,
  ldu    IN       POSITIVEN,
  vt     IN OUT   UTL_NLA_ARRAY_FLT,
  ldvt   IN       POSITIVEN,
  info   OUT      INTEGER,
  pack   IN       flag DEFAULT 'C');
```

Parameters

Table 293-44 LAPACK_GESVD Procedure Parameters

Parameter	Description
jobu	Specifies options for computing all or part of the matrix U: • 'A': All m columns of U are returned in array U. • 'S': The first min(m,n) columns of U (the left singular vectors) are returned in the array U. • 'O': The first min(m,n) columns of U (the left singular vectors) are overwritten on the array a. jobu and jobvt cannot both be 'O' • 'N': No columns of U (no left singular vectors) are computed.
jobvt	Specifies options for computing all or part of the matrix V**T: • 'A': All n rows of V**T are returned in the array vt. • 'S': The first min(m,n) rows of V**T (the right singular vectors) are returned in the array vt. • 'O': The first min(m,n) rows of V**T (the right singular vectors) are overwritten on the array a. jobvt and jobu cannot both be 'O'. • 'N': No rows of V**T (no right singular vectors) are computed.
m	The order of the matrix a. M >= 0.
n	The order of the matrix a. N >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). On entry, the n by n matrix A. On exit: • If jobu = 'O', A is overwritten with the first min(m,n) columns of U (the left singular vectors, stored columnwise); • If jobvt = 'O', A is overwritten with the first min(m,n) rows of V**T (the right singular vectors, stored rowwise); • If jobu.ne.'O' and jobvt.ne.'O', the contents of A are destroyed.
lda	The leading dimension of the array a. lda >= max(1,n).
s	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (min(m,n)). The singular values of A, sorted so that S(i) >= S(i+1).
u	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldu,ucol).(ldu,m)if jobu = 'A' or (ldu,min(m,n)) if jobu = 'S'. • If jobu = 'A', U contains the m by m orthogonal matrix U. • If jobu = 'S', U contains the first min(m,n) columns of U (the left singular vectors, stored columnwise). • If jobu = 'N' or 'O', U is not referenced.
ldu	The leading dimension of the array U. ldu >= 1. If jobu = 'S' or 'a', ldu >= m.

Table 293-44 (Cont.) LAPACK_GESVD Procedure Parameters

Parameter	Description
vt	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldvt, n). • If jobvt = 'A', vt contains the n by n orthogonal matrix V**T. • If jobvt = 'S', vt contains the first min(m,n) rows of V**T (the right singular vectors, stored rowwise). • If jobvt = 'N' or 'O', vt is not referenced.
ldvt	The leading dimension of the array vt. ldvt >= 1. • If jobvt = 'A', ldvt >= n. • If jobvt = 'S', ldvt >= min(m,n).
info	• = 0 : successful exit • < 0 : If info = -i, the i-th argument had an illegal value • > 0 : If SBDSQR did not converge, info specifies how many superdiagonals of an intermediate bidiagonal form B did not converge to zero.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.39 LAPACK_GEEV Procedures

This procedures computes for an n by n real nonsymmetric matrix A, the eigenvalues and, optionally, the left and/or right eigenvectors.

- The right eigenvector v(j) of A satisfies A * v(j) = lambda(j) * v(j) where lambda(j) is its eigenvalue.
- The left eigenvector u(j) of A satisfies u(j)**H * A = lambda(j) * u(j)**H where u(j)**H denotes the conjugate transpose of u(j).

The computed eigenvectors are normalized to have Euclidean norm equal to 1 and largest component real.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GEEV (
   jobvl    IN      flag,
   jobvr    IN      flag,
   n        IN      POSITIVEN,
   a        IN OUT  UTL_NLA_ARRAY_DBL,
   lda      IN      POSITIVEN,
   wr       IN OUT  UTL_NLA_ARRAY_DBL,
   wi       IN OUT  UTL_NLA_ARRAY_DBL,
   vl       IN OUT  UTL_NLA_ARRAY_DBL,
   ldvl     IN      POSITIVEN,
   vr       IN OUT  UTL_NLA_ARRAY_DBL,
```

```
   ldvr    IN      POSITIVEN,
   info    OUT     INTEGER,
   pack    IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_GEEV (
   jobvl   IN      flag,
   jobvr   IN      flag,
   n       IN      POSITIVEN,
   a       IN OUT  UTL_NLA_ARRAY_FLT,
   lda     IN      POSITIVEN,
   wr      IN OUT  UTL_NLA_ARRAY_FLT,
   wi      IN OUT  UTL_NLA_ARRAY_FLT,
   vl      IN OUT  UTL_NLA_ARRAY_FLT,
   ldvl    IN      POSITIVEN,
   vr      IN OUT  UTL_NLA_ARRAY_FLT,
   ldvr    IN      POSITIVEN,
   info    OUT     INTEGER,
   pack    IN      flag DEFAULT 'C');
```

Parameters

Table 293-45 LAPACK_GEEV Procedure Parameters

Parameter	Description
jobvl	• 'N': Left eigenvectors of A are not computed. • 'V': Left eigenvectors of A are computed.
jobvr	• 'N': Right eigenvectors of A are not computed. • 'V': Right eigenvectors of A are computed.
n	The order of the matrix a. N >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). • On entry, the n by n matrix A. • On exit, A has been overwritten.
lda	The leading dimension of the array a. lda >= max(1,n).
wr	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). wr and wi contain the real and imaginary parts respectively of the computed eigenvalues. Complex conjugate pairs of eigenvalues will appear consecutively with the eigenvalue having the positive imaginary part first.
wi	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz, n). wr and wi contain the real and imaginary parts respectively of the computed eigenvalues. Complex conjugate pairs of eigenvalues will appear consecutively with the eigenvalue having the positive imaginary part first.
vl	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • If jobvl = 'V', the left eigenvectors u(j) are stored one after another in the columns of vl, in the same order as their eigenvalues. • If jobvs = 'N', vl is not referenced. • If the j-th eigenvalue is real, then u(j) = VL(:,j), the j-th column of vl. • If the j-th and (j+1)-st eigenvalues form a complex conjugate pair, then u(j) = VL(:,j) + i*VL(:,j+1) and u(j+1) = VL(:,j) - i*VL(:,j+1).

Table 293-45 (Cont.) LAPACK_GEEV Procedure Parameters

Parameter	Description
ldv1	The leading dimension of the array vl. ldvl >= 1. If jobvl = 'v', ldvl >= n.
vr	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldvr, n). • If jobvr = 'V', the right eigenvectors v(j) are stored one after another in the columns of vr, in the same order as their eigenvalues.. • If jobvr = 'N', vr is not referenced. • If the j-th eigenvalue is real, then v(j) = VR(:,j), the j-th column of vr. • If the j-th and (j+1)-st eigenvalues form a complex conjugate pair, then v(j) = VR(:,j) + i*VR(:,j+1) and v(j+1) = VR(:,j) - i*VR(:,j+1).
ldvr	The leading dimension of the array vr. vr.ldvr >= 1. If jobvr = 'V', ldvr >= N
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, and i is <= N: the QR algorithm failed to compute all the eigenvalues, and no eigenvectors have been computed. Elements i+1:N of wr and wi contain eigenvalues which have converged..
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.40 LAPACK_GTSV Procedure

This procedure solves the equation a * x = b, where a is an n by n tridiagonal matrix, by Gaussian elimination with partial pivoting.

Note that the equation a'*x = b may be solved by interchanging the order of the arguments du and dl.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_GTSV (
   n      IN      POSITIVEN,
   nrhs   IN      POSITIVEN,
   dl     IN OUT  UTL_NLA_ARRAY_DBL,
   d      IN OUT  UTL_NLA_ARRAY_DBL,
   du     IN OUT  UTL_NLA_ARRAY_DBL,
   b      IN OUT  UTL_NLA_ARRAY_DBL,
   ldb    IN      POSITIVEN,
```

```
   info      OUT  INTEGER,
   pack   IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_GTSV (
   n      IN      POSITIVEN,
   nrhs   IN      POSITIVEN,
   dl     IN OUT  UTL_NLA_ARRAY_FLT,
   d      IN OUT  UTL_NLA_ARRAY_FLT,
   du     IN OUT  UTL_NLA_ARRAY_FLT,
   b      IN OUT  UTL_NLA_ARRAY_FLT,
   ldb    IN      POSITIVEN,
   info      OUT  INTEGER,
   pack   IN      flag DEFAULT 'C');
```

Parameters

Table 293-46 LAPACK_GTSV Procedure Parameters

Parameter	Description
n	The order of the matrix a .n >= 0
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
dl	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n-1). On entry, dl must contain the (n-1) sub-diagonal elements of a. On exit, dl is overwritten by the (n-2) elements of the second super-diagonal of the upper triangular matrix U from the LU factorization of a, in dl(1), ..., dl(n-2).
d	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). On entry, d must contain the diagonal elements of a. On exit, d is overwritten by the n diagonal elements of U.
du	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n-1). On entry, du must contain the (n-1) super-diagonal elements of a. On exit, du is overwritten by the (n-1) elements of the first super-diagonal of U.
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (LDB, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0, the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max (1, n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, U(i,i) is exactly zero, and the solution has not been computed. The factorization has not been completed unless i = n.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.41 LAPACK_PBSV Procedures

This procedure computes the solution to a real system of linear equations `a * x = b`, where `a` is an `n` by `n` symmetric positive definite band matrix and `x` and `b` are `n` by `nrhs` matrices.

The Cholesky decomposition is used to factor `A` as

`A = U**T*U` if `UPLO ='U'`

or

`A = L * L**T` if `UPLO = 'L'`

where `U` is an upper triangular matrix and `L` is a lower triangular matrix. The factored form of `A` is then used to solve the system of equations `A * X = B`.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_PBSV (
uplo        IN      flag,
n           IN      POSITIVEN,
kd          IN      NATURALN,
nrhs        IN      POSITIVEN,
ab          IN OUT  UTL_NLA_ARRAY_DBL,
ldab        IN      POSITIVEN,
b           IN OUT  UTL_NLA_ARRAY_DBL,
ldb         IN      POSITIVEN,
info        OUT     INTEGER,
pack        IN      flag DEFAULT 'C');


UTL_NLA.LAPACK_PBSV (
uplo       IN      flag,
n          IN      POSITIVEN,
kd         IN      NATURALN,
nrhs       IN      POSITIVEN,
ab         IN OUT  UTL_NLA_ARRAY_FLT,
ldab       IN      POSITIVEN,
b          IN OUT  UTL_NLA_ARRAY_FLT,
ldb        IN      POSITIVEN,
info       OUT     INTEGER,
pack       IN      flag DEFAULT 'C');
```

Parameters

Table 293-47 LAPACK_PBSV Procedure Parameters

Parameter	Description
`uplo`	• `uplo = 'U'`. Upper triangular of `A` is stored. • `uplo = 'L'`. Lower triangular of `A` is stored.

Table 293-47 (Cont.) LAPACK_PBSV Procedure Parameters

Parameter	Description
n	The number of linear equations, that is, the order of the matrix a .n >= 0
kd	The number of superdiagonals of the matrix A if uplo = 'U', or the number of subdiagonals if UPLO = 'L'. KD >= 0.
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
ab	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldab, n). On entry, the upper or lower triangle of the symmetric band matrix a, stored in the first kd+1 rows of the array. The j-th column of a is stored in the j-th column of the array ab is as follows: • if uplo = 'U', AB(KD+1+i-j,j) = A(i,j) for max(1,j-KD)<=i<=j; • if uplo = 'L', AB(1+i-j,j) = A(i,j) for j<=i<=min(N,j+KD) .See below for further details.On exit, if info = 0, the triangular factor U or L from the Cholesky factorization A = U**T*U or A = L*L**T of the bandmatrix A, in the same storage format as a.
ldab	The leading dimension of the array ab. ldb >= kd+1
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, the leading minor of order i of a is not positive definite, so the factorization could not be completed, and the solution has not been computed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.42 LAPACK_POSV Procedures

This procedure computes the solution to a real system of linear equations a * x = b, where a is an n by n symmetric positive definite matrix and x and b are n by nrhs matrices.

The Cholesky decomposition is used to factor A as

A = U**T* U if uplo = 'U'

or

A = L * L**T if UPLO = 'L'

where U is an upper triangular matrix and L is a lower triangular matrix. The factored form of A is then used to solve the system of equations A * X = B.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_POSV (
   uplo       IN      flag,
   n          IN      POSITIVEN,
   nrhs       IN      POSITIVEN,
   a          IN OUT  UTL_NLA_ARRAY_DBL,
   lda        IN      POSITIVEN,
   b          IN OUT  UTL_NLA_ARRAY_DBL,
   ldb        IN      POSITIVEN,
   info       OUT     INTEGER,
   pack       IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_POSV (
   uplo       IN      flag,
   n          IN      POSITIVEN,
   nrhs       IN      POSITIVEN,
   a          IN OUT  UTL_NLA_ARRAY_FLT,
   lda        IN      POSITIVEN,
   b          IN OUT  UTL_NLA_ARRAY_FLT,
   ldb        IN      POSITIVEN,
   info       OUT     INTEGER,
   pack       IN      flag DEFAULT 'C');
```

Parameters

Table 293-48 LAPACK_POSV Procedure Parameters

Parameter	Description
uplo	• uplo = 'U'. Upper triangular of A is stored. • uplo = 'L'. Lower triangular of A is stored.
n	The number of linear equations, that is, the order of the matrix a .n >= 0
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). If uplo = 'U', the leading NRHS n by n upper triangular part of a contains the upper NRHS triangular part of the matrix A, and the strictly lower NRHS triangular part of A is not referenced. If uplo = ''L', then rhs leading n by n lower triangular part of a contains the lower nrhs triangular part of the matrix a, and the strictly upper nrhs triangular part of a is not referenced. On exit, if info = 0, the factor U or L from the Cholesky factorization A = U**T*U or A = L*L**T.

Table 293-48 (Cont.) LAPACK_POSV Procedure Parameters

Parameter	Description
lda	The leading dimension of the array a. lda >= max (1, n)
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, the leading minor of order i of a is not positive definite, so the factorization could not be completed, and the solution has not been computed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.43 LAPACK_PPSV Procedures

This procedure computes the solution to a real system of linear equations a * x = b where a is an n by n symmetric positive definite matrix stored in packed format and x and b are n by nrhs matrices.

The Cholesky decomposition is used to factor A as

A = U**T* U if UPLO = 'U'

or

A = L * L**T if UPLO = 'L'

where **U** is an upper triangular matrix and L is a lower triangular matrix. The factored form of A is then used to solve the system of equations A * X = B.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_PPSV (
   uplo      IN      flag,
   n         IN      POSITIVEN,
   nrhs      IN      POSITIVEN,
   ap        IN OUT  UTL_NLA_ARRAY_DBL,
   b         IN OUT  UTL_NLA_ARRAY_DBL,
   ldb       IN      POSITIVEN,
   info      OUT     INTEGER,
```

```
   pack      IN      flag DEFAULT 'C');


UTL_NLA.LAPACK_PPSV (
   uplo      IN      flag,
   n         IN      POSITIVEN,
   nrhs      IN      POSITIVEN,
   ap        IN OUT  UTL_NLA_ARRAY_FLT,
   b         IN OUT  UTL_NLA_ARRAY_FLT,
   ldb       IN      POSITIVEN,
   info      OUT     INTEGER,
   pack      IN      flag DEFAULT 'C');
```

Parameters

Table 293-49 LAPACK_PPSV Procedure Parameters

Parameter	Description
uplo	• uplo = 'U' . Upper triangular of A is stored. • uplo = 'L'. Lower triangular of A is stored.
n	The number of linear equations, that is, the order of the matrix a .n >= 0
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
ap	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n*(n+1)/2). On entry, the upper or lower triangle of the symmetric matrix a, packed columnwise in a linear array. The j-th column of a is stored in the array ap as follows: If uplo = 'U', AP(i + (j-1)*j/2) = A(i,j) for 1<=i<=j; If uplo = 'L', AP(i + (j-1)*(2n-j)/2) = A(i,j) for j<=i<=n; On exit, if info = 0, the factor U or 'L' from the Cholesky factorization A = U**T*U or A = L*L**T in the same storage format as A.
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, the leading minor of order i of a is not positive definite, so the factorization could not be completed, and the solution has not been computed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.44 LAPACK_PTSV Procedures

This procedure computes the solution to a real system of linear equations `a * x = b`, where `a` is an `n` by `n` symmetric positive definite tridiagonal matrix, and `x` and `b` are `n` by `nrhs` matrices.

`a` is factored as `A = L*D*L**T`, and the factored form of a is then used to solve the system of equations.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_PTSV (
   n      IN       POSITIVEN,
   nrhs   IN       POSITIVEN,
   d      IN OUT   UTL_NLA_ARRAY_DBL,
   e      IN OUT   UTL_NLA_ARRAY_DBL,
   b      IN OUT   UTL_NLA_ARRAY_DBL,
   ldb    IN       POSITIVEN,
   info   OUT      INTEGER,
   pack   IN       flag  DEFAULT 'C');


UTL_NLA.LAPACK_PTSV (
   n      IN       POSITIVEN,
   nrhs   IN       POSITIVEN,
   d      IN OUT   UTL_NLA_ARRAY_FLT,
   e      IN OUT   UTL_NLA_ARRAY_FLT,
   b      IN OUT   UTL_NLA_ARRAY_FLT,
   ldb    IN       POSITIVEN,
   info   OUT      INTEGER,
   pack   IN       flag DEFAULT 'C');
```

Parameters

Table 293-50 LAPACK_PTSV Procedure Parameters

Parameter	Description
`n`	The order of the matrix `a`. `N >= 0`.
`nrhs`	The number of right-hand sides, which is the number of columns of the matrix `b`. `nrhs >= 0`.
`d`	`UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n)`. On entry, the `n` diagonal elements of the tridiagonal matrix `a`. On exit, the `n` diagonal elements of the diagonal matrix `d` from the factorization `A = L*D*L**T`.

Table 293-50 (Cont.) LAPACK_PTSV Procedure Parameters

Parameter	Description
e	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n-1). On entry, the (n-1) subdiagonal elements of the tridiagonal matrix a. On exit, the (n-1) diagonal elements of the unit bidiagonal factor L from the factorization A = L*D*L**T of a.(e can also be regarded as the superdiagonal of the unit bidiagonal factor U from the U**T*D*U factorization of a)
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, the leading minor of order i of a is not positive definite, so the factorization could not be completed, and the solution has not been computed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.45 LAPACK_SBEV Procedures

This procedure computes all the eigenvalues and, optionally, eigenvectors of a real symmetric band matrix A.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SBEV (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   kd       IN      NATURALN,
   ab       IN OUT  UTL_NLA_ARRAY_DBL,
   ldab     IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_DBL,
   z        IN OUT  UTL_NLA_ARRAY_DBL,
   ldz       IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SBEV (
   jobz     IN      flag,
```

```
  uplo     IN      flag,
  n        IN      POSITIVEN,
  kd       IN      NATURALN,
  ab       IN OUT  UTL_NLA_ARRAY_FLT,
  ldab     IN      POSITIVEN,
  w        IN OUT  UTL_NLA_ARRAY_FLT,
  z        IN OUT  UTL_NLA_ARRAY_FLT,
  ldz      IN      POSITIVEN,
  info     OUT     INTEGER,
  pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-51 LAPACK_SBEV Procedure Parameters

Parameter	Description
jobz	• 'N': Compute eigenvalues only. • 'V': Compute eigenvalues and eigenvectors.
uplo	• 'U': Upper triangle of A is stored. • 'L': Lower triangle of A is stored.
n	The order of the matrix a. N >= 0.
kd	The number of superdiagonals of the matrix A if uplo = 'U', or the number of subdiagonals if uplo = 'L'. kd >= 0.
ab	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldab, n). On entry, the upper or lower triangle of the symmetric band matrix A stored in the first kd+1 rows of the array. The j-th column of A is stored in the j-th column of the array ab: • If uplo = 'U', ab(kd+1+i-j,j) = a(i,j) for max(1,j-kd)<=i<=j. • If uplo = 'L', AB(1+i-j,j) = A(i,j) for j<=i<=min(n,j+kd). On exit, ab is overwritten by values generated during the reduction to tridiagonal form: • If uplo = 'U', the diagonal and first superdiagonal of the tridiagonal matrix T are returned in rows kd and kd+1 of ab. • If uplo = 'L', the diagonal and first subdiagonal of T are returned in the first two rows of ab.
ldab	The leading dimension of the array ab. ldab >= kd + 1.
w	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). If info = 0, the eigenvalues in ascending order.
z	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • If jobz = 'V', then if info = 0, z contains the orthonormal eigenvectors of the matrix A, with the i-th column of z holding the eigenvector associated with w(i). • If jobz = 'N', then z is not referenced.
ldz	The leading dimension of the array z. ldz >= 1, and if jobz = 'v', ldz >= max(1,n).
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero

Table 293-51 (Cont.) LAPACK_SBEV Procedure Parameters

Parameter	Description
`pack`	(Optional) Flags the packing of the matrices: • `'C'`: column-major (default) • `'R'`: row-major

293.5.46 LAPACK_SBEVD Procedures

This procedure computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SBEVD (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   kd       IN      NATURALN,
   ab       IN OUT  UTL_NLA_ARRAY_DBL,
   ldab     IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_DBL,
   z        IN OUT  UTL_NLA_ARRAY_DBL,
   ldz       IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SBEVD (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   kd       IN      NATURALN,
   ab       IN OUT  UTL_NLA_ARRAY_FLT,
   ldab     IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_FLT,
   z        IN OUT  UTL_NLA_ARRAY_FLT,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-52 LAPACK_SBEVD Procedure Parameters

Parameter	Description
jobz	• 'N': Compute eigenvalues only. • 'V': Compute eigenvalues and eigenvectors.
uplo	• 'U': Upper triangle of A is stored. • 'L': Lower triangle of A is stored.
n	The order of the matrix a. N >= 0.
kd	The number of superdiagonals of the matrix A if uplo = 'U', or the number of subdiagonals if uplo = 'L'. kd >= 0.
ab	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldab, n). On entry, the upper or lower triangle of the symmetric band matrix A stored in the first kd+1 rows of the array. The j-th column of A is stored in the j-th column of the array ab: • If uplo = 'U', ab(kd+1+i-j,j) = a(i,j) for max(1,j-kd)<=i<=j. • If uplo = 'L', AB(1+i-j,j) = A(i,j) for j<=i<=min(n,j+kd). On exit, ab is overwritten by values generated during the reduction to tridiagonal form: • If uplo = 'U', the diagonal and first superdiagonal of the tridiagonal matrix T are returned in rows kd and kd+1 of ab. • If uplo = 'L', the diagonal and first subdiagonal of T are returned in the first two rows of ab.
ldab	The leading dimension of the array ab. ldab >= kd + 1.
w	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz,n). If info = 0, the eigenvalues in ascending order.
z	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • If jobz = 'V', then if info = 0, z contains the orthonormal eigenvectors of the matrix A, with the i-th column of z holding the eigenvector associated with w(i). • If jobz = 'N', then z is not referenced.
ldz	The leading dimension of the array z. ldz >= 1, and if jobz = 'v', ldz >= max(1,n).
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.47 LAPACK_SPEV Procedures

This procedure computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A` in packed storage.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SPEV (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   ap       IN OUT  UTL_NLA_ARRAY_DBL,
   w        IN OUT  UTL_NLA_ARRAY_DBL,
   z        IN OUT  UTL_NLA_ARRAY_DBL,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SPEV (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   ap       IN OUT  UTL_NLA_ARRAY_FLT,
   w        IN OUT  UTL_NLA_ARRAY_FLT,
   z        IN OUT  UTL_NLA_ARRAY_FLT,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-53 LAPACK_SPEV Procedure Parameters

Parameter	Description
`jobz`	• `'N'`: Compute eigenvalues only. • `'V'`: Compute eigenvalues and eigenvectors.
`uplo`	• `'U'`: Upper triangle of `A` is stored. • `'L'`: Lower triangle of A is stored.
`n`	The order of the matrix `a. N >= 0`.

Table 293-53 (Cont.) LAPACK_SPEV Procedure Parameters

Parameter	Description
ap	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n*(n+1)/2). On entry, the upper or lower triangle of the symmetric matrix a packed columnwise in a linear array. The j-th column of a is stored in the array ap: • If uplo = 'U', ap(i + (j-1)*j/2) = a(i,j) for 1<=i<=j. • If uplo = 'L', ap(i + (j-1)*(2*n-j)/2) = a(i,j) for j<=i<=n. On exit, ap is overwritten by values generated during the reduction to tridiagonal form: • If uplo = 'U', the diagonal and first superdiagonal of the tridiagonal matrix T overwrite the corresponding elements of A. • If uplo = 'L', the diagonal and first subdiagonal of T overwrite the corresponding elements of A.
w	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). If info = 0, the eigenvalues in ascending order.
z	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz,n). • If jobz = 'V', then if info = 0, z contains the orthonormal eigenvectors of the matrix A, with the i-th column of z holding the eigenvector associated with w(i). • If jobz = 'N', then z is not referenced.
ldz	The leading dimension of the array z. ldz >= 1, and if jobz = 'v', ldz >= max(1,n).
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.48 LAPACK_SPEVD Procedures

This procedure computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix A in packed storage. If eigenvectors are desired, it uses a divide and conquer algorithm. The divide and conquer algorithm makes very mild assumptions about floating point arithmetic.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SPEVD (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   ap       IN OUT  UTL_NLA_ARRAY_DBL,
   w        IN OUT  UTL_NLA_ARRAY_DBL,
   z        IN OUT  UTL_NLA_ARRAY_DBL,
   ldz       IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SPEVD (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   ap       IN OUT  UTL_NLA_ARRAY_FLT,
   w        IN OUT  UTL_NLA_ARRAY_FLT,
   z        IN OUT  UTL_NLA_ARRAY_FLT,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-54 LAPACK_SPEVD Procedure Parameters

Parameter	Description
jobz	• 'N': Compute eigenvalues only. • 'V': Compute eigenvalues and eigenvectors.
uplo	• 'U': Upper triangle of A is stored. • 'L': Lower triangle of A is stored.
n	The order of the matrix a. N >= 0.
ap	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n*(n+1)/2). On entry, the upper or lower triangle of the symmetric matrix a packed columnwise in a linear array. The j-th column of a is stored in the array ap: • If uplo = 'U', ap(i + (j-1)*j/2) = a(i,j) for 1<=i<=j. • If uplo = 'L', ap(i + (j-1)*(2*n-j)/2) = a(i,j) for j<=i<=n. On exit, ap is overwritten by values generated during the reduction to tridiagonal form: • If uplo = 'U', the diagonal and first superdiagonal of the tridiagonal matrix T overwrite the corresponding elements of A. • If uplo = 'L', the diagonal and first subdiagonal of T overwrite the corresponding elements of A.
w	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). If info = 0, the eigenvalues in ascending order.

Table 293-54 (Cont.) LAPACK_SPEVD Procedure Parameters

Parameter	Description
z	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz,n). • If jobz = 'V', then if info = 0, z contains the orthonormal eigenvectors of the matrix A, with the i-th column of z holding the eigenvector associated with w(i). • If jobz = 'N', then z is not referenced.
ldz	The leading dimension of the array z. ldz >= 1, and if jobz = 'v', ldz >= max(1,n).
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.49 LAPACK_SPSV Procedures

This procedure computes the solution to a real system of linear equations a * x = b, where a is an n by n symmetric matrix stored in packed format, and x and b are n by nrhs matrices.

The diagonal pivoting method is used to factor A as

A = U * D * U**T, if UPLO = 'U'

or

A = L * D * L**T, if UPLO = 'L'

where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is symmetric and block diagonal with 1 by 1 and 2 by 2 diagonal blocks. The factored form of A is then used to solve the system of equations A * X = B.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SPSV (
   uplo     IN      flag,
   n        IN      POSITIVEN,
   nrhs     IN      POSITIVEN,
   ap       IN OUT  UTL_NLA_ARRAY_DBL,
   ipiv     IN OUT  UTL_NLA_ARRAY_INT,
   b        IN OUT  UTL_NLA_ARRAY_DBL,
   ldb      IN      POSITIVEN,
```

```
   info    OUT     INTEGER,
   pack    IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SPSV (
   uplo    IN      flag,
   n       IN      POSITIVEN,
   nrhs    IN      POSITIVEN,
   ap      IN OUT  UTL_NLA_ARRAY_FLT,
   ipiv    IN OUT  UTL_NLA_ARRAY_INT,
   b       IN OUT  UTL_NLA_ARRAY_FLT,
   ldb     IN      POSITIVEN,
   info    OUT     INTEGER,
   pack    IN      flag DEFAULT 'C');
```

Parameters

Table 293-55 LAPACK_SPSV Procedure Parameters

Parameter	Description
uplo	• uplo = 'U'. Upper triangular of A is stored. • uplo = 'L' . Lower triangular of A is stored.
n	The number of linear equations, which is the order of the matrix a. N >= 0.
nrhs	The number of right-hand sides, which is the number of columns of the matrix b. nrhs >= 0.
ap	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n*(n+1)/2). On entry, the upper or lower triangle of the symmetric matrix A, packed columnwise in a linear array. The j-th column of A is stored in the array ap as follows: • uplo = 'U': AP(i + (j-1)*j/2) = A(i,j) for 1<=i<=j • uplo = 'L' : AP(i + (j-1)*(2n-j)/2) = A(i,j) for j<=i<=n See below for further details. On exit, the block diagonal matrix D and the multipliers used to obtain the factor U or L from the factorization A = U*D*U**T or A = L*D*L**T as computed by SSPTRF, stored as a packed triangular matrix in the same storage format as A.
ipiv	INTEGER array, DIMENSION (n). Details of the interchanges and the block structure of d, as determined by SSPTRF. • If ipiv(k) > 0, then rows and columns k and ipiv(k) were interchanged, and d(k,k) is a 1 by1 diagonal block. • If uplo = 'U' and ipiv(k) = ipiv(k-1) < 0, then rows and columns k-1 and -ipiv(k) were interchanged and d(k-1:k,k-1:k) is a 2 by 2 diagonal block. • If uplo = 'L' and ipiv(k) = ipiv(k+1) < 0, then rows and columns k+1 and -ipiv(k) were interchanged and d(k:k+1,k:k+1) is a 2 by 2 diagonal block.
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)

Table 293-55 (Cont.) LAPACK_SPSV Procedure Parameters

Parameter	Description
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, d(i,i) is exactly zero. The factorization has been completed, but the block diagonal matrix d is exactly singular, so the solution could not be computed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.50 LAPACK_STEV Procedures

This procedure computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix A.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_STEV (
   jobz     IN      flag,
   n        IN      POSITIVEN,
   d        IN OUT  UTL_NLA_ARRAY_DBL,
   e        IN OUT  UTL_NLA_ARRAY_DBL,
   z        IN OUT  UTL_NLA_ARRAY_DBL,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_STEV (
   jobz     IN      flag,
   n        IN      POSITIVEN,
   d        IN OUT  UTL_NLA_ARRAY_FLT,
   e        IN OUT  UTL_NLA_ARRAY_FLT,
   z        IN OUT  UTL_NLA_ARRAY_FLT,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-56 LAPACK_STEV Procedure Parameters

Parameter	Description
jobz	• 'N': Compute eigenvalues only. • 'V': Compute eigenvalues and eigenvectors.

Table 293-56 (Cont.) LAPACK_STEV Procedure Parameters

Parameter	Description
n	The order of the matrix a. N >= 0.
d	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • On entry, the n diagonal elements of the tridiagonal matrix A. • On exit, if info = 0, the eigenvalues in ascending order.
e	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • On entry, the (n-1) subdiagonal elements of the tridiagonal matrix A, stored in elements 1 to n-1 of e. e(n) need not be set, but is used by the subprogram. • On exit, the contents of e are destroyed.
z	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz, n). • If jobz = 'V', then if info = 0, z contains the orthonormal eigenvectors of the matrix A, with the i-th column of z holding the eigenvector associated with d(i). • If jobz = 'N', then z is not referenced.
ldz	The leading dimension of the array z. ldz >= 1, and if jobz = 'v', ldz >= max(1,n).
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.51 LAPACK_STEVD Procedures

This procedure computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_STEVD (
   jobz     IN      flag,
   n        IN      POSITIVEN,
   d        IN OUT  UTL_NLA_ARRAY_DBL,
   e        IN OUT  UTL_NLA_ARRAY_DBL,
   z        IN OUT  UTL_NLA_ARRAY_DBL,
   ldz      IN      POSITIVEN,
```

```
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_STEVD(
   jobz     IN      flag,
   n        IN      POSITIVEN,
   d        IN OUT  UTL_NLA_ARRAY_FLT,
   e        IN OUT  UTL_NLA_ARRAY_FLT,
   z        IN OUT  UTL_NLA_ARRAY_FLT,
   ldz      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-57 LAPACK_STEVD Procedure Parameters

Parameter	Description
jobz	• 'N': Compute eigenvalues only. • 'V': Compute eigenvalues and eigenvectors.
n	The order of the matrix a. N >= 0.
d	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • On entry, the n diagonal elements of the tridiagonal matrix A. • On exit, if info = 0, the eigenvalues in ascending order.
e	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). • On entry, the (n-1) subdiagonal elements of the tridiagonal matrix A, stored in elements 1 to n-1 of e. e(n) need not be set, but is used by the subprogram. • On exit, the contents of e are destroyed.
z	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldz, n). • If jobz = 'V', then if info = 0, z contains the orthonormal eigenvectors of the matrix A, with the i-th column of z holding the eigenvector associated with d(i). • If jobz = 'N', then z is not referenced.
ldz	The leading dimension of the array z. ldz >= 1, and if jobz = 'v', ldz >= max(1,n).
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.52 LAPACK_SYEV Procedures

This procedure computes all eigenvalues and, optionally, eigenvectors of a real symmetric matrix `A`.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SYEV (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   a        IN OUT  UTL_NLA_ARRAY_DBL,
   lda      IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_DBL,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SYEV (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   a        IN OUT  UTL_NLA_ARRAY_FLT,
   lda      IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_FLT,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-58 LAPACK_SYEV Procedure Parameters

Paramete	Description
`jobz`	• `'N'`: Compute eigenvalues only. • `'V'`: Compute eigenvalues and eigenvectors.
`uplo`	• `'U'`: Upper triangle of A is stored. • `'L'`: Upper triangle of A is stored.
`n`	The order of the matrix `a`. `N >= 0`.

Table 293-58 (Cont.) LAPACK_SYEV Procedure Parameters

Paramete	Description
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). On entry, the symmetric matrix a: • If uplo = 'U', the leading n by n upper triangular part of a contains the upper triangular part of the matrix a. • If uplo = 'L', the leading n byn lower triangular part of a contains the lower triangular part of the matrix a. On exit: • If jobz = 'V', then if info = 0, a contains the orthonormal eigenvectors of the matrix a. • If jobz = 'N', then on exit the lower triangle (if uplo = 'L') or the upper triangle (if uplo='U') of a, including the diagonal, is destroyed.
lda	The leading dimension of the array a. lda >= max(1,n).
w	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). If info = 0, the eigenvalues in ascending order.
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.53 LAPACK_SYEVD Procedures

This procedure computes all eigenvalues and, optionally, eigenvectors of a real symmetric matrix A. If eigenvectors are desired, it uses a divide and conquer algorithm that makes mild assumptions about floating point arithmetic.

> **See Also:**
>
> LAPACK Driver Routines (LLS and Eigenvalue Problems) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SYEVD (

   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   a        IN OUT  UTL_NLA_ARRAY_DBL,
   lda      IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_DBL,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

```
UTL_NLA.LAPACK_SYEVD (
   jobz     IN      flag,
   uplo     IN      flag,
   n        IN      POSITIVEN,
   a        IN OUT  UTL_NLA_ARRAY_FLT,
   lda      IN      POSITIVEN,
   w        IN OUT  UTL_NLA_ARRAY_FLT,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');
```

Parameters

Table 293-59 LAPACK_SYEVD Procedure Parameters

Parameter	Description
jobz	• 'N': Compute eigenvalues only. • 'V': Compute eigenvalues and eigenvectors.
uplo	• 'U' : Upper triangle of A is stored. • 'L': Upper triangle of A is stored.
n	The order of the matrix a. N >= 0.
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (lda, n). On entry, the symmetric matrix a: • If uplo = 'U', the leading n by n upper triangular part of a contains the upper triangular part of the matrix a. • If uplo = 'L', the leading n byn lower triangular part of a contains the lower triangular part of the matrix a. On exit: • If jobz = 'V', then if info = 0, a contains the orthonormal eigenvectors of the matrix a. • If jobz = 'N', then on exit the lower triangle (if uplo = 'L') or the upper triangle (if uplo = 'U') of a, including the diagonal, is destroyed.
lda	The leading dimension of the array a. lda >= max(1,n).
w	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n). If info = 0, the eigenvalues in ascending order.
info	• = 0 : successful exit • < 0 : if info = -i, the i-th argument had an illegal value • > 0 : if info = i, the algorithm failed to converge; i off-diagonal elements of an intermediate tridiagonal form did not converge to zero
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

293.5.54 LAPACK_SYSV Procedures

This procedure computes the solution to a real system of linear equations a * x = b, where a is an n by n symmetric matrix, and x and b are n by nrhs matrices.

The diagonal pivoting method is used to factor A as

`A = U * D * U**T`, if `UPLO = 'U'`

or

`A = L * D * L**T`, if `UPLO = 'L'`

where `U` (or `L`) is a product of permutation and unit upper (lower) triangular matrices, and `D` is symmetric and block diagonal with 1 by 1 and 2 by 2 diagonal blocks. The factored form of `A` is then used to solve the system of equations `A * X = B`.

> **See Also:**
>
> LAPACK Driver Routines (Linear Equations) Subprograms for other subprograms in this group

Syntax

```
UTL_NLA.LAPACK_SYSV (
   uplo     IN      flag,
   n        IN      POSITIVEN,
   nrhs     IN      POSITIVEN,
   a        IN OUT  UTL_NLA_ARRAY_DBL,
   lda      IN      POSITIVEN,
   ipiv     IN OUT  UTL_NLA_ARRAY_INT,
   b        IN OUT  UTL_NLA_ARRAY_DBL,
   ldb      IN      POSITIVEN,
   info     OUT     INTEGER,
   pack     IN      flag DEFAULT 'C');

UTL_NLA.LAPACK_SYSV (
   uplo    IN       flag,
   n       IN       POSITIVEN,
   nrhs    IN       POSITIVEN,
   a       IN OUT   UTL_NLA_ARRAY_FLT,
   lda     IN       POSITIVEN,
   ipiv    IN OUT   UTL_NLA_ARRAY_INT,
   b       IN OUT   UTL_NLA_ARRAY_FLT,
   ldb     IN       POSITIVEN,
   info    OUT      INTEGER,
   pack    IN       flag DEFAULT 'C');
```

Parameters

Table 293-60 LAPACK_SYSV Procedure Parameters

Parameter	Description
`uplo`	• `uplo = 'U'`. Upper triangular of `A` is stored. • `uplo = 'L'` . Lower triangular of `A` is stored.
`n`	The number of linear equations, which is the order of the matrix `a`. `N >= 0`.
`nrhs`	The number of right-hand sides, which is the number of columns of the matrix `b`. `nrhs >= 0`.

Table 293-60 (Cont.) LAPACK_SYSV Procedure Parameters

Parameter	Description
a	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (n-1). On entry, the symmetric matrix a. If UPLO = 'U', the leading n by n upper triangular part of a contains the upper triangular part of the matrix a, and the strictly lower triangular part of a is not referenced. If uplo = 'L', the leading n by n lower triangular part of a contains the lower triangular part of the matrix a, and the strictly upper triangular part of a is not referenced. On exit, if info = 0, the block diagonal matrix d and the multipliers used to obtain the factor U or L from the factorization A = U*D*U**T or A = L*D*L**T as computed by SSYTRF.
lda	The leading dimension of the array a. lda >= max(1,n)
ipiv	INTEGER array, DIMENSION (ldb, nrhs). Details of the interchanges and the block structure of d, as determined by SSYTRF. • If ipiv(k) > 0, then rows and columns k and ipiv(k) were interchanged, and d(k,k) is a 1 by1 diagonal block. • If uplo = 'U' and ipiv(k) = ipiv(k-1) < 0, then rows and columns k-1 and -ipiv(k) were interchanged and d(k-1:k,k-1:k) is a 2 by 2 diagonal block. • If uplo = 'L' and ipiv(k) = ipiv(k+1) < 0, then rows and columns k+1 and -ipiv(k) were interchanged and d(k:k+1,k:k+1) is a 2 by 2 diagonal block.
b	UTL_NLA_ARRAY_FLT/DBL, DIMENSION (ldb, nrhs). On entry, the n by nrhs matrix of right hand side matrix b. On exit, if info = 0 , the n by nrhs solution matrix X.
ldb	The leading dimension of the array b. ldb >= max(1,n)
info	• = 0 : successful exit • < 0 : if info = -i , the i-th argument had an illegal value • > 0 : if info = i, d(i,i) is exactly zero. The factorization has been completed, but the block diagonal matrix d is exactly singular, so the solution could not be computed.
pack	(Optional) Flags the packing of the matrices: • 'C': column-major (default) • 'R': row-major

294
UTL_RAW

The `UTL_RAW` package provides SQL functions for manipulating `RAW` datatypes.

This chapter contains the following topics:

- Overview
- Operational Notes
- Summary of UTL_RAW Subprograms

294.1 UTL_RAW Overview

This package is necessary because normal SQL functions do not operate on `RAWs`, and PL/SQL does not allow overloading between a `RAW` and a `CHAR` datatype. `UTL_RAW` also includes subprograms that convert various COBOL number formats to, and from, `RAWs`.

`UTL_RAW` is not specific to the database environment, and it may actually be used in other environments. For this reason, the prefix UTL has been given to the package, instead of DBMS.

294.2 UTL_RAW Operational Notes

`UTL_RAW` allows a `RAW` "record" to be composed of many elements. By using the `RAW` datatype, character set conversion will not be performed, keeping the `RAW` in its original format when being transferred through remote procedure calls.

With the `RAW` functions, you can manipulate binary data that was previously limited to the `hextoraw` and `rawtohex` functions.

> **Note:**
>
> Notes on datatypes:
>
> - The `PLS_INTEGER` and `BINARY_INTEGER` datatypes are identical. This document uses `BINARY_INTEGER` to indicate datatypes in reference information (such as for table types, record types, subprogram parameters, or subprogram return values), but may use either in discussion and examples.
> - The `INTEGER` and `NUMBER(38)` datatypes are also identical. This document uses `INTEGER` throughout.

294.3 Summary of UTL_RAW Subprograms

This table lists the `UTL_RAW` subprograms and briefly describes them.

Table 294-1 ***UTL_RAW Package Subprograms***

Subprogram	Description
BIT_AND Function	Performs bitwise logical "and" of the values in `RAW r1` with `RAW r2` and returns the "anded" result `RAW`
BIT_COMPLEMENT Function	Performs bitwise logical "complement" of the values in `RAW r` and returns the "complement'ed" result `RAW`
BIT_OR Function	Performs bitwise logical "or" of the values in `RAW r1` with `RAW r2` and returns the "or'd" result `RAW`
BIT_XOR Function	Performs bitwise logical "exclusive or" of the values in `RAW r1` with `RAW r2` and returns the "xor'd" result `RAW`
CAST_FROM_BINARY_DOUBLE Function	Returns the `RAW` binary representation of a `BINARY_DOUBLE` value
CAST_FROM_BINARY_FLOAT Function	Returns the `RAW` binary representation of a `BINARY_FLOAT` value
CAST_FROM_BINARY_INTEGER Function	Returns the `RAW` binary representation of a `BINARY_INTEGER` value
CAST_FROM_NUMBER Function	Returns the `RAW` binary representation of a `NUMBER` value
CAST_TO_BINARY_DOUBLE Function	Casts the `RAW` binary representation of a `BINARY_DOUBLE` into a `BINARY_DOUBLE`
CAST_TO_BINARY_FLOAT Function	Casts the `RAW` binary representation of a `BINARY_FLOAT` into a `BINARY_FLOAT`
CAST_TO_BINARY_INTEGER Function	Casts the `RAW` binary representation of a `BINARY_INTEGER` into a `BINARY_INTEGER`
CAST_TO_NUMBER Function	Casts the `RAW` binary representation of a `NUMBER` into a `NUMBER`
CAST_TO_NVARCHAR2 Function	Converts a `RAW` value into a `VARCHAR2` value
CAST_TO_RAW Function	Converts a `VARCHAR2` value into a `RAW` value
CAST_TO_VARCHAR2 Function	Converts a RAW value into a `VARCHAR2` value
COMPARE Function	Compares `RAW r1` against `RAW r2`
CONCAT Function	Concatenates up to 12 `RAWs` into a single `RAW`
CONVERT Function	Converts `RAW r` from character set `from_charset` to character set `to_charset` and returns the resulting `RAW`
COPIES Function	Returns `n` copies of `r` concatenated together
LENGTH Function	Returns the length in bytes of a `RAW r`
OVERLAY Function	Overlays the specified portion of target `RAW` with overlay `RAW`, starting from byte position `pos` of target and proceeding for `len` bytes
REVERSE Function	Reverses a byte sequence in `RAW r` from end to end
SUBSTR Function	Returns `len` bytes, starting at `pos` from `RAW r`
TRANSLATE Function	Translates the bytes in the input `RAW r` according to the bytes in the translation `RAWs from_set` and `to_set`
TRANSLITERATE Function	Converts the bytes in the input `RAW r` according to the bytes in the transliteration `RAWs from_set` and `to_set`

Table 294-1 (Cont.) *UTL_RAW Package Subprograms*

Subprogram	Description
XRANGE Function	Returns a `RAW` containing all valid 1-byte encodings in succession, beginning with the value `start_byte` and ending with the value `end_byte`

294.3.1 BIT_AND Function

This function performs bitwise logical "and" of the values in `RAW r1` with `RAW r2` and returns the "anded" result `RAW`.

Syntax

```
UTL_RAW.BIT_AND (
   r1 IN RAW,
   r2 IN RAW)
RETURN RAW;
```

Pragmas

```
pragma restrict_references(bit_and, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-2 BIT_AND Function Parameters

Parameter	Description
`r1`	`RAW` to "and" with `r2`
`r2`	`RAW` to "and" with `r1`

Return Values

Table 294-3 BIT_AND Function Return Values

Return	Description
`RAW`	Containing the "and" of `r1` and `r2`
`NULL`	Either `r1` or `r2` input parameter was `NULL`

Usage Notes

If `r1` and r2 differ in length, the and operation is terminated after the last byte of the shorter of the two `RAWs`, and the unprocessed portion of the longer `RAW` is appended to the partial result. The result length equals the longer of the two input `RAWs`.

294.3.2 BIT_COMPLEMENT Function

This function performs bitwise logical "complement" of the values in `RAW r` and returns the complemented result `RAW`. The result length equals the input `RAW r` length.

Syntax

```
UTL_RAW.BIT_COMPLEMENT (
   r IN RAW)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(bit_complement, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-4 BIT_COMPLEMENT Function Parameters

Parameter	Description
`r`	`RAW` to perform "complement" operation

Return Values

Table 294-5 BIT_COMPLEMENT Function Return Values

Return	Description
`RAW`	The "complement" of `r1`
`NULL`	If `r` input parameter was `NULL`

294.3.3 BIT_OR Function

This function performs bitwise logical "or" of the values in `RAW r1` with `RAW r2` and returns the or'd result `RAW`.

Syntax

```
UTL_RAW.BIT_OR (
   r1 IN RAW,
   r2 IN RAW)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(bit_or, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-6 BIT_OR Function Parameters

Parameters	Description
`r1`	`RAW` to "or" with `r2`

Table 294-6 (Cont.) BIT_OR Function Parameters

Parameters	Description
r2	RAW to "or" with r1

Return Values

Table 294-7 BIT_OR Function Return Values

Return	Description
RAW	Containing the "or" of r1 and r2
NULL	Either r1 or r2 input parameter was NULL

Usage Notes

If r1 and r2 differ in length, then the "or" operation is terminated after the last byte of the shorter of the two RAWs, and the unprocessed portion of the longer RAW is appended to the partial result. The result length equals the longer of the two input RAWs.

294.3.4 BIT_XOR Function

This function performs bitwise logical "exclusive or" of the values in RAW r1 with RAW r2 and returns the xor'd result RAW.

Syntax

```
UTL_RAW.BIT_XOR (
   r1 IN RAW,
   r2 IN RAW)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(bit_xor, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-8 BIT_XOR Function Parameters

Parameter	Description
r1	RAW to "xor" with r2
r2	RAW to "xor" with r1

Return Values

Table 294-9 BIT_XOR Function Return Values

Return	Description
RAW	Containing the "xor" of r1 and r2

Table 294-9 (Cont.) BIT_XOR Function Return Values

Return	Description
NULL	If either r1 or r2 input parameter was NULL

Usage Notes

If r1 and r2 differ in length, then the "xor" operation is terminated after the last byte of the shorter of the two RAWs, and the unprocessed portion of the longer RAW is appended to the partial result. The result length equals the longer of the two input RAWs.

294.3.5 CAST_FROM_BINARY_DOUBLE Function

This function returns the RAW binary representation of a BINARY_DOUBLE value.

Syntax

```
UTL_RAW.CAST_FROM_BINARY_DOUBLE(
   n          IN BINARY_DOUBLE,
   endianess IN PLS_INTEGER DEFAULT 1)
RETURN RAW;
```

Pragmas

```
pragma restrict_references(cast_from_binary_double, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-10 CAST_FROM_BINARY_DOUBLE Function Parameters

Parameter	Description
n	BINARY_DOUBLE value
endianess	A BINARY_INTEGER value indicating the endianess. The function recognizes the defined constants big_endian (1), little_endian (2), and machine_endian (3). The default is big_endian. A setting of machine_endian has the same effect as big_endian on a big endian machine, or the same effect as little_endian on a little endian machine.

Return Values

The binary representation of the BINARY_DOUBLE value, or NULL if the input is NULL.

Usage Notes

- An 8-byte binary_double value maps to the IEEE 754 double-precision format as follows:

```
byte 0: bit 63 ~ bit 56
byte 1: bit 55 ~ bit 48
byte 2: bit 47 ~ bit 40
byte 3: bit 39 ~ bit 32
byte 4: bit 31 ~ bit 24
byte 5: bit 23 ~ bit 16
byte 6: bit 15 ~ bit  8
byte 7: bit  7 ~ bit  0
```

- The parameter endianess describes how the bytes of BINARY_DOUBLE are mapped to the bytes of RAW. In the following matrix, rb0 ~ rb7 refer to the bytes in raw and db0 ~ db7 refer to the bytes in BINARY_DOUBLE.

endianess	rb0	rb1	rb2	rb3	rb4	rb5	rb6	rb7
big_endian	db0	db1	db2	db3	db4	db5	db6	db7
little_endian	db7	db6	db5	db4	db3	db2	db1	db0

- In case of machine-endian, the 8 bytes of the BINARY_DOUBLE argument are copied straight across into the RAW return value. The effect is the same if the user has passed big_endian on a big-endian machine, or little_endian on a little-endian machine.

294.3.6 CAST_FROM_BINARY_FLOAT Function

This function returns the RAW binary representation of a BINARY_FLOAT value.

Syntax

```
UTL_RAW.CAST_FROM_BINARY_FLOAT(
   n          IN BINARY_FLOAT,
   endianess IN PLS_INTEGER DEFAULT 1)
RETURN RAW;
```

Pragmas

```
pragma restrict_references(cast_from_binary_float, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-11 CAST_FROM_BINARY_FLOAT Function Parameters

Parameter	Description
n	BINARY_FLOAT value
endianess	A BINARY_INTEGER value indicating the endianess. The function recognizes the defined constants big_endian (1), little_endian (2), and machine_endian (3). The default is big_endian. A setting of machine_endian has the same effect as big_endian on a big endian machine, or the same effect as little_endian on a little endian machine.

Return Values

The binary representation (RAW) of the BINARY_FLOAT value, or NULL if the input is NULL.

Usage Notes

- A 4-byte binary_float value maps to the IEEE 754 single-precision format as follows:

```
byte 0: bit 31 ~ bit 24
byte 1: bit 23 ~ bit 16
byte 2: bit 15 ~ bit  8
byte 3: bit 7 ~  bit  0
```

- The parameter endianess describes how the bytes of BINARY_FLOAT are mapped to the bytes of RAW. In the following matrix, rb0 ~ rb3 refer to the bytes in RAW and fb0 ~ fb3 refer to the bytes in BINARY_FLOAT.

Endianess	rb0	rb1	rb2	rb3
big_endian	fb0	fb1	fb2	fb3
little_endian	fb3	fb2	fb1	fb0

- In case of machine-endian, the 4 bytes of the BINARY_FLOAT argument are copied straight across into the RAW return value. The effect is the same if the user has passed big_endian on a big-endian machine, or little_endian on a little-endian machine.

294.3.7 CAST_FROM_BINARY_INTEGER Function

This function returns the RAW binary representation of a BINARY_INTEGER value.

Syntax

```
UTL_RAW.CAST_FROM_BINARY_INTEGER (
   n          IN BINARY_INTEGER
   endianess  IN PLS_INTEGER DEFAULT BIG_ENDIAN)
RETURN RAW;
```

Pragmas

```
pragma restrict_references(cast_from_binary_integer, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-12 CAST_FROM_BINARY_INTEGER Function Parameters

Parameter	Description
n	BINARY_INTEGER value.
endianess	A BINARY_INTEGER value indicating the endianess. The function recognizes the defined constants big_endian (1), little_endian (2), and machine_endian (3). The default is big_endian. A setting of machine_endian has the same effect as big_endian on a big endian machine, or the same effect as little_endian on a little endian machine.

Return Values

The binary representation of the BINARY_INTEGER value.

294.3.8 CAST_FROM_NUMBER Function

This function returns the RAW binary representation of a NUMBER value.

Syntax

```
UTL_RAW.CAST_FROM_NUMBER (
   n  IN NUMBER)
 RETURN RAW;
```

Pragmas

```
pragma restrict_references(cast_from_number, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-13 CAST_FROM_NUMBER Function Parameters

Parameter	Description
n	NUMBER value

Return Values

The binary representation of the NUMBER value.

294.3.9 CAST_TO_BINARY_DOUBLE Function

This function casts the RAW binary representation of a BINARY_DOUBLE into a BINARY_DOUBLE.

Syntax

```
UTL_RAW.CAST_TO_BINARY_DOUBLE (
   r          IN RAW
   endianess  IN PLS_INTEGER DEFAULT 1)
RETURN BINARY_DOUBLE;
```

Pragmas

```
pragma restrict_references(cast_to_binary_double, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-14 CAST_TO_BINARY_DOUBLE Function Parameters

Parameter	Description
r	Binary representation of a BINARY_DOUBLE
endianess	A PLS_INTEGER representing big-endian or little-endian architecture. The default is big-endian.

Return Values

The BINARY_DOUBLE value.

Usage Notes

- If the RAW argument is more than 8 bytes, only the first 8 bytes are used and the rest of the bytes are ignored. If the result is -0, +0 is returned. If the result is NaN, the value BINARY_DOUBLE_NAN is returned.
- If the RAW argument is less than 8 bytes, a VALUE_ERROR exception is raised.
- An 8-byte binary_double value maps to the IEEE 754 double-precision format as follows:

```
byte 0: bit 63 ~ bit 56
byte 1: bit 55 ~ bit 48
byte 2: bit 47 ~ bit 40
```

```
byte 3: bit 39 ~ bit 32
byte 4: bit 31 ~ bit 24
byte 5: bit 23 ~ bit 16
byte 6: bit 15 ~ bit  8
byte 7: bit  7 ~ bit  0
```

- The parameter endianess describes how the bytes of `BINARY_DOUBLE` are mapped to the bytes of `RAW`. In the following matrix, rb0 ~ rb7 refer to the bytes in raw and db0 ~ db7 refer to the bytes in `BINARY_DOUBLE`.

Architecture	rb0	rb1	rb2	rb3	rb4	rb5	rb6	rb7
big_endian	db0	db1	db2	db3	db4	db5	db6	db7
little_endian	db7	db6	db5	db4	db3	db2	db1	db0

- In case of machine-endian, the 8 bytes of the `RAW` argument are copied straight across into the `BINARY_DOUBLE` return value. The effect is the same if the user has passed `big_endian` on a big-endian machine, or `little_endian` on a little-endian machine.

294.3.10 CAST_TO_BINARY_FLOAT Function

This function casts the RAW binary representation of a `BINARY_FLOAT` into a `BINARY_FLOAT`.

Syntax

```
UTL_RAW.CAST_TO_BINARY_FLOAT (
   r          IN RAW
   endianess  IN PLS_INTEGER DEFAULT 1)
RETURN BINARY_FLOAT;
```

Pragmas

```
pragma restrict_references(cast_to_binary_float, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-15 CAST_TO_BINARY_FLOAT Function Parameters

Parameter	Description
`r`	Binary representation of a `BINARY_FLOAT`
`endianess`	A `PLS_INTEGER` representing big-endian or little-endian architecture. The default is big-endian.

Return Values

The `BINARY_FLOAT` value.

Usage Notes

- If the `RAW` argument is more than 4 bytes, only the first 4 bytes are used and the rest of the bytes are ignored. If the result is -0, +0 is returned. If the result is NaN, the value `BINARY_FLOAT_NAN` is returned.
- If the `RAW` argument is less than 4 bytes, a `VALUE_ERROR` exception is raised.
- A 4-byte `binary_float` value maps to the IEEE 754 single-precision format as follows:

```
byte 0: bit 31 ~ bit 24
byte 1: bit 23 ~ bit 16
byte 2: bit 15 ~ bit  8
byte 3: bit 7 ~  bit  0
```

- The parameter endianess describes how the bytes of `BINARY_FLOAT` are mapped to the bytes of `RAW`. In the following matrix, rb0 ~ rb3 refer to the bytes in `RAW` and fb0 ~ fb3 refer to the bytes in `BINARY_FLOAT`.

Endianness	rb0	rb1	rb2	rb3
big_endian	fbo	fb1	fb2	fb3
little_endian	fb3	fb2	fb1	fb0

- In case of machine-endian, the 4 bytes of the `RAW` argument are copied straight across into the `BINARY_FLOAT` return value. The effect is the same if the user has passed `big_endian` on a big-endian machine, or `little_endian` on a little-endian machine.

294.3.11 CAST_TO_BINARY_INTEGER Function

This function casts the RAW binary representation of a `BINARY_INTEGER` into a `BINARY_INTEGER`.

Syntax

```
UTL_RAW.CAST_TO_BINARY_INTEGER (
   r          IN RAW
   endianess  IN PLS_INTEGER DEFAULT BIG_ENDIAN)
RETURN BINARY_INTEGER;
```

Pragmas

```
pragma restrict_references(cast_to_binary_integer, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-16 CAST_TO_BINARY_INTEGER Function Parameters

Parameter	Description
`r`	Binary representation of a `BINARY_INTEGER`
`endianess`	A `PLS_INTEGER` representing big-endian or little-endian architecture. The default is big-endian.

Return Values

The `BINARY_INTEGER` value

294.3.12 CAST_TO_NUMBER Function

This function casts the RAW binary representation of a `NUMBER` into a `NUMBER`.

Syntax

```
UTL_RAW.CAST_TO_NUMBER (
   r  IN RAW)
 RETURN NUMBER;
```

Pragmas

```
pragma restrict_references(cast_to_number, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-17 CAST_TO_NUMBER function Parameters

Parameter	Description
r	Binary representation of a NUMBER

Return Values

The NUMBER value.

294.3.13 CAST_TO_NVARCHAR2 Function

This function converts a RAW value represented using some number of data bytes into an NVARCHAR2 value with that number of data bytes.

> **Note:**
>
> When casting to a NVARCHAR2, the current Globalization Support character set is used for the characters within that NVARCHAR2 value.

Syntax

```
UTL_RAW.CAST_TO_NVARCHAR2 (
   r IN RAW)
RETURN NVARCHAR2;
```

Pragmas

```
pragma restrict_references(cast_to_NVARCHAR2, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-18 CAST_TO_NVARCHAR2 Function Parameters

Parameter	Description
r	RAW (without leading length field) to be changed to a NVARCHAR2)

Return Values

Table 294-19 CAST_TO_NVARCHAR2 Function Return Values

Return	Description
NVARCHAR2	Containing having the same data as the input RAW
NULL	If r input parameter was NULL

294.3.14 CAST_TO_RAW Function

This function converts a VARCHAR2 value represented using some number of data bytes into a RAW value with that number of data bytes. The data itself is not modified in any way, but its datatype is recast to a RAW datatype.

Syntax

```
UTL_RAW.CAST_TO_RAW (
   c  IN VARCHAR2)
RETURN RAW;
```

Pragmas

```
pragma restrict_references(cast_to_raw, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-20 CAST_TO_RAW Function Parameters

Parameter	Description
c	VARCHAR2 to be changed to a RAW

Return Values

Table 294-21 CAST_TO_RAW Function Return Values

Return	Description
RAW	Containing the same data as the input VARCHAR2 and equal byte length as the input VARCHAR2 and without a leading length field
NULL	If c input parameter was NULL

294.3.15 CAST_TO_VARCHAR2 Function

This function converts a RAW value represented using some number of data bytes into a VARCHAR2 value with that number of data bytes.

> **Note:**
>
> When casting to a VARCHAR2, the current Globalization Support character set is used for the characters within that VARCHAR2.

Syntax

```
UTL_RAW.CAST_TO_VARCHAR2 (
   r IN RAW)
RETURN VARCHAR2;
```

Pragmas

```
pragma restrict_references(cast_to_VARCHAR2, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-22 CAST_TO_VARCHAR2 Function Parameters

Parameter	Description
r	RAW (without leading length field) to be changed to a VARCHAR2

Return Values

Table 294-23 CAST_TO_VARCHAR2 Function Return Values

Return	Description
VARCHAR2	Containing having the same data as the input RAW
NULL	If r input parameter was NULL

294.3.16 COMPARE Function

This function compares two RAW values. If they differ in length, then the shorter is extended on the right according to the optional pad parameter.

Syntax

```
UTL_RAW.COMPARE (
   r1  IN RAW,
   r2  IN RAW,
   pad IN RAW DEFAULT NULL)
  RETURN NUMBER;
```

Pragmas

```
pragma restrict_references(compare, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-24 COMPARE Function Parameters

Parameter	Description
r1	1st RAW to be compared, may be NULL or 0 length
r2	2nd RAW to be compared, may be NULL or 0 length
pad	This is an optional parameter. Byte to extend whichever of r1 or r2 is shorter. The default: x'00'

Return Values

Table 294-25 COMPARE Function Return Values

Return	Description
NUMBER	Equals 0 if RAW byte strings are both NULL or identical; or, Equals position (numbered from 1) of the first mismatched byte

294.3.17 CONCAT Function

This function concatenates up to 12 RAWs into a single RAW. If the concatenated size exceeds 32K, then an error is returned

Syntax

```
UTL_RAW.CONCAT (
   r1  IN RAW DEFAULT NULL,
   r2  IN RAW DEFAULT NULL,
   r3  IN RAW DEFAULT NULL,
   r4  IN RAW DEFAULT NULL,
   r5  IN RAW DEFAULT NULL,
   r6  IN RAW DEFAULT NULL,
   r7  IN RAW DEFAULT NULL,
   r8  IN RAW DEFAULT NULL,
   r9  IN RAW DEFAULT NULL,
   r10 IN RAW DEFAULT NULL,
   r11 IN RAW DEFAULT NULL,
   r12 IN RAW DEFAULT NULL)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(concat, WNDS, RNDS, WNPS, RNPS);
```

Parameters

r1....r12 are the RAW items to concatenate.

Return Values

Table 294-26 CONCAT Function Return Values

Return	Description
RAW	Containing the items concatenated

Exceptions

There is an error if the sum of the lengths of the inputs exceeds the maximum allowable length for a RAW, which is 32767 bytes.

294.3.18 CONVERT Function

This function converts `RAW r` from character set `from_charset` to character set `to_charset` and returns the resulting `RAW`.

Both `from_charset` and `to_charset` must be supported character sets defined to the Oracle server.

Syntax

```
UTL_RAW.CONVERT (
   r            IN RAW,
   to_charset   IN VARCHAR2,
   from_charset IN VARCHAR2)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(convert, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-27 CONVERT Function Parameters

Parameter	Description
`r`	`RAW` byte-string to be converted
`to_charset`	Name of the character set to which `r` is converted
`from_charset`	Name of the character set in which `r` is supplied

Return Values

Table 294-28 CONVERT Function Return Values

Return	Description
`RAW`	Byte string `r` converted according to the specified character sets.

Exceptions

Table 294-29 CONVERT Function Exceptions

Error	Description
`ORA-06502`	PL/SQL: numeric or value error
`ORA-12703`	This character set conversion is not supported
`ORA-12705`	Cannot access NLS data files or invalid environment specified

Usage Notes

- The NLS_LANG parameter form *language_territory.character set* is also accepted for `to_charset` and `from_charset`. However, this form is deprecated and should be avoided. Note that *language* and *territory* are ignored by this subprogram.

- The converted value is silently truncated if it exceeds the maximum length of a `RAW` value, which is 32767 bytes. Do not convert values longer than floor(32767/4) = 8191 bytes if you want to avoid this truncation for all possible combinations of `to_charset` and `from_charset`. You can use the maximum character width of the target character set `to_charset`, if known, to expand the limit to a less pessimistic value. For example, if the target character set is ZHS16GBK, the maximum safe source string length is floor(32767/2) = 16383 bytes. For single-byte target character sets, no truncation is ever necessary.

294.3.19 COPIES Function

This function returns `n` copies of `r` concatenated together.

Syntax

```
UTL_RAW.COPIES (
   r IN RAW,
   n IN NUMBER)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(copies, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-30 COPIES Function Parameters

Parameters	Description
`r`	`RAW` to be copied
`n`	Number of times to copy the `RAW` (must be positive)

Return Values

This returns the `RAW` copied `n` times.

Exceptions

Table 294-31 COPIES Function Exceptions

Error	Description
`VALUE_ERROR`	Either: - `r` is missing, `NULL` or 0 length - `n` < 1 - Length of result exceeds maximum length of a `RAW`

294.3.20 LENGTH Function

This function returns the length in bytes of a `RAW r`.

Syntax

```
UTL_RAW.LENGTH (
   r  IN RAW)
RETURN NUMBER;
```

Pragmas

```
pragma restrict_references(length, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-32 LENGTH Function Parameters

Parameter	Description
`r`	`RAW` byte stream to be measured

Return Values

Table 294-33 LENGTH Function Return Values

Return	Description
`NUMBER`	Current length of the `RAW`

294.3.21 OVERLAY Function

This function overlays the specified portion of target `RAW` with `overlay_str RAW`, starting from byte position `pos` of `target` and proceeding for `len` bytes.

Syntax

```
UTL_RAW.OVERLAY (
   overlay_str IN RAW,
   target      IN RAW,
   pos         IN BINARY_INTEGER DEFAULT 1,
   len         IN BINARY_INTEGER DEFAULT NULL,
   pad         IN RAW            DEFAULT NULL)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(overlay, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-34 OVERLAY Function Parameters

Parameters	Description
`overlay_str`	Byte-string used to overlay target

Table 294-34 (Cont.) OVERLAY Function Parameters

Parameters	Description
target	Byte-string which is to be overlaid
pos	Position in target (numbered from 1) to start overlay
len	The number of target bytes to overlay
pad	Pad byte used when overlay len exceeds overlay_str length or pos exceeds target length

Defaults and Optional Parameters

Table 294-35 OVERLAY Function Optional Parameters

Optional Parameter	Description
pos	1
len	To the length of overlay_str
pad	x'00'

Return Values

Table 294-36 OVERLAY Function Return Values

Return	Description
RAW	The target byte_string overlaid as specified.

Usage Notes

If overlay_str has less than len bytes, then it is extended to len bytes using the pad byte. If overlay_str exceeds len bytes, then the extra bytes in overlay_str are ignored. If len bytes beginning at position pos of target exceeds the length of target, then target is extended to contain the entire length of overlay_str.

If len is specified, it must be greater than or equal to 0. If pos is specified, it must be greater than or equal to 1. If pos exceeds the length of target, then target is padded with pad bytes to position pos, and target is further extended with overlay_str bytes.

Exceptions

Table 294-37 OVERLAY Function Exceptions

Error	Description
VALUE_ERROR	Either: - Overlay_str is NULL or has 0 length - Target is missing or undefined - Length of target exceeds maximum length of a RAW - len < 0 - pos < 1

294.3.22 REVERSE Function

This function reverses a byte sequence in `RAW r` from end to end.

For example, x'0102F3' would be reversed to x'F30201', and 'xyz' would be reversed to 'zyx'.The result length is the same as the input `RAW` length.

Syntax

```
UTL_RAW.REVERSE (
   r IN RAW)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(reverse, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-38 REVERSE Function Parameters

Parameter	Description
`r`	`RAW` to reverse

Return Values

Table 294-39 REVERSE Function Return Values

Return	Description
`RAW`	Containing the "reverse" of `r`

Exceptions

Table 294-40 REVERSE Function Exceptions

Error	Description
`VALUE_ERROR`	`R` is `NULL` or has 0 length

294.3.23 SUBSTR Function

This function returns `len` bytes, starting at `pos` from `RAW r`.

Syntax

```
UTL_RAW.SUBSTR (
   r   IN RAW,
   pos IN BINARY_INTEGER,
   len IN BINARY_INTEGER DEFAULT NULL)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(substr, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-41 SUBSTR Function Parameters

Parameter	Description
r	RAW byte-string from which a portion is extracted
pos	Byte position in r at which to begin extraction
len	Number of bytes from pos to extract from r (optional)

Defaults and Optional Parameters

Table 294-42 SUBSTR Function Optional Parameter

Optional Parameter	Description
len	Position pos through to the end of r

Return Values

Table 294-43 SUBSTR Function Return Values

Return	Description
portion of r	Beginning at pos for len bytes long
NULL	r input parameter was NULL

Usage Notes

- If pos is positive, then SUBSTR counts from the beginning of r to find the first byte. If pos is negative, then SUBSTR counts backward from the end of the r. The value pos cannot be 0.
- If len is omitted, then SUBSTR returns all bytes to the end of r. The value len cannot be less than 1.

Exceptions

Table 294-44 SUBSTR Function Exceptions

Error	Description
VALUE_ERROR	VALUE_ERROR is returned if: • pos = 0 or > length of r • len < 1 or > length of r - (pos-1)

294.3.24 TRANSLATE Function

This function translates the bytes in the input RAW r according to the bytes in the translation RAWs from_set and to_set.

If a byte in r has a matching byte in from_set, then it is replaced by the byte in the corresponding position in to_set, or deleted.

Bytes in `r`, but undefined in `from_set`, are copied to the result. Only the first (leftmost) occurrence of a byte in `from_set` is used. Subsequent duplicates are not scanned and are ignored.

Syntax

```
UTL_RAW.TRANSLATE (
   r        IN RAW,
   from_set IN RAW,
   to_set   IN RAW)
  RETURN RAW;
```

> **Note:**
>
> Be aware that *to_set* and *from_set* are reversed in the calling sequence compared to `TRANSLITERATE`.

Pragmas

```
pragma restrict_references(translate, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-45 TRANSLATE Function Parameters

Parameter	Description
`r`	`RAW` source byte-string to be translated
`from_set`	`RAW` byte-codes to be translated, if present in `r`
`to_set`	`RAW` byte-codes to which corresponding `from_str` bytes are translated

Return Values

Table 294-46 TRANSLATE Function Return Values

Return	Description
`RAW`	Translated byte-string

Usage Notes

- If `to_set` is shorter than `from_set`, the extra `from_set` bytes have no corresponding translation bytes. Bytes from the input `RAW` that match any such `from_set` bytes are not translated or included in the result. They are effectively translated to `NULL`.
- If `to_set` is longer than `from_set`, the extra `to_set` bytes are ignored.
- If a byte value is repeated in `from_set`, the repeated occurrence is ignored.

> **Note:**
>
> Differences from the TRANSLITERATE Function:
>
> - The `from_set` parameter comes before the `to_set` parameter in the calling sequence.
> - Bytes from `r` that appear in `from_set` but have no corresponding values in `to_set` are not translated or included in the result.
> - The resulting `RAW` value may be shorter than the input `RAW` value.
>
> Note that `TRANSLATE` and `TRANSLITERATE` only differ in functionality when `to_set` has fewer bytes than `from_set`.

Exceptions

Table 294-47 TRANSLATE Function Exceptions

Error	Description
`VALUE_ERROR`	Either: - `r` is `NULL` or has 0 length - `from_set` is `NULL` or has 0 length - `to_set` is `NULL` or has 0 length

294.3.25 TRANSLITERATE Function

This function converts the bytes in the input `RAW r` according to the bytes in the transliteration `RAWs from_set` and `to_set`.

Successive bytes in `r` are looked up in the `from_set`, and, if not found, copied unaltered to the result `RAW`. If found, then they are replaced in the result `RAW` by either corresponding bytes in the `to_set`, or the `pad` byte when no correspondence exists.

Bytes in `r`, but undefined in `from_set`, are copied to the result. Only the first (leftmost) occurrence of a byte in `from_set` is used. Subsequent duplicates are not scanned and are ignored. The result `RAW` is always the same length as `r`.

Syntax

```
UTL_RAW.TRANSLITERATE (
   r        IN RAW,
   to_set   IN RAW DEFAULT NULL,
   from_set IN RAW DEFAULT NULL,
   pad      IN RAW DEFAULT NULL)
  RETURN RAW;
```

> **Note:**
>
> Be aware that *`to_set`* and *`from_set`* are reversed in the calling sequence compared to `TRANSLATE`.

Pragmas

```
pragma restrict_references(transliterate, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-48 TRANSLITERATE Function Parameters

Parameter	Description
r	RAW input byte-string to be converted
to_set	RAW byte-codes to which corresponding from_set bytes are converted (any length)
from_set	RAW byte-codes to be converted, if presenting r (any length)
pad	1 byte used when to-set is shorter than the from_set

Defaults and Optional Parameters

Table 294-49 TRANSLITERATE Function Optional Parameters

Optional Parameter	Description
to_set	To the NULL string and effectively extended with pad to the length of from_set as necessary
from_set	x'00' through x'fff'
pad	x'00'

Return Values

Table 294-50 TRANSLITERATE Function Return Values

Return	Description
RAW	Converted byte-string.

Usage Notes

- If to_set is shorter than from_set, the extra from_set bytes have no corresponding conversion bytes. Bytes from the input RAW that match any such from_set bytes are converted in the result to the pad byte instead.
- If to_set is longer than from_set, the extra to_set bytes are ignored.
- If a byte value is repeated in from_set, the repeated occurrence is ignored.

> **Note:**
>
> Differences from the TRANSLATE Function:
>
> - The `to_set` parameter comes before the `from_set` parameter in the calling sequence.
> - Bytes from `r` that appear in `from_set` but have no corresponding values in `to_set` are replaced by pad in the result.
> - The resulting `RAW` value always has the same length as the input `RAW` value.
>
> Note that `TRANSLATE` and `TRANSLITERATE` only differ in functionality when `to_set` has fewer bytes than `from_set`.

Exceptions

Table 294-51 TRANSLITERATE Function Exceptions

Error	Description
`VALUE_ERROR`	`R` is `NULL` or has 0 length

294.3.26 XRANGE Function

This function returns a RAW `value` containing the succession of one-byte encodings beginning and ending with the specified byte-codes. The specified byte-codes must be single-byte `RAW` values. If the `start_byte` value is greater than the `end_byte` value, then the succession of resulting bytes begins with `start_byte`, wraps through `x'FF'` back to `x'00'`, then ends at `end_byte`.

Syntax

```
UTL_RAW.XRANGE (
   start_byte IN RAW DEFAULT NULL,
   end_byte   IN RAW DEFAULT NULL)
  RETURN RAW;
```

Pragmas

```
pragma restrict_references(xrange, WNDS, RNDS, WNPS, RNPS);
```

Parameters

Table 294-52 XRANGE Function Parameters

Parameters	Description
`start_byte`	Beginning byte-code value of resulting sequence. The default is `x'00'`.
`end_byte`	Ending byte-code value of resulting sequence. The default is `x'FF'`.

Return Values

Table 294-53 XRANGE Function Return Values

Return	Description
`RAW`	Containing succession of 1-byte hexadecimal encodings

295
UTL_RECOMP

The UTL_RECOMP package recompiles invalid PL/SQL modules, invalid views, Java classes, indextypes and operators in a database, either sequentially or in parallel.

This chapter contains the following topics:

- Overview
- Operational Notes
- Examples
- Summary of UTL_RECOMP Subprograms

295.1 UTL_RECOMP Overview

This script is particularly useful after a major-version upgrade that typically invalidates all PL/SQL and Java objects. Although invalid objects are recompiled automatically on use, it is useful to run this script prior to operation because this will either eliminate or minimize subsequent latencies due to on-demand automatic recompilation at runtime.

Parallel recompilation can exploit multiple CPUs to reduce the time taken to recompile invalid objects. The degree of parallelism is specified by the first argument to RECOMP_PARALLEL Procedure.

In general, a parallelism setting of one thread for each available CPU provides a good initial setting. However, please note that the process of recompiling an invalid object writes a significant amount of data to system tables and is fairly I/O intensive. A slow disk system may be a significant bottleneck and limit speedups available from a higher degree of parallelism.

295.2 UTL_RECOMP Operational Notes

`UTL_RECOMP` has several operational notes.

- This package uses the job queue for parallel recompilation.
- This package must be run using `SQL*PLUS`.
- You must be connected `AS SYSDBA` to run this script.
- This package expects the following packages to have been created with VALID status:
  - `STANDARD` (`standard.sql`)
  - `DBMS_STANDARD` (`dbmsstdx.sql`)
  - `DBMS_JOB` (`dbmsjob.sql`)
  - `DBMS_RANDOM` (`dbmsrand.sql`)
- There should be no other DDL on the database while running entries in this package. Not following this recommendation may lead to deadlocks.

295.3 UTL_RECOMP Examples

These examples show various ways that UTL_RECOMP can recompile objects.

- Recompile all objects sequentially:

```
EXECUTE UTL_RECOMP.RECOMP_SERIAL();
```

- Recompile objects in schema SCOTT sequentially:

```
EXECUTE UTL_RECOMP.RECOMP_SERIAL('SCOTT');
```

- Recompile all objects using 4 parallel threads:

```
EXECUTE UTL_RECOMP.RECOMP_PARALLEL(4);
```

- Recompile objects in schema JOE using the number of threads specified in the parameter JOB_QUEUE_PROCESSES:

```
EXECUTE UTL_RECOMP.RECOMP_PARALLEL(NULL, 'JOE');
```

295.4 Summary of UTL_RECOMP Subprograms

This table lists the UTL_RECOMP subprograms and briefly describes them.

Table 295-1 *UTL_RECOMP Package Subprograms*

Subprogram	Description
RECOMP_PARALLEL Procedure	Recompiles invalid objects in a given schema, or all invalid objects in the database, in parallel
RECOMP_SERIAL Procedure	Recompiles invalid objects in a given schema or all invalid objects in the database

295.4.1 RECOMP_PARALLEL Procedure

This procedure uses the information exposed in the DBA_Dependencies view to recompile invalid objects in the database, or in a given schema, in parallel.

Syntax

```
UTL_RECOMP.RECOMP_PARALLEL(
   threads  IN   PLS_INTEGER DEFAULT NULL,
   schema   IN   VARCHAR2    DEFAULT NULL,
   flags    IN   PLS_INTEGER DEFAULT 0);
```

Parameters

Table 295-2 RECOMP_PARALLEL Procedure Parameters

Parameter	Description
threads	The number of recompile threads to run in parallel. If NULL, use the value of 'job_queue_processes'.
schema	The schema in which to recompile invalid objects. If NULL, all invalid objects in the database are recompiled.

Table 295-2 (Cont.) RECOMP_PARALLEL Procedure Parameters

Parameter	Description
flags	Flag values are intended for internal testing and diagnosability only.

Usage Notes

The UTL_RECOMP_COMPILED table is updated with objects that RECOMP_PARALLEL has attempted to compile, or that it has compiled so far. If you monitor the progress of the recompilation, then the number of objects should increase with time. These tables are used internally by RECOMP_PARALLEL to prevent an object from being compiled more than once.

Table 295-3 Tables Generated by UTL_RECOMP.RECOMP_PARALLEL

Table	Purpose
UTL_RECOMP_ERRORS	A table of objects that are sorted into groups, based on dependencies. It is used internally by RECOMP_PARALLEL to guide parallel recompilation.
UTL_RECOMP_COMPILED	A table of objects that RECOMP_PARALLEL has attempted to compile. It is used internally by recomp_parallel to prevent an object from being compiled more than once.
UTL_RECOMP_COMP_IDX1	An index created on UTL_RECOMP_COMPILED, which uses object number as the primary key.
UTL_RECOMP_SKIP_LIST	A table of objects that RECOMP_PARALLEL will not attempt to compile. Typical usage is for RECOMP_PARALLEL to call UTL_RECOMP.POPULATE_UTL_RECOMP_SKIP_LIST before the update starts so that this table is filled with previously identified invalid objects. Afterward, when RECOMP_PARALLEL is run after upgrade, it will not attempt to compile those previously identified invalid objects.

Table 295-4 Views Generated by UTL_RECOMP.RECOMP_PARALLEL

View	Purpose
UTL_RECOMP_ALL_OBJECTS	A view that includes all objects that can be compiled.
UTL_RECOMP_INVALID_ALL	A view that selects invalid objects from UTL_RECOMP_ALL_OBJECTS that have not been compiled and are not on the skip list.
UTL_RECOMP_INVALID_SEQ	A view that selects objects from UTL_RECOMP_INVALID_ALL that must be compiled sequentially.
UTL_RECOMP_INVALID_PARALLEL	A view that selects objects from UTL_RECOMP_INVALID_ALL that can be compiled in parallel.
UTL_RECOMP_INVALID_JAVA_SYN	A view that selects from UTL_RECOMP_ALL_OBJECTS objects that are invalid JAVA synonyms.
UTL_RECOMP_CIRCULAR_MV	A view that selects from UTL_RECOMP_ALL_OBJECTS materialized view objects that have a circular dependency chain.
UTL_RECOMP_INVALID_MV	A view that selects from UTL_RECOMP_ALL_OBJECTS materialized view objects that are invalid but don't have a circular dependency chain.

Table 295-4 (Cont.) Views Generated by UTL_RECOMP.RECOMP_PARALLEL

View	Purpose
UTL_RECOMP_ALL_STUBS	A view that includes all stub objects that have not been compiled, and are not on the skip list. When this view is used in RECOMP_PARALLEL, it is restricted to show only stubs in the current edition.

295.4.2 RECOMP_SERIAL Procedure

This procedure recompiles invalid objects in a given schema or all invalid objects in the database.

Syntax

```
UTL_RECOMP.RECOMP_SERIAL(
   schema    IN   VARCHAR2    DEFAULT NULL,
   flags     IN   PLS_INTEGER DEFAULT 0);
```

Parameters

Table 295-5 RECOMP_SERIAL Procedure Parameters

Parameter	Description
schema	The schema in which to recompile invalid objects. If NULL, all invalid objects in the database are recompiled.
flags	Flag values are intended for internal testing and diagnosability only.

296
UTL_REF

The `UTL_REF` package provides PL/SQL procedures to support reference-based operations. Unlike SQL, `UTL_REF` procedures enable you to write generic type methods without knowing the object table name.

This chapter contains the following topics:

- Overview
- Security Model
- Types
- Exceptions
- Summary of UTL_REF Subprograms

296.1 UTL_REF Overview

`UTL_REF` procedures enable you to write generic type methods without knowing the object table name.

Oracle supports user-defined composite type or object type. Any instance of an object type is called an object. An object type can be used as the type of a column or as the type of a table.

In an object table, each row of the table stores an object. You can uniquely identify an object in an object table with an object identifier.

A reference is a persistent pointer to an object, and each reference can contain an object identifier. The reference can be an attribute of an object type, or it can be stored in a column of a table. Given a reference, an object can be retrieved.

296.2 UTL_REF Security Model

The procedural option is needed to use this package. This package must be created under `SYS` (`CONNECT /AS SYSDBA`). Operations provided by this package are performed under the current calling user, not under the package owner `SYS`.

You can use the `UTL_REF` package from stored PL/SQL procedures/packages on the server, as well as from client/side PL/SQL code.

When invoked from PL/SQL procedures/packages on the server, `UTL_REF` verifies that the invoker has the appropriate privileges to access the object pointed to by the `REF`.

> **Note:**
>
> This is in contrast to PL/SQL packages/procedures on the server which operate with definer's privileges, where the package owner must have the appropriate privileges to perform the desired operations.

Thus, if UTL_REF is defined under user SYS, and user A invokes UTL_REF.SELECT to select an object from a reference, then user A (the invoker) requires the privileges to check.

When invoked from client-side PL/SQL code, UTL_REF operates with the privileges of the client session under which the PL/SQL execution is being done.

296.3 UTL_REF Types

An object type is a composite datatype defined by the user or supplied as a library type.

You can create the object type employee_type using the following syntax:

```
CREATE TYPE employee_type AS OBJECT (
   name    VARCHAR2(20),
   id      NUMBER,

member function GET_ID
    (name VARCHAR2)
   RETURN MEMBER);
```

The object type employee_type is a user-defined type that contains two attributes, name and id, and a member function, GET_ID().

You can create an object table using the following SQL syntax:

```
CREATE TABLE employee_table OF employee_type;
```

296.4 UTL_REF Exceptions

Exceptions can be returned during execution of UTL_REF functions for various reasons.

For example, the following scenarios would result in exceptions:

- The object selected does not exist. This could be because either:
  1. The object has been deleted, or the given reference is dangling (invalid).
  2. The object table was dropped or does not exist.
- The object cannot be modified or locked in a serializable transaction. The object was modified by another transaction after the serializable transaction started.
- You do not have the privilege to select or modify the object. The caller of the UTL_REF subprogram must have the proper privilege on the object that is being selected or modified.

Table 296-1 *UTL_REF Exceptions*

Exceptions	Description
errnum == 942	Insufficient privileges.
errnum == 1031	Insufficient privileges.
errnum == 8177	Unable to serialize, if in a serializable transaction.
errnum == 60	Deadlock detected.
errnum == 1403	No data found (if the REF is NULL, and so on.).

The UTL_REF package does not define any named exceptions. You may define exception handling blocks to catch specific exceptions and to handle them appropriately.

296.5 Summary of UTL_REF Subprograms

This table lists the `UTL_REF` subprograms and briefly describes them.

Table 296-2 UTL_REF Package Subprograms

Subprogram	Description
DELETE_OBJECT Procedure	Deletes an object given a reference
LOCK_OBJECT Procedure	Locks an object given a reference
SELECT_OBJECT Procedure	Selects an object given a reference
UPDATE_OBJECT Procedure	Updates an object given a reference

296.5.1 DELETE_OBJECT Procedure

This procedure deletes an object given a reference.

The semantic of this subprogram is similar to the following SQL statement:

```
DELETE FROM object_table
WHERE REF(t) = reference;
```

Unlike the preceding SQL statement, this subprogram does not require you to specify the object table name where the object resides.

Syntax

```
UTL_REF.DELETE_OBJECT (
   reference IN REF "<typename>");
```

Parameters

Table 296-3 DELETE_OBJECT Procedure Parameters

Parameter	Description
`reference`	Reference of the object to delete.

Exceptions

May be raised.

Examples

The following example illustrates usage of the `UTL_REF` package to implement this scenario: if an employee of a company changes their address, their manager should be notified.

... declarations of `Address_t` and others...

```
CREATE OR REPLACE TYPE Person_t (
   name    VARCHAR2(64),
   gender  CHAR(1),
   address Address_t,
   MEMBER PROCEDURE setAddress(addr IN Address_t)
);
```

```
CREATE OR REPLACE TYPE BODY Person_t (
   MEMBER PROCEDURE setAddress(addr IN Address_t) IS
   BEGIN
      address := addr;
   END;
);

CREATE OR REPLACE TYPE Employee_t (
```

Under `Person_t`: Simulate implementation of inheritance using a `REF` to `Person_t` and delegation of `setAddress` to it.

```
   thePerson  REF Person_t,
   empno      NUMBER(5),
   deptREF    Department_t,
   mgrREF     Employee_t,
   reminders  StringArray_t,
   MEMBER PROCEDURE setAddress(addr IN Address_t),
   MEMBER procedure addReminder(reminder VARCHAR2);
);

CREATE TYPE BODY Employee_t (
   MEMBER PROCEDURE setAddress(addr IN Address_t) IS
      myMgr Employee_t;
      meAsPerson Person_t;
   BEGIN
```

Update the address by delegating the responsibility to `thePerson`. Lock the Person object from the reference, and also select it:

```
      UTL_REF.LOCK_OBJECT(thePerson, meAsPerson);
      meAsPerson.setAddress(addr);
```

Delegate to `thePerson`:

```
        UTL_REF.UPDATE_OBJECT(thePerson, meAsPerson);
        if mgr is NOT NULL THEN
```

Give the manager a reminder:

```
         UTL_REF.LOCK_OBJECT(mgr);
         UTL_REF.SELECT_OBJECT(mgr, myMgr);
         myMgr.addReminder
         ('Update address in the employee directory for' ||
         thePerson.name || ', new address: ' || addr.asString);
         UTL_REF.UPDATE_OBJECT(mgr, myMgr);
      END IF;
   EXCEPTION
      WHEN OTHERS THEN
      errnum := SQLCODE;
      errmsg := SUBSTR(SQLERRM, 1, 200);
```

296.5.2 LOCK_OBJECT Procedure

This procedure locks an object given a reference. In addition, this procedure lets the program select the locked object.

The semantic of this subprogram is similar to the following SQL statement:

```
SELECT VALUE(t)
  INTO object
  FROM object_table t
```

```
  WHERE REF(t) = reference
  FOR UPDATE;
```

Unlike the preceding SQL statement, this subprogram does not require you to specify the object table name where the object resides. It is not necessary to lock an object before updating/deleting it.

Syntax

```
UTL_REF.LOCK_OBJECT (
   reference IN REF "<typename>");

UTL_REF.LOCK_OBJECT (
   reference IN REF "<typename>",
   object    IN OUT "<typename>");
```

Parameters

Table 296-4 LOCK_OBJECT Procedure Parameters

Parameter	Description
`reference`	Reference of the object to lock.
`object`	The PL/SQL variable that stores the locked object. This variable should be of the same object type as the locked object.

Exceptions

May be raised.

296.5.3 SELECT_OBJECT Procedure

This procedure selects an object given its reference. The selected object is retrieved from the database and its value is put into the PL/SQL variable 'object'.

The semantic of this subprogram is similar to the following SQL statement:

```
SELECT VALUE(t)
INTO object
FROM object_table t
WHERE REF(t) = reference;
```

Unlike the preceding SQL statement, this subprogram does not require you to specify the object table name where the object resides.

Syntax

```
UTL_REF.SELECT_OBJECT (
   reference IN REF "<typename>",
   object    IN OUT "<typename>");
```

Parameters

Table 296-5 SELECT_OBJECT Procedure Parameters

Parameter	Description
`reference`	Reference to the object to select or retrieve.

Table 296-5 (Cont.) SELECT_OBJECT Procedure Parameters

Parameter	Description
object	The PL/SQL variable that stores the selected object; this variable should be of the same object type as the referenced object.

Exceptions

May be raised.

296.5.4 UPDATE_OBJECT Procedure

This procedure updates an object given a reference. The referenced object is updated with the value contained in the PL/SQL variable 'object'.

The semantic of this subprogram is similar to the following SQL statement:

```
UPDATE object_table t
SET VALUE(t) = object
WHERE REF(t) = reference;
```

Unlike the preceding SQL statement, this subprogram does not require you to specify the object table name where the object resides.

Syntax

```
UTL_REF.UPDATE_OBJECT (
   reference IN REF "<typename>",
   object    IN     "<typename>");
```

Parameters

Table 296-6 UPDATE_OBJECT Procedure Parameters

Parameter	Description
reference	Reference of the object to update.
object	The PL/SQL variable that contains the new value of the object. This variable should be of the same object type as the object to update.

Exceptions

May be raised.

297
UTL_RPADV

The `UTL_RPADV` package provides subprograms to collect and analyze statistics for the Oracle Replication components in a distributed database environment. This package uses the Oracle Replication Performance Advisor to gather statistics.

This chapter contains the following topic:

- UTL_RPADV Overview
- DBMS_COMPARISON Security Model
- UTL_RPADV Operational Notes
- Summary of UTL_RPADV Subprograms

297.1 UTL_RPADV Overview

This package enables you to collect and analyze statistics about the performance or Oracle Replication components. You can either collect statistics on demand or you can create a monitoring job that continually monitors Oracle Replication performance.

When this package is used on an Oracle Database 11*g* Release 2 (11.2) database, it can monitor Oracle Database 10*g* Release 2 (10.2) and later databases. It cannot monitor databases before release 10.2.

297.2 DBMS_COMPARISON Security Model

Security on this package can be controlled by either granting `EXECUTE` on this package to selected users or roles, or by granting `EXECUTE_CATALOG_ROLE` to selected users or roles.

If subprograms in the package are run from within a stored procedure, then the user who runs the subprograms must be granted `EXECUTE` privilege on the package directly. It cannot be granted through a role.

To ensure that the user who runs the subprograms in this package has the necessary privileges, configure an Oracle Replication administrator and connect as the Oracle Replication administrator when using this package.

297.3 UTL_RPADV Operational Notes

To use this package, you must connect to an Oracle database as an Oracle Replication administrator and run the `utlrpadv.sql` script in the `rdbms/admin` directory in `ORACLE_HOME`.

The `utlrpadv.sql` script creates the following tables:

- STREAMS$_PA_COMPONENT Table
- STREAMS$_PA_COMPONENT_LINK Table
- STREAMS$_PA_COMPONENT_PROP Table
- STREAMS$_PA_COMPONENT_STAT Table

- STREAMS$_PA_CONTROL Table
- STREAMS$_PA_DATABASE Table
- STREAMS$_PA_DATABASE_PROP Table
- STREAMS$_PA_MONITORING Table
- STREAMS$_PA_PATH_BOTTLENECK Table
- STREAMS$_PA_PATH_STAT Table
- STREAMS$_PA_SHOW_COMP_STAT Table
- STREAMS$_PA_SHOW_PATH_STAT Table

The Oracle Replication Performance Advisor populates these tables when it is run.

STREAMS$_PA_COMPONENT Table

The `STREAMS$_PA_COMPONENT` table displays information about the Oracle Replication components at each database.

Table 297-1 STREAMS$_PA_COMPONENT Table

Column	Datatype	NULL	Description
`COMPONENT_ID`	`NUMBER`	`NOT NULL`	Identification number assigned to the component by the Oracle Replication Performance Advisor
`COMPONENT_NAME`	`VARCHAR2(194)`		Name of the component
`COMPONENT_DB`	`VARCHAR2(128)`		Name of the database that contains the component
`COMPONENT_TYPE`	`VARCHAR2(20)`		Type of the component The following types are possible: • `CAPTURE` for a capture process • `PROPAGATION SENDER` for a propagation sender • `PROPAGATION RECEIVER` for a propagation receiver • `APPLY` for an apply process • `QUEUE` for a queue
`COMPONENT_CHANGED_TIME`	`DATE`		Time when the component was last changed

STREAMS$_PA_COMPONENT_LINK Table

The `STREAMS$_PA_COMPONENT_LINK` table displays information about how information flows between Oracle Replication components.

Table 297-2 STREAMS$_PA_COMPONENT_LINK Table

Column	Datatype	NULL	Description
`PATH_ID`	`NUMBER`	`NOT NULL`	Identification number assigned to the path by the Oracle Replication Performance Advisor

Table 297-2 (Cont.) STREAMS$_PA_COMPONENT_LINK Table

Column	Datatype	NULL	Description
PATH_KEY	VARCHAR2(4000)		Unique key assigned to the path by the Oracle Replication Performance Advisor
SOURCE_COMPONENT_ID	NUMBER	NOT NULL	Source component ID for the path The path starts with this component.
DESTINATION_COMPONENT_ID	NUMBER	NOT NULL	Destination component ID for the path The path ends with this component.
POSITION	NUMBER		Position of the component in the path

STREAMS$_PA_COMPONENT_PROP Table

The STREAMS$_PA_COMPONENT_PROP table displays information about capture processes and apply processes necessary for analysis by the Replication Performance Advisor.

Table 297-3 STREAMS$_PA_COMPONENT_PROP Table

Column	Datatype	NULL	Description
COMPONENT_ID	NUMBER	NOT NULL	Identification number assigned to the component by the Oracle Replication Performance Advisor

Table 297-3 (Cont.) STREAMS$_PA_COMPONENT_PROP Table

Column	Datatype	NULL	Description
PROP_NAME	VARCHAR2(30)		Property name For a capture process, the component properties include the following: • SOURCE_DATABASE - The source database for the changes captured by the capture process • PARALLELISM - The setting for the parallelism capture process parameter • OPTIMIZATION_MODE - Indicates whether the capture process uses combined capture and apply (greater than zero) or does not use combined capture and apply (0) For an apply process, the component properties include the following: • SOURCE_DATABASE - The source database for the messages applied by the apply process • PARALLELISM - The setting for the parallelism apply process parameter • APPLY_CAPTURED - Indicates whether the apply process applies captured messages (YES) persistent messages (NO) • MESSAGE_DELIVERY_MODE - Either buffered or persistent
PROP_VALUE	VARCHAR2(30)		Property value

STREAMS$_PA_COMPONENT_STAT Table

The STREAMS$_PA_COMPONENT_STAT table displays performance statistics and session statistics about each Oracle Replication component.

Table 297-4 STREAMS$_PA_COMPONENT_STAT Table

Column	Datatype	NULL	Description
ADVISOR_RUN_ID	NUMBER		Identification number of the Oracle Replication Performance Advisor run
ADVISOR_RUN_TIME	DATE		Time when the Oracle Replication Performance Advisor was run for the advisor run ID

Table 297-4 (Cont.) STREAMS$_PA_COMPONENT_STAT Table

Column	Datatype	NULL	Description
COMPONENT_ID	NUMBER		Identification number assigned to the component by the Oracle Replication Performance Advisor
STATISTIC_TIME	DATE		Time when the statistic was recorded
STATISTIC_NAME	VARCHAR2(64)		Name of the statistic
STATISTIC_VALUE	NUMBER		Value recorded for the statistic
STATISTIC_UNIT	VARCHAR2(64)		Unit of measurement for the statistic
SUB_COMPONENT_TYPE	VARCHAR2(64)		Type of the subcomponent Only capture processes and apply processes have subcomponents. The following capture process subcomponent types are possible: • LOGMINER READER for a builder server of a capture process • LOGMINER PREPARER for a preparer server of a capture process • LOGMINER BUILDER for a reader server of a capture process • CAPTURE SESSION for a capture process session The following apply process subcomponent types are possible: • PROPAGATION SENDER+RECEIVER for sending LCRs from a capture process directly to an apply process in a combined capture and apply configuration in which both the capture process and apply process run on a single database • APPLY READER for a reader server of an apply process • APPLY COORDINATOR for a coordinator process of an apply process • APPLY SERVER for a reader server of an apply process
SESSION_ID	NUMBER		Identification number of the session for the component. Query the V$SESSION view for information about the session.
SESSION_SERIAL#	NUMBER		Session serial number of the session for the component. Query the V$SESSION view for information about the session.

STREAMS$_PA_CONTROL Table

The STREAMS$_PA_CONTROL table displays the parameters set for the COLLECT_STATS procedure in this package. The parameters control the monitoring behavior.

Table 297-5 STREAMS$_PA_CONTROL Table

Column	Datatype	NULL	Description
ADVISOR_RUN_ID	NUMBER		Identification number of the Oracle Replication Performance Advisor run
ADVISOR_RUN_TIME	DATE		Time when the Oracle Replication Performance Advisor was last run
PARAM_NAME	VARCHAR2(30)		The name of the parameter
PARAM_VALUE	VARCHAR2(4000)		The value set for the parameter
PARAM_UNIT	VARCHAR2(30)		The unit of the parameter

STREAMS$_PA_DATABASE Table

The STREAMS$_PA_DATABASE table displays information about each database that contains Oracle Replication components.

Table 297-6 STREAMS$_PA_DATABASE Table

Column	Datatype	NULL	Description
GLOBAL_NAME	VARCHAR2(128)	NOT NULL	Global name of the database analyzed by the Oracle Replication Performance Advisor
LAST_QUERIED	DATE		The time when the Performance Advisor successfully collected information from a database in its last run
ERROR_NUMBER	NUMBER		The error number of the error encountered when the database was last queried
ERROR_MESSAGE	VARCHAR2(4000)		The error message of the error encountered when the database was last queried

STREAMS$_PA_DATABASE_PROP Table

The STREAMS$_PA_DATABASE_PROP table displays Oracle Replication database property information necessary for analysis by the Replication Performance Advisor.

Table 297-7 STREAMS$_PA_DATABASE_PROP Table

Column	Datatype	NULL	Description
GLOBAL_NAME	VARCHAR2(128)	NOT NULL	Global name of the database analyzed by the Oracle Replication Performance Advisor
PROP_NAME	VARCHAR2(30)		Property name The database properties include the following: • VERSION • COMPATIBILITY • MANAGEMENT_PACK_ACCESS • DB_UNIQUE_NAME
PROP_VALUE	VARCHAR2(30)		Property value

STREAMS$_PA_MONITORING Table

The STREAMS$_PA_MONITORING table displays information about each monitoring job running in a database.

Table 297-8 STREAMS$_PA_MONITORING Table

Column	Datatype	NULL	Description
JOB_NAME	VARCHAR2(30)	NOT NULL	Name of the monitoring job
CLIENT_NAME	VARCHAR2(30)		Name of the client that submitted the job **See Also:** "Full Monitoring Job Names"
QUERY_USER_NAME	VARCHAR2(30)		User granted privileges to view the monitoring results
SHOW_STATS_TABLE	VARCHAR2(30)		Name of the table used by the SHOW_STATS procedure to display statistics
STARTED_TIME	TIMESTAMP		Time the monitoring job started
STOPPED_TIME	TIMESTAMP		Time the monitoring job last stopped
ALTERED_TIME	TIMESTAMP		Time the monitoring job was last altered
STATE	VARCHAR2(30)		State of the monitoring job, either ENABLED or STOPPED

STREAMS$_PA_PATH_BOTTLENECK Table

The STREAMS$_PA_PATH_BOTTLENECK table displays information about Oracle Replication components that might be slowing down the flow of messages.

Table 297-9 STREAMS$_PA_PATH_BOTTLENECK Table

Column	Datatype	NULL	Description
ADVISOR_RUN_ID	NUMBER		Identification number of the Oracle Replication Performance Advisor run
ADVISOR_RUN_TIME	DATE		Time when the Oracle Replication Performance Advisor was last run
ADVISOR_RUN_REASON	VARCHAR2(4000)		Reason for the bottleneck
PATH_ID	NUMBER		Identification number assigned to the path by the Oracle Replication Performance Advisor
PATH_KEY	VARCHAR2(4000)		Unique key assigned to the path by the Oracle Replication Performance Advisor
COMPONENT_ID	NUMBER		Identification number assigned to the component by the Oracle Replication Performance Advisor
TOP_SESSION_ID	NUMBER		Session ID of the top component. Query the V$SESSION view for information about the session.
TOP_SESSION_SERIAL#	NUMBER		Session serial number of the top component. Query the V$SESSION view for information about the session.
ACTION_NAME	VARCHAR2(32)		Action name for the top session

Table 297-9 (Cont.) STREAMS$_PA_PATH_BOTTLENECK Table

Column	Datatype	NULL	Description
BOTTLENECK_IDENTIFIED	VARCHAR2(30)		Whether a bottleneck was identified

STREAMS$_PA_PATH_STAT Table

The STREAMS$_PA_PATH_STAT table displays performance statistics about each stream path.

Table 297-10 STREAMS$_PA_PATH_STAT Table

Column	Datatype	NULL	Description
ADVISOR_RUN_ID	NUMBER		Identification number of the Oracle Replication Performance Advisor run
ADVISOR_RUN_TIME	DATE		Time when the Oracle Replication Performance Advisor was run for the advisor run ID
PATH_ID	NUMBER		Identification number assigned to the path by the Oracle Replication Performance Advisor
PATH_KEY	VARCHAR2(4000)		Unique key assigned to the path by the Oracle Replication Performance Advisor
STATISTIC_TIME	DATE		Time when the statistic was recorded
STATISTIC_NAME	VARCHAR2(64)		Name of the statistic
STATISTIC_VALUE	NUMBER		Value recorded for the statistic
STATISTIC_UNIT	VARCHAR2(64)		Unit of measurement for the statistic

STREAMS$_PA_SHOW_COMP_STAT Table

The STREAMS$_PA_SHOW_COMP_STAT table displays statistics for Oracle Replication components.

Table 297-11 STREAMS$_PA_SHOW_COMP_STAT Table

Column	Datatype	NULL	Description
ADVISOR_RUN_ID	NUMBER		Identification number of the Oracle Replication Performance Advisor run
ADVISOR_RUN_TIME	DATE		Time when the Oracle Replication Performance Advisor was last run
PATH_ID	NUMBER		Identification number assigned to the path by the Oracle Replication Performance Advisor
POSITION	NUMBER		Position of the component in the path
COMPONENT_ID	NUMBER		Identification number assigned to the component by the Oracle Replication Performance Advisor
COMPONENT_NAME	VARCHAR2(194)		Name of the component

Table 297-11 (Cont.) STREAMS$_PA_SHOW_COMP_STAT Table

Column	Datatype	NULL	Description
COMPONENT_TYPE	VARCHAR2(30)		Type of the component The following types are possible: • CAPTURE for a capture process • PROPAGATION SENDER for a propagation sender • PROPAGATION RECEIVER for a propagation receiver • APPLY for an apply process • QUEUE for a queue
SUB_COMPONENT_TYPE	VARCHAR2(30)		Type of the subcomponent Only capture processes and apply processes have subcomponents. The following capture process subcomponent types are possible: • LOGMINER READER for a builder server of a capture process • LOGMINER PREPARER for a preparer server of a capture process • LOGMINER BUILDER for a reader server of a capture process • CAPTURE SESSION for a capture process session The following apply process subcomponent types are possible: • PROPAGATION SENDER+RECEIVER for sending LCRs from a capture process directly to an apply process in a combined capture and apply configuration in which both the capture process and apply process run on a single database • APPLY READER for a reader server of an apply process • APPLY COORDINATOR for a coordinator process of an apply process • APPLY SERVER for a reader server of an apply process
SESSION_ID	NUMBER		Identification number of the session for the component. Query the V$SESSION view for information about the session.
SESSION_SERIAL#	NUMBER		Session serial number of the session for the component. Query the V$SESSION view for information about the session.
STATISTIC_ALIAS	VARCHAR2(30)		Name of the statistic
STATISTIC_NAME	VARCHAR2(128)		Name of the statistic
STATISTIC_VALUE	NUMBER		Value recorded for the statistic
STATISTIC_UNIT	VARCHAR2(128)		Unit of measurement for the statistic

STREAMS$_PA_SHOW_PATH_STAT Table

The STREAMS$_PA_SHOW_PATH_STAT table displays statistics for the stream paths in an Oracle Replication configuration. A monitoring job uses this table as the default table for the statistics collected for stream paths.

Table 297-12 STREAMS$_PA_SHOW_PATH_STAT Table

Column	Datatype	NULL	Description
PATH_ID	NUMBER		Identification number assigned to the path by the Oracle Replication Performance Advisor
ADVISOR_RUN_ID	NUMBER		Identification number of the Oracle Replication Performance Advisor run
ADVISOR_RUN_TIME	DATE		Time when the Oracle Replication Performance Advisor was last run
SETTING	VARCHAR2(2000)		Setting for the Oracle Replication Performance Advisor Run
STATISTICS	VARCHAR2(4000)		Component-level statistics
SESSION_STATISTICS	VARCHAR2(4000)		Session-level statistics
OPTIMIZATION	NUMBER		Whether the path uses the combined capture and apply optimization 0 (zero) means that the path does not use the combined capture and apply optimization. 1 means that the path uses the combined capture and apply optimization.

297.4 Summary of UTL_RPADV Subprograms

This table lists the UTL_RPADV subprograms and briefly describes them.

Table 297-13 UTL_RPADV Package Subprograms

Subprogram	Description
ALTER_MONITORING Procedure	Alters the monitoring job submitted by the current user.
COLLECT_STATS Procedure	Uses the Oracle Replication Performance Advisor to gather statistics about the Oracle Replication components and subcomponents in a distributed database environment.
IS_MONITORING Function	Checks whether a monitoring job is currently running.
SHOW_STATS Procedure	Generates output that includes the statistics gathered by the COLLECT_STATS procedure.
SHOW_STATS_HTML Procedure	Generates HTML output that includes the statistics gathered by the COLLECT_STATS procedure.
START_MONITORING Procedure	Starts a monitoring job.
STOP_MONITORING Procedure	Stops a monitoring job.

297.4.1 ALTER_MONITORING Procedure

This procedure alters the monitoring job submitted by the current user.

Syntax

```
UTL_RPADV.ALTER_MONITORING(
   interval                     IN NUMBER  DEFAULT NULL,
   top_event_threshold          IN NUMBER  DEFAULT NULL,
   bottleneck_idle_threshold    IN NUMBER  DEFAULT NULL,
   bottleneck_flowctrl_threshold IN NUMBER  DEFAULT NULL,
   retention_time               IN NUMBER  DEFAULT NULL);
```

Parameters

Table 297-14 ALTER_MONITORING Procedure Parameters

Parameter	Description
`interval`	The amount of time, in seconds, between each Performance Advisor run. The maximum is 3600 seconds. If `NULL`, then the current value is not changed.
`top_event_threshold`	A percentage that determines whether a top wait event statistic is collected. The percentage for a wait event must be greater than the value specified in this parameter for the procedure to collect the wait event statistic. For example, if `15` is specified, then only wait events with a value larger than 15% are collected. If `NULL`, then the current value is not changed.
`bottleneck_idle_threshold`	A percentage that determines whether an Oracle Replication component session is eligible for bottleneck analysis based on its `IDLE` percentage. The `IDLE` percentage must be less than or equal to the value specified in this parameter for the Oracle Replication component session to be eligible for bottleneck analysis. For example, if `50` is specified, then only components that are idle 50% of the time or less are eligible for bottleneck analysis. If `NULL`, then the current value is not changed.
`bottleneck_flowctrl_threshold`	A percentage that determines whether an Oracle Replication component session is eligible for bottleneck analysis based on its `FLOW CONTROL` percentage. The `FLOW CONTROL` percentage must be less than or equal to the value specified in this parameter for the Oracle Replication component session to be eligible for bottleneck analysis. For example, if `50` is specified, then only components that are paused for flow control 50% of the time or less are eligible for bottleneck analysis. If `NULL`, then the current value is not changed.
`retention_time`	The number of hours to retain monitoring results. If `NULL`, then the current value is not changed.

Exceptions

Table 297-15 ALTER_MONITORING Procedure Exceptions

Exception	Description
ORA-20113	no active monitoring job found

297.4.2 COLLECT_STATS Procedure

This procedure uses the Oracle Replication Performance Advisor to gather statistics about the Oracle Replication components and subcomponents in a distributed database environment.

> **Note:**
>
> This procedure commits.

Syntax

```
UTL_RPADV.COLLECT_STATS(
   interval                      IN NUMBER  DEFAULT 60,
   num_runs                      IN NUMBER  DEFAULT 10,
   comp_stat_table              IN VARCHAR2  DEFAULT 'STREAMS$_ADVISOR_COMP_STAT',
   path_stat_table              IN VARCHAR2  DEFAULT 'STREAMS$_ADVISOR_PATH_STAT',
   top_event_threshold           IN NUMBER  DEFAULT 15,
   bottleneck_idle_threshold     IN NUMBER  DEFAULT 50,
   bottleneck_flowctrl_threshold IN NUMBER  DEFAULT 50);
```

Parameters

Table 297-16 *COLLECT_STATS Procedure Parameters*

Parameter	Description
interval	The amount of time, in seconds, between each Performance Advisor run. The maximum is 3600 seconds.
num_runs	The number of times that the Oracle Replication Performance Advisor is run by the procedure.
comp_stat_table	The name of the table that stores the statistics collected for Oracle Replication components and subcomponents. Specify the table name as *[schema_name.]object_name*. If the schema is not specified, then the current user is the default. The procedure creates the specified table if it does not exist. Oracle recommends that you use the default table STREAMS$_ADVISOR_COMP_STAT. See "Usage Notes" for more information about this parameter.

Table 297-16 (Cont.) *COLLECT_STATS Procedure Parameters*

Parameter	Description
path_stat_table	The name of the table that stores the statistics collected for stream paths. Specify the table name as *[schema_name.]object_name*. If the schema is not specified, then the current user is the default. The procedure creates the specified table if it does not exist. Oracle recommends that you use the default table STREAMS$_ADVISOR_PATH_STAT. See "Usage Notes" for more information about this parameter.
top_event_threshold	A percentage that determines whether a top wait event statistic is collected. The percentage for a wait event must be greater than the value specified in this parameter for the procedure to collect the wait event statistic. For example, if 15 is specified, then only wait events with a value larger than 15% are collected.
bottleneck_idle_threshold	A percentage that determines whether an Oracle Replication component session is eligible for bottleneck analysis based on its IDLE percentage. The IDLE percentage must be less than or equal to the value specified in this parameter for the Oracle Replication component session to be eligible for bottleneck analysis. For example, if 50 is specified, then only components that are idle 50% of the time or less are eligible for bottleneck analysis.
bottleneck_flowctrl_threshold	A percentage that determines whether an Oracle Replication component session is eligible for bottleneck analysis based on its FLOW CONTROL percentage. The FLOW CONTROL percentage must be less than or equal to the value specified in this parameter for the Oracle Replication component session to be eligible for bottleneck analysis. For example, if 50 is specified, then only components that are paused for flow control 50% of the time or less are eligible for bottleneck analysis.

Usage Notes

The table specified in the path_stat_table parameter stores stream path statistics. This table also concatenates the component and subcomponent statistics stored in the table specified in the comp_stat_table parameter. The SHOW_STATS procedure in this package shows only the statistics stored in the table specified in the path_stat_table parameter.

297.4.3 IS_MONITORING Function

This function checks whether a monitoring job is currently running. This function either returns TRUE if a monitoring job is currently running or FALSE if a monitoring job is not currently running.

A monitoring job is submitted using the START_MONITORING procedure.

> **See Also:**
>
> "START_MONITORING Procedure"

Syntax

```
UTL_RPADV.IS_MONITORING(
   job_name    IN VARCHAR2  DEFAULT 'STREAMS$_MONITORING_JOB',
   client_name IN VARCHAR2  DEFAULT NULL)
RETURN BOOLEAN;
```

Parameters

Table 297-17 IS_MONITORING Function Parameters

Parameter	Description
job_name	The name of the job for which to check.
client_name	The name of the client that submitted the job.

297.4.4 SHOW_STATS Procedure

This procedure generates output that includes the statistics gathered by the COLLECT_STATS and START_MONITORING procedures.

The output is formatted so that it can be imported into a spreadsheet for analysis.

> **Note:**
>
> This procedure does not commit.

> **See Also:**
>
> - "COLLECT_STATS Procedure"
> - "START_MONITORING Procedure"

Syntax

```
UTL_RPADV.SHOW_STATS(
   path_stat_table IN VARCHAR2  DEFAULT 'STREAMS$_ADVISOR_PATH_STAT',
   path_id           IN NUMBER    DEFAULT NULL,
   bgn_run_id        IN NUMBER    DEFAULT -1,
   end_run_id        IN NUMBER    DEFAULT -10,
   show_path_id      IN BOOLEAN   DEFAULT TRUE,
   show_run_id       IN BOOLEAN   DEFAULT TRUE,
   show_run_time     IN BOOLEAN   DEFAULT TRUE,
   show_optimization IN BOOLEAN   DEFAULT TRUE,
   show_setting      IN BOOLEAN   DEFAULT FALSE,
   show_stat         IN BOOLEAN   DEFAULT TRUE,
```

```
  show_sess         IN BOOLEAN   DEFAULT FALSE,
  show_legend       IN BOOLEAN   DEFAULT TRUE);
```

Parameters

Table 297-18 *SHOW_STATS Procedure Parameters*

Parameter	Description
path_stat_table	The name of the table that contains the stream path statistics. Specify the table name as *[schema_name.]object_name*. If the schema is not specified, then the current user is the default. When you gather statistics using the COLLECT_STATS procedure, this table is specified in the path_stat_table parameter in the COLLECT_STATS procedure. The default table is STREAMS$_ADVISOR_PATH_STAT. When you gather statistics using the START_MONITORING procedure, you can determine the name for this table by querying the SHOW_STATS_TABLE column in the STREAMS$_PA_MONITORING view. The default table for a monitoring job is STREAMS$_PA_SHOW_PATH_STAT.
path_id	A stream path ID. If non-NULL, then the procedure shows output for the specified stream path only. If NULL, then the procedure shows output for all active stream paths.
bgn_run_id	The first Oracle Replication Performance Advisor run ID to show in the range of runs. See "Usage Notes" for more information about this parameter.
end_run_id	The last Oracle Replication Performance Advisor run ID to show in the range of runs. See "Usage Notes" for more information about this parameter.
show_path_id	If TRUE, then the path ID for each stream path is included in the output. If FALSE, then the path ID for each stream path is not included in the output.
show_run_id	If TRUE, then the Oracle Replication Performance Advisor run ID is included in the output. If FALSE, then the Oracle Replication Performance Advisor run ID is not included in the output.
show_run_time	If TRUE, then the Oracle Replication Performance Advisor run time is included in the output. If FALSE, then the Oracle Replication Performance Advisor run time is not included in the output.
show_optimization	If TRUE, then path output includes information pertaining to the combined capture and apply optimization. If FALSE, then path output does not include information pertaining to the combined capture and apply optimization.
show_setting	If TRUE, then the settings for the threshold parameters are included in the output. The threshold parameters are the top_event_threshold, bottleneck_idle_threshold, and bottleneck_flowctrl_threshold parameters in the COLLECT_STATS procedure. If FALSE, then the settings for the threshold parameters are not included in the output.

Table 297-18 (Cont.) *SHOW_STATS Procedure Parameters*

Parameter	Description
show_stat	If TRUE, then the component-level and subcomponent-level statistics are included in the output. These components include capture processes, queues, propagation senders, propagation receivers, and apply processes. The subcomponents are the subcomponents for capture processes and apply processes. If FALSE, then the component-level and subcomponent-level statistics are not included in the output.
show_sess	If TRUE, then the session-level statistics are included in the output. Session-level statistics include IDLE, FLOW CONTROL, and EVENT statistics. If FALSE, then the session-level statistics are not included in the output.
show_legend	If TRUE, then the legend is included in the output. The legend describes the abbreviations used in the output. If FALSE, then the legend is not included in the output.

Usage Notes

Use the bgn_run_id and end_run_id together to specify the range of Oracle Replication Performance Advisor runs to display. Positive numbers show statistics from an earlier run forward. Negative numbers show statistics from a later run backward.

For example, if bgn_run_id is set to 1 and end_run_id is set to 10, then the procedure shows statistics for the first ten Oracle Replication Performance Advisor runs.

However, if bgn_run_id is set to -1 and end_run_id is set to -10, then the procedure shows statistics for the last ten Oracle Replication Performance Advisor runs.

297.4.5 SHOW_STATS_HTML Procedure

This procedure generates HTML output that includes the statistics gathered by the COLLECT_STATS and START_MONITORING procedures.

> **Note:**
>
> This procedure does not commit.

> **See Also:**
>
> - "COLLECT_STATS Procedure"
> - "START_MONITORING Procedure"

Syntax

```
UTL_RPADV.SHOW_STATS_HTML(
   directory       IN VARCHAR2,
   reportname      IN VARCHAR2  DEFAULT 'RPADVREPORT.HTML',
```

```
  comp_stat_table IN VARCHAR2  DEFAULT 'STREAMS$_ADVISOR_COMP_STAT',
  path_id         IN NUMBER    DEFAULT NULL,
  bgn_run_id      IN NUMBER    DEFAULT -1,
  end_run_id      IN NUMBER    DEFAULT -10,
  detailed        IN BOOLEAN   DEFAULT TRUE);
```

Parameters

Table 297-19 SHOW_STATS_HTML Procedure Parameters

Parameter	Description
`directory`	The directory object for the directory on the local computer system into which the generated HTML report is placed The specified directory object must be created using the SQL statement `CREATE DIRECTORY`, and the user who invokes the procedure must have `READ` and `WRITE` privilege on each one.
`reportname`	The name of the HTML report
`comp_stat_table`	The name of the table that stores the statistics collected for Oracle Replication components and subcomponents. Specify the table name as *`[schema_name.]object_name`*. If the schema is not specified, then the current user is the default. When you gather statistics using the `COLLECT_STATS` procedure, this table is specified in the `comp_stat_table` parameter in the `COLLECT_STATS` procedure. The default table is `STREAMS$_ADVISOR_COMP_STAT`. When you gather statistics using the `START_MONITORING` procedure, you can determine the name for this table by querying the `SHOW_STATS_TABLE` column in the `STREAMS$_PA_MONITORING` view. The default table for a monitoring job is `STREAMS$_PA_SHOW_PATH_STAT`. Oracle recommends that you start a monitoring job with the `START_MONITORING` procedure in this package and use the appropriate the `STREAMS$_PA_SHOW_PATH_STAT` table.
`path_id`	A stream path ID. If non-`NULL`, then the procedure shows output for the specified stream path only. If `NULL`, then the procedure shows output for all active stream paths.
`bgn_run_id`	The first Oracle Replication Performance Advisor run ID to show in the range of runs. See "Usage Notes" for more information about this parameter.
`end_run_id`	The last Oracle Replication Performance Advisor run ID to show in the range of runs. See "Usage Notes" for more information about this parameter.
`detailed`	If `TRUE`, then the procedure generates component-level statistics. If `FALSE`, then the procedure does not generate component-level statistics.

Usage Notes

Use the `bgn_run_id` and `end_run_id` together to specify the range of Oracle Replication Performance Advisor runs to display. Positive numbers show statistics from an earlier run forward. Negative numbers show statistics from a later run backward.

For example, if `bgn_run_id` is set to `1` and `end_run_id` is set to `10`, then the procedure shows statistics for the first ten Oracle Replication Performance Advisor runs.

However, if `bgn_run_id` is set to `-1` and `end_run_id` is set to `-10`, then the procedure shows statistics for the last ten Oracle Replication Performance Advisor runs.

297.4.6 START_MONITORING Procedure

This procedure starts a monitoring job.

This procedure runs the `COLLECT_STATS` procedure to gather statistics about the Oracle Replication components and subcomponents in a distributed database environment.

> **Note:**
>
> This procedure commits.

> **See Also:**
>
> - "COLLECT_STATS Procedure"

Syntax

```
UTL_RPADV.START_MONITORING(
   job_name                      IN VARCHAR2  DEFAULT 'STREAMS$_MONITORING_JOB',
   client_name                   IN VARCHAR2  DEFAULT NULL,
   query_user_name               IN VARCHAR2  DEFAULT NULL,
   interval                      IN NUMBER    DEFAULT 60,
   top_event_threshold           IN NUMBER    DEFAULT 15,
   bottleneck_idle_threshold     IN NUMBER    DEFAULT 50,
   bottleneck_flowctrl_threshold IN NUMBER    DEFAULT 50,
   retention_time                IN NUMBER    DEFAULT 24);
```

Parameters

Table 297-20 START_MONITORING Procedure Parameters

Parameter	Description
`job_name`	The name of the monitoring job to create.
`client_name`	The name of the client.
`query_user_name`	The user who will query the result tables. This procedure grants privileges to the specified user to enable the user to query the result tables.
`interval`	The amount of time, in seconds, between each Performance Advisor run. The maximum is 3600 seconds. The specified interval is used for the interval parameter in the `COLLECT_STATS` procedure.

Table 297-20 (Cont.) START_MONITORING Procedure Parameters

Parameter	Description
top_event_threshold	A percentage that determines whether a top wait event statistic is collected. The percentage for a wait event must be greater than the value specified in this parameter for the procedure to collect the wait event statistic. For example, if 15 is specified, then only wait events with a value larger than 15% are collected.
bottleneck_idle_threshold	A percentage that determines whether an Oracle Replication component session is eligible for bottleneck analysis based on its IDLE percentage. The IDLE percentage must be less than or equal to the value specified in this parameter for the Oracle Replication component session to be eligible for bottleneck analysis. For example, if 50 is specified, then only components that are idle 50% of the time or less are eligible for bottleneck analysis.
bottleneck_flowctrl_threshold	A percentage that determines whether an Oracle Replication component session is eligible for bottleneck analysis based on its FLOW CONTROL percentage. The FLOW CONTROL percentage must be less than or equal to the value specified in this parameter for the Oracle Replication component session to be eligible for bottleneck analysis. For example, if 50 is specified, then only components that are paused for flow control 50% of the time or less are eligible for bottleneck analysis.
retention_time	The number of hours to retain monitoring results.

Exceptions

Table 297-21 START_MONITORING Procedure Exceptions

Exception	Description
ORA-20111	cannot start monitoring due to active EM monitoring job Stop the Oracle Enterprise Manager (EM) monitoring job, and run the START_MONITORING procedure again.
ORA-20112	cannot start monitoring due to active Replication monitoring job Stop the Replication monitoring job, and run the START_MONITORING procedure again.

Usage Notes

The following are usage notes for the START_MONITORING procedure:

- Requirements for the User Running the Procedure
- Full Monitoring Job Names
- Restrictions on Monitoring Jobs

Requirements for the User Running the Procedure

The user who runs the START_MONITORING procedure must meet the following requirements:

- The user must have access to a database link to each database that contains Oracle Replication components.
- The user must have been granted privileges using the DBMS_XSTREAM_AUTH.GRANT_ADMIN_PRIVILEGE procedure, and each database link must connect to a user at the remote database that has been granted privileges using the DBMS_XSTREAM_AUTH.GRANT_ADMIN_PRIVILEGE procedure.

Full Monitoring Job Names

When you submit a monitoring job, the client name and job name are concatenated to form the full monitoring job name. You specify the client name using the client_name parameter and the job name using the job_name parameter when you run the START_MONITORING procedure. The client name for a monitoring job submitted by Oracle Enterprise Manager is always EM.

The following table show examples of full monitoring job names:

Setting for client_name Parameter	Setting for job_name parameter	Full Monitoring Job Name
NULL	STREAMS$_MONITORING_JOB	STREAMS$_MONITORING_JOB
EM	STREAMS$_MONITORING_JOB	EMSTREAMS$_MONITORING_JOB
strm	STREAMS$_MONITORING_JOB	strmSTREAMS$_MONITORING_JOB
strm	mjob1	strmmjob1

Restrictions on Monitoring Jobs

The following restrictions apply to monitoring jobs:

- The limit for the length of the full monitoring job name is 30 bytes.
- Two monitoring jobs cannot have the same full monitoring job name, even if the monitoring jobs were submitted by different schemas. The name check is not case-sensitive. For example, strmSTREAMS$_MONITORING_JOB and STRMSTREAMS$_MONITORING_JOB are considered to be the same name.
- Oracle Enterprise Manager can have at most one monitoring job for each database.
- Each schema can have at most one monitoring job.

297.4.7 STOP_MONITORING Procedure

This procedure stops a monitoring job that was submitted by the current user.

Syntax

```
UTL_RPADV.STOP_MONITORING(
   purge IN BOOLEAN  DEFAULT FALSE);
```

Parameters

Table 297-22 STOP_MONITORING Procedure Parameters

Parameter	Description
`purge`	If `TRUE`, then the procedure purges information about the monitoring job from the result tables. If `FALSE`, then the procedure retains information about the monitoring job in the result tables.

298
UTL_SMTP

The `UTL_SMTP` package is designed for sending electronic mails (e-mails) over Simple Mail Transfer Protocol (SMTP) as specified by RFC821.

This chapter contains the following topics:

- Overview
- Security Model
- Constants
- Types
- Reply Codes
- UTL_SMTP Operational Notes
- Exceptions
- Rules and Limits
- Examples
- Summary of UTL_SMTP Subprograms

> **See Also:**
>
> *Oracle Database Development Guide*

298.1 UTL_SMTP Overview

The `UTL_SMTP` protocol consists of a set of commands for an e-mail client to dispatch e-mails to an SMTP server. The `UTL_SMTP` package provides interfaces to the SMTP commands.

For many of the commands, the package provides both a procedural and a functional interface. The functional form returns the reply from the server for processing by the client. The procedural form checks the reply and raises an exception if the reply indicates a transient (400-range reply code) or permanent error (500-range reply code). Otherwise, it discards the reply.

Note that the original SMTP protocol communicates using 7-bit ASCII. Using `UTL_SMTP`, all text data (in other words, those in `VARCHAR2`) is converted to US7ASCII before it is sent to the server. Some implementations of SMTP servers that support SMTP extension 8BITMIME [RFC1652] support full 8-bit communication between client and server. The body of the DATA command can be transferred in full 8 bits, but the rest of the SMTP command and response must be in 7 bits. When the target SMTP server supports 8BITMIME extension, users of multibyte databases may convert their non-US7ASCII, multibyte `VARCHAR2` data to `RAW` and use the `WRITE_RAW_DATA` subprogram to send multibyte data using 8-bit MIME encoding.

`UTL_SMTP` provides for SMTP communication as specified in RFC821, but does not provide an API to format the content of the message according to RFC 822 (for example, setting the subject of an electronic mail). You must format the message appropriately. In addition,

UTL_SMTP does not have the functionality to implement an SMTP server for an e-mail clients to send e-mails using SMTP.

298.2 UTL_SMTP Security Model

This package is an invoker's rights package. The invoking user must have the connect privilege granted in the access control list assigned to the remote network host to which the user must connect.

> **Note:**
>
> For more information on managing fine-grained access, see *Oracle Database Security Guide*

298.3 UTL_SMTP Constants

UTL_SMTP defines several constants to use when specifying parameter values.

These are shown in the following table.

Table 298-1 UTL_SMTP Constants

Name	Type	Value	Description
ALL_SCHEMES	VARCHAR2(256)	'CRAM-MD5 PLAIN LOGIN'	List of all authentication schemes UTL_SMTP supports, in order of their relative security strength. The subset of the schemes in ALL_SCHEMES (namely, PLAIN and LOGIN) in which cleartext passwords are sent over SMTP must be used only in SMTP connections that are secured by Secure Socket Layer / Transport Layer Security (SSL/TLS).
NON_CLEARTEXT_PASSWORD_SCHEMES	VARCHAR2(256)	'CRAM-MD5'	List of authentication schemes that UTL_SMTP supports and in which no cleartext passwords are sent over SMTP. They can be used in SMTP connections that are not secured by SSL/TLS. Note that these schemes may still be weak when used in an insecure SMTP connection.

298.4 UTL_SMTP Types

UTL_SMTP uses a CONNECTION record type and REPLY_REPLIES record types.

CONNECTION Record Type

This is a PL/SQL record type used to represent an SMTP connection.

Syntax

```
TYPE connection IS RECORD (
    host              VARCHAR2(255),
    port              PLS_INTEGER,
    tx_timeout        PLS_INTEGER,
    private_tcp_con   utl_tcp.connection,
    private_state     PLS_INTEGER);
```

Fields

Table 298-2 CONNECTION Record Type Fields

Field	Description
`host`	Name of the remote host when connection is established. `NULL` when no connection is established.
`port`	Port number of the remote SMTP server connected. `NULL` when no connection is established.
`tx_timeout`	Time in seconds that the `UTL_SMTP` package waits before timing out in a read or write operation in this connection. In read operations, this package times out if no data is available for reading immediately. In write operations, this package times out if the output buffer is full and no data is to be sent into the network without being blocked. `0` indicates not to wait at all. `NULL` indicates to wait forever.
`private_tcp_con`	Private, for implementation use only. You should not modify this field.
`private_state`	Private, for implementation use only. You should not modify this field.

Usage Notes

The read-only fields in a connection record are used to return information about the SMTP connection after the connection is successfully made with the OPEN_CONNECTION Functions. Changing the values of these fields has no effect on the connection. The fields `private_tcp_con` and `private_state` for implementation use only. You should not modify these fields.

REPLY_REPLIES Record Types

These are PL/SQL record types used to represent an SMTP reply line. Each SMTP reply line consists of a reply code followed by a text message. While a single reply line is expected for most SMTP commands, some SMTP commands expect multiple reply lines. For those situations, a PL/SQL table of reply records is used to represent multiple reply lines.

Syntax

```
TYPE reply IS RECORD (
  code    PLS_INTEGER,
  text    VARCHAR2(508));

TYPE replies IS TABLE OF reply INDEX BY BINARY_INTEGER;
```

Fields

Table 298-3 REPLY, REPLIES Record Type Fields

Field	Description
code	3-digit reply code
text	Text message of the reply

298.5 UTL_SMTP Reply Codes

SMTP servers send reply codes that indicate message or server status.

The following is a list of the SMTP reply codes.

Table 298-4 SMTP Reply Codes

Reply Code	Meaning
211	System status, or system help reply
214	Help message [Information on how to use the receiver or the meaning of a particular non-standard command; this reply is useful only to the human user]
220	<domain> Service ready
221	<domain> Service closing transmission channel
250	Requested mail action okay, completed
251	User not local; forwards to <forward-path>
252	OK, pending messages for node <node> started. Cannot VRFY user (for example, info is not local), but takes message for this user and attempts delivery.
253	OK, <messages> pending messages for node <node> started
354	Start mail input; end with <CRLF.CRLF>
355	Octet-offset is the transaction offset
421	<domain> Service not available, closing transmission channel (This can be a reply to any command if the service knows it must shut down.)
450	Requested mail action not taken: mailbox unavailable [for example, mailbox busy]
451	Requested action terminated: local error in processing
452	Requested action not taken: insufficient system storage
453	You have no mail.
454	TLS not available due to temporary reason. Encryption required for requested authentication mechanism.
458	Unable to queue messages for node <node>
459	Node <node> not allowed: reason
500	Syntax error, command unrecognized (This may include errors such as command line too long.)
501	Syntax error in parameters or arguments
502	Command not implemented

Table 298-4 (Cont.) SMTP Reply Codes

Reply Code	Meaning
503	Bad sequence of commands
504	Command parameter not implemented
521	<Machine> does not accept mail.
530	Must issue a STARTTLS command first. Encryption required for requested authentication mechanism.
534	Authentication mechanism is too weak.
538	Encryption required for requested authentication mechanism.
550	Requested action not taken: mailbox unavailable [for, mailbox not found, no access]
551	User not local; please try <forward-path>
552	Requested mail action terminated: exceeded storage allocation
553	Requested action not taken: mailbox name not allowed [for example, mailbox syntax incorrect]
554	Transaction failed

298.6 UTL_SMTP Operational Notes

An SMTP connection is initiated by a call to OPEN_CONNECTION Functions which returns a SMTP connection.

After a connection is established, the following calls are required to send a mail:

HELO Function and Procedure - identify the domain of the sender

MAIL Function and Procedure- start a mail, specify the sender

RCPT Function - specify the recipient

OPEN_DATA Function and Procedure- start the mail body

WRITE_RAW_DATA Procedure - write the mail body (multiple calls allowed)

CLOSE_DATA Function and Procedure - close the mail body and send the mail

The SMTP connection is closed by calling the QUIT Function and Procedure.

Related Topics

- OPEN_CONNECTION Functions
  These functions open a connection to an SMTP server.
- QUIT Function and Procedure
  This subprogram terminates an SMTP session and disconnects from the server.

298.7 UTL_SMTP Exceptions

This table lists the exceptions that can be raised by the interface of the UTL_SMTP package.

The network error is transferred to a reply code of 421- service not available.

Table 298-5 UTL_SMTP Exceptions

Exception	Description
INVALID_OPERATION	Raised when an invalid operation is made. In other words, calling API other than the WRITE_DATA Procedure, the WRITE_RAW_DATA Procedure or the CLOSE_DATA Function and Procedure after the OPEN_DATA Function and Procedure is called, or calling WRITE_DATA, WRITE_RAW_DATA or CLOSE_DATA without first calling OPEN_DATA.
TRANSIENT_ERROR	Raised when receiving a reply code in 400 range
PERMANENT_ERROR	Raised when receiving a reply code in 500 range

298.8 UTL_SMTP Rules and Limits

The API imposes no imitation or range-checking. However, you must be aware of the limitations on various elements of SMTP. Sending data that exceed these limits may result in errors returned by the server.

The following table describes the size limitations encountered by the UTL_SMTP subprograms.

Table 298-6 SMTP Size Limitation

Element	Size Limitation
user	Maximum total length of a user name is 64 characters
domain	Maximum total length of a domain name or number is 64 characters
path	Maximum total length of a reverse-path or forward-path is 256 characters (including the punctuation and element separators)
command line	Maximum total length of a command line including the command word and the <CRLF> is 512 characters
reply line	Maximum total length of a reply line including the reply code and the <CRLF> is 512 characters
text line	Maximum total length of a text line including the <CRLF> is 1000 characters (but not counting the leading dot duplicated for transparency)
recipients buffer	Maximum total number of recipients that must be buffered is 100 recipients

298.9 UTL_SMTP Examples

This example illustrates how UTL_SMTP is used by an application to send e-mail. The application connects to an SMTP server at port 25 and sends a simple text message.

```
DECLARE
  c UTL_SMTP.CONNECTION;

  PROCEDURE send_header(name IN VARCHAR2, header IN VARCHAR2) AS
  BEGIN
    UTL_SMTP.WRITE_DATA(c, name || ': ' || header || UTL_TCP.CRLF);
  END;
```

```
BEGIN
  c := UTL_SMTP.OPEN_CONNECTION('smtp-server.acme.com');
  UTL_SMTP.HELO(c, 'foo.com');
  UTL_SMTP.MAIL(c, 'sender@foo.com');
  UTL_SMTP.RCPT(c, 'recipient@foo.com');
  UTL_SMTP.OPEN_DATA(c);
  send_header('From',    '"Sender" <sender@foo.com>');
  send_header('To',      '"Recipient" <recipient@foo.com>');
  send_header('Subject', 'Hello');
  UTL_SMTP.WRITE_DATA(c, UTL_TCP.CRLF || 'Hello, world!');
  UTL_SMTP.CLOSE_DATA(c);
  UTL_SMTP.QUIT(c);
EXCEPTION
  WHEN utl_smtp.transient_error OR utl_smtp.permanent_error THEN
    BEGIN
      UTL_SMTP.QUIT(c);
    EXCEPTION
      WHEN UTL_SMTP.TRANSIENT_ERROR OR UTL_SMTP.PERMANENT_ERROR THEN
        NULL; -- When the SMTP server is down or unavailable, we don't have
              -- a connection to the server. The QUIT call raises an
              -- exception that we can ignore.
    END;
    raise_application_error(-20000,
      'Failed to send mail due to the following error: ' || sqlerrm);
END;
```

298.10 Summary of UTL_SMTP Subprograms

This table lists the `UTL_SMTP` subprograms and briefly describes them.

Table 298-7 UTL_SMTP Package Subprograms

Subprogram	Description
AUTH Function and Procedure	Sends the `AUTH` command to authenticate to the SMTP server
CLOSE_CONNECTION Procedure	Closes the SMTP connection, causing the current SMTP operation to terminate
CLOSE_DATA Function and Procedure	Closes the data session
COMMAND Function and Procedure	Performs a generic SMTP command
COMMAND_REPLIES Function	Performs a generic SMTP command and retrieves multiple reply lines
DATA Function and Procedure	Sends the e-mail body
EHLO Function and Procedure	Performs the initial handshake with SMTP server using the `EHLO` command
HELO Function and Procedure	Performs the initial handshake with SMTP server using the `HELO` command
HELP Function	Sends `HELP` command
MAIL Function and Procedure	Initiates an e-mail transaction with the server, the destination is a mailbox
NOOP Function and Procedure	`NULL` command
OPEN_CONNECTION Functions	Opens a connection to an SMTP server

Table 298-7 (Cont.) UTL_SMTP Package Subprograms

Subprogram	Description
OPEN_DATA Function and Procedure	Sends the DATA command
QUIT Function and Procedure	Terminates an SMTP session and disconnects from the server
RCPT Function	Specifies the recipient of an e-mail message
RSET Function and Procedure	Terminates the current e-mail transaction
STARTTLS Function and Procedure	Sends STARTTLS command to secure the SMTP connection using SSL/TLS
VRFY Function	Verifies the validity of a destination e-mail address
WRITE_DATA Procedure	Writes a portion of the e-mail message
WRITE_RAW_DATA Procedure	Writes a portion of the e-mail message with RAW data

298.10.1 AUTH Function and Procedure

This subprogram sends the AUTH command to authenticate to the SMTP server. The UTL_SMTP package goes through the user's choices of authentication schemes, skips any that is not supported by the SMTP server and uses the first supported.

To determine the schemes the SMTP server supports from its EHLO reply, the user must call the EHLO Function and Procedure. Otherwise, UTL_SMTP uses the first scheme in the list.

Syntax

```
UTL_SMTP.AUTH (
   c          IN OUT NOCOPY connection,
   username   IN            VARCHAR2,
   password   IN            VARCHAR2,
   schemes    IN            VARCHAR2 DEFAULT NON_CLEARTEXT_PASSWORD_SCHEMES)
 RETURN reply;

UTL_SMTP.AUTH (
   c          IN OUT NOCOPY connection,
   username   IN            VARCHAR2,
   password   IN            VARCHAR2,
   schemes    IN            VARCHAR2 DEFAULT NON_CLEARTEXT_PASSWORD_SCHEMES);
```

Parameters

Table 298-8 AUTH Function and Procedure Parameters

Parameter	Description
c	SMTP connection
username	Username
password	Password
schemes	Space-separated list of authentication schemes UTL_SMTP is allowed to use in the preferred order. See the ALL_SCHEMES and NON_CLEARTEXT_PASSWORD_SCHEMES constants for suggestions.

Return Values

Table 298-9 AUTH Function and Procedure Function Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

- Currently only PLAIN, LOGIN and CRAM-MD5 authentication schemes are supported by UTL_SMTP.
- Since the SMTP server may change the authentication schemes it supports after the SMTP connection is secured by SSL/TLS after the STARTTLS command (for example, adding PLAIN and LOGIN), the caller must call the EHLO Function and Procedure again for UTL_SMTP to update the list after the STARTTLS Function and Procedure is called.

Examples

```
DECLARE
  c utl_smtp.connection;
BEGIN
  c := utl_smtp.open_connection(
      host => 'smtp.example.com',
      port => 25,
      wallet_path => 'file:/oracle/wallets/smtp_wallet',
      wallet_password => 'password',
      secure_connection_before_smtp => FALSE);
  UTL_SMTP.STARTTLS(c);
  UTL_SMTP.AUTH(
      c => c,
      username => 'scott',
      password => 'password'
      schemes  => utl_smtp.all_schemes);
END;
```

298.10.2 CLOSE_CONNECTION Procedure

This procedure closes the SMTP connection, causing the current SMTP operation to terminate. Use this procedure only to cancel an e-mail in the middle of the data session.

To end the SMTP connection properly, use the QUIT Function and Procedure.

Syntax

```
UTL_SMTP.CLOSE_CONNECTION (
   c     IN OUT NOCOPY connection);
```

Parameters

Table 298-10 CLOSE_CONNECTION Procedure Parameters

Parameter	Description
c	SMTP connection

298.10.3 CLOSE_DATA Function and Procedure

This subprogram ends the e-mail message by sending the sequence <CR><LF>.<CR><LF> (a single period at the beginning of a line).

Syntax

```
UTL_SMTP.CLOSE_DATA (
   c     IN OUT NOCOPY connection)
RETURN reply;

UTL_SMTP.CLOSE_DATA (
   c     IN OUT NOCOPY connection);
```

Parameters

Table 298-11 CLOSE_DATA Function and Procedure Parameters

Parameter	Description
c	SMTP connection

Return Values

Table 298-12 CLOSE_DATA Function and Procedure Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

The calls to OPEN_DATA, WRITE_DATA, WRITE_RAW_DATA and CLOSE_DATA must be made in the right order. A program calls OPEN_DATA to send the DATA command to the SMTP server. After that, it can call WRITE_DATA or WRITE_RAW_DATA repeatedly to send the actual data. The data is terminated by calling CLOSE_DATA. After OPEN_DATA is called, the only subprograms that can be called are WRITE_DATA, WRITE_RAW_DATA, or CLOSE_DATA. A call to other subprograms results in an INVALID_OPERATION exception being raised.

CLOSE_DATA must be called only after OPEN_CONNECTION, HELO or EHLO, MAIL, and RCPT have been called. The connection to the SMTP server must be open and a mail transaction must be active when this routine is called.

Note that there is no function form of WRITE_DATA because the SMTP server does not respond until the data-terminator is sent during the call to CLOSE_DATA.

298.10.4 COMMAND Function and Procedure

This subprogram performs a generic SMTP command.

Syntax

```
UTL_SMTP.COMMAND (
   c     IN OUT NOCOPY    connection,
```

```
   cmd   IN              VARCHAR2,
   arg   IN              VARCHAR2 DEFAULT NULL)
RETURN reply;

UTL_SMTP.COMMAND (
   c     IN OUT NOCOPY    connection,
   cmd   IN              VARCHAR2,
   arg   IN              VARCHAR2 DEFAULT NULL);
```

Parameters

Table 298-13 COMMAND Function and Procedure Parameters

Parameter	Description
`c`	SMTP connection
`cmd`	SMTP command to send to the server
`arg`	Optional argument to the SMTP argument. A space is inserted between `cmd` and `arg`.

Return Values

Table 298-14 COMMAND Function and Procedure Return Values

Return Value	Description
`reply`	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

This function is used to invoke generic SMTP commands. Use `COMMAND` if only a single reply line is expected. Use `COMMAND_REPLIES` if multiple reply lines are expected.

For `COMMAND`, if multiple reply lines are returned from the SMTP server, it returns the last reply line only.

298.10.5 COMMAND_REPLIES Function

This function performs a generic SMTP command and retrieves multiple reply lines.

Syntax

```
UTL_SMTP.COMMAND_REPLIES (
   c     IN OUT NOCOPY    connection,
   cmd   IN              VARCHAR2,
   arg   IN              VARCHAR2 DEFAULT NULL)
RETURN replies;
```

Parameters

Table 298-15 COMMAND_REPLIES Function Parameters

Parameter	Description
c	SMTP connection
cmd	SMTP command to send to the server
arg	Optional argument to the SMTP argument. A space is inserted between cmd and arg.

Return Values

Table 298-16 COMMAND_REPLIES Function Return Values

Return Value	Description
replies	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types)

Usage Notes

This function is used to invoke generic SMTP commands. Use COMMAND if only a single reply line is expected. Use COMMAND_REPLIES if multiple reply lines are expected.

For COMMAND, if multiple reply lines are returned from the SMTP server, it returns the last reply line only.

298.10.6 DATA Function and Procedure

This subprogram specifies the body of an e-mail message.

Syntax

```
UTL_SMTP.DATA (
   c     IN OUT NOCOPY connection
   body  IN  VARCHAR2 CHARACTER SET ANY_CS)
RETURN reply;

UTL_SMTP.DATA (
   c     IN OUT NOCOPY connection
   body  IN VARCHAR2 CHARACTER SET ANY_CS);
```

Parameters

Table 298-17 DATA Function and Procedure Parameters

Parameter	Description
c	SMTP Connection
body	Text of the message to be sent, including headers, in [RFC822] format

Return Values

Table 298-18 DATA Function and Procedure Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

The application must ensure that the contents of the body parameter conform to the MIME(RFC822) specification. The DATA routine terminates the message with a <CR><LF>.<CR><LF> sequence (a single period at the beginning of a line), as required by RFC821. It also translates any sequence of <CR><LF>.<CR><LF> (single period) in body to <CR><LF>..<CR><LF> (double period). This conversion provides the transparency as described in Section 4.5.2 of RFC821.

The DATA subprogram must be called only after OPEN_CONNECTION, HELO or EHLO, MAIL and RCPT have been called. The connection to the SMTP server must be open, and a mail transaction must be active when this routine is called.

The expected response from the server is a message beginning with status code 250. The 354 response received from the initial DATA command is not returned to the caller.

298.10.7 EHLO Function and Procedure

This subprogram performs the initial handshake with SMTP server using the EHLO command.

Syntax

```
UTL_SMTP.EHLO (
   c       IN OUT NOCOPY connection,
   domain  IN)
RETURN replies;

UTL_SMTP.EHLO (
   c       IN OUT NOCOPY connection,
   domain  IN);
```

Parameters

Table 298-19 EHLO Function and Procedure Parameters

Parameter	Description
c	SMTP connection
domain	Domain name of the local (sending) host. Used for identification purposes.

Return Values

Table 298-20 EHLO Function and Procedure Return Values

Return Value	Description
replies	Reply of the command (see REPLY_ REPLIES Record Types in UTI_SMTP Types).

Usage Notes

The EHLO interface is identical to HELO except that it allows the server to return more descriptive information about its configuration. [RFC1869] specifies the format of the information returned, which the PL/SQL application can retrieve using the functional form of this call. For compatibility with HELO, each line of text returned by the server begins with status code 250.

Related Functions

HELO Function and Procedure

298.10.8 HELO Function and Procedure

This subprogram performs the initial handshake with SMTP server using the HELO command.

Syntax

```
UTL_SMTP.HELO (
   c       IN OUT NOCOPY   connection,
   domain  IN              VARCHAR2)
RETURN reply;

UTL_SMTP.HELO (
   c       IN OUT NOCOPY   connection,
   domain  IN              VARCHAR2);
```

Parameters

Table 298-21 HELO Function and Procedure Parameters

Parameter	Description
c	SMTP connection
domain	Domain name of the local (sending) host. Used for identification purposes.

Return Values

Table 298-22 HELO Function and Procedure Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTI_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

RFC 821 specifies that the client must identify itself to the server after connecting. This routine performs that identification. The connection must have been opened through a call to OPEN_CONNECTION Functions before calling this routine.

The expected response from the server is a message beginning with status code 250.

Related Functions

EHLO Function and Procedure

298.10.9 HELP Function

This function sends the `HELP` command.

Syntax

```
UTL_SMTP.HELP (
   c         IN OUT NOCOPY   connection,
   command   IN              VARCHAR2 DEFAULT NULL)
RETURN replies;
```

Parameters

Table 298-23 HELP Function Parameters

Parameter	Description
`c`	SMTP connection
`command`	Command to get the help message

Return Values

Table 298-24 HELP Function Return Values

Return Value	Description
`replies`	Reply of the command (see REPLY_ REPLIES Record Types in UTI_SMTP Types)

298.10.10 MAIL Function and Procedure

This subprogram initiate a mail transaction with the server. The destination is a mailbox.

Syntax

```
UTL_SMTP.MAIL (
   c           IN OUT NOCOPY   connection,
   sender      IN              VARCHAR2,
   parameters  IN              VARCHAR2 DEFAULT NULL)
RETURN reply;

UTL_SMTP.MAIL (
   c           IN OUT NOCOPY   connection,
```

```
   sender      IN              VARCHAR2,
   parameters  IN              VARCHAR2 DEFAULT NULL);
```

Parameters

Table 298-25 MAIL Function and Procedure Parameters

Parameter	Description
c	SMTP connection
sender	E-mail address of the user sending the message.
parameters	Additional parameters to mail command as defined in Section 6 of [RFC1869]. It must follow the format of "XXX=XXX (XXX=XXX ....)".

Return Values

Table 298-26 MAIL Function and Procedure Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

This command does not send the message; it simply begins its preparation. It must be followed by calls to RCPT and DATA to complete the transaction. The connection to the SMTP server must be open and a HELO or EHLO command must have already been sent.

The expected response from the server is a message beginning with status code 250.

298.10.11 NOOP Function and Procedure

This subprogram issues the NULL command.

Syntax

```
UTL_SMTP.NOOP (
   c  IN OUT NOCOPY connection)
RETURN reply;

UTL_SMTP.NOOP (
   c  IN OUT NOCOPY connection);
```

Parameter

Table 298-27 NOOP Function and Procedure Parameters

Parameter	Description
c	SMTP connection

Return Values

Table 298-28 NOOP Function and Procedure Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

- This command has no effect except to elicit a successful reply from the server. It can be issued at any time after the connection to the server has been established with OPEN_CONNECTION. The NOOP command can be used to verify that the server is still connected and is listening properly.
- This command replies with a single line beginning with status code 250.

298.10.12 OPEN_CONNECTION Functions

These functions open a connection to an SMTP server.

Syntax

```
UTL_SMTP.OPEN_CONNECTION (
   host                           IN  VARCHAR2,
   port                           IN  PLS_INTEGER DEFAULT 25,
   c                              OUT connection,
   tx_timeout                     IN  PLS_INTEGER DEFAULT NULL,
   wallet_path                    IN  VARCHAR2 DEFAULT NULL,
   wallet_password                IN  VARCHAR2 DEFAULT NULL,
   secure_connection_before_smtp  IN  BOOLEAN DEFAULT FALSE,
   secure_host                    IN  VARCHAR2 DEFAULT NULL)
 RETURN reply;

UTL_SMTP.OPEN_CONNECTION (
   host                           IN  VARCHAR2,
   port                           IN  PLS_INTEGER DEFAULT 25,
   tx_timeout                     IN  PLS_INTEGER DEFAULT NULL,
   wallet_path                    IN  VARCHAR2 DEFAULT NULL,
   wallet_password                IN  VARCHAR2 DEFAULT NULL,
   secure_connection_before_smtp  IN  BOOLEAN DEFAULT FALSE,
   secure_host                    IN  VARCHAR2 DEFAULT NULL)
RETURN connection;
```

Parameters

Table 298-29 OPEN_CONNECTION Functions Parameters

Parameter	Description
host	Name of the SMTP server host
port	Port number on which SMTP server is listening (usually 25)
c	SMTP connection

Table 298-29 (Cont.) OPEN_CONNECTION Functions Parameters

Parameter	Description
tx_timeout	Time in seconds that the UTL_SMTP package waits before timing out in a read or write operation for this connection. In read operations, this package times out if no data is available for reading immediately. In write operations, this package times out if the output buffer is full and no data is to be sent into the network without being blocked. 0 indicates not to wait at all. NULL indicates to wait forever.
wallet_path	Directory path that contains the Oracle wallet for SSL/TLS. The format is file: <directory-path> If you want to use the operating system certificate store to act in place of the Oracle wallet, then set the path parameter to system: (include the colon). Doing so greatly improves performance in the database.
wallet_password	Password to open the wallet. When the wallet is auto-login enabled, the password can be set to NULL. If you set path to system:, then omit the password by setting it to NULL.
secure_connection_before_smtp	If TRUE, a secure connection with SSL/TLS is made before SMTP communication. If FALSE, no connection is made.
secure_host	The host name to be matched against the common name (CN) of the SMTP server's certificate when a secure connection is used. It can also be a domain name like "*.example.com". If NULL, the SMTP host name to connect to will be used.

Return Values

Table 298-30 OPEN_CONNECTION Functions Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

- The expected response from the server is a message beginning with status code 220.
- The version of OPEN_CONNECTION that returns UTL_SMTP.CONNECTION record checks the reply code returned by an SMTP server when the connection is first established. It raises an exception when the reply indicates an error. Otherwise, it discards the reply. If you want to examine the reply, invoke the version of OPEN_CONNECTION that returns REPLY.
- tx_timeout is intended to govern both the read operations and the write operations. However, an implementation restriction prevents tx_timeout from governing write operations in the current release.

Examples

```
DECLARE
  c utl_smtp.connection;
BEGIN
  c := UTL_SMTP.OPEN_CONNECTION(
```

```
   host => 'smtp.example.com',
   port => 465,
   wallet_path => 'file:/oracle/wallets/smtp_wallet',
   wallet_password => 'password',
   secure_connection_before_smtp => TRUE);
END;
```

298.10.13 OPEN_DATA Function and Procedure

This subprogram sends the `DATA` command after which you can use `WRITE_DATA` and `WRITE_RAW_DATA` to write a portion of the e-mail message.

Syntax

```
UTL_SMTP.OPEN_DATA (
   c     IN OUT NOCOPY connection)
RETURN reply;

UTL_SMTP.OPEN_DATA (
   c     IN OUT NOCOPY connection);
```

Parameters

Table 298-31 OPEN_DATA Function and Procedure Parameters

Parameter	Description
`c`	SMTP connection
`data`	Portion of the text of the message to be sent, including headers, in RFC822 format.

Return Values

Table 298-32 OPEN_DATA Function and Procedure Function Return Values

Return Value	Description
`reply`	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

- The calls to `OPEN_DATA`, `WRITE_DATA`, `WRITE_RAW_DATA` and `CLOSE_DATA` must be made in the right order. A program calls `OPEN_DATA` to send the `DATA` command to the SMTP server. After that, it can call `WRITE_DATA` or `WRITE_RAW_DATA` repeatedly to send the actual data. The data is terminated by calling `CLOSE_DATA`. After `OPEN_DATA` is called, the only subprograms that can be called are `WRITE_DATA`, `WRITE_RAW_DATA`, or `CLOSE_DATA`. A call to other subprograms results in an `INVALID_OPERATION` exception being raised.
- `OPEN_DATA` must be called only after `OPEN_CONNECTION`, `HELO` or `EHLO`, `MAIL`, and `RCPT` have been called. The connection to the SMTP server must be open and a mail transaction must be active when this routine is called.

298.10.14 QUIT Function and Procedure

This subprogram terminates an SMTP session and disconnects from the server.

Syntax

```
UTL_SMTP.QUIT (
   c  IN OUT NOCOPY connection)
RETURN reply;

UTL_SMTP.QUIT (
   c  IN OUT NOCOPY connection);
```

Parameter

Table 298-33 QUIT Function and Procedure Parameters

Parameter	Description
c	SMTP connection

Return Values

Table 298-34 QUIT Function and Procedure Function Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

The QUIT command informs the SMTP server of the client's intent to terminate the session. It then closes the connection established by OPEN_CONNECTION which must have been called before executing this command. If a mail transaction is in progress when QUIT is issued, it is canceled in the same manner as RSET.

The function form of this command returns a single line beginning with the status code 221 on successful termination. In all cases, the connection to the SMTP server is closed. The fields REMOTE_HOST and REMOTE_PORT of c are reset.

Related Functions

RSET Function and Procedure

298.10.15 RCPT Function

This subprogram specifies the recipient of an e-mail message.

Syntax

```
UTL_SMTP.RCPT (
   c           IN OUT NOCOPY     connection,
   recipient   IN                VARCHAR2,
   parameters  IN                VARCHAR2 DEFAULT NULL)
```

```
RETURN reply;

UTL_SMTP.RCPT (
   c           IN OUT NOCOPY     connection,
   recipient   IN                VARCHAR2,
   parameters  IN                VARCHAR2 DEFAULT NULL);
```

Table 298-35 RCPT Function and Procedure Parameters

Parameter	Description
c	SMTP connection
recipient	E-mail address of the user to which the message is being sent
parameters	Additional parameters to RCPT command as defined in Section 6 of [RFC1869]. It must follow the format of "XXX=XXX (XXX=XXX ....)".

Return Values

Table 298-36 RCPT Function and Procedure Function Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTI_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

To send a message to multiple recipients, call this routine multiple times. Each invocation schedules delivery to a single e-mail address. The message transaction must have been begun by a prior call to `MAIL`, and the connection to the mail server must have been opened and initialized by prior calls to `OPEN_CONNECTION` and `HELO` or `EHLO` respectively.

The expected response from the server is a message beginning with status code 250 or 251.

298.10.16 RSET Function and Procedure

This subprogram terminates the current mail transaction.

Syntax

```
UTL_SMTP.RSET (
   c  IN OUT NOCOPY connection)
RETURN reply;

UTL_SMTP.RSET (
   c  IN OUT NOCOPY connection);
```

Parameters

Table 298-37 RSET Function and Procedure Parameters

Parameter	Description
c	SMTP connection

Return Values

Table 298-38 RSET Function and Procedure Return Values

Return Value	Description
`reply`	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

- This command allows the client to cancel an e-mail message it was in the process of composing. No mail is sent. The client can call `RSET` at any time after the connection to the SMTP server has been opened by means of `OPEN_CONNECTION` until `DATA` or `OPEN_DATA` is called. Once the e-mail data has been sent, it is too late to prevent the e-mail from being sent.
- The server responds to `RSET` with a message beginning with status code 250.

Related Functions

QUIT Function and Procedure

298.10.17 STARTTLS Function and Procedure

This subprogram sends the `STARTTLS` command to secure the SMTP connection using SSL/TLS.

SSL/TLS requires an Oracle wallet which must be specified when the connection was opened by the OPEN_CONNECTION Functions.

Syntax

```
UTL_SMTP.STARTTLS (
   c            IN OUT NOCOPY  connection,
   secure_host  IN     VARCHAR2 DEFAULT NULL)
RETURN reply;

UTL_SMTP.STARTTLS (
   c            IN OUT NOCOPY connection,
   secure_host  IN     VARCHAR2 DEFAULT NULL);
```

Parameters

Table 298-39 STARTTLS Function and Procedure Parameters

Parameter	Description
`c`	SMTP connection
`secure_host`	The host name to be matched against the common name (CN) of the SMTP server's certificate. It can also be a domain name like "`*.example.com`". If `NULL`, the SMTP host name to connect to will be used.

Return Values

Table 298-40 STARTTLS Function and Procedure Return Values

Return Value	Description
reply	SMTP reply

Usage Notes

The STARTTLS command must only be issued on an unencrypted connection and when the SMTP server indicates the support of the command in the reply of the EHLO command. The wallet to be used for encryption must have been specified when the initial SMTP connection was opened by the OPEN_CONNECTION function.

Examples

```
DECLARE
  c utl_smtp.connection;
BEGIN
  c := utl_smtp.open_connection(
     host => 'smtp.example.com',
     port => 25,
     wallet_path => 'file:/oracle/wallets/smtp_wallet',
     wallet_password => 'password',
     secure_connection_before_smtp => FALSE);
  utl_smtp.starttls(c);
END
```

298.10.18 VRFY Function

This function verifies the validity of a destination e-mail address.

Syntax

```
UTL_SMTP.VRFY (
   c          IN OUT NOCOPY connection
   recipient  IN VARCHAR2)
RETURN reply;
```

Parameters

Table 298-41 VRFY Function Parameters

Parameter	Description
c	SMTP connection
recipient	E-mail address to be verified

Return Values

Table 298-42 VRFY Function Return Values

Return Value	Description
reply	Reply of the command (see REPLY_ REPLIES Record Types in UTl_SMTP Types). In cases where there are multiple replies, the last reply is returned.

Usage Notes

The server attempts to resolve the destination address `recipient`. If successful, it returns the recipient's full name and fully qualified mailbox path. The connection to the server must have already been established by means of `OPEN_CONNECTION` and `HELO` or `EHLO` before making this request.

Successful verification returns one or more lines beginning with status code 250 or 251.

298.10.19 WRITE_DATA Procedure

This procedure writes a portion of the e-mail message. A repeat call to `WRITE_DATA` appends data to the e-mail message.

Syntax

```
UTL_SMTP.WRITE_DATA (
   c     IN OUT NOCOPY connection,
   data  IN VARCHAR2 CHARACTER SET ANY_CS);
```

Parameters

Table 298-43 WRITE_DATA Procedure Parameters

Parameter	Description
c	SMTP connection
data	Portion of the text of the message to be sent, including headers, in [RFC822] format

Usage Notes

- The calls to the OPEN_DATA Function and Procedure, WRITE_DATA Procedure, WRITE_RAW_DATA Procedure and CLOSE_DATA Function and Procedure must be made in the correct order. A program calls `OPEN_DATA` to send the `DATA` command to the SMTP server. After that, it can call `WRITE_DATA` or `WRITE_RAW_DATA` repeatedly to send the actual data. The data is terminated by calling `CLOSE_DATA`. After `OPEN_DATA` is called, the only subprograms that can be called are `WRITE_DATA`, `WRITE_RAW_DATA`, or `CLOSE_DATA`. A call to other subprograms results in an `INVALID_OPERATION` exception being raised.
- The application must ensure that the contents of the body parameter conform to the MIME(RFC822) specification. The `DATA` routine terminates the message with a `<CR><LF>.<CR><LF>` sequence (a single period at the beginning of a line), as required by RFC821. It also translates any sequence of `<CR><LF>.<CR><LF>` (single period) in the body

to <CR><LF>..<CR><LF> (double period). This conversion provides the transparency as described in Section 4.5.2 of RFC821.

- The OPEN_DATA Function and Procedure, WRITE_DATA Procedure, WRITE_RAW_DATA Procedure and CLOSE_DATA Function and Procedure must be called only after OPEN_CONNECTION Functions, HELO Function and Procedure, or EHLO Function and Procedure, MAIL Function and Procedure, and RCPT Function have been called. The connection to the SMTP server must be open and a mail transaction must be active when this routine is called.
- Note that there is no function form of the WRITE_DATA Procedure because the SMTP server does not respond until the data-terminator is sent during the call to CLOSE_DATA Function and Procedure.
- Text (VARCHAR2) data sent using WRITE_DATA is converted to US7ASCII before it is sent. If the text contains multibyte characters, each multibyte character in the text that cannot be converted to US7ASCII is replaced by a '?' character. If 8BITMIME extension is negotiated with the SMTP server using the EHLO subprogram, multibyte VARCHAR2 data can be sent by first converting the text to RAW using the UTL_RAW package, and then sending the RAW data using WRITE_RAW_DATA.

298.10.20 WRITE_RAW_DATA Procedure

This procedure writes a portion of the e-mail message. A repeat call to WRITE_RAW_DATA appends data to the e-mail message.

Syntax

```
UTL_SMTP.WRITE_RAW_DATA (
   c     IN OUT NOCOPY connection
   data  IN RAW);
```

Parameters

Table 298-44 WRITE_RAW_DATA Procedure Parameters

Parameter	Description
c	SMTP connection
data	Portion of the text of the message to be sent, including headers, in [RFC822] format

Usage Notes

- The calls to the OPEN_DATA Function and Procedure, WRITE_DATA Procedure, WRITE_RAW_DATA Procedure and CLOSE_DATA Function and Procedure must be made in the correct order. A program calls OPEN_DATA to send the DATA command to the SMTP server. After that, it can call WRITE_DATA or WRITE_RAW_DATA repeatedly to send the actual data. The data is terminated by calling CLOSE_DATA. After OPEN_DATA is called, the only subprograms that can be called are WRITE_DATA, WRITE_RAW_DATA, or CLOSE_DATA. A call to other subprograms results in an INVALID_OPERATION exception being raised.
- The application must ensure that the contents of the body parameter conform to the MIME(RFC822) specification. The DATA routine terminates the message with a <CR><LF>.<CR><LF> sequence (a single period at the beginning of a line), as required by RFC821. It also translates any sequence of <CR><LF>.<CR><LF> (single period) in the body

to <CR><LF>..<CR><LF> (double period). This conversion provides the transparency as described in Section 4.5.2 of RFC821.

- The OPEN_DATA Function and Procedure, WRITE_DATA Procedure, WRITE_RAW_DATA Procedure and CLOSE_DATA Function and Procedure must be called only after OPEN_CONNECTION Functions, HELO Function and Procedure, or EHLO Function and Procedure, MAIL Function and Procedure, and RCPT Function have been called. The connection to the SMTP server must be open and a mail transaction must be active when this routine is called.
- Note that there is no function form of the WRITE_DATA Procedure because the SMTP server does not respond until the data-terminator is sent during the call to CLOSE_DATA Function and Procedure.

299
UTL_TCP

With the `UTL_TCP` package and its procedures and functions, PL/SQL applications can communicate with external TCP/IP-based servers using TCP/IP. Because many Internet application protocols are based on TCP/IP, this package is useful to PL/SQL applications that use Internet protocols and e-mail.

This chapter contains the following topics:

- Overview
- Security Model
- Types
- Exceptions
- Rules and Limits
- Examples
- Summary of UTL_TCP Subprograms

299.1 UTL_TCP Overview

The `UTL_TCP` package provides TCP/IP client-side access functionality in PL/SQL.

299.2 UTL_TCP Security Model

This package is an invoker's rights package and the invoking user needs the connect privilege granted in the access control list assigned to the remote network host to which he wants to connect.

> **Note:**
>
> For more information about managing fine-grained access, see *Oracle Database Security Guide*

299.3 UTL_TCP Types

The `UTL_TCP` package includes a CONNECTION type and a carriage-return line-feed (CRLF) type.

CONNECTION Type

This is a PL/SQL record type used to represent a TCP/IP connection.

Syntax

```
TYPE connection IS RECORD (
    remote_host    VARCHAR2(255),
    remote_port    PLS_INTEGER,
    local_host     VARCHAR2(255),
    local_port     PLS_INTEGER,
    charset        VARCHAR2(30),
    newline        VARCHAR2(2),
    tx_timeout     PLS_INTEGER,
    private_sd     PLS_INTEGER);
```

Fields

Table 299-1 Connection Record Type Fields

Field	Description
`remote_host`	Name of the remote host when connection is established. `NULL` when no connection is established.
`remote_port`	Port number of the remote host connected. `NULL` when no connection is established.
`local_host`	Name of the local host used to establish the connection. `NULL` when no connection is established.
`local_port`	Port number of the local host used to establish the connection. `NULL` when no connection is established.
`charset`	The on-the-wire character set. Since text messages in the database may be encoded in a character set that is different from the one expected on the wire (that is, the character set specified by the communication protocol, or the one stipulated by the other end of the communication), text messages in the database are converted to and from the on-the-wire character set as they are sent and received on the network.
`newline`	Newline character sequence. This newline character sequence is appended to the text line sent by `WRITE_LINE` API.
`tx_timeout`	Time in seconds that the UTL_TCP package waits before giving up in a read or write operation in this connection. In read operations, this package gives up if no data is available for reading immediately. In write operations, this package gives up if the output buffer is full and no data is to be sent in the network without being blocked. Zero (0) indicates not to wait at all. `NULL` indicates to wait forever.

Usage Notes

The fields in a connection record are used to return information about the connection, which is often made using `OPEN_CONNECTION`. Changing the values of those fields has no effect on the connection. The fields `private_XXXX` are for implementation use only. You should not modify the values.

In the current release of the `UTL_TCP` package, the parameters `local_host` and `local_port` are ignored when `open_connection` makes a TCP/IP connection. It does not attempt to use the specified local host and port number when the connection is made. The `local_host` and `local_port` fields are not set in the connection record returned by the function.

Time out on write operations is not supported in the current release of the `UTL_TCP` package.

CRLF

The character sequence carriage-return line-feed. It is the newline sequence commonly used by many communication standards.

Syntax

```
CRLF CONSTANT VARCHAR2(2 CHAR);
```

Usage Notes

This package variable defines the newline character sequence commonly used in many Internet protocols. This is the default value of the newline character sequence for `WRITE_LINE`, specified when a connection is opened. While such protocols use `<CR><LF>` to denote a new line, some implementations may choose to use just line-feed to denote a new line. In such cases, users can specify a different newline character sequence when a connection is opened.

299.4 UTL_TCP Exceptions

`UTL_TCP` will raise an exception when it encounters a processing issue.

The exceptions raised by the TCP/IP package are listed in the following table.

Table 299-2 ***TCP/IP Exceptions***

Exception	Description
`BUFFER_TOO_SMALL`	Buffer is too small for input that requires look-ahead
`END_OF_INPUT`	Raised when no more data is available to read from the connection
`NETWORK_ERROR`	Generic network error
`BAD_ARGUMENT`	Bad argument passed in an API call (for example, a negative buffer size)
`TRANSFER_TIMEOUT`	No data is read and a read time out occurred
`PARTIAL_MULTIBYTE_CHAR`	No complete character is read and a partial multibyte character is found at the end of the input

299.5 UTL_TCP Rules and Limits

The interface provided in the package only allows connections to be initiated by the PL/SQL program. It does not allow the PL/SQL program to accept connections initiated outside the program.

299.6 UTL_TCP Examples

Some possible uses for `UTL_TCP` include retrieving a Web page over HTTP or sending an e-mail.

The following code example illustrates how the TCP/IP package can be used to retrieve a Web page over HTTP. It connects to a Web server listening at port 80 (standard port for HTTP) and requests the root document.

```
DECLARE
  c  utl_tcp.connection;  -- TCP/IP connection to the Web server
```

```
  ret_val pls_integer;
BEGIN
  c := utl_tcp.open_connection(remote_host => 'www.acme.com',
                               remote_port =>  80,
                               charset     => 'US7ASCII');  -- open connection
  ret_val := utl_tcp.write_line(c, 'GET / HTTP/1.0');    -- send HTTP request
  ret_val := utl_tcp.write_line(c);
  BEGIN
    LOOP
      dbms_output.put_line(utl_tcp.get_line(c, TRUE));  -- read result
    END LOOP;
  EXCEPTION
    WHEN utl_tcp.end_of_input THEN
      NULL; -- end of input
  END;
  utl_tcp.close_connection(c);
END;
```

The following code example illustrates how the TCP/IP package can be used by an application to send e-mail (also known as email from PL/SQL). The application connects to an SMTP server at port 25 and sends a simple text message.

```
PROCEDURE send_mail (sender    IN VARCHAR2,
                     recipient IN VARCHAR2,
                   message   IN VARCHAR2) IS
    mailhost   VARCHAR2(30) := 'mailhost.mydomain.com';
    smtp_error  EXCEPTION;
    mail_conn   utl_tcp.connection;
    PROCEDURE smtp_command(command IN VARCHAR2,
                           ok      IN VARCHAR2 DEFAULT '250')
    IS
        response varchar2(3);
        len pls_integer;
    BEGIN
        len := utl_tcp.write_line(mail_conn, command);
        response := substr(utl_tcp.get_line(mail_conn), 1, 3);
        IF (response <> ok) THEN
            RAISE smtp_error;
        END IF;
    END;

BEGIN
    mail_conn := utl_tcp.open_connection(remote_host => mailhost,
                                         remote_port => 25,
                                         charset     => 'US7ASCII');
    smtp_command('HELO ' || mailhost);
    smtp_command('MAIL FROM: ' || sender);
    smtp_command('RCPT TO: ' || recipient);
    smtp_command('DATA', '354');
    smtp_command(message);
    smtp_command('QUIT', '221');
    utl_tcp.close_connection(mail_conn);
EXCEPTION
    WHEN OTHERS THEN
        -- Handle the error
END;
```

299.7 Summary of UTL_TCP Subprograms

This table lists the `UTL_TCP` subprograms and briefly describes them.

Table 299-3 UTL_TCP Package Subprograms

Subprogram	Description
AVAILABLE Function	Determines the number of bytes available for reading from a TCP/IP connection
CLOSE_ALL_CONNECTIONS Procedure	Closes all open TCP/IP connections
CLOSE_CONNECTION Procedure	Closes an open TCP/IP connection
FLUSH Procedure	Transmits immediately to the server all data in the output buffer, if a buffer is used
GET_LINE Function	Returns the line of data read
GET_LINE_NCHAR Function	Returns the line of data read in NCHAR form
GET_RAW Function	Return the data read instead of the amount of data read
GET_TEXT Function	Returns the text data read
GET_TEXT_NCHAR Function	Returns the text data read in NCHAR form
OPEN_CONNECTION Function	Opens a TCP/IP connection to a specified service
READ_LINE Function	Receives a text line from a service on an open connection
READ_RAW Function	Receives binary data from a service on an open connection
READ_TEXT Function	Receives text data from a service on an open connection
SECURE_CONNECTION Procedure	Secures a TCP/IP connection using SSL/TLS
WRITE_LINE Function	Transmits a text line to a service on an open connection
WRITE_RAW Function	Transmits a binary message to a service on an open connection
WRITE_TEXT Function	Transmits a text message to a service on an open connection

299.7.1 AVAILABLE Function

This function determines the number of bytes available for reading from a TCP/IP connection. It is the number of bytes that can be read immediately without blocking. Determines if data is ready to be read from the connection.

Syntax

```
UTL_TCP.AVAILABLE (
   c         IN OUT NOCOPY connection,
   timeout   IN PLS_INTEGER DEFAULT 0)
RETURN PLS_INTEGER;
```

Parameters

Table 299-4 AVAILABLE Function Parameters

Parameter	Description
c	TCP connection to determine the amount of data that is available to be read

Table 299-4 (Cont.) AVAILABLE Function Parameters

Parameter	Description
timeout	Time in seconds to wait before giving up and reporting that no data is available. Zero (0) indicates not to wait at all. NULL indicates to wait forever.

Return Values

The number of bytes available for reading without blocking

Usage Notes

The connection must have already been opened through a call to OPEN_CONNECTION. Users may use this API to determine if data is available to be read before calling the read API so that the program are not blocked because data is not ready to be read from the input.

The number of bytes available for reading returned by this function may be less than what is actually available. On some platforms, this function may only return 1, to indicate that some data is available. If you are concerned about the portability of your application, then assume that this function returns a positive value when data is available for reading, and 0 when no data is available. This function returns a positive value when all the data at a particular connection has been read and the next read result in the END_OF_INPUT exception.

The following example illustrates using this function in a portable manner:

```
DECLARE
   c   utl_tcp.connection
   data VARCHAR2(256);
   len  PLS_INTEGER;
BEGIN
   c := utl_tcp.open_connection(...);
   LOOP
      IF (utl_tcp.available(c) > 0) THEN
         len := utl_tcp.read_text(c, data, 256);
      ELSE
         ---do some other things
        . . . .
      END IF
   END LOOP;
END;
```

299.7.2 CLOSE_ALL_CONNECTIONS Procedure

This procedure closes all open TCP/IP connections.

Syntax

```
UTL_TCP.CLOSE_ALL_CONNECTIONS;
```

Usage Notes

This call is provided to close all connections before a PL/SQL program ends to avoid dangling connections.

299.7.3 CLOSE_CONNECTION Procedure

This procedure closes an open TCP/IP connection.

Syntax

```
UTL_TCP.CLOSE_CONNECTION (
   c IN OUT NOCOPY connection);
```

Parameters

Table 299-5 CLOSE_CONNECTION Procedure Parameters

Parameter	Description
c	TCP connection to close

Usage Notes

Connection must have been opened by a previous call to `OPEN_CONNECTION`. The fields `remote_host`, `remote_port`, `local_host`, `local_port` and `charset` of `c` are reset after the connection is closed.

An open connection must be closed explicitly. An open connection remains open when the PL/SQL record variable that stores the connection goes out-of-scope in the PL/SQL program. Failing to close unwanted connections may result in unnecessary tying up of local and remote system resources.

299.7.4 FLUSH Procedure

This procedure transfers immediately to the server all data in the output buffer, if a buffer is used.

Syntax

```
UTL_TCP.FLUSH (
   c IN OUT NOCOPY connection);
```

Parameters

Table 299-6 FLUSH Procedure Parameters

Parameter	Description
c	TCP connection to which to send data

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`.

299.7.5 GET_LINE Function

This function returns the line of data read.

Syntax

```
UTL_TCP.GET_LINE (
   c           IN OUT NOCOPY connection,
   remove_crlf IN            BOOLEAN DEFAULT FALSE,
   peek        IN            BOOLEAN DEFAULT FALSE)
 RETURN VARCHAR2;
```

Parameters

Table 299-7 GET_LINE Function Parameters

Parameter	Description
c	TCP connection from which to receive data
remove_crlf	If TRUE, then one ore more trailing CRLF characters are removed from the received message.
peek	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to TRUE and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.

Return Values

The text line read

Usage Notes

- The connection must have already been opened through a call to OPEN_CONNECTION.
- See READ_LINE for the read time out, character set conversion, buffer size, and multibyte character issues.

299.7.6 GET_LINE_NCHAR Function

This function returns the line of data read in NCHAR form.

Syntax

```
UTL_TCP.GET_LINE_NCHAR (
   c           IN OUT NOCOPY connection,
   remove_crlf IN            BOOLEAN DEFAULT FALSE,
   peek        IN            BOOLEAN DEFAULT FALSE)
 RETURN NVARCHAR2;
```

Parameters

Table 299-8 GET_LINE_NCHAR Function Parameters

Parameter	Description
c	TCP connection from which to receive data
remove_crlf	If TRUE, then one ore more trailing CRLF characters are removed from the received message.
peek	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to TRUE and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.

Return Values

The text line read

Usage Notes

- The connection must have already been opened through a call to OPEN_CONNECTION.
- See READ_LINE for the read time out, character set conversion, buffer size, and multibyte character issues.

299.7.7 GET_RAW Function

This function returns the data read instead of the amount of data read.

Syntax

```
UTL_TCP.GET_RAW (
   c     IN OUT NOCOPY connection,
   len   IN            PLS_INTEGER DEFAULT 1,
   peek  IN            BOOLEAN     DEFAULT FALSE)
 RETURN RAW;
```

Parameters

Table 299-9 GET_RAW Function Parameters

Parameter	Description
c	TCP connection from which to receive data
len	The number of bytes (or characters for VARCHAR2) of data to receive. Default is 1.

Table 299-9 (Cont.) GET_RAW Function Parameters

Parameter	Description
peek	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to TRUE and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.
remove_crlf	If TRUE, then one ore more trailing CRLF characters are removed from the received message.

Return Values

The binary data read

Usage Notes

The connection must have already been opened through a call to OPEN_CONNECTION.

For all the get_* APIs described in this section, see the corresponding READ_* API for the read time out issue. For GET_TEXT and GET_LINE, see the corresponding READ_* API for character set conversion, buffer size, and multibyte character issues.

299.7.8 GET_TEXT Function

This function returns the text data read.

Syntax

```
UTL_TCP.GET_TEXT (
   c    IN OUT NOCOPY connection,
   len  IN            PLS_INTEGER DEFAULT 1,
   peek IN            BOOLEAN     DEFAULT FALSE)
 RETURN VARCHAR2;
```

Parameters

Table 299-10 GET_TEXT Function Parameters

Parameter	Description
c	TCP connection from which to receive data
len	Number of bytes (or characters for VARCHAR2) of data to receive. Default is 1.
peek	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to TRUE and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.

Table 299-10 (Cont.) GET_TEXT Function Parameters

Parameter	Description
remove_crlf	If TRUE, then one ore more trailing CRLF characters are removed from the received message.

Return Values

The text data read

Usage Notes

The connection must have already been opened through a call to OPEN_CONNECTION.

For all the get_* APIs described in this section, see the corresponding read_* API for the read time out issue. For GET_TEXT and GET_LINE, see the corresponding READ_* API for character set conversion, buffer size, and multibyte character issues.

299.7.9 GET_TEXT_NCHAR Function

This function returns the text data read in NCHAR form.

Syntax

```
UTL_TCP.GET_TEXT_NCHAR (
   c    IN OUT NOCOPY connection,
   len  IN            PLS_INTEGER DEFAULT 1,
   peek IN            BOOLEAN     DEFAULT FALSE)
 RETURN NVARCHAR2;
```

Parameters

Table 299-11 GET_TEXT_NCHAR Function Parameters

Parameter	Description
c	TCP connection from which to receive data
len	The number of bytes (or characters for VARCHAR2) of data to receive. Default is 1.
peek	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to TRUE and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.
remove_crlf	If TRUE, then one ore more trailing CRLF characters are removed from the received message.

Return Values

The text data read

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`.

For all the `get_*` APIs described in this section, see the corresponding `read_*` API for the read time out issue. For `GET_TEXT` and `GET_LINE`, see the corresponding `READ_*` API for character set conversion, buffer size, and multibyte character issues.

299.7.10 OPEN_CONNECTION Function

This function opens a TCP/IP connection to a specified service.

Syntax

```
UTL_TCP.OPEN_CONNECTION  (
   remote_host          IN VARCHAR2,
   remote_port          IN PLS_INTEGER,
   local_host           IN VARCHAR2 DEFAULT NULL,
   local_port           IN PLS_INTEGER DEFAULT NULL,
   in_buffer_size       IN PLS_INTEGER DEFAULT NULL,
   out_buffer_size      IN PLS_INTEGER DEFAULT NULL,
   charset              IN VARCHAR2 DEFAULT NULL,
   newline              IN VARCHAR2 DEFAULT CRLF,
   tx_timeout           IN PLS_INTEGER DEFAULT NULL,
   wallet_path          IN  VARCHAR2 DEFAULT NULL,
   wallet_password      IN  VARCHAR2 DEFAULT NULL,
  RETURN connection;
```

Parameters

Table 299-12 OPEN_CONNECTION Function Parameters

Parameter	Description
`remote_host`	Name of the host providing the service. When `remote_host` is `NULL`, it connects to the local host.
`remote_port`	Port number on which the service is listening for connections
`local_host`	Name of the host providing the service. `NULL` means does not care.
`local_port`	Port number on which the service is listening for connections. `NULL` means don't care.
`in_buffer_size`	The size of input buffer. The use of an input buffer can speed up execution performance in receiving data from the server. The appropriate size of the buffer depends on the flow of data between the client and the server, and the traffic/latency on the network. A zero value means no buffer should be used. A `NULL` value means the caller does not care if a buffer is used or not. The maximum size of the input buffer is 32767 bytes.
`out_buffer_size`	The size of output buffer. The use of an output buffer can speed up execution performance in sending data to the server. The appropriate size of buffer depends on the flow of data between the client and the server, and the network condition. A zero value means no buffer should be used. A `NULL` value means the caller does not care if a buffer is used or not. The maximum size of the output buffer is 32767 bytes.

Table 299-12 (Cont.) OPEN_CONNECTION Function Parameters

Parameter	Description
`charset`	The on-the-wire character set. Since text messages in the database may be encoded in a character set that is different from the one expected on the wire (that is, the character set specified by the communication protocol, or the one stipulated by the other end of the communication), text messages in the database are converted to and from the on-the-wire character set as they are sent and received on the network using `READ_TEXT`, `READ_LINE`, `WRITE_TEXT` and `WRITE_LINE`. Set this parameter to `NULL` when no conversion is needed.
`newline`	Newline character sequence. This newline character sequence is appended to the text line sent by `WRITE_LINE` API.
`tx_timeout`	Time in seconds that the `UTL_TCP` package should wait before giving up in a read or write operations in this connection. In read operations, this package gives up if no data is available for reading immediately. In write operations, this package gives up if the output buffer is full and no data is to be sent in the network without being blocked. Zero (0) indicates not to wait at all. `NULL` indicates to wait forever.
`wallet_path`	Directory path that contains the Oracle wallet for SSL/TLS. The format is `file:`*`directory-path`*. If you want to use the operating system certificate store to act in place of the Oracle wallet, then set the `path` parameter to `system:` (include the colon). Doing so greatly improves performance in the database.
`wallet_password`	Password to open the wallet. When the wallet is auto-login enabled, the password may be set to `NULL`. If you set `path` to `system:`, then omit the password by setting it to `NULL`.

Return Values

A connection to the targeted TCP/IP service

Usage Notes

- Note that connections opened by this `UTL_TCP` package can remain open and be passed from one database call to another in a shared server configuration. However, the connection must be closed explicitly. The connection remains open when the PL/SQL record variable that stores the connection goes out-of-scope in the PL/SQL program. Failing to close unwanted connections may result in unnecessary tying up of local and remote system resources.
- In the current release of the `UTL_TCP` package, the parameters `local_host` and `local_port` are ignored when `open_connection` makes a TCP/IP connection. It does not attempt to use the specified local host and port number when the connection is made. The `local_host` and `local_port` fields is not set in the connection record returned by the function.
- `tx_timeout` is intended to govern both the read operations and the write operations. However, an implementation restriction prevents `tx_timeout` from governing write operations in the current release.

Examples

```
DECLARE
  c UTL_TCP.CONNECTION;
BEGIN
  c := UTL_TCP.OPEN_CONNECTION(
   host           => 'www.example.com',
   port           => 443,
   wallet_path    => 'file:/oracle/wallets/smtp_wallet',
   wallet_password => '****');
  UTL_TCP.SECURE_CONNECTION (c => c);
END;
```

299.7.11 READ_LINE Function

This function receives a text line from a service on an open connection.

A line is terminated by a line-feed, a carriage-return or a carriage-return followed by a line-feed.

Syntax

```
UTL_TCP.READ_LINE (
   c           IN OUT NOCOPY connection,
   data        IN OUT NOCOPY VARCHAR2 CHARACTER SET ANY_CS,
   peek        IN            BOOLEAN DEFAULT FALSE)
 RETURN PLS_INTEGER;
```

Parameters

Table 299-13 READ_LINE Function Parameters

Parameter	Description
`c`	TCP connection from which to receive data
`data`	Data received.
`remove_crlf`	If `TRUE`, then one ore more trailing CRLF characters are removed from the received message.
`peek`	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to `TRUE` and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.

Return Values

The number of characters of data received

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`. This function does not return until the end-of-line have been reached, or the end of input has been reached. Text messages is converted from the on-the-wire character set, specified when the connection was opened, to the database character set before they are returned to the caller.

If transfer time out is set when the connection is opened, then this function waits for each data packet to be ready to read until time out occurs. If it occurs, then this function stops reading and returns all the data read successfully. If no data is read successfully, then the `transfer_timeout` exception is raised. The exception can be handled and the read operation can be retried later.

If a partial multibyte character is found at the end of input, then this function stops reading and returns all the complete multibyte characters read successfully. If no complete character is read successfully, then the `partial_multibyte_char` exception is raised. The exception can be handled and the bytes of that partial multibyte character can be read as binary by the `READ_RAW` function. If a partial multibyte character is seen in the middle of the input because the remaining bytes of the character have not arrived and read time out occurs, then the `transfer_timeout` exception is raised instead. The exception can be handled and the read operation can be retried later.

299.7.12 READ_RAW Function

This function receives binary data from a service on an open connection.

Syntax

```
UTL_TCP.READ_RAW (
   c     IN OUT NOCOPY connection,
   data  IN OUT NOCOPY RAW,
   len   IN            PLS_INTEGER DEFAULT 1,
   peek  IN            BOOLEAN     DEFAULT FALSE)
 RETURN PLS_INTEGER;
```

Parameters

Table 299-14 READ_RAW Function Parameters

Parameter	Description
`c`	TCP connection from which to receive data
`data (IN OUT COPY)`	Data received
`len`	Number of bytes of data to receive
`peek`	Normally, you want to read the data and remove it from the input queue, that is, consume it. In some situations, you may just want to look ahead at the data, that is, peek at it, without removing it from the input queue, so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to `TRUE` and set up an input buffer before the connection is opened. The amount of data you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.

Return Values

The number of bytes of data received

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`. This function does not return until the specified number of bytes have been read, or the end of input has been reached.

If transfer time out is set when the connection is opened, then this function waits for each data packet to be ready to read until time out occurs. If it occurs, then this function stops reading and returns all the data read successfully. If no data is read successfully, then the `transfer_timeout` exception is raised. The exception can be handled and the read operation can be retried later.

299.7.13 READ_TEXT Function

This function receives text data from a service on an open connection.

Syntax

```
UTL_TCP.READ_TEXT (
   c    IN OUT NOCOPY connection,
   data IN OUT NOCOPY VARCHAR2 CHARACTER SET ANY_CS,
   len  IN            PLS_INTEGER DEFAULT 1,
   peek IN            BOOLEAN     DEFAULT FALSE)
RETURN PLS_INTEGER;
```

Parameters

Table 299-15 READ_TEXT Function Parameters

Parameter	Description
`c`	TCP connection from which to receive data
`data`	Data received
`len`	Number of characters of data to receive
`peek`	Normally, users want to read the data and remove it from the input queue, that is, consume it. In some situations, users may just want to look ahead at the data without removing it from the input queue so that it is still available for reading (or even peeking) in the next call. To keep the data in the input queue, set this flag to `TRUE` and an input buffer must be set up when the connection is opened. The amount of data that you can peek at (that is, read but keep in the input queue) must be less than the size of input buffer.

Return Values

The number of characters of data received

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`. This function does not return until the specified number of characters has been read, or the end of input has been reached. Text messages is converted from the on-the-wire character set, specified when the connection was opened, to the database character set before they are returned to the caller.

Unless explicitly overridden, the size of a `VARCHAR2` buffer is specified in terms of bytes, while the parameter `len` refers to the maximum number of characters to be read. When the database character set is multibyte, where a single character may consist of more than 1 byte, you should ensure that the buffer can hold the maximum of characters. In general, the size of the `VARCHAR2` buffer should equal the number of characters to be read, multiplied by the maximum number of bytes of a character of the database character set.

If transfer time out is set when the connection is opened, then this function waits for each data packet to be ready to read until time out occurs. If it occurs, then this function stops reading and returns all the data read successfully. If no data is read successfully, then the `transfer_timeout` exception is raised. The exception can be handled and the read operation can be retried later.

If a partial multibyte character is found at the end of input, then this function stops reading and returns all the complete multibyte characters read successfully. If no complete character is read successfully, then the `partial_multibyte_char` exception is raised. The exception can be handled and the bytes of that partial multibyte character can be read as binary by the `READ_RAW` function. If a partial multibyte character is seen in the middle of the input because the remaining bytes of the character have not arrived and read time out occurs, then the `transfer_timeout` exception is raised instead. The exception can be handled and the read operation can be retried later.

299.7.14 SECURE_CONNECTION Procedure

This procedure secures a TCP/IP connection using SSL/TLS.

SSL/TLS requires an Oracle wallet which must be specified when the connection was opened by the OPEN_CONNECTION Function.

Syntax

```
UTL_TCP.SECURE_CONNECTION (
   c    IN OUT NOCOPY connection);
```

Parameters

Table 299-16 SECURE_CONNECTION Procedure Parameters

Parameter	Description
`c`	TCP connection from which to receive data

299.7.15 WRITE_LINE Function

This function transmits a text line to a service on an open connection. The `newline` character sequence is appended to the message before it is transmitted.

Syntax

```
UTL_TCP.WRITE_LINE (
   c    IN OUT NOCOPY connection,
   data IN            VARCHAR2 DEFAULT NULL CHARACTER SET ANY_CS)
  RETURN PLS_INTEGER;
```

Parameters

Table 299-17 WRITE_LINE Function Parameters

Parameter	Description
`c`	TCP connection to which to send data
`data`	Buffer containing the data to be sent

Return Values

The actual number of characters of data transmitted

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`. Text messages are converted to the on-the-wire character set, specified when the connection was opened, before they are transmitted on the wire.

299.7.16 WRITE_RAW Function

This function transmits a binary message to a service on an open connection. The function does not return until the specified number of bytes have been written.

Syntax

```
UTL_TCP.WRITE_RAW (
   c    IN OUT NOCOPY connection,
   data IN            RAW,
   len  IN            PLS_INTEGER DEFAULT NULL)
 RETURN PLS_INTEGER;
```

Parameters

Table 299-18 WRITE_RAW Function Parameters

Parameter	Description
`c`	TCP connection to which to send data
`data`	Buffer containing the data to be sent
`len`	The number of bytes of data to transmit. When `len` is `NULL`, the whole length of data is written.

Return Values

The number of bytes of data transmitted

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`.

299.7.17 WRITE_TEXT Function

This function transmits a text message to a service on an open connection.

Syntax

```
UTL_TCP.WRITE_TEXT (
   c    IN OUT NOCOPY connection,
   data IN            VARCHAR2 CHARACTER SET ANY_CS,
   len  IN            PLS_INTEGER DEFAULT NULL)
 RETURN num_chars PLS_INTEGER;
```

Parameters

Table 299-19 WRITE_TEXT Function Parameters

Parameter	Description
`c`	TCP connection to which to send data
`data`	Buffer containing the data to be sent
`len`	The number of characters of data to transmit. When `len` is `NULL`, the whole length of data is written. The actual amount of data written may be less because of network condition.

Return Values

The actual number of characters of data transmitted

Usage Notes

The connection must have already been opened through a call to `OPEN_CONNECTION`. Text messages are converted to the on-the-wire character set, specified when the connection was opened, before they are transmitted on the wire.

300
UTL_URL

The `UTL_URL` package has two functions: `ESCAPE` and `UNESCAPE`.

This chapter contains the following topics:

- Overview
- Exceptions
- Examples
- Summary of UTL_URL Subprograms

> **See Also:**
>
> UTL_HTTP

300.1 UTL_URL Overview

A Uniform Resource Locator (URL) is a string that identifies a Web resource, such as a page or a picture. Use a URL to access such resources by way of the HyperText Transfer Protocol (HTTP).

For example, the URL for Oracle's Web site is:

```
http://www.oracle.com
```

Normally, a URL contains English alphabetic characters, digits, and punctuation symbols. These characters are known as the *unreserved characters*. Any other characters in URLs, including multibyte characters or binary octet codes, must be escaped to be accurately processed by Web browsers or Web servers. Some punctuation characters, such as dollar sign (`$`), question mark (`?`), colon (`:`), and equals sign (`=`), are reserved as delimiters in a URL. They are known as the *reserved characters*. To literally process these characters, instead of treating them as delimiters, they must be escaped.

The unreserved characters are:

- `A` through `Z`, `a` through `z`, and `0` through `9`
- Hyphen (-), underscore (_), period (.), exclamation point (!), tilde (~), asterisk (*), accent (`'`), left parenthesis (( ), right parenthesis ( ))

The reserved characters are:

- Semi-colon (;) slash (/), question mark (?), colon (:), at sign (@), ampersand (&), equals sign (=), plus sign (+), dollar sign ($), percentage sign (%), and comma (,)

The `UTL_URL` package has two functions that provide escape and unescape mechanisms for URL characters. Use the escape function to escape a URL before the URL is used fetch a Web page by way of the `UTL_HTTP` package. Use the unescape function to unescape an escaped URL before information is extracted from the URL.

For more information, refer to the Request For Comments (RFC) document RFC2396. Note that this URL escape and unescape mechanism is different from the `x-www-form-urlencoded` encoding mechanism described in the HTML specification:

`http://www.w3.org/TR/html`

300.2 UTL_URL Exceptions

`UTL_URL` raises an exception when it encounter a processing issue.

The following table lists the exceptions that can be raised when the `UTL_URL` package API is invoked.

Table 300-1 UTL_URL Exceptions

Exception	Error Code	Reason
BAD_URL	29262	The URL contains badly formed escape code sequences
BAD_FIXED_WIDTH_CHARSET	29274	Fixed-width multibyte character set is not allowed as a URL character set.

300.3 UTL_URL Examples

`UTL_URL` can be used for encoding and decoding.

You can implement the `x-www-form-urlencoded` encoding using the `UTL_URL.ESCAPE` function as follows:

```
CREATE OR REPLACE FUNCTION form_url_encode (
   data    IN VARCHAR2,
   charset IN VARCHAR2) RETURN VARCHAR2 AS
BEGIN
  RETURN utl_url.escape(data, TRUE, charset); -- note use of TRUE
END;
```

For decoding data encoded with the `form-URL-encode scheme`, the following function implements the decording scheme:

```
CREATE OR REPLACE FUNCTION form_url_decode(
   data    IN VARCHAR2,
   charset IN VARCHAR2) RETURN VARCHAR2 AS
BEGIN
  RETURN utl_url.unescape(
     replace(data, '+', ' '),
     charset);
END;
```

300.4 Summary of UTL_URL Subprograms

This table lists and briefly describes the `UTL_URL` subprograms.

Table 300-2 UTL_URL Package Subprograms

Subprogram	Description
ESCAPE Function	Returns a URL with illegal characters (and optionally reserved characters) escaped using the `%2-digit-hex-code` format
UNESCAPE Function	Unescapes the escape character sequences to their original forms in a URL. Convert the `%XX` escape character sequences to the original characters

300.4.1 ESCAPE Function

This function returns a URL with illegal characters (and optionally reserved characters) escaped using the `%2-digit-hex-code` format.

Syntax

```
UTL_URL.ESCAPE (
   url                  IN VARCHAR2 CHARACTER SET ANY_CS,
   escape_reserved_chars IN BOOLEAN  DEFAULT  FALSE,
   url_charset          IN VARCHAR2 DEFAULT  utl_http.body_charset)
 RETURN VARCHAR2;
```

Parameters

Table 300-3 ESCAPE Function Parameters

Parameter	Description
`url`	The original URL
`escape_reserved_chars`	Indicates whether the URL reserved characters should be escaped. If set to `TRUE`, both the reserved and illegal URL characters are escaped. Otherwise, only the illegal URL characters are escaped. The default value is `FALSE`.
`url_charset`	When escaping a character (single-byte or multibyte), determine the target character set that character should be converted to before the character is escaped in %hex-code format. If `url_charset` is `NULL`, the database charset is assumed and no character set conversion will occur. The default value is the current default body character set of the `UTL_HTTP` package, whose default value is `ISO-8859-1`. The character set can be named in Internet Assigned Numbers Authority (IANA) or in the Oracle naming convention.

Usage Notes

Use this function to escape URLs that contain illegal characters as defined in the URL specification RFC 2396. The legal characters in URLs are:

- `A` through `Z`, `a` through `z`, and `0` through `9`
- Hyphen (-), underscore (_), period (.), exclamation point (!), tilde (~), asterisk (*), accent (`'`), left parenthesis (( ), right parenthesis ( ))

The reserved characters consist of:

- Semi-colon (;) slash (/), question mark (?), colon (:), at sign (@), ampersand (&), equals sign (=), plus sign (+), dollar sign ($), and comma (,)

Many of the reserved characters are used as delimiters in the URL. You should escape characters beyond those listed here by using escape_url. Also, to use the reserved characters in the name-value pairs of the query string of a URL, those characters must be escaped separately. An escape_url cannot recognize the need to escape those characters because once inside a URL, those characters become indistinguishable from the actual delimiters. For example, to pass a name-value pair `$logon=HR/<password>` into the query string of a URL, escape the `$` and `/` separately as `%24logon=HR%2F<password>` and use it in the URL.

Normally, you will escape the entire URL, which contains the reserved characters (delimiters) that should not be escaped. For example:

```
utl_url.escape('http://www.acme.com/a url with space.html')
```

Returns:

```
http://www.acme.com/a%20url%20with%20space.html
```

In other situations, you may want to send a query string with a value that contains reserved characters. In that case, escape only the value fully (with `escape_reserved_chars` set to `TRUE`) and then concatenate it with the rest of the URL. For example:

```
url := 'http://www.acme.com/search?check=' || utl_url.escape
('Is the use of the "$" sign okay?', TRUE);
```

This expression escapes the question mark (?), dollar sign ($), and space characters in `'Is the use of the "$" sign okay?'` but not the `?` after `search` in the URL that denotes the use of a query string.

The Web server that you intend to fetch Web pages from may use a character set that is different from that of your database. In that case, specify the url_charset as the Web server character set so that the characters that need to be escaped are escaped in the target character set. For example, a user of an EBCDIC database who wants to access an ASCII Web server should escape the URL using `US7ASCII` so that a space is escaped as `%20` (hex code of a space in ASCII) instead of `%40` (hex code of a space in EBCDIC).

This function does not validate a URL for the proper URL format.

300.4.2 UNESCAPE Function

This function unescapes the escape character sequences to its original form in a URL, to convert the `%XX` escape character sequences to the original characters.

Syntax

```
UTL_URL.UNESCAPE (
   url             IN VARCHAR2 CHARACTER SET ANY_CS,
   url_charset     IN VARCHAR2 DEFAULT utl_http.body_charset)
                   RETURN VARCHAR2;
```

Parameters

Table 300-4 UNESCAPE Function Parameters

Parameter	Description
`url`	The URL to unescape

Table 300-4 (Cont.) UNESCAPE Function Parameters

Parameter	Description
`url_charset`	After a character is unescaped, the character is assumed to be in the `source_charset` character set and it will be converted from the `source_charset` to the database character set before the URL is returned. If `source_charset` is `NULL`, the database charset is assumed and no character set conversion occurred. The default value is the current default body character set of the `UTL_HTTP` package, whose default value is "ISO-8859-1". The character set can be named in Internet Assigned Numbers Authority (IANA) or Oracle naming convention.

Usage Notes

The Web server that you receive the URL from may use a character set that is different from that of your database. In that case, specify the url_charset as the Web server character set so that the characters that need to be unescaped are unescaped in the source character set. For example, a user of an EBCDIC database who receives a URL from an ASCII Web server should unescape the URL using `US7ASCII` so that `%20` is unescaped as a space (0x20 is the hex code of a space in ASCII) instead of a `?` (because 0x20 is not a valid character in EBCDIC).

This function does not validate a URL for the proper URL format.

301
WPG_DOCLOAD

The WPG_DOCLOAD package provides an interface to download files, BLOBs and BFILEs.

> **See Also:**
>
> For more information about implementation of this package:
>
> - *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
> - *Oracle Fusion Middleware User's Guide for mod_plsql*

The chapter contains the following topics:

- Constants
- Summary of WPG_DOCLOAD Subprograms

301.1 WPG_DOCLOAD Constants

WPG_DOCLOAD defines several constants to use when specifying parameter values.

The WPG_DOCLOAD constants are listed below:

- NAME_COL_LEN
- MIMET_COL_LEN
- MAX_DOCTABLE_NAME_LEN

NAME_COL_LEN

The NAME column in your document table must be the same as the value of name_col_len.

```
name_col_len CONSTANT pls_integer := 64;
```

MIMET_COL_LEN

The MIME_TYPE column in your document table must be the same as the value of mimet_col_len.

```
mimet_col_len CONSTANT pls_integer := 48;
```

MAX_DOCTABLE_NAME_LEN

The name length of your document table must be less than max_doctable_name_len.

```
max_doctable_name_len CONSTANT pls_integer := 256;
```

301.2 Summary of WPG_DOCLOAD Subprograms

The WPG_DOCLOAD package uses one subprogram, the DOWNLOAD_FILE procedure.

Table 301-1 WPG_DOCLOAD Package Subprograms

Subprogram	Description
DOWNLOAD_FILE Procedures	Downloads files, BLOBS and BFILES

301.2.1 DOWNLOAD_FILE Procedures

There are three versions of this download file procedure.

- The first version downloads files and is invoked from within a document download procedure to signal the PL/SQL Gateway that p_filename is to be downloaded from the document table to the client's browser.
- The second version can be called from within any procedure to signal the PL/SQL Gateway that p_blob is to be downloaded to the client's browser.
- The third version can be called from within any procedure to signal the PL/SQL Gateway that p_bfile is to be downloaded to the client's browser.

Syntax

```
WPG_DOCLOAD.DOWNLOAD_FILE(
   p_filename      IN             VARCHAR2,
   p_bcaching      IN             BOOLEAN DEFAULT TRUE);


WPG_DOCLOAD.DOWNLOAD_FILE(
   p_blob          IN OUT NOCOPY  BLOB);

WPG_DOCLOAD.DOWNLOAD_FILE(
   p_bfile         IN OUT         BFILE);
```

Parameters

Table 301-2 DOWNLOAD_FILE Procedure Parameters

Parameter	Description
p_filename	The file to download from the document table.
p_blob	The BLOB to download.
p_bfile	The BFILE to download (see Usage Notes).
p_bcaching	Whether browser caching is enabled (see Usage Notes).

Usage Notes

- Normally, a document will be downloaded to the browser unless the browser sends an 'If-Modified-Since' header to the gateway indicating that it has the requested document in its cache. In that case, the gateway will determine if the browser's cached copy is up to date, and if it is, it will send an HTTP 304 status message to the browser indicating that the browser should display the cached copy. However, because a document URL and a document do not necessarily have a one-to-one relationship in the PL/SQL Web Gateway, in some cases it may be undesirable to have the cached copy of a document displayed. In those cases, the p_bcaching parameter should be set to FALSE to indicate to the gateway to ignore the 'If-Modified-Since' header, and download the document.

- `p_bfile` and `p_blob` are declared as `IN OUT` because the locator is initially opened to check for file accessibility and existence. The open operation can only be performed if the locator is writable and readable.

302
ANYDATA TYPE

An ANYDATA TYPE contains an instance of a given type, plus a description of the type. In this sense, an ANYDATA is self-describing. An ANYDATA can be persistently stored in the database.

This chapter contains the following topics:

- Restrictions
- Operational Notes
- Summary of ANYDATA Subprograms

302.1 ANYDATA TYPE Restrictions

Persistent storage of ANYDATA instances whose type contains embedded LOBs other than BFILEs is not currently supported.

302.2 ANYDATA TYPE Operational Notes

This section contains notes related to ANYDATA TYPE construction and access.

Construction

There are 2 ways to construct an ANYDATA. The CONVERT* calls enable construction of the ANYDATA in its entirety with a single call. They serve as explicit CAST functions from any type in the Oracle ORDBMS to ANYDATA.

```
STATIC FUNCTION ConvertBDouble(dbl IN BINARY_DOUBLE) return ANYDATA,
STATIC FUNCTION ConvertBfile(b IN BFILE) RETURN ANYDATA,
STATIC FUNCTION ConvertBFloat(fl IN BINARY_FLOAT) return ANYDATA,
STATIC FUNCTION ConvertBlob(b IN BLOB) RETURN ANYDATA,
STATIC FUNCTION ConvertChar(c IN CHAR) RETURN ANYDATA,
STATIC FUNCTION ConvertClob(c IN CLOB) RETURN ANYDATA,
STATIC FUNCTION ConvertCollection(col IN "collection_type") RETURN ANYDATA,
STATIC FUNCTION ConvertDate(dat IN DATE) RETURN ANYDATA,
STATIC FUNCTION ConvertIntervalDS(inv IN INTERVAL DAY TO SECOND) return ANYDATA,
STATIC FUNCTION ConvertIntervalYM(invIN INTERVAL YEAR TO MONTH) return ANYDATA,
STATIC FUNCTION ConvertNchar(nc IN NCHAR) return ANYDATA,
STATIC FUNCTION ConvertNClob(nc IN NCLOB) return ANYDATA,
STATIC FUNCTION ConvertNumber(num IN NUMBER) RETURN ANYDATA,
STATIC FUNCTION ConvertNVarchar2(nc IN NVARCHAR2) return ANYDATA,
STATIC FUNCTION ConvertObject(obj IN "<object_type>") RETURN ANYDATA,
STATIC FUNCTION ConvertRaw(r IN RAW) RETURN ANYDATA,
STATIC FUNCTION ConvertRef(rf IN REF "<object_type>") RETURN ANYDATA,
STATIC FUNCTION ConvertTimestamp(ts IN TIMESTAMP) return ANYDATA,
STATIC FUNCTION ConvertTimestampTZ(ts IN TIMESTAMP WITH TIMEZONE) return ANYDATA,
STATIC FUNCTION ConvertTimestampLTZ(ts IN TIMESTAMP WITH LOCAL TIMEZONE) return ANYDATA,
STATIC FUNCTION ConvertURowid(rid IN UROWID) return ANYDATA,
STATIC FUNCTION ConvertVarchar(c IN VARCHAR) RETURN ANYDATA,
STATIC FUNCTION ConvertVarchar2(c IN VARCHAR2) RETURN ANYDATA,
```

The second way to construct an ANYDATA is a piece by piece approach. The BEGINCREATE Static Procedure call begins the construction process and ENDCREATE Member Procedure call finishes the construction process. In between these two calls, the individual attributes of an object type or the elements of a collection can be set using SET* calls. For piece by piece access of the attributes of objects and elements of collections, the PIECEWISE Member Procedure should be invoked prior to GET* calls.

Note: The ANYDATA has to be constructed or accessed sequentially starting from its first attribute (or collection element). The BEGINCREATE call automatically begins the construction in a piece-wise mode. There is no need to call PIECEWISE immediately after BEGINCREATE. ENDCREATE should be called to finish the construction process (before which any access calls can be made).

Access

Access functions are available based on SQL. These functions do not throw exceptions on type-mismatch. Instead, they return NULL if the type of the ANYDATA does not correspond to the type of access. If you wish to use only ANYDATA functions of the appropriate types returned in a query, you should use a WHERE clause which uses GETTYPENAME and choose the type you are interested in (say "SYS.NUMBER"). Each of these functions returns the value of a specified datatype inside a SYS.ANYDATA wrapper.

```
MEMBER FUNCTION AccessBDouble(self IN ANYDATA) return BINARY_DOUBLE
   DETERMINISTIC,
MEMBER FUNCTION AccessBfile(self IN ANYDATA) return BFILE,
MEMBER FUNCTION AccessBFloat(self IN ANYDATA) return BINARY_FLOAT
   DETERMINISTIC,
MEMBER FUNCTION AccessBlob(self IN ANYDATA) return BLOB,
MEMBER FUNCTION AccessChar(self IN ANYDATA) return CHAR,
MEMBER FUNCTION AccessClob(self IN ANYDATA) return CLOB,
MEMBER FUNCTION AccessDate(self IN ANYDATA) return DATE,
MEMBER FUNCTION AccessIntervalYM(self IN ANYDATA) return INTERVAL YEAR TO MONTH,
MEMBER FUNCTION AccessIntervalDS(self IN ANYDATA) return INTERVAL DAY TO SECOND,
MEMBER FUNCTION AccessNchar(self IN ANYDATA) return NCHAR,
MEMBER FUNCTION AccessNClob(self IN ANYDATA) return NCLOB
MEMBER FUNCTION AccessNumber(self IN ANYDATA) return NUMBER,
MEMBER FUNCTION AccessNVarchar2(self IN ANYDATA) return NVARCHAR2,
MEMBER FUNCTION AccessRaw(self IN ANYDATA) return RAW,
MEMBER FUNCTION AccessTimestamp(self IN ANYDATA) return TIMESTAMP,
MEMBER FUNCTION AccessTimestampLTZ(self IN ANYDATA) return TIMESTAMP WITH LOCAL
   TIMEZONE,
MEMBER FUNCTION AccessTimestampTZ(self IN ANYDATA) return TIMESTAMP WITH
   TIMEZONE,
MEMBER FUNCTION AccessURowid(self IN ANYDATA) return UROWID DETERMINISTIC
MEMBER FUNCTION AccessVarchar(self IN ANYDATA) return VARCHAR,
MEMBER FUNCTION AccessVarchar2(self IN ANYDATA) return VARCHAR2,
```

302.3 Summary of ANYDATA Subprograms

This table lists the ANYDATA subprograms in alphabetical order and briefly describes them.

Table 302-1 ***ANYDATA Type Subprograms***

Subprogram	Description
BEGINCREATE Static Procedure	Begins creation process on a new ANYDATA

Table 302-1 (Cont.) *ANYDATA Type Subprograms*

Subprogram	Description
ENDCREATE Member Procedure	Ends creation of an `ANYDATA`
GET* Member Functions	Gets the current data value (which should be of appropriate type)
GETTYPE Member Function	Gets the Type of the `ANYDATA`
GETTYPENAME Member Function	Get the fully qualified type name for the `ANYDATA`
PIECEWISE Member Procedure	Sets the `MODE` of access of the current data value to be an attribute at a time (if the data value is of `TYPECODE_OBJECT`)
SET* Member Procedures	Sets the current data value.

302.3.1 BEGINCREATE Static Procedure

This procedure begins the creation process on a new `ANYDATA`.

Syntax

```
STATIC PROCEDURE BeginCreate(
   dtype          IN OUT NOCOPY AnyType,
   adata          OUT NOCOPY ANYDATA);
```

Parameters

Table 302-2 BEGINCREATE Procedure Parameters

Parameter	Description
`dtype`	The type of the `ANYDATA`. (Should correspond to `OCI_TYPECODE_OBJECT` or a Collection typecode.)
`adata`	ANYDATA being constructed.

Exception

`DBMS_TYPES.INVALID_PARAMETERS`: dtype is invalid (not fully constructed, and similar deficits.)

Usage Notes

There is no need to call `PIECEWISE` immediately after this call. The construction process begins in a piece-wise manner automatically.

302.3.2 ENDCREATE Member Procedure

This procedure ends creation of an `ANYDATA`. Other creation functions cannot be called after this call.

Syntax

```
MEMBER PROCEDURE EndCreate(
   self         IN OUT NOCOPY ANYDATA);
```

Parameters

Table 302-3 ENDCREATE Procedure Parameter

Parameter	Description
self	An ANYDATA.

302.3.3 GET* Member Functions

These functions get the current data value (which should be of appropriate type).

The type of the current data value depends on the MODE by which it is accessed (depending on whether the PIECEWISE call is invoked).

If PIECEWISE has NOT been called, the ANYDATA is accessed in its entirety and the type of the data value should match the type of the ANYDATA.

If PIECEWISE has been called, the ANYDATA is accessed piece-wise. The type of the data value should match the type of the attribute (or collection element) at the current position.

Syntax

```
MEMBER FUNCTION GetBDouble(
   self         IN ANYDATA,
   dbl          OUT NOCOPY BINARY_DOUBLE)
RETURN PLS_INTEGER;

MEMBER FUNCTION GetBfile(
   self         IN ANYDATA,
   b            OUT NOCOPY BFILE)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetBFloat(
   self         IN ANYDATA,
   fl           OUT NOCOPY BINARY_FLOAT)
RETURN PLS_INTEGER;

MEMBER FUNCTION GetBlob(
   self         IN ANYDATA,
   b            OUT NOCOPY BLOB)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetChar(
   self         IN ANYDATA,
   c            OUT NOCOPY CHAR)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetClob(
   self         IN ANYDATA,
   c            OUT NOCOPY CLOB)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetCollection(
   self         IN ANYDATA,
   col          OUT NOCOPY "<collection_type>")
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetDate(
```

```
   self         IN ANYDATA,
   dat          OUT NOCOPY DATE)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetIntervalDS(
   self         IN ANYDATA,
   inv          OUT NOCOPY INTERVAL DAY TO SECOND)
 RETURN         PLS_INTEGER;

MEMBER FUNCTION GetIntervalYM(
   self         IN ANYDATA,
   inv          OUT NOCOPY INTERVAL YEAR TO MONTH)
 RETURN         PLS_INTEGER;

MEMBER FUNCTION GetNchar(
   self         IN ANYDATA,
   nc           OUT NOCOPY NCHAR)
 RETURN         PLS_INTEGER;

MEMBER FUNCTION GetNClob(
   self         IN ANYDATA,
   nc           OUT NOCOPY NCLOB)
 RETURN         PLS_INTEGER;

MEMBER FUNCTION GetNumber(
   self         IN ANYDATA,
   num          OUT NOCOPY NUMBER)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetNVarchar2(
   self         IN ANYDATA,
   nc           OUT NOCOPY NVARCHAR2)
 RETURN         PLS_INTEGER;

MEMBER FUNCTION GetObject(
   self         IN ANYDATA,
   obj          OUT NOCOPY "<object_type>")
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetRaw(
   self         IN ANYDATA,
   r            OUT NOCOPY RAW)
   RETURN       PLS_INTEGER;

MMEMBER FUNCTION GetRef(
   self         IN ANYDATA,
   rf           OUT NOCOPY REF "<object_type>")
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetTimestamp(
   self         IN ANYDATA,
   ts           OUT NOCOPY TIMESTAMP)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GetTimestampTZ(
   self         IN ANYDATA,
   ts           OUT NOCOPY TIMESTAMP WITH TIME ZONE)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GetTimestampLTZ(
   self         IN ANYDATA,
   ts           OUT NOCOPY TIMESTAMP WITH LOCAL TIME ZONE)
```

```
 RETURN PLS_INTEGER;

MEMBER FUNCTION GetVarchar(
   self         IN ANYDATA,
   c            OUT NOCOPY VARCHAR)
   RETURN       PLS_INTEGER;

MEMBER FUNCTION GetVarchar2(
   self         IN ANYDATA,
   c            OUT NOCOPY VARCHAR2)
   RETURN       PLS_INTEGER;
```

Parameters

Table 302-4 GET* Function Parameter

Parameter	Description
`self`	An ANYDATA.
`num`	The number to be obtained.

Return Values

`DBMS_TYPES.SUCCESS` or `DBMS_TYPES.NO_DATA`

The return value is relevant only if `PIECEWISE` has been already called (for a collection). In such a case, DBMS_TYPES.NO_DATA signifies the end of the collection when all elements have been accessed.

Exceptions

`DBMS_TYPES.TYPE_MISMATCH`: When the expected type is different from the passed in type.

`DBMS_TYPES.INVALID_PARAMETERS`: Invalid Parameters (if it is not appropriate to add a number at this point in the creation process).

`DBMS_TYPES.INCORRECT_USAGE`: Incorrect usage.

302.3.4 GETTYPE Member Function

This function gets the typecode of the `ANYDATA`.

Syntax

```
MEMBER FUNCTION GETTYPE(
   self          IN ANYDATA,
   typ           OUT NOCOPY AnyType)
   RETURN        PLS_INTEGER;
```

Parameters

Table 302-5 GETTYPE Function Parameter

Parameter	Description
`self`	An `ANYDATA`.

Table 302-5 (Cont.) GETTYPE Function Parameter

Parameter	Description
typ	The AnyType corresponding to the ANYDATA. May be NULL if it does not represent a user-defined type.

Return Values

The typecode corresponding to the type of the ANYDATA.

302.3.5 GETTYPENAME Member Function

This function gets the fully qualified type name for the ANYDATA.

If the ANYDATA is based on a built-in type, this function will return NUMBER and other relevant information.

If it is based on a user defined type, this function will return *schema_name.type_name*, for example, SCOTT.FOO.

If it is based on a transient anonymous type, this function will return NULL.

Syntax

```
MEMBER FUNCTION GETTYPENAME(
   self         IN ANYDATA)
   RETURN       VARCHAR2;
```

Parameters

Table 302-6 GETTYPENAME Function Parameter

Parameter	Description
self	An ANYDATA.

Return Values

Type name of the ANYDATA.

302.3.6 PIECEWISE Member Procedure

This procedure sets the MODE of access of the current data value to be an attribute at a time (if the data value is of TYPECODE_OBJECT).

It sets the MODE of access of the data value to be a collection element at a time (if the data value is of collection type). Once this call has been made, subsequent calls to SET* and GET* will sequentially obtain individual attributes or collection elements.

Syntax

```
MEMBER PROCEDURE PIECEWISE(
   self         IN OUT NOCOPY ANYDATA);
```

Parameters

Table 302-7 PIECEWISE Procedure Parameters

Parameter	Description
self	The current data value.

Exceptions

- DBMS_TYPES.INVALID_PARAMETERS
- DBMS_TYPES.INCORRECT_USAGE: On incorrect usage.

Usage Notes

The current data value must be of an OBJECT or COLLECTION type before this call can be made.

Piece-wise construction and access of nested attributes that are of object or collection types is not supported.

302.3.7 SET* Member Procedures

This procedure sets the current data value.

This is a list of procedures that should be called depending on the type of the current data value. The type of the data value should be the type of the attribute at the current position during the piece-wise construction process.

Syntax

```
MEMBER PROCEDURE SETBDOUBLE(
   self        IN OUT NOCOPY ANYDATA,
   dbl         IN BINARY_DOUBLE,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETBFILE(
   self        IN OUT NOCOPY ANYDATA,
   b           IN BFILE,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETBFLOAT(
   self        IN OUT NOCOPY ANYDATA,
   fl          IN            BINARY_FLOAT,
   last_elem   IN            boolean DEFAULT FALSE);

MEMBER PROCEDURE SETBLOB(
   self        IN OUT NOCOPY ANYDATA,
   b           IN BLOB,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETCHAR(
   self        IN OUT NOCOPY ANYDATA,
   c           IN CHAR,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETCLOB(
   self        IN OUT NOCOPY ANYDATA,
   c           IN CLOB,
```

```
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETCOLLECTION(
   self        IN OUT NOCOPY ANYDATA,
   col         IN "<collectyion_type>",
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETDATE(
   self        IN OUT NOCOPY ANYDATA,
   dat         IN DATE,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETINTERVALDS(
   self        IN OUT NOCOPY ANYDATA,
   inv         IN INTERVAL DAY TO SECOND,
   last_elem IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETINTERVALYM(
   self        IN OUT NOCOPY ANYDATA,
   inv         IN INTERVAL YEAR TO MONTH,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETNCHAR(
   self        IN OUT NOCOPY ANYDATA,
   nc          IN NCHAR,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETNCLOB(
   self        IN OUT NOCOPY ANYDATA,
   nc          IN NClob,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETNUMBER(
   self        IN OUT NOCOPY ANYDATA,
   num         IN NUMBER,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETNVARCHAR2(
   self        IN OUT NOCOPY ANYDATA,
   nc          IN NVarchar2,
   last_elem   IN boolean DEFAULT FALSE),

MEMBER PROCEDURE SETOBJECT(
   self        IN OUT NOCOPY ANYDATA,
   obj         IN "<object_type>",
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETRAW(
   self        IN OUT NOCOPY ANYDATA,
   r           IN RAW,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETREF(
   self        IN OUT NOCOPY ANYDATA,
   rf          IN REF "<object_type>",
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETTIMESTAMP(
   self       IN OUT NOCOPY ANYDATA,
   ts         IN TIMESTAMP,
   last_elem  IN BOOLEAN DEFAULT FALSE);
```

```
MEMBER PROCEDURE SETTIMESTAMPTZ(self IN OUT NOCOPY ANYDATA,
   ts          IN TIMESTAMP WITH TIME ZONE,
   last_elem   IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETTIMESTAMPLTZ(
   self IN OUT NOCOPY ANYDATA,
   ts IN TIMESTAMP WITH LOCAL TIME ZONE,
   last_elem IN boolean DEFAULT FALSE),

MEMBER PROCEDURE SETVARCHAR(
   self        IN OUT NOCOPY ANYDATA,
   c           IN VARCHAR,
   last_elem   IN boolean DEFAULT FALSE);

MEMBER PROCEDURE SETVARCHAR2(
   self        IN OUT NOCOPY ANYDATA,
   c           IN VARCHAR2,
   last_elem   IN boolean DEFAULT FALSE);
```

Parameters

Table 302-8 SET* Procedure Parameters

Parameter	Description
`self`	An `ANYDATA`.
`num`	The number, and associated information, that is to be set.
`last_elem`	Relevant only if `ANYDATA` represents a collection. Set to `TRUE` if it is the last element of the collection, `FALSE` otherwise.

Exceptions

- `DBMS_TYPES.INVALID_PARAMETERS`: Invalid Parameters (if it is not appropriate to add a number at this point in the creation process).
- `DBMS_TYPES.INCORRECT_USAGE`: Incorrect usage.
- `DBMS_TYPES.TYPE_MISMATCH`: When the expected type is different from the passed in type.

Usage Notes

When `BEGINCREATE` is called, construction has already begun in a piece-wise fashion. Subsequent calls to `SET*` will set the successive attribute values.

If the `ANYDATA` is a standalone collection, the `SET*` call will set the successive collection elements.

303
ANYDATASET TYPE

An ANYDATASET TYPE contains a description of a given type plus a set of data instances of that type. An ANYDATASET can be persistently stored in the database if desired, or it can be used as interface parameters to communicate self-descriptive sets of data, all of which belong to a certain type.

This chapter contains the following topics:

- Construction
- Summary of ANYDATASET TYPE Subprograms

303.1 ANYDATASET TYPE Construction

The ANYDATASET needs to be constructed value by value, sequentially.

For each data instance (of the type of the ANYDATASET), the ADDINSTANCE function must be invoked. This adds a new data instance to the ANYDATASET. Subsequently, SET* can be called to set each value in its entirety.

The MODE of construction/access can be changed to attribute/collection element wise by making calls to PIECEWISE.

- If the type of the ANYDATASET is TYPECODE_OBJECT, individual attributes will be set with subsequent SET* calls. Likewise on access.
- If the type of the current data value is a collection type individual collection elements will be set with subsequent SET* calls. Likewise on access. This call is very similar to ANYDATA.PIECEWISE call defined for the type ANYDATA.

Note that there is no support for piece-wise construction and access of nested (not top level) attributes that are of object types or collection types.

ENDCREATE should be called to finish the construction process (before which no access calls can be made).

303.2 Summary of ANYDATASET TYPE Subprograms

This table lists the ANYDATASET Type subprograms and briefly describes them.

Table 303-1 ***ANYDATASET Type Subprograms***

Subprogram	Description
ADDINSTANCE Member Procedure	Adds a new data instance to an ANYDATASET.
BEGINCREATE Static Procedure	Creates a new ANYDATASET which can be used to create a set of data values of the given ANYTYPE.
ENDCREATE Member Procedure	Ends Creation of a ANYDATASET. Other creation functions cannot be called after this call.

Table 303-1 (Cont.) *ANYDATASET Type Subprograms*

Subprogram	Description
GET* Member Functions	Gets the current data value (which should be of appropriate type).
GETCOUNT Member Function	Gets the number of data instances in an `ANYDATASET`.
GETINSTANCE Member Function	Gets the next instance in an `ANYDATASET`.
GETTYPE Member Function	Gets the `ANYTYPE` describing the type of the data instances in an `ANYDATASET`. current data value (which should be of appropriate type).
GETTYPENAME Member Function	Gets the AnyType describing the type of the data instances in an `ANYDATASET`.
PIECEWISE Member Procedure	Sets the `MODE` of construction, access of the data value to be an attribute at a time (if the data value is of `TYPECODE_OBJECT`).
SET* Member Procedures	Sets the current data value.

303.2.1 ADDINSTANCE Member Procedure

This procedure adds a new data instance to an `ANYDATASET`.

Syntax

```
MEMBER PROCEDURE AddInstance(
   self          IN OUT NOCOPY ANYDATASET);
```

Parameters

Table 303-2 ADDINSTANCE Procedure Parameter

Parameter	Description
`self`	The `ANYDATASET` being constructed.

Exceptions

```
DBMS_TYPES.invalid_parameters: Invalid parameters.
DBMS_TYPES.incorrect_usage: On incorrect usage.
```

Usage Notes

The data instances have to be added sequentially. The previous data instance must be fully constructed (or set to NULL) before a new one can be added.

This call DOES NOT automatically set the mode of construction to be piece-wise. The user has to explicitly call `PIECEWISE` if a piece-wise construction of the instance is intended.

303.2.2 BEGINCREATE Static Procedure

This procedure creates a new ANYDATASET which can be used to create a set of data values of the given ANYTYPE.

Syntax

```
STATIC PROCEDURE BeginCreate(
   typecode     IN PLS_INTEGER,
   rtype        IN OUT NOCOPY AnyType,
   aset         OUT NOCOPY ANYDATASET);
```

Parameters

Table 303-3 BEGINCREATE Procedure Parameter

Parameter	Description
typecode	The typecode for the type of the ANYDATASET.
dtype	The type of the data values. This parameter is a must for user-defined types like TYPECODE_OBJECT, Collection typecodes, and similar others.
aset	The ANYDATASET being constructed.

Exceptions

DBMS_TYPES.invalid_parameters: dtype is invalid (not fully constructed, and like errors.)

303.2.3 ENDCREATE Member Procedure

This procedure ends creation of a ANYDATASET. Other creation functions cannot be called after this call.

Syntax

```
MEMBER PROCEDURE ENDCREATE(
   self              IN OUT NOCOPY ANYDATASET);
```

Parameters

Table 303-4 ENDCREATE Procedure Parameter

Parameter	Description
self	The ANYDATASET being constructed.

303.2.4 GET* Member Functions

These functions get the current data value (which should be of the appropriate type).

The type of the current data value depends on the MODE used for accessing it (depending on how the PIECEWISE call is invoked). If PIECEWISE has not been called, the instance is accessed in its entirety, and the type of the data value should match the type of the ANYDATASET.

If `PIECEWISE` has been called, the instance is accessed piece-wise. The type of the data value should match the type of the attribute (or collection element) at the current position.

Syntax

```
MEMBER FUNCTION GETBDOUBLE(
   self        IN ANYDATASET,
   dbl         OUT NOCOPY BINARY_DOUBLE)
   RETURN PLS_INTEGER;

MEMBER FUNCTION GETBFLOAT(
   self        IN ANYDATASET,
   fl          OUT NOCOPY BINARY_FLOAT)
   RETURN PLS_INTEGER;

MEMBER FUNCTION GETBFILE(
   self        IN ANYDATASET,
   b           OUT NOCOPY BFILE)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETBLOB(
   self        IN ANYDATASET,
   b           OUT NOCOPY BLOB)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETCHAR(
   self        IN ANYDATASET,
   c           OUT NOCOPY CHAR)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETCLOB(
   self        IN ANYDATASET,
   c           OUT NOCOPY CLOB)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETCOLLECTION(
   self        IN ANYDATASET,
   col         OUT NOCOPY "<collection_type>")
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETDATE(
   self        IN ANYDATASET,
   dat         OUT NOCOPY DATE)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETINTERVALDS(
   self        IN ANYDATASET,
   inv         IN OUT NOCOPY INTERVAL DAY TO SECOND)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETINTERVALYM(
   self        IN ANYDATASET,
   inv IN OUT NOCOPY INTERVAL YEAR TO MONTH)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETNCHAR(
   self        IN ANYDATASET,
   nc          OUT NOCOPY NCHAR)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETNCLOB(
   self        IN ANYDATASET,
```

```
   nc          OUT NOCOPY NCLOB)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETNUMBER(
   self        IN ANYDATASET,
   num         OUT NOCOPY NUMBER)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETNVARCHAR2(
   self        IN ANYDATASET,
   nc          OUT NOCOPY NVARCHAR2)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETOBJECT(
   self        IN ANYDATASET,
   obj         OUT NOCOPY "<object_type>")
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETRAW(
   self        IN ANYDATASET,
   r           OUT NOCOPY RAW)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETREF(
   self        IN ANYDATASET,
   rf          OUT NOCOPY REF "<object_type>")
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETTIMESTAMP(
   self        IN ANYDATASET,
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETTIMESTAMPLTZ(
   self        IN ANYDATASET,
   ts          OUT NOCOPY TIMESTAMP WITH LOCAL TIME ZONE)
 RETURN PLS_INTEGER;

MEMBER FUNCTION GETTIMESTAMPTZ(
   self        IN ANYDATASET,
   ts          OUT NOCOPY TIMESTAMP WITH TIME ZONE)
 RETURN PLS_INTEGER,

MEMBER FUNCTION GETUROWID(
   self        IN ANYDATASET,
   rid         OUT NOCOPY UROWID)
 RETURN PLS_INTEGER

MEMBER FUNCTION GETVARCHAR(
   self        IN ANYDATASET,
   c           OUT NOCOPY VARCHAR)
   RETURN      PLS_INTEGER;

MEMBER FUNCTION GETVARCHAR2(
   self        IN ANYDATASET,
   c           OUT NOCOPY VARCHAR2)
   RETURN      PLS_INTEGER;
```

Parameters

Table 303-5 GET* Function Parameters

Parameter	Description
self	The ANYDATASET being accessed.
num	The number, and associated information., that is to be obtained.

Return Values

DBMS_TYPES.SUCCESS or DBMS_TYPES.NO_DATA

The return value is relevant only if PIECEWISE has been already called (for a collection). In such a case, DBMS_TYPES.NO_DATA signifies the end of the collection when all elements have been accessed.

Exceptions

DBMS_TYPES.INVALID_PARAMETERs: Invalid Parameters (if it is not appropriate to add a number at this point in the creation process).

DBMS_TYPES.INCORRECT_USAGE: Incorrect usage

DBMS_TYPES.TYPE_MISMATCH: When the expected type is different from the passed in type.

303.2.5 GETCOUNT Member Function

This function gets the number of data instances in an ANYDATASET.

Syntax

```
MEMBER FUNCTION GetCount(
   self         IN ANYDATASET)
   RETURN       PLS_INTEGER;
```

Parameter

Table 303-6 GETCOUNT Function Parameter

Parameter	Description
self	The ANYDATASET being accessed.

Return Values

The number of data instances.

303.2.6 GETINSTANCE Member Function

This function gets the next instance in an ANYDATASET. Only sequential access to the instances in an ANYDATASET is allowed.

After this function has been called, the GET* functions can be invoked on the ANYDATASET to access the current instance. If PIECEWISE is called before doing the GET* calls, the individual attributes (or collection elements) can be accessed.

It is an error to invoke this function before the ANYDATASET is fully created.

Syntax

```
MEMBER FUNCTION GETINSTANCE(
   self           IN OUT NOCOPY ANYDATASET)
   RETURN         PLS_INTEGER;
```

Parameters

Table 303-7 GETINSTANCE Function Parameter

Parameter	Description
self	The ANYDATASET being accessed.

Return Values

DBMS_TYPES.SUCCESS or DBMS_TYPES.NO_DATA

DBMS_TYPES.NO_DATA signifies the end of the ANYDATASET (all instances have been accessed).

Usage Notes

This function should be called even before accessing the first instance.

303.2.7 GETTYPE Member Function

This function gets the AnyType describing the type of the data instances in an ANYDATASET.

Syntax

```
MEMBER FUNCTION GETTYPE(
   self           IN ANYDATASET,
   typ            OUT NOCOPY AnyType)
   RETURN         PLS_INTEGER;
```

Parameters

Table 303-8 GETTYPE Function Parameter

Parameter	Description
self	The ANYDATASET.
typ	The ANYTYPE corresponding to the AnyData. May be NULL if it does not represent a user-defined function.

Return Values

The typecode corresponding to the type of the ANYDATA.

303.2.8 GETTYPENAME Member Function

This procedure gets the fully qualified type name for the ANYDATASET.

If the ANYDATASET is based on a built-in, this function will return NUMBER and associated information.

If it is based on a user defined type, this function will return *schema_name.type_name*. for example, SCOTT.FOO.

If it is based on a transient anonymous type, this function will return `NULL`.

Syntax

```
MEMBER FUNCTION GETTYPENAME(
   self            IN ANYDATASET)
   RETURN          VARCHAR2;
```

Parameter

Table 303-9 GETTYPENAME Function Parameter

Parameter	Description
`self`	The ANYDATASET being constructed.

Return Values

Type name of the `ANYDATASET`.

303.2.9 PIECEWISE Member Procedure

This procedure sets the `MODE` of construction, access of the data value to be an attribute at a time (if the data value is of `TYPECODE_OBJECT`).

It sets the `MODE` of construction, access of the data value to be a collection element at a time (if the data value is of a collection `TYPE`). Once this call has been made, subsequent `SET*` and `GET*` calls will sequentially obtain individual attributes or collection elements.

Syntax

```
MEMBER PROCEDURE PIECEWISE(
   self          IN OUT NOCOPY ANYDATASET);
```

Parameters

Table 303-10 PIECEWISE Procedure Parameter

Parameter	Description
`self`	The ANYDATASET being constructed.

Exceptions

`DBMS_TYPES.INVALID_PARAMETERS`: Invalid parameters.

`DBMS_TYPES.INCORRECT_USAGE`: On incorrect usage.

Usage Notes

The current data value must be of an object or collection type before this call can be made. There is no support for piece-wise construction or access of embedded object type attributes or nested collections.

303.2.10 SET* Member Procedures

This procedure sets the current data value.

The type of the current data value depends on the `MODE` with which we are constructing (depending on how we have invoked the `PIECEWISE` call). The type of the current data should be the type of the `ANYDATASET` if `PIECEWISE` has NOT been called. The type should be the type of the attribute at the current position if `PIECEWISE` has been called.

Syntax

```
MEMBER PROCEDURE SETBDOUBLE(
   self              IN OUT NOCOPY ANYDATASET,
   dbl               IN BINARY_DOUBLE,
   last_elem         IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETBFLOAT(
   self              IN OUT NOCOPY ANYDATASET,
   fl                IN BINARY_FLOAT,
   last_elem         IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETBFILE(
   self              IN OUT NOCOPY ANYDATASET,
   b                 IN BFILE,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETBLOB(
   self              IN OUT NOCOPY ANYDATASET,
   b                 IN BLOB,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETCHAR(
   self              IN OUT NOCOPY ANYDATASET,
   c                 IN CHAR,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETCLOB(
   self              IN OUT NOCOPY ANYDATASET,
   c                 IN CLOB,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETCOLLECTION(
   self              IN OUT NOCOPY ANYDATASET,
   col               IN "<collection_type>",
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETDATE(
   self              IN OUT NOCOPY ANYDATASET,
   dat               IN DATE,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETINTERVALDS(
   self              IN OUT NOCOPY ANYDATASET,
   inv               IN INTERVAL DAY TO SECOND,
   last_elem         IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETINTERVALYM(
   self              IN OUT NOCOPY ANYDATASET,
   inv               IN INTERVAL YEAR TO MONTH,
   last_elem         IN BOOLEAN DEFAULT FALSE);
```

```
MEMBER PROCEDURE SETNCHAR(
   self              IN OUT NOCOPY ANYDATASET,
   nc                IN NCHAR,
   last_elem IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETNCLOB(
  self               IN OUT NOCOPY ANYDATASET,
  nc                 IN NClob,
  last_elem          IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETNUMBER(
   self              IN OUT NOCOPY ANYDATASET,
   num               IN NUMBER,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETNVARCHAR2(
   self              IN OUT NOCOPY ANYDATASET,
   nc                IN NVarchar2,
   last_elem         IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETOBJECT(
   self              IN OUT NOCOPY ANYDATASET,
   obj               IN "<object_type>",
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETRAW(
   self              IN OUT NOCOPY ANYDATASET,
   r                 IN RAW,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETREF(
   self              IN OUT NOCOPY ANYDATASET,
   rf                IN REF "<object_type>",
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETTIMESTAMP(
   self              IN OUT NOCOPY ANYDATASET,
   ts                IN TIMESTAMP,
   last_elem IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETTIMESTAMPLTZ(
   self              IN OUT NOCOPY ANYDATASET,
   ts                IN TIMESTAMP WITH LOCAL TIME ZONE,
   last_elem         IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETTIMESTAMPTZ(
   self              IN OUT NOCOPY ANYDATASET,
   ts                IN TIMESTAMP WITH TIME ZONE,
   last_elem         IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETUROWID(
   self              IN OUT NOCOPY ANYDATASET,
   rid               IN UROWID,
   last_elem IN BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETVARCHAR(
   self              IN OUT NOCOPY ANYDATASET,
   c                 IN VARCHAR,
   last_elem BOOLEAN DEFAULT FALSE);

MEMBER PROCEDURE SETVARCHAR2(
```

```
self             IN OUT NOCOPY ANYDATASET,
c                IN VARCHAR2,
last_elem BOOLEAN DEFAULT FALSE);
```

Parameters

Table 303-11 SET* Procedure Parameters

Parameter	Description
`self`	The `ANYDATASET` being accessed.
`num`	The number, and associated information, that is to be set.
`last_elem`	Relevant only if `PIECEWISE` has been already called (for a collection). Set to TRUE if it is the last element of the collection, FALSE otherwise.

Exceptions

- `DBMS_TYPES.INVALID_PARAMETERS`: Invalid parameters (if it is not appropriate to add a number at this point in the creation process).
- `DBMS_TYPES.INCORRECT_USAGE`: Incorrect usage.
- `DBMS_TYPES.TYPE_MISMATCH`: When the expected type is different from the passed in type.

304
ANYTYPE TYPE

An ANYTYPE TYPE can contain a type description of any persistent SQL type, named or unnamed, including object types and collection types. It can also be used to construct new transient type descriptions.

New persistent types can only be created using the CREATE TYPE statement. Only new transient types can be constructed using the ANYTYPE interfaces.

Starting with Oracle Database 19c, the GETANYTYPEFROMPERSISTENT function replaces the GETPERISTENT static function.

This chapter discusses the following:

- Summary of ANYTYPE Subprograms

304.1 Summary of ANYTYPE Subprograms

This table lists the ANYTYPE subprograms in alphabetical order and briefly describes them.

Table 304-1 ***ANYTYPE Subprograms***

Subprogram	Description
BEGINCREATE Static Procedure	Creates a new instance of ANYTYPE which can be used to create a transient type description.
SETINFO Member Procedure	Sets any additional information required for constructing a COLLECTION or builtin type.
ADDATTR Member Procedure	Adds an attribute to an ANYTYPE (of typecode DBMS_TYPES.TYPECODE_OBJECT).
ENDCREATE Member Procedure	Ends creation of a transient ANYTYPE. Other creation functions cannot be called after this call.
GETINFO Member Function	Gets the type information for the ANYTYPE.
GETATTRELEMINFO Member Function	Gets the type information for an attribute of the type (if it is of TYPECODE_OBJECT). Gets the type information for a collection's element type if the *self* parameter is of a collection type.
GETANYTYPEFROMPERSISTENT Function	This standalone function replaces the old static function GETPERSISTENT. This returns an ANYTYPE corresponding to a persistent type created earlier using the CREATE TYPE SQL statement.

304.1.1 BEGINCREATE Static Procedure

This procedure creates a new instance of `ANYTYPE` which can be used to create a transient type description.

Syntax

```
STATIC PROCEDURE BEGINCREATE(
   typecode       IN          PLS_INTEGER,
   atype          OUT NOCOPY  ANYTYPE);
```

Parameters

Table 304-2 BEGINCREATE Procedure Parameters

Parameter	Description
`typecode`	Use a constant from `DBMS_TYPES` package. Typecodes for user-defined type: • `DBMS_TYPES.TYPECODE_OBJECT` • `DBMS_TYPES.TYPECODE_VARRAY` or • `DBMS_TYPES.TYPECODE_TABLE` Typecodes for builtin types: • `DBMS_TYPES.TYPECODE_NUMBER`, and similar types.
`atype`	`ANYTYPE` for a transient type

304.1.2 SETINFO Member Procedure

This procedure sets any additional information required for constructing a `COLLECTION` or builtin type.

Syntax

```
MEMBER PROCEDURE SETINFO(
   self           IN OUT NOCOPY ANYTYPE,
   prec           IN PLS_INTEGER,
   scale          IN PLS_INTEGER,
   len            IN PLS_INTEGER,
   csid           IN PLS_INTEGER,
   csfrm          IN PLS_INTEGER,
   atype          IN ANYTYPE DEFAULT NULL,
   elem_tc        IN PLS_INTEGER DEFAULT NULL,
   elem_count     IN PLS_INTEGER DEFAULT 0);
```

Parameters

Table 304-3 SETINFO Procedure Parameters

Parameter	Description
`self`	The transient `ANYTYPE` that is being constructed.
`prec`	Optional.Required if typecode represents a `NUMBER`. Give precision and scale. Ignored otherwise.

Table 304-3 (Cont.) SETINFO Procedure Parameters

Parameter	Description
scale	Optional.Required if typecode represents a NUMBER. Give precision and scale. Ignored otherwise.
len	Optional. Required if typecode represents a RAW, CHAR, VARCHAR, or VARCHAR2 type. Gives length.
csid	Required if typecode represents types requiring character information such as CHAR, VARCHAR, or VARCHAR2.
csfrm	Required if typecode represents types requiring character information such as CHAR, VARCHAR, or VARCHAR2.
atype	Optional. Required if collection element typecode is a user-defined type such as TYPECODE_OBJECT, and similar others. It is also required for a built-in type that needs user-defined type information such as TYPECODE_REF. This parameter is not needed otherwise.

The Following Parameters Are Required For Collection Types

Table 304-4 SETINFO Procedure Parameters - Collection Types

Parameter	Description
elem_tc	Must be of the collection element's typecode (from DBMS_TYPES package).
elem_count	Pass 0 for elem_count if the self represents a nested table (TYPECODE_TABLE). Otherwise pass the collection count if *self* represents a VARRAY.

Exceptions

- DBMS_TYPES.INVALID_PARAMETER: Invalid Parameters (typecode, typeinfo)
- DBMS_TYPES.INCORRECT_USAGE: Incorrect usage (cannot call after calling ENDCREATE, and similar actions.)

Usage Notes

It is an error to call this function on an ANYTYPE that represents a persistent user defined type.

304.1.3 ADDATTR Member Procedure

This procedure adds an attribute to an ANYTYPE (of typecode DBMS_TYPES.TYPECODE_OBJECT).

Syntax

```
MEMBER PROCEDURE ADDATTR(
   self          IN OUT NOCOPY ANYTYPE,
   aname         IN VARCHAR2,
   typecode      IN PLS_INTEGER,
   prec          IN PLS_INTEGER,
   scale         IN PLS_INTEGER,
   len           IN PLS_INTEGER,
   csid          IN PLS_INTEGER,
```

```
   csfrm         IN PLS_INTEGER,
   attr_type     IN ANYTYPE DEFAULT NULL);
```

Parameters

Table 304-5 ADDATTR Procedure Parameters

Parameter	Description
self	The transient ANYTYPE that is being constructed. Must be of type DBMS_TYPES.TYPECODE_OBJECT.
aname	Optional. Attribute's name. Could be NULL.
typecode	Attribute's typecode. Can be built-in or user-defined typecode (from DBMS_TYPES package).
prec	Optional. Required if typecode represents a NUMBER. Give precision and scale. Ignored otherwise.
scale	Optional. Required if typecode represents a NUMBER. Give precision and scale. Ignored otherwise.
len	Optional. Required if typecode represents a RAW, CHAR, VARCHAR, or VARCHAR2 type. Give length.
csid	Optional. Required if typecode represents a type requiring character information, such as CHAR, VARCHAR, or VARCHAR2.
csfrm	Optional. Required if typecode represents a type requiring character information, such as CHAR, VARCHAR, or VARCHAR2.
attr_type	Optional. ANYTYPE corresponding to a user-defined type. This parameter is required if the attribute is a user defined type.

Exceptions

- DBMS_TYPES.INVALID_PARAMETERS: Invalid Parameters (typecode, typeinfo)
- DBMS_TYPES.INCORRECT_USAGE: Incorrect usage (cannot call after calling EndCreate, and similar actions.)

304.1.4 ENDCREATE Member Procedure

This procedure ends creation of a transient ANYTYPE. Other creation functions cannot be called after this call.

Syntax

```
MEMBER PROCEDURE ENDCREATE(
   self           IN OUT NOCOPY ANYTYPE);
```

Parameter

Table 304-6 ENDCREATE Procedure Parameter

Parameter	Description
self	The transient ANYTYPE that is being constructed.

304.1.5 GETINFO Member Function

This function gets the type information for the ANYTYPE.

Syntax

```
MEMBER FUNCTION GETINFO (
   self        IN ANYTYPE,
   prec        OUT PLS_INTEGER,
   scale       OUT PLS_INTEGER,
   len         OUT PLS_INTEGER,
   csid        OUT PLS_INTEGER,
   csfrm       OUT PLS_INTEGER,
   schema_name OUT VARCHAR2,
   type_name   OUT VARCHAR2,
   version     OUT varchar2,
   numelems    OUT PLS_INTEGER)
   RETURN      PLS_INTEGER;
```

Parameters

Table 304-7 GETINFO Function Parameters

Parameter	Description
self	The ANYTYPE.
prec	If typecode represents a number. Gives precision and scale. Ignored otherwise.
scale	If typecode represents a number. Gives precision and scale. Ignored otherwise.
len	If typecode represents a RAW, CHAR, VARCHAR, or VARCHAR2 type. Gives length.
csid	If typecode represents a type requiring character information such as: CHAR, VARCHAR, or VARCHAR2.
csid	If typecode represents a type requiring character information such as: CHAR, VARCHAR, or VARCHAR2.
schema_name	Type's schema (if persistent).
type_name	Type's typename.
version	Type's version.
numelems	If *self* is a TYPECODE_VARRAY, this gives the VARRAY count. If self is of TYPECODE_OBJECT, this gives the number of attributes.

Return Values

The typecode of *self*.

Exceptions

- DBMS_TYPES.INVALID_PARAMETERS: Invalid Parameters (position is beyond bounds or the ANYTYPE is not properly Constructed).

304.1.6 GETATTRELEMINFO Member Function

This function gets the type information for an attribute of the type (if it is of `TYPECODE_OBJECT`). Gets the type information for a collection's element type if the *self* parameter is of a collection type.

Syntax

```
MEMBER FUNCTION GETATTRELEMINFO (
   self           IN ANYTYPE,
   pos            IN PLS_INTEGER,
   prec           OUT PLS_INTEGER,
   scale          OUT PLS_INTEGER,
   len            OUT PLS_INTEGER,
   csid           OUT PLS_INTEGER,
   csfrm          OUT PLS_INTEGER,
   attr_elt_type  OUT ANYTYPE
   aname          OUT VARRCHAR2)
   RETURN         PLS_INTEGER;
```

Parameters

Table 304-8 GETATTRELEMINFO Function Parameters

Parameter	Description
`self`	The `ANYTYPE`.
`pos`	If self is of `TYPECODE_OBJECT`, this gives the attribute position (starting at 1). It is ignored otherwise.
`prec`	If attribute/collection element typecode represents a `NUMBER`. Gives precision and scale. Ignored otherwise.
`scale`	If attribute/collection element typecode represents a `NUMBER`. Gives precision and scale. Ignored otherwise.
`len`	If typecode represents a `RAW`, `CHAR`, `VARCHAR`, or `VARCHAR2` type. Gives length.
`csid, csfrm`	If typecode represents a type requiring character information such as: `CHAR`, `VARCHAR`, or `VARCHAR2`. Gives character set `ID`, character set form.
`attr_elt_type`	If attribute/collection element typecode represents a user-defined type, this returns the `ANYTYPE` corresponding to it. User can subsequently describe the *attr_elt_type*.
`aname`	Attribute name (if it is an attribute of an object type, `NULL` otherwise).

Return Values

The typecode of the attribute or collection element.

Exceptions

`DBMS_TYPES.INVALID_PARAMETERS`: Invalid Parameters (position is beyond bounds or the ANYTYPE is not properly constructed).

304.1.7 GETANYTYPEFROMPERSISTENT Function

This standalone function returns an `ANYTYPE` corresponding to a persistent type created earlier using the `CREATE TYPE` SQL statement. Starting with Oracle Database 19c, the `GETANYTYPEFROMPERSISTENT` function replaces the `GETPERISTENT` static function.

Syntax

```
GETANYTYPEFROMPERSISTENT(
   schema_name      IN VARCHAR2,
   type_name        IN VARCHAR2)
RETURN           ANYTYPE;
```

Parameters

Table 304-9 GETANYTYPEFROMPERSISTENT Function Parameters

Parameter	Description
`schema_name`	Schema name of the type.
`type_name`	Type name.

Return Values

An `ANYTYPE` corresponding to a persistent type created earlier using the `CREATE TYPE` SQL statement.

305

Oracle Database Advanced Queuing (AQ) Types

This chapter describes the types used with Oracle Database Advanced Queuing (AQ) packages for PL/SQL, `DBMS_AQ`, and `DBMS_AQADM`.

This chapter contains the following topics:

- Security Model
- Summary of Types

> **See Also:**
>
> *Oracle Database Advanced Queuing User's Guide* for information about using Oracle Database Advanced Queuing.

305.1 Advanced Queuing (AQ)Types Security Model

`PUBLIC` is granted `EXECUTE` privilege on the types described in this chapter.

305.2 Oracle Database Advanced Queuing (AQ) Types — Summary of Types

Oracle Database Advanced Queuing (AQ) uses types to specify certain types of information.

- AQ$_AGENT Type
- AQ$_AGENT_LIST_T Type
- AQ$_DESCRIPTOR Type
- AQ$_NTFN_DESCRIPTOR Type
- AQ$_NTFN_MSGID_ARRAY Type
- AQ$_POST_INFO Type
- AQ$_POST_INFO_LIST Type
- AQ$_PURGE_OPTIONS_T Type
- AQ$_RECIPIENT_LIST_T Type
- AQ$_REG_INFO Type
- AQ$_REG_INFO_LIST Type
- AQ$_SUBSCRIBER_LIST_T Type
- DEQUEUE_OPTIONS_T Type

- ENQUEUE_OPTIONS_T Type
- QUEUE_PROPS_T Type
- SEEK_INPUT_T Type
- SEEK_OUTPUT_T Type
- SYS.MSG_PROP_T Type
- MESSAGE_PROPERTIES_T Type
- MESSAGE_PROPERTIES_ARRAY_T Type
- MSGID_ARRAY_T Type

305.2.1 AQ$_AGENT Type

This type identifies a producer or a consumer of a message.

Syntax

```
TYPE SYS.AQ$_AGENT IS OBJECT (
   name       VARCHAR2(512),
   address    VARCHAR2(1024),
   protocol   NUMBER);
```

Attributes

Table 305-1 AQ$_AGENT Attributes

Attribute	Description
`name`	Name of a producer or consumer of a message. The name must follow object name guidelines in the *Oracle Database SQL Language Reference* with regard to reserved characters.
`address`	Protocol-specific address of the recipient. If the protocol is `0`, then the address is of the form [*schema.*]*queue*[*@dblink*]. For example, a queue named `emp_messages` in the `HR` queue at the site `dbs1.net` has the address: hr.emp_messages@dbs1.net
`protocol`	Protocol to interpret the address and propagate the message. Protocols 1-127 are reserved for internal use. If the protocol number is in the range 128 - 255, the address of the recipient is not interpreted by Oracle Database Advanced Queuing.

305.2.2 AQ$_AGENT_LIST_T Type

This type identifies the list of agents for which `DBMS_AQ.LISTEN` listens.

> **See Also:**
>
> "AQ$_AGENT Type"

Syntax

```
TYPE SYS.AQ$_AGENT_LIST_T IS TABLE OF SYS.AQ$_AGENT
   INDEX BY BINARY INTEGER;
```

305.2.3 AQ$_DESCRIPTOR Type

This type specifies the Oracle Database Advanced Queuing descriptor received by the AQ PL/SQL callbacks upon notification.

> **See Also:**
>
> "MESSAGE_PROPERTIES_T Type"

Syntax

```
TYPE SYS.AQ$_DESCRIPTOR IS OBJECT (
   queue_name      VARCHAR2(261),
   consumer_name   VARCHAR2(512),
   msg_id          RAW(16),
   msg_prop        MSG_PROP_T,
   gen_desc        AQ$_NTFN_DESCRIPTOR,
   msgid_array     SYS.AQ$_NTFN_MSGID_ARRAY,
   ntfnsRecdInGrp  NUMBER);
```

Attributes

Table 305-2 AQ$_DESCRIPTOR Attributes

Attribute	Description
queue_name	Name of the queue in which the message was enqueued which resulted in the notification
consumer_name	Name of the consumer for the multiconsumer queue
msg_id	Identification number of the message
msg_prop	Message properties specified by the MSG_PROP_T type
gen_desc	Indicates the timeout specifications
msgid_array	Group notification message ID list
ntfnsRecdInGrp	Notifications received in group

305.2.4 AQ$_NTFN_DESCRIPTOR Type

This type is for storing a generic notification descriptor regarding PL/SQL notification flags.

Syntax

```
TYPE SYS.AQ$_NTFN_DESCRIPTOR IS OBJECT(   ntfn_flags   NUMBER)
```

Attributes

Table 305-3 AQ$_NTFN_DESCRIPTOR Attributes

Attribute	Description
ntfn_flags	Set to 1 if the notifications are already removed after a stipulated timeout. Set 2 to denote grouping. Default is 0.

305.2.5 AQ$_NTFN_MSGID_ARRAY Type

This type is for storing grouping notification data for AQ namespace, value 2^{30} which is the max varray size.

Syntax

```
TYPE SYS.AQ$_NTFN_MSGID_ARRAY   AS VARRAY(1073741824)OF RAW(16);
```

305.2.6 AQ$_POST_INFO Type

This type specifies anonymous subscriptions to which you want to post messages.

Syntax

```
TYPE SYS.AQ$_POST_INFO IS OBJECT (
   name        VARCHAR2(512),
   namespace   NUMBER,
   payload     RAW(32767));
```

Attributes

Table 305-4 AQ$_POST_INFO Attributes

Attribute	Description
name	Name of the anonymous subscription to which you want to post
namespace	To receive notifications from other applications through DBMS_AQ.POST or OCISubscriptionPost(), the namespace must be DBMS_AQ.NAMESPACE_ANONYMOUS
payload	The payload to be posted to the anonymous subscription

305.2.7 AQ$_POST_INFO_LIST Type

This type identifies the list of anonymous subscriptions to which you want to post messages.

> **See Also:**
>
> AQ$_POST_INFO Type

Syntax

```
TYPE SYS.AQ$_POST_INFO_LIST AS VARRAY(1024) OF SYS.AQ$_POST_INFO;
```

305.2.8 AQ$_PURGE_OPTIONS_T Type

This type specifies the options available for purging a queue table.

> **See Also:**
>
> PURGE_QUEUE_TABLE Procedure.

Syntax

```
TYPE AQ$_PURGE_OPTIONS_T is RECORD (
   block           BOOLEAN       DEFAULT FALSE
   delivery_mode   PLS_INTEGER   DEFAULT PERSISTENT);
```

Table 305-5 AQ$_PURGE_OPTIONS_T Type Attributes

Attribute	Description
`block`	`TRUE/FALSE`. • If `block` is `TRUE`, then an exclusive lock on all the queues in the queue table is held while purging the queue table. This will cause concurrent enqueuers and dequeuers to block while the queue table is purged. The purge call always succeeds if `block` is `TRUE`. • The default for `block` is `FALSE`. This will not block enqueuers and dequeuers, but it can cause the purge to fail with an error during high concurrency times.
`delivery_mode`	Kind of messages to purge, either `DBMS_AQ.BUFFERED` or `DBMS_AQ.PERSISTENT`

305.2.9 AQ$_RECIPIENT_LIST_T Type

Identifies the list of agents that receive the message. This type can be used only when the queue is enabled for multiple dequeues.

> **See Also:**
>
> "AQ$_AGENT Type"

Syntax

```
TYPE SYS.AQ$_RECIPIENT_LIST_T IS TABLE OF SYS.AQ$_AGENT
   INDEX BY BINARY_INTEGER;
```

305.2.10 AQ$_REG_INFO Type

This type identifies a producer or a consumer of a message.

Syntax

```
TYPE SYS.AQ$_REG_INFO IS OBJECT (
   name                              VARCHAR2(512),
   namespace                         NUMBER,
   callback                          VARCHAR2(4000),
   context                           RAW(2000)  DEFAULT NULL,
   anyctx                            ANYDATA,
   ctxtype                           NUMBER,
   payloadcbk                        VARCHAR2(4000),
   timeout                           NUMBER,
   ntfn_grouping_class               NUMBER,
   ntfn_grouping_value               NUMBER    DEFAULT 600,
   ntfn_grouping_type                NUMBER,
   ntfn_grouping_start_time          TIMESTAMP(6) WITH TIME ZONE,
   ntfn_grouping_repeat_count        NUMBER);
```

Attributes

Table 305-6 AQ$_REG_INFO Type Attributes

Attribute	Description
`name`	Specifies the name of the subscription. The subscription name is of the form *`schema.queue`* if the registration is for a single consumer queue or *`schema.queue:consumer_name`* if the registration is for a multiconsumer queues.
`namespace`	Specifies the namespace of the subscription. To receive notification from Oracle Database Advanced Queuing queues, the namespace must be `DBMS_AQ.NAMESPACE_AQ`. To receive notifications from other applications through `DBMS_AQ.POST` or `OCISubscriptionPost()`, the namespace must be `DBMS_AQ.NAMESPACE_ANONYMOUS`.
`callback`	Specifies the action to be performed on message notification. For HTTP notifications, use `http://www.company.com:8080`. For e-mail notifications, use `mailto://xyz@company.com`. For raw message payload for the `PLSQLCALLBACK` procedure, use `plsql://`*`schema.procedure?`*`PR=0`. For user-defined type message payload converted to XML for the `PLSQLCALLBACK` procedure, use `plsql://`*`schema.procedure?`*`PR=1`
`context`	Specifies the context that is to be passed to the callback function
`anyctx`	Specifies the `ANYDATA` type context that is passed to the callback function
`ctxtype`	Specifies the context type. Valid values are `0 (RAW)` or `1 (ANYDATA)`

Table 305-6 (Cont.) AQ$_REG_INFO Type Attributes

Attribute	Description
payloadcbk	Specifies payload for the callback function
timeout	Specifies the time in seconds
ntfn_grouping_class	Currently, only the following flag can be set to specify criterion for grouping. The default value will be 0. If ntfn_grouping_class is 0, all other notification grouping attributes must be 0. • NTFN_GROUPING_CLASS_TIME - Notifications grouped by time, that is, the user specifies a time value and a single notification gets published at the end of that time.
ntfn_grouping_value	Time-period of grouping notifications specified in seconds, meaning the time after which grouping notification would be sent periodically until ntfn_grouping_repeat_count is exhausted.
ntfn_grouping_type	• NTFN_GROUPING_TYPE_SUMMARY - Summary of all notifications that occurred in the time interval. (Default) • NTFN_GROUPING_TYPE_LAST - Last notification that occurred in the interval.
ntfn_grouping_start_time	Notification grouping start time. Notification grouping can start from a user-specified time that should a valid timestamp with time zone. If ntfn_grouping_start_time is not specified when using grouping, the default is to current timestamp with time zone
ntfn_grouping_repeat_count	Grouping notifications will be sent as many times as specified by the notification grouping repeat count and after that revert to regular notifications. The ntfn_grouping_repeat_count, if not specified, will default to • NTFN_GROUPING_FOREVER - Keep sending grouping notifications forever.

Usage Notes

You can use the following notification mechanisms:

- OCI callback
- e-mail callback
- PL/SQL callback

Table 305-7 shows the actions performed for nonpersistent queues for different notification mechanisms when RAW presentation is specified. Table 305-8 shows the actions performed when XML presentation is specified.

Table 305-7 Actions Performed for Nonpersistent Queues When RAW Presentation Specified

Queue Payload Type	OCI Callback	E-mail	PL/SQL Callback
`RAW`	OCI callback receives the RAW data in the payload.	Not supported	PL/SQL callback receives the RAW data in the payload.
`Oracle object type`	Not supported	Not supported	Not supported

Table 305-8 Actions Performed for Nonpersistent Queues When XML Presentation Specified

Queue Payload Type	OCI Callback	E-mail	PL/SQL Callback
`RAW`	OCI callback receives the XML data in the payload.	XML data is formatted as a SOAP message and e-mailed to the registered e-mail address.	PL/SQL callback receives the XML data in the payload.
`Oracle object type`	OCI callback receives the XML data in the payload.	XML data is formatted as a SOAP message and e-mailed to the registered e-mail address.	PL/SQL callback receives the XML data in the payload.

305.2.11 AQ$_REG_INFO_LIST Type

Identifies the list of registrations to a queue.

> **See Also:**
>
> "AQ$_REG_INFO Type"

Syntax

```
TYPE SYS.AQ$_REG_INFO_LIST AS VARRAY(1024) OF SYS.AQ$_REG_INFO;
```

305.2.12 AQ$_SUBSCRIBER_LIST_T Type

This type identifies the list of subscribers that subscribe to a queue.

> **See Also:**
>
> "AQ$_AGENT Type"

Syntax

```
TYPE SYS.AQ$_SUBSCRIBER_LIST_T IS TABLE OF SYS.AQ$_AGENT
   INDEX BY BINARY_INTEGER;
```

305.2.13 DEQUEUE_OPTIONS_T Type

`DEQUEUE_OPTIONS_T` specifies the options available for the dequeue operation.

Syntax

```
TYPE DEQUEUE_OPTIONS_T IS RECORD (
   consumer_name     VARCHAR2(30)    DEFAULT NULL,
   dequeue_mode      BINARY_INTEGER  DEFAULT REMOVE,
   navigation        BINARY_INTEGER  DEFAULT NEXT_MESSAGE,
   visibility        BINARY_INTEGER  DEFAULT ON_COMMIT,
   wait              BINARY_INTEGER  DEFAULT FOREVER,
   msgid             RAW(16)         DEFAULT NULL,
   correlation       VARCHAR2(128)   DEFAULT NULL,
   deq_condition     VARCHAR2(4000)  DEFAULT NULL,
   signature         aq$_sig_prop    DEFAULT NULL,
   transformation    VARCHAR2(61)    DEFAULT NULL,
   delivery_mode     PLS_INTEGER     DEFAULT PERSISTENT);
```

Attributes

Table 305-9 DEQUEUE_OPTIONS_T Attributes

Attribute	Description
`consumer_name`	Name of the consumer. Only those messages matching the consumer name are accessed. If a queue is not set up for multiple consumers, then this field should be set to `NULL`. For secure queues, consumer_name must be a valid AQ agent name, mapped to the database user performing the dequeue operation, through `dbms_aqadm.enable_db_access` procedure call.
`dequeue_mode`	Specifies the locking behavior associated with the dequeue. Possible settings are: `BROWSE`: Read the message without acquiring any lock on the message. This specification is equivalent to a select statement. `LOCKED`: Read and obtain a write lock on the message. The lock lasts for the duration of the transaction. This setting is equivalent to a select for update statement. `REMOVE`: Read the message and delete it. This setting is the default. The message can be retained in the queue table based on the retention properties. `REMOVE_NODATA`: Mark the message as updated or deleted. The message can be retained in the queue table based on the retention properties.

Table 305-9 (Cont.) DEQUEUE_OPTIONS_T Attributes

Attribute	Description
navigation	Specifies the position of the message that will be retrieved. First, the position is determined. Second, the search criterion is applied. Finally, the message is retrieved. Possible settings are: NEXT_MESSAGE: Retrieve the next message that is available and matches the search criteria. If the previous message belongs to a message group, then AQ retrieves the next available message that matches the search criteria and belongs to the message group. This setting is the default. NEXT_TRANSACTION: Skip the remainder of the current transaction group (if any) and retrieve the first message of the next transaction group. This setting can only be used if message grouping is enabled for the current queue. FIRST_MESSAGE: Retrieves the first message which is available and matches the search criteria. This setting resets the position to the beginning of the queue. FIRST_MESSAGE_MULTI_GROUP: indicates that a call to DBMS_AQ.DEQUEUE_ARRAY will reset the position to the beginning of the queue and dequeue messages (possibly across different transaction groups) that are available and match the search criteria, until reaching the ARRAY_SIZE limit. Refer to the TRANSACTION_GROUP attribute for the message to distinguish between transaction groups. NEXT_MESSAGE_MULTI_GROUP: indicates that a call to DBMS_AQ.DEQUEUE_ARRAY will dequeue the next set of messages (possibly across different transaction groups) that are available and match the search criteria, until reaching the ARRAY_SIZE limit. Refer to the TRANSACTION_GROUP attribute for the message to distinguish between transaction groups.
visibility	Specifies whether the new message is dequeued as part of the current transaction.The visibility parameter is ignored when using the BROWSE dequeue mode. Possible settings are: ON_COMMIT: The dequeue will be part of the current transaction. This setting is the default. IMMEDIATE: The dequeue operation is not part of the current transaction, but an autonomous transaction which commits at the end of the operation.
wait	Specifies the wait time if there is currently no message available which matches the search criteria. Possible settings are: FOREVER: Wait forever. This setting is the default. NO_WAIT: Do not wait. NUMBER: Wait time in seconds.
msgid	Specifies the message identifier of the message to be dequeued.
correlation	Specifies the correlation identifier of the message to be dequeued. Special pattern matching characters, such as the percent sign (%) and the underscore (_) can be used. If more than one message satisfies the pattern, then the order of dequeuing is undetermined.

Table 305-9 (Cont.) DEQUEUE_OPTIONS_T Attributes

Attribute	Description
deq_condition	A conditional expression based on the message properties, the message data properties, and PL/SQL functions. A deq_condition is specified as a Boolean expression using syntax similar to the WHERE clause of a SQL query. This Boolean expression can include conditions on message properties, user data properties (object payloads only), and PL/SQL or SQL functions (as specified in the WHERE clause of a SQL query). Message properties include priority, corrid and other columns in the queue table To specify dequeue conditions on a message payload (object payload), use attributes of the object type in clauses. You must prefix each attribute with tab.user_data as a qualifier to indicate the specific column of the queue table that stores the payload. The deq_condition parameter cannot exceed 4000 characters. If more than one message satisfies the dequeue condition, then the order of dequeuing is undetermined.
signature	Currently not implemented
transformation	Specifies a transformation that will be applied after dequeuing the message. The source type of the transformation must match the type of the queue.
delivery_mode	The dequeuer specifies the delivery mode of the messages it wishes to dequeue in the dequeue options. It can be BUFFERED or PERSISTENT or PERSISTENT_OR_BUFFERED. The message properties of the dequeued message indicate the delivery mode of the dequeued message. Array dequeue is only supported for buffered messages with an array size of '1'.

305.2.14 ENQUEUE_OPTIONS_T Type

ENQUEUE_OPTIONS_T Type specifies the options available for the enqueue operation.

Syntax

```
TYPE SYS.ENQUEUE_OPTIONS_T IS RECORD (
   visibility            BINARY_INTEGER  DEFAULT ON_COMMIT,
   relative_msgid        RAW(16)         DEFAULT NULL,
   sequence_deviation    BINARY_INTEGER  DEFAULT NULL,
   transformation        VARCHAR2(61)    DEFAULT NULL,
   delivery_mode         PLS_INTEGER     NOT NULL DEFAULT PERSISTENT);
```

Attributes

Table 305-10 ENQUEUE_OPTIONS_T Attributes

Attribute	Description
visibility	Specifies the transactional behavior of the enqueue request. Possible settings are: ON_COMMIT: The enqueue is part of the current transaction. The operation is complete when the transaction commits. This setting is the default. IMMEDIATE: The enqueue operation is not part of the current transaction, but an autonomous transaction which commits at the end of the operation. This is the only value allowed when enqueuing to a non-persistent queue.

Table 305-10 (Cont.) ENQUEUE_OPTIONS_T Attributes

Attribute	Description
relative_msgid	Specifies the message identifier of the message which is referenced in the sequence deviation operation. This field is valid only if BEFORE is specified in sequence_deviation. This parameter is ignored if sequence deviation is not specified.
sequence_deviation	Specifies whether the message being enqueued should be dequeued before other messages already in the queue. Possible settings are: BEFORE: The message is enqueued ahead of the message specified by relative_msgid. TOP: The message is enqueued ahead of any other messages.
transformation	Specifies a transformation that will be applied before enqueuing the message. The return type of the transformation function must match the type of the queue.
delivery_mode	The enqueuer specifies the delivery mode of the messages it wishes to enqueue in the enqueue options. It can be BUFFERED or PERSISTENT. The message properties of the enqueued message indicate the delivery mode of the enqueued message. Array enqueue is only supported for buffered messages with an array size of '1'.

305.2.15 QUEUE_PROPS_T Type

This type specifies the Oracle Database Advanced Queuing descriptor received by the AQ PL/SQL callbacks upon notification.

Syntax

```
TYPE QUEUE_PROPS_T IS RECORD (
  retry_delay     NUMBER         DEFAULT 0,
  retention_time  NUMBER         DEFAULT 0,
  sort_list       VARCHAR2(30)   DEFAULT NULL,
  cache_hint      BINARY_INTEGER DEFAULT AUTO,
  retention_type  BINARY_INTEGER DEFAULT DBMS_AQ.DEQUEUE_TIME
  );
```

Attributes

Table 305-11 QUEUE_PROPS_T Attributes

Attribute	Description
retry_delay	Specifies the number of seconds after which this message is scheduled for processing again after an application rollback
retention_time	Specifies the number of seconds a message is retained in the queue table after being dequeued from the queue
sort_list	Sharded queues can be sorted by ENQ_TIME or PRIORITY, ENQ_TIME .
cache_hint	Specifies a hint to sharded queue whether to cache messages or not. User can specify following: • UNCACHED: Message cache will not cache messages • CACHED: Message cache will cache messages and give preference to those queues with CACHED hint. • AUTO: Message cache does best effort based on available memory

Table 305-11 (Cont.) QUEUE_PROPS_T Attributes

Attribute	Description
retention_type	Specifies the type of retention. DBMS_AQ.DEQUEUE_TIME (default): The retention time starts after a subshard is dequeued by all the subscribers. If a new subscriber seeks back before or to this subshard, then the retention clock is reset.

305.2.16 SEEK_INPUT_T Type

This type specifies the seek input per shard.

Syntax

```
TYPE SEEK_INPUT_T IS RECORD(
               shard             binary_integer DEFAULT DBMS_AQ.ALL_SHARDS,
               priority          binary_integer DEFAULT DBMS_AQ.ALL_PRIORITIES,
               seek_msgid        raw(16)        DEFAULT NULL,
               seek_time         TIMESTAMP WITH TIME ZONE DEFAULT NULL
             );
```

Attributes

Table 305-12 SEEK_INPUT_T Attributes

Attribute	Description
shard	Shard identifier on which seek operation needs to be performed. Possible values are DBMS_AQ.ALL_SHARDS or a specific shard on which seek is to be performed If value is DBMS_AQ.ALL_SHARDS, then seek is performed on requested shards, otherwise it is performed on the specific shard mentioned in this field.
priority	Seek is performed on all priorities of the shard individually. If value is DBMS_AQ.ALL_PRIORITIES, then seek is performed on all priorities, otherwise it is performed on the specific priority mentioned in this field
seek_msgid	Input message id when seek_type is DBMS_AQ.MESSAGE.
seek_time	Specifies the input time when seek_type is DBMS_AQ.TIME.

305.2.17 SEEK_OUTPUT_T Type

This type specifies the seek output of a seek call.

Syntax

```
TYPE SEEK_OUTPUT_T IS RECORD(
          shard             binary_integer,
          priority          binary_integer,
          seeked_from       raw(16),
          seeked_to         raw(16));
```

Attributes

Table 305-13 SEEK_OUTPUT_T Attributes

Attribute	Description
shard	Specifies the shard identifier on which the seek operation is performed for the subscriber.
priority	Specifies the priority of the shard, which is applicable only for priority queues, else value is null.
seeked_from	Secifies the message id of dequeue position before seek. NULL value signifies seeking from end of shard.
seeked_to	Specifies the message id of dequeue position after seek. NULL value signifies seeking to end of shard.

305.2.18 SYS.MSG_PROP_T Type

This type is used in PL/SQL notification, as one field in aq$_descriptor, to pass message properties of an AQ message to the PL/SQL notification client callback.

Syntax

```
CREATE or replace TYPE sys.msg_prop_t AS OBJECT (
   priority         NUMBER,
   delay            NUMBER,
   expiration       NUMBER,
   correlation      VARCHAR2(128),
   attempts         NUMBER,
   exception_queue  VARCHAR2(51),
   enqueue_time     DATE,
   state            NUMBER,
   sender_id        aq$_agent,
   original_msgid   RAW(16),
   delivery_mode    NUMBER);
```

Parameters

Table 305-14 SYS.MSG_PROP_T Type Attributes

Parameter	Description
priority	Specifies the priority of the message. A smaller number indicates higher priority. The priority can be any number, including negative numbers.
delay	Specifies the delay of the enqueued message. The delay represents the number of seconds after which a message is available for dequeuing. Dequeuing by msgid overrides the delay specification. A message enqueued with delay set is in the WAITING state, and when the delay expires, the message goes to the READY state. DELAY processing requires the queue monitor to be started. However the queue monitor is started automatically by the system if needed. Delay is set by the producer who enqueues the message. The possible settings follow: NO_DELAY: The message is available for immediate dequeuing number: The number of seconds to delay the message

Table 305-14 (Cont.) SYS.MSG_PROP_T Type Attributes

Parameter	Description
expiration	Specifies the expiration of the message. It determines, in seconds, the duration the message is available for dequeuing. This parameter is an offset from the time the message is ready for dequeue. Expiration processing requires the queue monitor to be running. However the queue monitor is started automatically by the system if needed. The possible settings follow: NEVER: The message does not expire number: The number of seconds message remains in READY state. If the message is not dequeued before it expires, then it is moved to the exception queue in the EXPIRED state.
correlation	Returns the identifier supplied by the producer of the message at enqueue time.
attempts	Returns the number of attempts that have been made to dequeue the message. This parameter cannot be set at enqueue time.
exception_queue	Specifies the name of the queue into which the message is moved if it cannot be processed successfully. Messages are moved automatically into the exception queue. Messages are moved into the exception queue in the following cases: • RETRY_COUNT, the number of unsuccessful dequeue attempts, has exceeded the specification for the MAX_RETRIES parameter in the DBMS_AQADM.CREATE_QUEUE procedure during queue creation. For multiconsumer queues, the message becomes eligible to be moved to the exception queue even if failed dequeue attempts exceeds the MAX_RETRIES parameter for only one of the consumers. But the message will not be moved until either all other consumers have successfully consumed the message or failed more than MAX_RETRIES. You can view MAX_RETRIES for a queue in the ALL_QUEUES data dictionary view. If a dequeue transaction fails because the server process dies (including ALTER SYSTEM KILL SESSION) or SHUTDOWN ABORT on the instance, then RETRY_COUNT is not incremented. • A message was not dequeued before the expiration time elapsed. • Message propagation to the specified destination queue failed with one of the following errors: * There were no recipients for the multiconsumer destination queue. * Recipients were specified for a single-consumer destination queue. * Destination queue was an exception queue * There was an error when applying transformation. The default is the exception queue associated with the queue table. If the exception queue specified does not exist at the time of the move, then the message is moved to the default exception queue associated with the queue table, and a warning is logged in the alert log. If the default exception queue is specified, then the parameter returns a NULL value at dequeue time.
enqueue_time	Specifies the time the message was enqueued. This value is determined by the system and cannot be set by the user at enqueue time.

Table 305-14 (Cont.) SYS.MSG_PROP_T Type Attributes

Parameter	Description
state	Specifies the state of the message at the time of the dequeue. This parameter cannot be set at enqueue time. The possible states follow: • DBMS_AQ.READY: The message is ready to be processed. • DBMS_AQ.WAITING: The message delay has not yet been reached. • DBMS_AQ.PROCESSED: The message has been processed and is retained. • DBMSAQ.EXPIRED: The message has been moved to the exception queue.
sender_id	The application-sender identification specified at enqueue time by the message producer. Sender id is of type aq$_agent. Sender name is required for secure queues at enqueue time. This must be a valid AQ agent name, mapped to the database user performing the enqueue operation, through dbms_aqadm.enable_db_access procedure call. Sender address and protocol should not be specified. The Sender id in the message properties returned at dequeue time may have a sender address if the message was propagated from another queue. The value of the address is the source_queue, source database name if it was a remote database [format source_queue@source_database_name]
original_msgid	This parameter is used by Oracle Database Advanced Queuing for propagating messages.
delivery_mode	DBMS_AQ.BUFFERED or DBMS_AQ.PERSISTENT.

305.2.19 MESSAGE_PROPERTIES_T Type

This type is defined inside the DBMS_AQ package, and describes the information that AQ uses to convey the state of individual messages. These are set at enqueue time, and their values are returned at dequeue time.

> **See Also:**
>
> AQ$_RECIPIENT_LIST_T Type

Syntax

```
TYPE message_properties_t IS RECORD (
   priority               BINARY_INTEGER  NOT NULL DEFAULT 1,
   delay                  BINARY_INTEGER  NOT NULL DEFAULT NO_DELAY,
   expiration             BINARY_INTEGER  NOT NULL DEFAULT NEVER,
   correlation            VARCHAR2(128)   DEFAULT NULL,
   attempts               BINARY_INTEGER,
   recipient_list         AQ$_RECIPIENT_LIST_T,
   exception_queue        VARCHAR2(61)    DEFAULT NULL,
   enqueue_time           DATE,
   state                  BINARY_INTEGER,
   sender_id              SYS.AQ$_AGENT   DEFAULT NULL,
   original_msgid         RAW(16)         DEFAULT NULL,
   signature              aq$_sig_prop    DEFAULT NULL,
```

```
  transaction_group       VARCHAR2(30)     DEFAULT NULL,
  user_property           SYS.ANYDATA      DEFAULT NULL
  delivery_mode           PLS_INTEGER      NOT NULL DEFAULT DBMS_AQ.PERSISTENT);
```

Attributes

Table 305-15 MESSAGE_PROPERTIES_T Attributes

Attribute	Description
`priority`	Specifies the priority of the message. A smaller number indicates higher priority. The priority can be any number, including negative numbers.
`delay`	Specifies the delay of the enqueued message. The delay represents the number of seconds after which a message is available for dequeuing. Dequeuing by `msgid` overrides the delay specification. A message enqueued with delay set is in the `WAITING` state, and when the delay expires, the message goes to the `READY` state. `DELAY` processing requires the queue monitor to be started. However the queue monitor is started automatically by the system if needed. Delay is set by the producer who enqueues the message. The possible settings follow: `NO_DELAY`: The message is available for immediate dequeuing number: The number of seconds to delay the message
`expiration`	Specifies the expiration of the message. It determines, in seconds, the duration the message is available for dequeuing. This parameter is an offset from the time the message is ready for dequeue. Expiration processing requires the queue monitor to be running. However the queue monitor is started automatically by the system if needed. The possible settings follow: `NEVER`: The message does not expire number: The number of seconds message remains in `READY` state. If the message is not dequeued before it expires, then it is moved to the exception queue in the `EXPIRED` state.
`correlation`	Returns the identifier supplied by the producer of the message at enqueue time.
`attempts`	Returns the number of attempts that have been made to dequeue the message. This parameter cannot be set at enqueue time.
`recipient_list`	This parameter is only valid for queues that allow multiple consumers. The default recipients are the queue subscribers. This parameter is not returned to a consumer at dequeue time. For type definition, see the "AQ$_AGENT Type".

Table 305-15 (Cont.) MESSAGE_PROPERTIES_T Attributes

Attribute	Description
exception_queue	Specifies the name of the queue into which the message is moved if it cannot be processed successfully. Messages are moved automatically into the exception queue. Messages are moved into the exception queue in the following cases: • RETRY_COUNT, the number of unsuccessful dequeue attempts, has exceeded the specification for the MAX_RETRIES parameter in the DBMS_AQADM.CREATE_QUEUE procedure during queue creation. For multiconsumer queues, the message becomes eligible to be moved to the exception queue even if failed dequeue attempts exceeds the MAX_RETRIES parameter for only one of the consumers. But the message will not be moved until either all other consumers have successfully consumed the message or failed more than MAX_RETRIES. You can view MAX_RETRIES for a queue in the ALL_QUEUES data dictionary view. If a dequeue transaction fails because the server process dies (including ALTER SYSTEM KILL SESSION) or SHUTDOWN ABORT on the instance, then RETRY_COUNT is not incremented. • A message was not dequeued before the expiration time elapsed. • Message propagation to the specified destination queue failed with one of the following errors: * There were no recipients for the multiconsumer destination queue. * Recipients were specified for a single-consumer destination queue. * Destination queue was an exception queue * There was an error when applying transformation. The default is the exception queue associated with the queue table. If the exception queue specified does not exist at the time of the move, then the message is moved to the default exception queue associated with the queue table, and a warning is logged in the alert log. If the default exception queue is specified, then the parameter returns a NULL value at dequeue time.
enqueue_time	The enqueue_time attribute specifies the time the message was enqueued. This value is always in Universal Coordinated Time (UTC), and is determined by the system and cannot be set by the user at enqueue time.
state	Specifies the state of the message at the time of the dequeue. This parameter cannot be set at enqueue time. The possible states follow: • DBMS_AQ.READY: The message is ready to be processed. • DBMS_AQ.WAITING: The message delay has not yet been reached. • DBMS_AQ.PROCESSED: The message has been processed and is retained. • DBMSAQ.EXPIRED: The message has been moved to the exception queue.

Table 305-15 (Cont.) MESSAGE_PROPERTIES_T Attributes

Attribute	Description
sender_id	The application-sender identification specified at enqueue time by the message producer. Sender id is of type aq$_agent. Sender name is required for secure queues at enqueue time. This must be a valid AQ agent name, mapped to the database user performing the enqueue operation, through dbms_aqadm.enable_db_access procedure call. Sender address and protocol should not be specified. The Sender id in the message properties returned at dequeue time may have a sender address if the message was propagated from another queue. The value of the address is the source_queue, source database name if it was a remote database [format source_queue@source_database_name]
original_msgid	This parameter is used by Oracle Database Advanced Queuing for propagating messages.
signature	Currently not implemented
transaction_group	Specifies the transaction_group for the dequeued message. Messages belonging to the same transaction group will have the same value for this attribute. This attribute is only set by the DBMS_AQ.DEQUEUE_ARRAY. This attribute cannot be used to set the transaction group of a message through DBMS_AQ.ENQUEUE or DBMS_AQ.ENQUEUE_ARRAY calls.
user_property	This optional attribute is used to store additional information about the payload.
delivery_mode	The message publisher specifies the delivery mode in the message_properties. This can be DBMS_AQ.BUFFERED or DBMS_AQ.PERSISTENT. Array enqueue is only supported for buffered messages with an array size of '1'.

305.2.20 MESSAGE_PROPERTIES_ARRAY_T Type

This type is used by dbms_aq.enqueue_array and dbms_aq.dequeue_array calls to hold the set of message properties.

Each element in the payload_array should have a corresponding element in the MESSAGE_PROPERTIES_ARRAY_T VARRAY.

> **See Also:**
>
> "MESSAGE_PROPERTIES_T Type"

Syntax

```
TYPE MESSAGE_PROPERTIES_ARRAY_T IS VARRAY (2147483647)
   OF MESSAGE_PROPERTIES_T;
```

305.2.21 MSGID_ARRAY_T Type

The `msgid_array_t` type is used in `dbms_aq.enqueue_array` and `dbms_aq.dequeue_array` calls to hold the set of message IDs that correspond to the enqueued or dequeued messages.

Syntax

```
TYPE MSGID_ARRAY_T IS TABLE OF RAW(16) INDEX BY BINARY_INTEGER
```

306

DBFS Content Interface Types

DBFS Content Interface Types are public types that support the `DBMS_DBFS_CONTENT` interface.

This chapter contains the following topics:

- Overview
- Security Model
- Data Structures

Related Topics

- DBMS_DBFS_CONTENT
  The `DBMS_DBFS_CONTENT` package provides an interface comprising a file system-like abstraction backed by one or more Store Providers.

306.1 DDBFS Content Interface Types Overview

The type definitions described in this chapter support the DBMS_DBFS_CONTENT interface in implementing metadata tables, packages, views, dependent application-side entities, and service-provider entities.

306.2 DBFS Content Interface Types Security Model

The user can access the content operational and administrative interfaces (packages, types, tables, and so on) with the `DBFS_ROLE`. This role can be granted to users as needed.

306.3 DBFS Content Interface Types —Data Structures

Types that support the DBMS_DBFS_CONTENT interface include both Object and Table types.

Object Types

- DBMS_DBFS_CONTENT_CONTEXT_T Object Type
- DBMS_DBFS_CONTENT_LIST_ITEM_T Object Type
- DBMS_DBFS_CONTENT_PROPERTY_T Object Type

Table Types

- DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type
- DBMS_DBFS_CONTENT_PROPERTIES_T Table Type
- DBMS_DBFS_CONTENT_RAW_T Table Type

306.3.1 DBMS_DBFS_CONTENT_CONTEXT_T Object Type

This type describes the execution context for the providers. It provides the user performing the operation with the Access Control List, the owner of the item(s), a timestamp for doing `asof` queries, and whether or not the item(s) are `read_only`. This type can be used both as input, in the case of path item creation functions, and output, in the case of path item query, or both.

Syntax

```
CREATE OR REPLACE TYPE dbms_dbfs_content_context_t
    AUTHID DEFINER
AS OBJECT (
    principal     VARCHAR2(32),
    acl           VARCHAR2(1024),
    owner         VARCHAR2(32),
    asof          TIMESTAMP,
    read_only     INTEGER);
```

Fields

Table 306-1 DBMS_DBFS_CONTENT_CONTEXT_T Fields

Field	Description
`principal`	File system user
`acl`	Access control list
`owner`	Path item owner
`asof`	Timestamp
`read_only`	Nonzero if the path item is read-only

306.3.2 DBMS_DBFS_CONTENT_LIST_ITEM_T Object Type

This type describes a type to assist in listing the contents of a directory.

Syntax

```
CREATE OR REPLACE TYPE dbms_dbfs_content_list_item_t
    AUTHID DEFINER
AS OBJECT (
    path          VARCHAR2(1024),
    item_name     VARCHAR2(256),
    item_type     INTEGER);
```

Fields

Table 306-2 DBMS_DBFS_CONTENT_LIST_ITEM_T Fields

Field	Description
`path`	Path to the path item
`item_name`	Name of the path item
`item_type`	Type of path item. (See Table 62-4

306.3.3 DBMS_DBFS_CONTENT_PROPERTY_T Object Type

This type describes a single (`name`, `value`, `typecode`) property tuple. All properties (standard, optional, and user-defined) are described using such tuples.

The type is used by both the client-facing interfaces and by store providers for the `DBMS_DBFS_CONTENT` interface.

Syntax

```
CREATE OR REPLACE TYPE dbms_dbfs_content_property_t
    AUTHID DEFINER
AS OBJECT (
    propname     VARCHAR2(32),
    propvalue    VARCHAR2(1024),
    typecode     INTEGER);
```

Fields

Table 306-3 DBMS_DBFS_CONTENT_PROPERTY_T Fields

Field	Description
`prop_name`	Name of property
`prop_value`	Value of property
`typecode`	Property type (See Constants in DBMS_TYPES)

306.3.4 DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type

This type is a variable-sized array of DBMS_DBFS_CONTENT_LIST_ITEM_T Object Type. It is used by both the client-facing interfaces and by store providers for the DBMS_DBFS_CONTENT interface.

Syntax

```
CREATE OR REPLACE TYPE dbms_dbfs_content_list_items_t AS
    TABLE OF dbms_dbfs_content_list_item_t;
```

Related Topics

- DBMS_DBFS_CONTENT_LIST_ITEM_T Object Type
  This type describes a type to assist in listing the contents of a directory.
- DBMS_DBFS_CONTENT
  The `DBMS_DBFS_CONTENT` package provides an interface comprising a file system-like abstraction backed by one or more Store Providers.

306.3.5 DBMS_DBFS_CONTENT_PROPERTIES_T Table Type

This type is a variable-sized array of property tuples of DBMS_DBFS_CONTENT_PROPERTY_T Object Type. It is used by both the client-facing interfaces and by store providers for the `DBMS_DBFS_CONTENT` interface.

Syntax

```
CREATE OR REPLACE TYPE dbms_dbfs_content_properties_t AS
    TABLE OF dbms_dbfs_content_property_t;
```

Related Topics

- DBMS_DBFS_CONTENT_PROPERTY_T Object Type
  This type describes a single (`name`, `value`, `typecode`) property tuple. All properties (standard, optional, and user-defined) are described using such tuples.

306.3.6 DBMS_DBFS_CONTENT_RAW_T Table Type

This type is an array of `RAW`. It is to enable `RAW` data transport for batch interfaces in the DBMS_DBFS_CONTENT interface.

Syntax

```
CREATE OR REPLACE TYPE dbms_dbfs_content_raw_t AS
    TABLE OF RAW(32767);
```

Related Topics

- DBMS_DBFS_CONTENT
  The `DBMS_DBFS_CONTENT` package provides an interface comprising a file system-like abstraction backed by one or more Store Providers.

307

Database URI TYPEs

Oracle supports the `UriType` family of types that can be used to store and query Uri-refs inside the database. The `UriType` itself is an abstract object type and the `HTTPURITYPE`, `XDBURITYPE` and `DBURITYPE` are subtypes of it.

You can create a `UriType` column and store instances of the `DBURITYPE`, `XDBURITYPE` or the `HTTPURITYPE` inside of it. You can also define your own subtypes of the `UriType` to handle different URL protocols.

Oracle also provides a `UriFactory` package that can be used as a factory method to automatically generate various instances of these `UriTypes` by scanning the prefix, such as `http://` or `/oradb`. You can also register your subtype and provide the prefix that you support. For instance, if you have written a subtype to handle the gopher protocol, you can register the prefix `gopher://` to be handled by your subtype. The `UriFactory` will then generate your subtype instance for any URL starting with that prefix.

This chapter contains the following topics:

- Summary of URITYPE Supertype Subprograms
- Summary of HTTPURITYPE Subtype Subprograms
- Summary of DBURITYPE Subtype Subprogams
- Summary of XDBURITYPE Subtype Subprograms
- Summary of URIFACTORY Package Subprograms

> **See Also:**
>
> – *Oracle XML DB Developer's Guide*

307.1 Summary of URITYPE Supertype Subprograms

The `UriType` is the abstract super type. It provides a standard set of functions to get the value pointed to by the URI. The actual implementation of the protocol must be defined by the subtypes of this type.

Instances of this type cannot be created directly. However, you can create columns of this type and store subtype instances in it, and also select from columns without knowing the instance of the URL stored.

Table 307-1 URITYPE Type Subprograms

Method	Description
GETBLOB	Returns the `BLOB` located at the address specified by the URL.
GETCLOB	Returns the `CLOB` located at the address specified by the URL.

Table 307-1 (Cont.) URITYPE Type Subprograms

Method	Description
GETCONTENTTYPE	Returns the URL, in escaped format, stored inside the `UriType` instance.
GETEXTERNALURL	Returns the URL, in escaped format, stored inside the `UriType` instance.
GETURL	Returns the URL, in non-escaped format, stored inside the `UriType` instance.
GETXML	Returns the `XMLType` located at the address specified by the URL.

307.1.1 UriType Supertype GETBLOB

This function returns the `BLOB` located at the address specified by the URL.

This function can be overridden in the subtype instances. The options are described below.

This function returns the `BLOB` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getBlob()
  RETURN BLOB;
```

This function returns the `BLOB` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getBlob(content OUT VARCHAR2)
  RETURN BLOB;
```

This function returns the `BLOB` located at the address specified by the URL in the specified character set.

Syntax

```
FUNCTION getBlob(csid IN NUMBER)
  RETURN BLOB;
```

Table 307-2 UriType Supertype GETBLOB Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.
`csid`	`(IN)`	Character set id of the document. Must be a valid Oracle id and greater than `0`; otherwise returns an error

307.1.2 UriType Supertype GETCLOB

This function returns the `CLOB` located at the address specified by the URL.

This function can be overridden in the subtype instances. This function returns either a permanent `CLOB` or a temporary `CLOB`. If a temporary `CLOB` is returned, it must be freed. The options are described below.

This function returns the `CLOB` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getClob()
  RETURN CLOB;
```

This function returns the `CLOB` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getClob(content OUT VARCHAR2)
  RETURN CLOB;
```

Table 307-3 UriType Supertype GETCLOB Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.

307.1.3 UriType Supertype GETCONTENTTYPE

This function returns the content type of the document pointed to by the URI.

This function can be overridden in the subtype instances. This function returns the content type as `VARCHAR2`.

Syntax

```
MEMBER FUNCTION getContentType()
RETURN VARCHAR2;
```

307.1.4 UriType Supertype GETEXTERNALURL

This function returns the URL, in escaped format, stored inside the `UriType` instance.

The subtype instances override this member function to provide additional semantics. For instance, the `HTTPURITYPE` function does not store the prefix `http://` in the URL itself. When generating the external URL, it appends the prefix and generates it. For this reason, use the `getExternalUrl` function or the `getUrl` function to get to the URL value instead of using the attribute present in the `UriType` instance.

Syntax

```
MEMBER FUNCTION getExternalUrl()
RETURN varchar2;
```

307.1.5 UriType Supertype GETURL

This function returns the URL, in non-escaped format, stored inside the `UriType` instance.

The subtype instances override this member function to provide additional semantics. For instance, the `HTTPURITYPE` function does not store the prefix `http://` in the URL itself. When generating the external URL, it appends the prefix and generates it. For this reason, use the `getExternalUrl` function or the `getUrl` function to get to the URL value instead of using the attribute present in the `UriType` instance.

Syntax

```
MEMBER FUNCTION getUrl()
RETURN varchar2;
```

307.1.6 UriType Supertype GETXML

This function returns the `XMLType` located at the address specified by the URL.

This function can be overridden in the subtype instances. The options are described below.

This function returns the `XMLType` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getXML()
  RETURN XMLType;
```

This function returns the `XMLType` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getXML(content OUT VARCHAR2)
  RETURN XMLType;
```

Table 307-4 UriType Supertype GETXML Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.

307.2 Summary of HTTPURITYPE Subtype Subprograms

The `HTTPURITYPE` is a subtype of the `UriType` that provides support for the HTTP protocol. This uses the `UTL_HTTP` package underneath to access the HTTP URLs. Proxy and secure wallets are not supported in this release.

Table 307-5 HTTPURITYPE Type Subprorgams

Method	Description
CREATEURI	Creates an instance of `HTTPURITYPE` from the given URI.
GETBLOB	Returns the `BLOB` located at the address specified by the URL.

Table 307-5 (Cont.) HTTPURITYPE Type Subprorgams

Method	Description
GETCLOB	Returns the CLOB located at the address specified by the URL.
GETCONTENTTYPE	Returns the content type of the document pointed to by the URI.
GETEXTERNALURL	Returns the URL, in escaped format, stored inside the UriType instance.
GETURL	Returns the URL, in non-escaped format, stored inside the UriType instance.
GETXML	Returns the XMLType located at the address specified by the URL
HTTPURITYPE	Creates an instance of HTTPURITYPE from the given URI.

307.2.1 HttpUriType Subtype CREATEURI

This static function constructs a HTTPURITYPE instance. The HTTPURITYPE instance does not contain the prefix http:// in the stored URL.

Syntax

```
STATIC FUNCTION createUri(
   url IN varchar2)
RETURN HTTPURITYPE;
```

Table 307-6 HttpUriType Subtype CREATEURI Parameters

Parameter	IN / OUT	Description
url	(IN)	The URL string containing a valid HTTP URL; escaped format.

307.2.2 HttpUriType Subtype GETBLOB

This function returns the BLOB located at the address specified by the HTTP URL.

The subprograms of the URI Subtype ATTPURITYPE member subprogram GETBLOB are described below.

This function returns the BLOB located at the address specified by the HTTP URL.

Syntax

```
MEMBER FUNCTION getBlob()
  RETURN BLOB;
```

This function returns the BLOB located at the address specified by the HTTP URL and the content type.

Syntax

```
MEMBER FUNCTION getBlob(content OUT VARCHAR2)
  RETURN BLOB;
```

This function returns the BLOB located at the address specified by the URL in the specified character set.

Syntax

```
FUNCTION getBlob(csid IN NUMBER)
  RETURN BLOB;
```

Table 307-7 HttpUriType Subtype GETBLOB Parameters

Parameter	IN / OUT	Description
content	(OUT)	Content type of the document to which URI is pointing.
csid	(IN)	Character set id of the document. Must be a valid Oracle id and greater than 0; otherwise returns an error.

307.2.3 HttpUriType Subtype GETCLOB

This function returns the CLOB located by the HTTP URL address. If a temporary CLOB is returned, it must be freed.

The subprograms of the URI Subtype HTTPURITYPE member subprogram GETCLOB are described below.

Returns the CLOB located at the address specified by the HTTP URL.

Syntax

```
MEMBER FUNCTION getClob()
  RETURN CLOB;
```

Returns the CLOB located at the address specified by the HTTP URL and the content type.

Syntax

```
MEMBER FUNCTION getClob(content OUT VARCHAR2)
  RETURN CLOB;
```

Table 307-8 HttpUriType Subtype GETCLOB Parameters

Parameter	IN / OUT	Description
content	(OUT)	Content type of the document to which URI is pointing.

307.2.4 HttpUriType Subtype GETCONTENTTYPE

This function returns the content type of the document pointed to by the URI.

Syntax

```
MEMBER FUNCTION getContentType()
RETURN VARCHAR2;
```

307.2.5 HttpUriType Subtype GETEXTERNALURL

This function returns the URL, in escaped format, stored inside the `HTTPURITYPE` instance. The subtype instances override this member function.

The `HTTPURITYPE` function does not store the prefix `http://`, but generates it for the external URL.

Syntax

```
MEMBER FUNCTION getExternalUrl()
RETURN varchar2;
```

307.2.6 HttpUriType Subtype GETURL

This function returns the URL, in non-escaped format, stored inside the `HTTPURITYPE` instance.

Syntax

```
MEMBER FUNCTION getUrl()
RETURN varchar2;
```

307.2.7 HttpUriType Subtype GETXML

This function returns the `XMLType` located at the address specified by the URL. An error is thrown if the address does not point to a valid XML document.

The subprograms of the URI Subtype HttpUriType member subprogram GETXML are described below.

This function returns the `XMLType` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getXML()
  RETURN XMLType;
```

This function returns the `XMLType` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getXML(content OUT VARCHAR2)
  RETURN XMLType;
```

Table 307-9 HttpUriType Subtype GETXML Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.

307.2.8 HttpUriType Subtype HTTPURITYPE

This constructs a `HTTPURITYPE` instance. The `HTTPURITYPE` instance does not contain the prefix `http://` in the stored URL.

Syntax

```
CONSTRUCTOR FUNCTION HTTPURITYPE(
   url IN VARCHAR2);
```

Table 307-10 HttpUriType Subtype HTTPURITYPE Parameters

Parameter	IN / OUT	Description
`url`	`(IN)`	The URL string containing a valid HTTP URL. The URL string is expected in escaped format. For example, non-url characters are represented as the hexadecimal value for the UTF-8 encoding of those characters.

307.3 Summary of DBURITYPE Subtype Subprogams

The `DBURITYPE` is a subtype of the UriType that provides support for DBUri-refs. A DBUri-ref is an intra-database URL that can be used to reference any row or row-column data in the database.

The URL is specified as an XPath expression over a XML visualization of the database. The schemas become elements which contain tables and views. These tables and views further contain the rows and columns inside them.

Table 307-11 DBURITYPE Type Subprograms

Method	Description
CREATEURI	Constructs a `DBURITYPE` instance.
DBURITYPE	Creates an instance of `DBURITYPE` from the given URI.
GETBLOB	Returns the `BLOB` located at the address specified by the `DBURITYPE` instance.
GETCLOB	Returns the `CLOB` located at the address specified by the `DBURITYPE` instance.
GETCONTENTTYPE	Returns the content type of the document pointed to by the URI.
GETEXTERNALURL	Returns the URL, in escaped format, stored inside the `DBURITYPE` instance.
GETURL	Returns the URL, in non-escaped format, stored inside the `DBURITYPE` instance.
GETXML	Returns the `XMLType` located at the address specified by the URL

307.3.1 DBUriType Subtype CREATEURI

This static function constructs a DBURITYPE instance. Parses the URL given and creates a DBURITYPE instance.

Syntax

```
STATIC FUNCTION createUri(
   url IN varchar2)
RETURN DBURITYPE;
```

Table 307-12 DBUriType Subtype CREATEURI Parameters

Parameter	IN / OUT	Description
url	(IN)	The URL string, in escaped format, containing a valid DBURITYPE.

307.3.2 DBUriType Subtype DBURITYPE

This constructs a DBURITYPE instance.

Syntax

```
CONSTRUCTOR FUNCTION DBURITYPE(
   url IN varchar2);
```

Table 307-13 DBUriType Subtype DBURITYPE Parameters

Parameter	IN / OUT	Description
url	(IN)	The URL string containing a valid DBURITYPE. The URL string is expected in escaped format. For example, non-URL characters are represented as the hexadecimal value for the UTF-8 encoding of those characters.

307.3.3 DBUriType Subtype GETBLOB

This function returns the BLOB located at the address specified by the URL.

The subprograms of the URI Subtype DBURITYPE member subprogram GETBLOB are described below. The options are described in the following table.

This function returns the BLOB located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getBlob()
  RETURN BLOB;
```

This function returns the BLOB located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getBlob(content OUT VARCHAR2)
  RETURN BLOB;
```

This function returns the BLOB located at the address specified by the URL in the specified character set.

Syntax

```
FUNCTION getBlob(csid IN NUMBER)
  RETURN BLOB;
```

Table 307-14 DBUriType Subtype GETBLOB Parameters

Parameter	IN / OUT	Description
content	(OUT)	Content type of the document to which URI is pointing.
csid	(IN)	Character set id of the document. Must be a valid Oracle id and greater than 0; otherwise returns an error.

307.3.4 DBUriType Subtype GETCLOB

This function returns the CLOB located at the address specified by the DBURITYPE instance.

If a temporary CLOB is returned, it must be freed. The document returned may be an XML document or a text document. When the DBUri-ref identifies an element in the XPath, the result is a well-formed XML document. On the other hand, if it identifies a text node, then what is returned is only the text content of the column or attribute. The options are described below.

The following function returns the CLOB located at the address specified by the DBURITYPE instance.

Syntax

```
MEMBER FUNCTION getClob()
  RETURN CLOB;
```

The following function returns the CLOB located at the address specified by the DBURITYPE instance and the content type.

Syntax

```
MEMBER FUNCTION getClob(content OUT VARCHAR2)
RETURN CLOB;
```

Table 307-15 DBUriType Subtype GETCLOB Parameters

Parameter	IN / OUT	Description
content	(OUT)	Content type of the document to which URI is pointing.

307.3.5 DBUriType Subtype GETCONTENTTYPE

This function returns the content type of the document pointed to by the URI.

Syntax

```
MEMBER FUNCTION getContentType()
RETURN VARCHAR2;
```

307.3.6 DBUriType Subtype GETEXTERNALURL

This function returns the URL, in escaped format, stored inside the `DBURITYPE` instance. The `DBUri` servlet URL that processes the `DBURITYPE` has to be appended before using the escaped URL in Web pages.

Syntax

```
MEMBER FUNCTION getExternalUrl()
RETURN varchar2;
```

307.3.7 DBUriType Subtype GETURL

This function returns the URL, in non-escaped format, stored inside the `DBURITYPE` instance.

Syntax

```
MEMBER FUNCTION getUrl()
RETURN varchar2;
```

307.3.8 DBUriType Subtype GETXML

This function returns the `XMLType` located at the address specified by the URL.

The subprograms of the URI Subtype DBPURITYPE member subprogram GETXML are described below. The options are described in the following table.

This function returns the `XMLType` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getXML()
  RETURN XMLType;
```

This function returns the `XMLType` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getXML(content OUT VARCHAR2)
  RETURN XMLType;
```

Table 307-16 DBUriType Subtype GETXML Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.

307.4 Summary of XDBURITYPE Subtype Subprograms

`XDBURITYPE` is a new subtype of URIType. It provides a way to expose documents in the Oracle XML DB hierarchy as URIs that can be embedded in any URIType column in a table.

The URL part of the URI is the hierarchical name of the XML document it refers to. The optional fragment part uses the XPath syntax, and is separated from the URL part by '#'. The more general XPointer syntax for specifying a fragment is not currently supported.

Table 307-17 XDBURITYPE Type Subprograms

Method	Description
CREATEURI	Returns the `UriType` corresponding to the specified URL.
GETBLOB	Returns the `BLOB` corresponding to the contents of the document specified by the `XDBURITYPE` instance.
GETCLOB	Returns the `CLOB` corresponding to the contents of the document specified by the `XDBURITYPE` instance.
GETCONTENTTYPE	Returns the content type of the document pointed to by the URI.
GETEXTERNALURL	Returns the URL, in escaped format, stored inside the `XDBURITYPE` instance.
GETURL	Returns the URL, in non-escaped format, stored inside the `XDBURITYPE` instance.
GETXML	Returns the `XMLType` corresponding to the contents of the document specified by the URL.
XDBURITYPE	Creates an instance of `XDBURITYPE` from the given URI.

307.4.1 XDBUriType Subtype CREATEURI

This static function constructs a `XDBURITYPE` instance. It parses the given URL and creates a `XDBURITYPE` instance.

Syntax

```
STATIC FUNCTION createUri(
   url IN varchar2)
RETURN XDBURITYPE
```

Parameter	IN / OUT	Description
`url`	`(IN)`	The URL string, in escaped format, containing a valid `XDBURITYPE`.

307.4.2 XDBUriType Subtype GETBLOB

This function returns the `BLOB` located at the address specified by the `XDBURITYPE` instance.

The subprograms of the URI Subtype XDBURITYPE member subprogram GETBLOB are described below. The options are described in the following table.

This function returns the `BLOB` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getBlob()
  RETURN BLOB;
```

This function returns the `BLOB` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getBlob(content OUT VARCHAR2)
  RETURN BLOB;
```

This function returns the `BLOB` located at the address specified by the URL in the specified character set.

Syntax

```
FUNCTION getBlob(csid IN NUMBER)
  RETURN BLOB;
```

Table 307-18 XDBUriType Subtype GETBLOB Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.
`csid`	`(IN)`	Character set id of the document. Must be a valid Oracle id and greater than `0`; otherwise returns an error.

307.4.3 XDBUriType Subtype GETCLOB

This function returns the `CLOB` located at the address specified by the `XDBURITYPE` instance. If a temporary `CLOB` is returned, it must be freed.

The subprograms of the URI Subtype XDBURITYPE member subprogram GETCLOB are described below. The options are described in the following table.

Returns the `CLOB` located at the address specified by the `XDBUriType` instance.

Syntax

```
MEMBER FUNCTION getClob()
  RETURN CLOB;
```

Returns the `CLOB` located at the address specified by the `XDBUriType` instance and the content type.

Syntax

```
MEMBER FUNCTION getClob(content OUT VARCHAR2)
  RETURN CLOB;
```

Table 307-19 XDBUriType Subtype GETCLOB Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.

307.4.4 XDBUriType Subtype GETCONTENTTYPE

This function returns the content type of the document pointed to by the URI. This function returns the content type as `VARCHAR2`.

Syntax

```
MEMBER FUNCTION getContentType()
RETURN VARCHAR2;
```

307.4.5 XDBUriType Subtype GETEXTERNALURL

This function returns the URL, in escaped format, stored inside the `XDBURITYPE` instance.

Syntax

```
MEMBER FUNCTION getExternalUrl()
RETURN varchar2;
```

307.4.6 XDBUriType Subtype GETURL

This function returns the URL, in non-escaped format, stored inside the `XDBURITYPE` instance.

Syntax

```
MEMBER FUNCTION getUrl()
RETURN varchar2;
```

307.4.7 XDBUriType Subtype GETXML

This function returns the `XMLType` located at the address specified by the URL.

The subprograms of the URI Subtype XDBURITYPE member subprogram GETXML are described below.The options are described in the following table.

This function returns the `XMLType` located at the address specified by the URL.

Syntax

```
MEMBER FUNCTION getXML()
  RETURN XMLType;
```

This function returns the `XMLType` located at the address specified by the URL and the content type.

Syntax

```
MEMBER FUNCTION getXML(content OUT VARCHAR2)
  RETURN XMLType;
```

Table 307-20 XDBUriType Subtype GETXML Parameters

Parameter	IN / OUT	Description
`content`	`(OUT)`	Content type of the document to which URI is pointing.

307.4.8 XDBUriType subtype XDBURITYPE

This constructs a `XDBURITYPE` instance.

Syntax

```
CONSTRUCTOR FUNCTION XDBURITYPE(
   url      IN   VARCHAR2,
   flags    IN   RAW := NULL)
 RETURN self AS RESULT;
```

Table 307-21 XDBUriType subtype XDBURITYPE Parameters

Parameter	IN / OUT	Description
`url`	`(IN)`	The URL string containing a valid `XDBUriType`. The URL string is expected in escaped format. For example, non-URL characters are represented as the hexadecimal value for the UTF-8 encoding of those characters.
`flags`	`(IN)`	Possible values are: • `1` - Expand all `XInclude` elements before returning the result contents. If any `XInclude` element cannot be successfully resolved according to the `XInclude` fallback semantics, then an error is raised. • `2` - Indicates that any errors during document retrieval should be suppressed • `3` - Both flag bits (`1`, `2`) are enabled

307.5 Summary of URIFACTORY Package Subprograms

The `UriFactory` package contains factory methods that can be used to generate the appropriate instance of the URI types without having to hard code the implementation in the program.

The `UriFactory` package also provides the ability to register new subtypes of the `UriType` to handle various other protocols. For example, you can invent a new protocol `ecom://` and define a subtype of the `UriType` to handle that protocol and register it with `UriFactory`. After that any factory method would generate the new subtype instance if it sees the `ecom://` prefix.

Table 307-22 URIFACTORY Type Subprograms

Method	Description
GETURI	Returns the correct URL handler for the given URL string.
ESCAPEURI	Returns a URL in escaped format.
UNESCAPEURI	Returns a URL in unescaped format.
REGISTERURLHANDLER	Registers a particular type name for handling a particular URL.
UNREGISTERURLHANDLER	Unregisters a URL handler.

307.5.1 UriFactory Package GETURI

This factory method returns the correct URI handler for the given URI string.

It returns a subtype instance of the `UriType` that can handle the protocol. By default, it always creates an `XDBURITYPE` instance, if it cannot resolve the URL. A URL handler can be registered for a particular prefix using the REGISTERURLHANDLER function. If the prefix matches, GETURI would then use that subtype.

Syntax

```
FUNCTION getUri(
   url IN Varchar2)
RETURN UriType;
```

Table 307-23 UriFactory Package GETURI Parameters

Parameter	IN / OUT	Description
`uri`	`(IN)`	The URL string, in escaped format, containing a valid HTTP URL.

307.5.2 UriFactory Package ESCAPEURI

This function returns a URL in escaped format.

The subtype instances override this member function to provide additional semantics. For instance, the `HTTPURITYPE` does not store the prefix `http://` in the URL itself. When generating the external URL, it appends the prefix and generates it. For this reason, use the GETEXTERNALURL function or the GETURI function to get to the URL value instead of using the attribute present in the `UriType`.

Syntax

```
MEMBER FUNCTION escapeUri()
RETURN varchar2;
```

Table 307-24 UriFactory Package ESCAPEURI Parameters

Parameter	IN / OUT	Description
`url`	`(IN)`	The URL string to be returned in escaped format.

307.5.3 UriFactory Package UNESCAPEURI

This function returns a URL in unescaped format.

This function is the reverse of the ESCAPEURI function. This function scans the string and converts any non-URL hexadecimal characters into the equivalent UTF-8 characters. Since the return type is a `VARCHAR2`, the characters would be converted into the equivalent characters as defined by the database character set.

Syntax

```
FUNCTION unescapeUri()
RETURN varchar2;
```

Table 307-25 UriFactory Package UNESCAPEURI Parameters

Parameter	IN / OUT	Description
`url`	(IN)	The URL string to be returned in unescaped format.

307.5.4 UriFactory Package REGISTERURLHANDLER

This package registers a particular type name for handling a particular URL.

The type specified must be valid and must be a subtype of the `UriType` or one of its subtypes. It must also implement the `createUri` static member function. This function is called by the GETURI function to generate an instance of the type. The `stripprefix` parameter indicates that the prefix must be stripped off before calling this function.

Syntax

```
PROCEDURE registerUrlHandler(
   prefix IN varchar2,
   schemaName IN varchar2,
   typename IN varchar2,
   ignoreCase IN boolean := true,
   stripprefix IN boolean := true);
```

Table 307-26 UriFactory Package REGISTERURLHANDLER Parameters

Parameter	IN / OUT	Description
`prefix`	(IN)	The prefix to handle; for example, `http://`.
`schemaName`	(IN)	Name of the schema where the type resides; case sensitive.
`typename`	(IN)	The name of the type to handle the URL; case sensitive.
`ignoreCase`	(IN)	Ignore case when matching prefixes.
`stripprefix`	(IN)	Strip prefix before generating the instance of the type.

307.5.5 UriFactory Package UNREGISTERURLHANDLER

This procedure unregisters a URL handler. This only unregisters user registered handler prefixes and not predefined system prefixes such as `http://`.

Syntax

```
PROCEDURE unregisterUrlHandler(
   prefix IN VARCHAR2);
```

Table 307-27 UriFactory Package UNREGISTERURLHANDLER Parameters

Parameter	IN / OUT	Description
`prefix`	(IN)	The prefix to be unregistered.

308
JMS Types

PL/SQL users can use the `DBMS_AQ` package to enqueue and dequeue messages from JMS queues.

The JMS types member and static functions and procedures in this chapter are needed to populate JMS messages for enqueuing or to interpret a dequeued JMS message.

This chapter contains these topics:

- Overview
- Security Model
- Java Versus PL/SQL Datatypes
- More on Bytes_ Stream and Map Messages
- Upcasting and Downcasting Between General and Specific Messages
- JMS Types Error Reporting
- Oracle JMS Type Constants
- JMS Types Error Reporting
- Oracle JMS Type Constants
- CONVERT_JMS_SELECTOR
- Summary of JMS Types

308.1 JMS Types Overview

Java Message Service (JMS) is a well known public standard interface for accessing messaging systems. Oracle JMS (OJMS) implements JMS based on Oracle Advanced Queuing (AQ) and a relational database system (RDBMS). Messages are stored in queues as OJMS specific ADTs. Java clients use OJMS packages to enqueue, dequeue, and manipulate these messages.

PL/SQL users, on the other hand, use the `DBMS_AQ` package to enqueue and dequeue JMS messages and the member functions in this chapter to populate and interpret them. Oracle Database Advanced Queuing offers such member functions for the following JMS ADTs:

- `aq$_jms_header`
- `aq$_jms_message`
- `aq$_jms_text_message`
- `aq$_jms_bytes_message`
- `aq$_jms_map_message`
- `aq$_jms_stream_message`

In addition to these populating and interpreting member functions, Oracle Database Advanced Queuing offers:

- Casting between `aq$_jms_message` and other message ADTs.
- PL/SQL stored procedures for converting JMS selectors to equivalent Oracle Database Advanced Queuing rules

308.2 JMS Types Security Model

`PUBLIC` is granted `EXECUTE` privilege in these JMS types.

- SYS.AQ$_JMS_MESSAGE Type
- SYS.AQ$_JMS_TEXT_MESSAGE Type
- SYS.AQ$_JMS_BYTES_MESSAGE Type
- SYS.AQ$_JMS_MAP_MESSAGE Type
- SYS.AQ$_JMS_STREAM_MESSAGE Type
- SYS.AQ$_JMS_OBJECT_MESSAGE Type
- SYS.AQ$_JMS_NAMEARRAY Type
- SYS.AQ$_JMS_VALUE Type
- SYS.AQ$_JMS_EXCEPTION Type

308.3 Java Versus PL/SQL Datatypes

Datatypes do not map one-to-one between PL/SQL and Java.

Some Java types, such as `BYTE` and `SHORT`, are not present in PL/SQL. PL/SQL type `INT` was chosen to represent these types. If a PL/SQL `INT` value intended to hold a Java `BYTE` or `SHORT` value exceeds the corresponding range Java enforces, an out-of-range error is thrown.

Other Java types have more than one counterpart in PL/SQL with different capabilities. A Java String can be represented by both `VARCHAR2` and `CLOB`, but `VARCHAR2` has a maximum limit of 4000 bytes. When retrieving `TEXT` data from map, stream, and bytes message types, a `CLOB` is always returned. When updating the map, stream and bytes message types, users can submit either a `VARCHAR2` or `CLOB`.

Similarly, a Java `BYTE ARRAY` can be represented by both `RAW` and `BLOB`, with `RAW` having a maximum size of 32767. When retrieving `BYTE ARRAY` data from map, stream, and bytes message types, a `BLOB` is always returned. When updating the map, stream and bytes message types, users can submit either a `RAW` or `BLOB`.

> **See Also:**
>
> JMS specification 3.11.3, Conversion Provided by StreamMessage and MapMessage

New JMS Support in Oracle Database 10*g*

In Oracle Database 10*g*, a new `AQ$_JMS_VALUE` ADT has been added in the `SYS` schema for OJMS PL/SQL users. It is specifically used to implement the `read_object` procedure of `aq$_jms_stream_message` and `get_object` procedure of `aq$_jms_map_message`, to mimic the Java general object class `Object`. `AQ$_JMS_VALUE` ADT can represent any datatype that JMS `StreamMessage` and `MapMessage` can hold.

The collection ADT `AQ$_JMS_NAMEARRAY` was added for the `getNames` method of `MapMessage`. It holds an array of names.

In this release the ADT `AQ$_JMS_EXCEPTION` was added to represent a Java exception thrown in an OJMS JAVA stored procedure on the PL/SQL side. Now you can retrieve a Java exception thrown by an OJMS stored procedure and analyze it on the PL/SQL side.

308.4 More on Bytes, Stream and Map Messages

Oracle uses Java stored procedure to implement some of the procedures of `AQ$_MAP_MESSAGE`, `AQ$_JMS_STREAM_MESSAGE`, and `AQ$_JMS_BYTES_MESSAGE` types. These types have some common functionality that are different from `AQ$_JMS_TEXT_MESSAGE` type. This section discusses this common functionality.

This section contains these topics:

- Using Java Stored Procedures to Encode and Decode Oracle Database Advanced Queuing Messages
- Initialize the Jserv Static Variable
- Get the Payload Data Back to PL/SQL
- Garbage Collect the Static Variable
- Use a Message Store: A Static Variable Collection
- Typical Calling Sequences
- Read-Only and Write-Only Modes Enforced for Stream and Bytes Messages
- Differences Between Bytes and Stream Messages
- Getting and Setting Bytes, Map, and Stream Messages as RAW Bytes

Using Java Stored Procedures to Encode and Decode Oracle Database Advanced Queuing Messages

The major difference between map, stream, bytes, and other messages is that the message payload is encoded as a byte stream by JAVA. Retrieving and updating these payloads in PL/SQL therefore requires Oracle JAVA stored procedures.

A message payload is stored in two places during processing. On the PL/SQL side it is stored as the data members of a JMS message ADT, and on the Jserv side it is stored as a static variable. (Jserv is the JVM inside Oracle Database.) When the payload is processed, the payload data is first transformed to a static variable on the Jserv side. Once the static variable is initialized, all later updates on the message payload are performed on this static variable. At the end of processing, payload data is flushed back to the PL/SQL side.

Oracle provides member procedures that maintain the status of the Jserv static variable and enforce rules when calling these member procedures. These procedures are in the following ADTs:

- `aq$_jms_bytes_message`
- `aq$_jms_map_message`
- `aq$_jms_stream_message`

Initialize the Jserv Static Variable

Before you make any other calls to manipulate the payload data, the Jserv static variable must be properly initialized. This is done by calling the `prepare` or `clear_body` procedure. The `prepare` procedure uses the payload data in PL/SQL ADTs to initialize the static variable, while `clear_body` initializes the static variable to an empty payload (empty hashtable or stream).

> **Note:**
>
> It is important to call the `prepare` or `clear_body` procedure before any other calls to properly initialize the Jserv static variables. Usually these two methods are called once at the beginning. But they can be called multiple times for one message. Any call of these two methods without first calling the `flush` procedure wipes out all updates made to the messages.

Get the Payload Data Back to PL/SQL

Calling the `flush` procedure synchronizes changes made to the Jserv static variable back to the PL/SQL ADTs. The `flush` call is required when you want the changes made to be reflected in the ADT payload. It is important to synchronize the changes back to the ADT, because it is the ADT payload that matters.

Garbage Collect the Static Variable

The `clean` procedure forces garbage collection of the static variable. It is there to do cleanup and free JVM memory. You can avoid memory leaks by doing it immediately after finishing processing the message.

Use a Message Store: A Static Variable Collection

Instead of a single static variable, Oracle uses a collection of static variables to process the message payload on the Jserv side. This collection is called the message store. Each map, bytes, or stream message type has its own message store within one session.

Oracle uses the operation ID parameter to locate the correct static variable to work on within the message store. Initialization calls such as `prepare` and `clear_body` give users an operation ID, which is used in later message access.

After users complete message processing, they must call the `clean` procedure with the operation ID to clean up the message store. This avoids possible memory leaks. The `clean_all` static procedures of message ADTs `aq$_jms_bytes_message`, `aq$_jms_map_message`, and `aq$_jms_stream_message` clean up all static variables of their corresponding message stores.

Typical Calling Sequences

This section describes typical procedures for retrieving and populating messages.

Here is a typical procedure for retrieving messages

1. Call `prepare` for a message.

   This call also gives you an operation ID if you do not specify one.
2. Call multiple retrieving procedures with the provided operation ID.

3. Call the `clean` procedure with the provided operation ID.

Here is a typical procedure for populating messages:

1. Call `clear_body` for a message.

   For `aq$_jms_map_message`, you can also call `prepare` to update the message based on the existing payload. This call also gives you an operation ID if you do not specify one.

2. Call multiple updating procedures with the provided operation ID.
3. Call the `flush` method with the provided operation ID.
4. Call the `clean` procedure with the provided operation ID.

Read-Only and Write-Only Modes Enforced for Stream and Bytes Messages

According to the JMS specification, when a message is received, its body is read-only. Users can call the `clear_body` method to make the body writable. This method erases the current message body and sets the message body to be empty.

The OJMS JAVA API follows the rule set by JMS specification. In updating the JMS message ADTs in PL/SQL, however, Oracle enforces the rule selectively:

- Map messages

  The restriction is relaxed, because adding more entries on top of a existing map payload is a convenient way for users to update the payload. Therefore there are no read-only or write-only modes for map messages.

- Stream and bytes messages

  The restriction is not relaxed, because these payloads use a stream when reading and writing data. It is difficult to update the payload while in the middle of a stream. Oracle enforces read-only and write-only modes in processing stream and bytes message payloads. Calling the `prepare` procedure initializes the message payload in read-only mode. Calling the `clear_body` procedure initializes the message payload in write-only mode.

  Calling the `reset` procedure resets the pointer to the beginning of the stream and switches the mode from write-only to read-only. The `reset` procedure keeps the updates made to the message payload in the Jserv static variable.

  The `prepare` procedure, on the other hand, overwrites the message payload in the Jserv static variable with the payload in the PL/SQL ADT.

  Oracle provides member function `get_mode` for users to query the mode.

Differences Between Bytes and Stream Messages

Member functions of bytes messages are not exactly the same as those of stream messages. Stream messages are encoded using Java `ObjectOutputStream` and bytes messages are encoded using Java `DataOutputStream`. In stream messages each primitive type is written and read as a Java Object, but in a bytes message they are written and read as raw bytes according to the encoding mechanism of `DataOutputStream`.

For stream messages, the `read_bytes` method works on a stream of bytes to the end of the byte array field written by the corresponding `write_bytes` method. The `read_bytes` method of bytes message works on a stream of bytes to the end of the whole byte stream. This is why the `read_bytes` member procedure of `aq$_bytes_message` also requires a `length` parameter to tell how long it is to read.

You will not see a type conversion error raised by bytes message, because bytes messages do not support type conversion.

Methods `get_unsigned_byte` and `get_unsigned_short` are available for bytes messages, but not for stream messages. This is because stream messages read Java objects, and there are no Java objects as unsigned bytes or unsigned shorts.

Methods `read_string` and `write_string` methods are not available for bytes messages. The bytes message ADT must enforce some character encoding. It has methods `read_utf` and `write_utf` which support `utf-8` encoding.

> **Note:**
>
> All data written by bytes messages use `DataOutputStream` as the basis. See JDK API documentation JavaSoft.com for details on how the data is encoded into bytes.

Getting and Setting Bytes, Map, and Stream Messages as RAW Bytes

The payloads of bytes, map, and stream message types are stored as either `RAW` or `BLOB` in the database. In this release Oracle Database Advanced Queuing provides the following member functions to set and get these payloads as raw bytes without interpreting them:

```
set_bytes(payload IN BLOB)
set_bytes(payload IN RAW)
get_bytes(payload OUT BLOB)
get_bytes(payload OUT RAW)
```

These functions were provided for bytes messages in Oracle9*i* Release 2 (9.2).

308.5 Upcasting and Downcasting Between General and Specific Messages

OJMS ADT `aq$_jms_message` is used to represent a general message, so that different types of messages can reside on the same Oracle Database Advanced Queuing queue. Oracle Database Advanced Queuing supports retrieving and populating of `aq$_jms_message` by supporting upcasting and downcasting between this ADT and ADTs of specific message types.

To read an `aq$_jms_message`, you must first downcast it to a specific message type according to its `message_type` field

To populate an `aq$_jms_message`, you must first populate a specific message and upcast it to `aq$_jms_message`. This avoids copying all member functions of other specific message ADTs to this ADT. It also guarantees that the manipulation of this ADT is consistent with other specific message ADTs.

308.6 JMS Types Error Reporting

This table lists Oracle JMS types related errors.

Table 308-1 Oracle JMS Types Errors

ORA error number	dbms_jms_plsql package constants	Explanation
ORA-24190	ERROR_DATA_OVERFLOW	The payload data exceeds the size that an out parameter can hold. For example, the get_text procedure with a VARCHAR2 parameter of aq$_jms_text_message or get_bytes procedure with a RAW parameter of aq$_jms_bytes_message.
ORA-24191	ERROR_PROP_NAME_EXIST	Setting a property that is previous set
ORA-24192	ERROR_PROP_NAME_NULL	Occurs when setting a property with null property name.
ORA-24193	ERROR_EXCEED_RANGE	PL/SQL number type exceeds the valid range of the respective Java type. For example set_byte_property, set_short_property of aq$_jms_head ADT; set_byte and set_short of aq$_jms_map_message ADT; write_byte and write_short of aq$_jms_stream_message and aq$_jms_bytes_message ADT.
ORA-24194	ERROR_TYPE_MISMATCH	The type conversion between the Java type of the retrieving method and the Java type of a field of the payload is not valid.
ORA-24195	ERROR_MAP_TOO_LARGE	The size of the map exceeds the aq$_jms_namearray ADT capacity. The current size limit is 1024. You can use the get_names function with offset and length parameters to retrieve the name array in multiple small chunks.
ORA-24196	ERROR_WRONG_MODE	The message payload is being accessed with a wrong access mode. For example, trying to read a message payload with write-only mode or trying to write a message payload with the read-only mode.
ORA-24197	ERROR_JAVA_EXCEPTION	ORA-24197 error is raised when a Java exception is raised that does not fit in any of the other error categories. You can use the get_exception static procedure of aq$_jms_map_message, aq$_jms_bytes_message, and aq$_jms_stream_message to retrieve the exception information last thrown by the Java stored procedure. A single static variable is used to store the last exception and is overwritten if another exception is thrown before you retrieve it. A new ADT aq$_jms_exception is created to represent the exception information on the PL/SQL side.
ORA-24198	ERROR_INVALID_ID	An invalid operation ID is being provided to access a message.
ORA-24199	ERROR_STORE_OVERFLOW	The number of messages (with the same type) that users are trying to manipulate exceeds the size of the message store on the Java stored procedure side. The current size of the store is 20. It unusual to need to manipulate more than 20 messages at the same time. A common mistake is to forget to call the clean procedure after using one message. The clean procedure frees the message slot for use by other messages attempting access.

308.7 Oracle JMS Type Constants

These constants can be useful when dealing with message type functions.

DBMS_AQ Package Constants

DBMS_AQ package constants specify different types of JMS messages. They are useful when dealing with general message types during upcasting and downcasting or constructing a general message with a specific message type:

```
JMS_TEXT_MESSAGE   CONSTANT BINARY_INTEGER;
JMS_BYTES_MESSAGE  CONSTANT BINARY_INTEGER;
JMS_STREAM_MESSAGE CONSTANT BINARY_INTEGER;
JMS_MAP_MESSAGE    CONSTANT BINARY_INTEGER;
JMS_OBJECT_MESSAGE CONSTANT BINARY_INTEGER;
```

SYS.DBMS_JMS_PLSQL Package Constants

`SYS.DBMS_JMS_PLSQL` package constants are new in Oracle Database 10*g*.

These constants specify the mode of message payload. They are useful when interpreting the mode of the message payload returned from the `get_mode` function:

```
MESSAGE_ACCESS_READONLY  CONSTANT PLS_INTEGER;
MESSAGE_ACCESS_WRITEONLY CONSTANT PLS_INTEGER;
```

These constants specify the ADT type of an Oracle Database Advanced Queuing queue. They are useful during the conversion of JMS selectors to Oracle Database Advanced Queuing rules:

```
DESTPLOAD_JMSTYPE CONSTANT PLS_INTEGER;
DESTPLOAD_USERADT CONSTANT PLS_INTEGER;
DESTPLOAD_ANYDATA CONSTANT PLS_INTEGER;
```

These constants specify the type of data that can be held by a `aq$_jms_value` type. They are useful when interpreting the `aq$_jms_value` returned by the `get_object` method of `AQ$_JMS_MAP_MESSAGE` or `read_object` method of `AQ$_JMS_STREAM_MESSAGE`:

```
DATA_TYPE_BYTE           CONSTANT PLS_INTEGER;
DATA_TYPE_SHORT          CONSTANT PLS_INTEGER;
DATA_TYPE_INTEGER        CONSTANT PLS_INTEGER;
DATA_TYPE_LONG           CONSTANT PLS_INTEGER;
DATA_TYPE_FLOAT          CONSTANT PLS_INTEGER;
DATA_TYPE_DOUBLE         CONSTANT PLS_INTEGER;
DATA_TYPE_BOOLEAN        CONSTANT PLS_INTEGER;
DATA_TYPE_CHARACTER      CONSTANT PLS_INTEGER;
DATA_TYPE_STRING         CONSTANT PLS_INTEGER;
DATA_TYPE_BYTES          CONSTANT PLS_INTEGER;
DATA_TYPE_UNSIGNED_BYTE  CONSTANT PLS_INTEGER;
DATA_TYPE_UNSIGNED_SHORT CONSTANT PLS_INTEGER;
```

These constants specify the error number of the ORA errors that can be raised by the functions of message type ADTs. They are useful in user error handlers:

```
ERROR_DATA_OVERFLOW   CONSTANT PLS_INTEGER := -24190;
ERROR_PROP_NAME_EXIST CONSTANT PLS_INTEGER := -24191;
ERROR_PROP_NAME_NULL  CONSTANT PLS_INTEGER := -24192;
ERROR_EXCEED_RANGE    CONSTANT PLS_INTEGER := -24193;
ERROR_TYPE_MISMATCH   CONSTANT PLS_INTEGER := -24194;
ERROR_MAP_TOO_LARGE   CONSTANT PLS_INTEGER := -24195;
ERROR_WRONG_MODE      CONSTANT PLS_INTEGER := -24196;
ERROR_JAVA_EXCEPTION  CONSTANT PLS_INTEGER := -24197;
ERROR_INVALID_ID      CONSTANT PLS_INTEGER := -24198;
ERROR_STORE_OVERFLOW  CONSTANT PLS_INTEGER := -24199;
```

308.8 CONVERT_JMS_SELECTOR

Oracle Database includes three stored procedures to help users convert JMS selectors into Oracle Database Advanced Queuing rules. These rules can be used in `ADD_SUBSCRIBER`

operations as subscriber rules or in DEQUEUE operations as dequeue conditions. These procedures are in the SYS.dbms_jms_plsql package.

Convert with Minimal Specification

The first procedure assumes the destination payload type is one of the JMS ADTs whose corresponding constant is dbms_jms_plsql.DESTPLOAD_JMSTYPE and also assumes that the J2EE compliant mode is true.

Syntax

```
Function convert_jms_selector(selector IN VARCHAR2) RETURN VARCHAR2
```

Returns

The converted Oracle Database Advanced Queuing rule or null if there is any conversion error.

Exceptions

ORA-24197 if the Java stored procedure throws an exception during execution.

Convert with Destination Payload Type Specified

The second procedure takes one more parameter: dest_pload_type. The conversion of a JMS selector to an Oracle Database Advanced Queuing rule happens only if this parameter is SYS.dbms_jms_plsql.DESTPLOAD_JMSTYPE or SYS.dbms_jms_plsql.DESTPLOAD_ANYDATA. The function returns exactly the same VARCHAR2 value as the selector parameter if the dest_pload_type parameter is SYS.dbms_jms_plsql.DESTPLOAD_USERADT. The function returns null if dest_pload_type parameter is none of these three constants.

This function assumes that the J2EE compliant mode is true.

Syntax

```
Function convert_jms_selector(
   selector IN VARCHAR2,
   dest_pload_type IN PLS_INTEGER)
RETURN VARCHAR2
```

Returns

The converted Oracle Database Advanced Queuing rule or null if there is any conversion error.

Exceptions

ORA-24197 if the Java stored procedure throws an exception during execution.

Convert with Destination Payload Type and Compliant Mode Specified

The third procedure takes a dest_pload_type parameter and a compliant parameter. The conversion of a JMS selector to an Oracle Database Advanced Queuing rule happens only if the dest_pload_type parameter is SYS.dbms_jms_plsql.DESTPLOAD_JMSTYPE or SYS.dbms_jms_plsql.DESTPLOAD_ANYDATA. The function returns exactly the same VARCHAR2 value as the selector parameter if the dest_pload_type parameter is SYS.dbms_jms_plsql.DESTPLOAD_USERADT. The function returns null if the dest_pload_type parameter is none of these three constants.

The compliant parameter controls if the conversion is in J2EE compliant mode or not. The noncompliant conversion of a JMS selector is for backward compatibility.

Syntax

```
Function convert_jms_selector(
   selector         IN  VARCHAR2,
   dest_pload_type  IN  PLS_INTEGER,
   compliant        IN  BOOLEAN )
```

Returns

The converted Oracle Database Advanced Queuing rule or null if there is any conversion error.

Exceptions

ORA-24197 if the Java stored procedure throws an exception during execution.

308.9 Summary of JMS Types

This lists shows the JMS types.

- SYS.AQ$_JMS_MESSAGE Type
- SYS.AQ$_JMS_TEXT_MESSAGE Type
- SYS.AQ$_JMS_BYTES_MESSAGE Type
- SYS.AQ$_JMS_MAP_MESSAGE Type
- SYS.AQ$_JMS_STREAM_MESSAGE Type
- SYS.AQ$_JMS_OBJECT_MESSAGE Type
- SYS.AQ$_JMS_NAMEARRAY Type
- SYS.AQ$_JMS_VALUE Type
- SYS.AQ$_JMS_EXCEPTION Type

308.9.1 SYS.AQ$_JMS_MESSAGE Type

This ADT type can represent any of five different JMS message types: text message, bytes message, stream message, map message, or object message. Queues created using this ADT can therefore store all five types of JMS messages.

This section contains these topics:

- CONSTRUCT Static Functions
- Cast Methods
- JMS Header Methods
- System Properties Methods
- User Properties Methods
- Payload Methods

Syntax

```
TYPE AQ$_JMS_MESSAGE AS OBJECT(
 header        aq$_jms_header,
 senderid      varchar2(100),
 message_type  INT,
 text_len      INT,
```

```
bytes_len     INT,
text_vc       varchar2(4000),
bytes_raw     raw(2000),
text_lob      clob,
bytes_lob     blob,
STATIC FUNCTION  construct (mtype      IN  INT)
  RETURN aq$_jms_message,
STATIC FUNCTION  construct (text_msg   IN  aq$_jms_text_message)
  RETURN aq$_jms_message,
STATIC FUNCTION  construct (bytes_msg  IN  aq$_jms_bytes_message)
  RETURN aq$_jms_message,
STATIC FUNCTION  construct (stream_msg IN  aq$_jms_stream_message)
  RETURN aq$_jms_message,
STATIC FUNCTION  construct (map_msg    IN  aq$_jms_map_message)
  RETURN aq$_jms_message,
STATIC FUNCTION  construct (object_msg IN  aq$_jms_object_message)
  RETURN aq$_jms_message,
MEMBER FUNCTION  cast_to_bytes_msg  RETURN aq$_jms_bytes_message,
MEMBER FUNCTION  cast_to_map_msg    RETURN aq$_jms_map_message,
MEMBER FUNCTION  cast_to_object_msg RETURN aq$_jms_object_message,
MEMBER FUNCTION  cast_to_stream_msg RETURN aq$_jms_stream_message,
MEMBER FUNCTION  cast_to_text_msg   RETURN aq$_jms_text_message,
MEMBER PROCEDURE set_replyto  (replyto IN sys.aq$_agent),
MEMBER PROCEDURE set_type     (type     IN  VARCHAR),
MEMBER PROCEDURE set_userid   (userid   IN  VARCHAR),
MEMBER PROCEDURE set_appid    (appid    IN  VARCHAR),
MEMBER PROCEDURE set_groupid  (groupid  IN  VARCHAR),
MEMBER PROCEDURE set_groupseq (groupseq IN  INT),
MEMBER FUNCTION  get_replyto  RETURN sys.aq$_agent,
MEMBER FUNCTION  get_type     RETURN VARCHAR,
MEMBER FUNCTION  get_userid   RETURN VARCHAR,
MEMBER FUNCTION  get_appid    RETURN VARCHAR,
MEMBER FUNCTION  get_groupid  RETURN VARCHAR,
MEMBER FUNCTION  get_groupseq RETURN INT,
MEMBER PROCEDURE clear_properties,
MEMBER PROCEDURE set_boolean_property (property_name IN VARCHAR,
  property_value IN BOOLEAN),
MEMBER PROCEDURE set_byte_property    (property_name IN VARCHAR,
  property_value IN INT),
MEMBER PROCEDURE set_double_property  (property_name IN VARCHAR,
  property_value IN DOUBLE PRECISION),
MEMBER PROCEDURE set_float_property   (property_name IN VARCHAR,
  property_value IN FLOAT),
MEMBER PROCEDURE set_int_property     (property_name IN VARCHAR,
  property_value IN INT),
MEMBER PROCEDURE set_long_property    (property_name IN VARCHAR,
  property_value IN NUMBER),
MEMBER PROCEDURE set_short_property   (property_name IN VARCHAR,
  property_value IN INT),
MEMBER PROCEDURE set_string_property  (property_name IN VARCHAR,
  property_value IN VARCHAR),
MEMBER FUNCTION get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN,
MEMBER FUNCTION get_byte_property    (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_double_property  (property_name IN VARCHAR)
  RETURN DOUBLE PRECISION,
MEMBER FUNCTION get_float_property   (property_name IN VARCHAR) RETURN FLOAT,
MEMBER FUNCTION get_int_property     (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_long_property    (property_name IN VARCHAR) RETURN NUMBER,
MEMBER FUNCTION get_short_property   (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_string_property  (property_name IN VARCHAR) RETURN VARCHAR,
MEMBER PROCEDURE set_text     (payload  IN  VARCHAR2),
MEMBER PROCEDURE set_text     (payload  IN  CLOB),
```

```
MEMBER PROCEDURE set_bytes    (payload  IN  RAW),
MEMBER PROCEDURE set_bytes    (payload  IN  BLOB),
MEMBER PROCEDURE get_text     (payload  OUT VARCHAR2),
MEMBER PROCEDURE get_text     (payload  OUT CLOB),
MEMBER PROCEDURE get_bytes    (payload  OUT RAW),
MEMBER PROCEDURE get_bytes    (payload  OUT BLOB));
```

CONSTRUCT Static Functions

There are six `CONSTRUCT` static functions in this type.

STATIC FUNCTION construct (mtype IN INT) RETURN aq$_jms_message
Creates an instance of `aq$_jms_message`, which can hold a specific type of JMS message (TextMessage, BytesMessage, MapMessage, StreamMessage or ObjectMessage). The message type of the created `aq$_jms_message` instance depends on the `mtype` parameter passed to the construct method. Once a message has been constructed, it can be used to store JMS messages of the type it has been constructed to hold.
The `mtype` parameter must be one of the following constants described in "Oracle JMS Type Constants":

```
DBMS_AQ.JMS_TEXT_MESSAGE
DBMS_AQ.JMS_BYTES_MESSAGE
DBMS_AQ.JMS_STREAM_MESSAGE
DBMS_AQ.JMS_MAP_MESSAGE
DBMS_AQ.JMS_OBJECT_MESSAGE
```

STATIC FUNCTION construct (text_msg IN aq$_jms_text_message) RETURN aq$_jms_message
Creates an `aq$_jms_message` from an `aq$_jms_text_message`.

STATIC FUNCTION construct (bytes_msg IN aq$_jms_bytes_message) RETURN aq$_jms_message;
Creates an `aq$_jms_message` from an `aq$_jms_bytes_message`.

STATIC FUNCTION construct (stream_msg IN aq$_jms_stream_message) RETURN aq$_jms_message;
Creates an `aq$_jms_message` from an `aq$_jms_stream_message`.

STATIC FUNCTION construct (map_msg IN aq$_jms_map_message) RETURN aq$_jms_message;
Creates an `aq$_jms_message` from an `aq$_jms_map_message`.

STATIC FUNCTION construct (object_msg IN aq$_jms_object_message) RETURN aq$_jms_message;
Creates an `aq$_jms_message` from an `aq$_jms_object_message`.

Cast Methods

cast_to_bytes_msg RETURN aq$_jms_bytes_message
Casts an `aq$_jms_message` to an `aq$_jms_bytes_message`. Returns an `aq$_jms_bytes_message` or null if the `message_type` attribute of the `aq$_jms_message` is not `DBMS_AQ.JMS_BYTES_MESSAGE`. This function raises ORA-24198 if the `message_type` field of the `aq$_jms_message` is not `DBMS_AQJMS.JMS_BYTES_MESSAGE`.

cast_to_map_msg RETURN aq$_jms_map_message
Casts an `aq$_jms_message` to an `aq$_jms_map_message`. Returns an `aq$_jms_map_message` or null if the `message_type` attribute of the `aq$_jms_message` is not `DBMS_AQ.JMS_MAP_MESSAGE`.

This function raises ORA-24198 if the `message_type` field of the `aq$_jms_message` is not `DBMS_AQJMS.JMS_MAP_MESSAGE`.

cast_to_object_msg RETURN aq$_jms_object_message
Casts an `aq$_jms_message` to an `aq$_jms_object_message`. Returns an `aq$_jms_object_message` or null if the `message_type` attribute of the `aq$_jms_message` is not `DBMS_AQ.JMS_OBJECT_MESSAGE`. This function raises ORA-24198 if the `message_type` field of the `aq$_jms_message` is not `DBMS_AQJMS.JMS_OBJECT_MESSAGE`.

cast_to_stream_msg RETURN aq$_jms_stream_message
Casts an `aq$_jms_message` to an `aq$_jms_stream_message`. Returns an `aq$_jms_stream_message` or null if the `message_type` attribute of the `aq$_jms_message` is not `DBMS_AQ.JMS_STREAM_MESSAGE`. This function raises ORA-24198 if the `message_type` field of the `aq$_jms_message` is not `DBMS_AQJMS.JMS_STREAM_MESSAGE`.

cast_to_text_msg RETURN aq$_jms_text_message
Casts an `aq$_jms_message` to an `aq$_jms_text_message`. Returns an `aq$_jms_text_message` or null if the `message_type` attribute of the `aq$_jms_message` is not `DBMS_AQ.JMS_TEXT_MESSAGE`. This function raises ORA-24198 if the `message_type` field of the `aq$_jms_message` is not `DBMS_AQJMS.JMS_TEXT_MESSAGE`.

JMS Header Methods

set_replyto (replyto IN sys.aq$_agent)
Sets the `replyto` parameter, which corresponds to `JMSReplyTo`.

get_replyto RETURN sys.aq$_agent
Returns `replyto`, which corresponds to `JMSReplyTo`.

set_type (type IN VARCHAR)
Sets the JMS type, which can be any text and corresponds to `JMSType`.

get_type RETURN VARCHAR
Returns `type`, which corresponds to `JMSType`.

System Properties Methods

set_userid (userid IN VARCHAR)
Sets `userid`, which corresponds to `JMSXUserID`.

set_appid (appid IN VARCHAR)
Sets `appid`, which corresponds to `JMSXAppID`.

set_groupid (groupid IN VARCHAR)
Sets `groupid`, which corresponds to `JMSXGroupID`.

set_groupseq (groupseq IN INT)
Sets `groupseq`, which corresponds to `JMSXGroupSeq`.

get_userid RETURN VARCHAR
Returns `userid`, which corresponds to `JMSXUserID`.

get_appid RETURN VARCHAR
Returns `appid`, which corresponds to `JMSXAppID`.

get_groupid RETURN VARCHAR
Returns `groupid`, which corresponds to `JMSXGroupID`.

get_groupseq RETURN VARCHAR
Returns `groupseq`, which corresponds to `JMSXGroupSeq`.

User Properties Methods

clear_properties
Clears all user properties. This procedure does not affect system properties.

set_boolean_property (property_name IN VARCHAR, property_value IN BOOLEAN)
Checks whether `property_name` is null or exists. If it is not null, the procedure stores `property_value` in an internal representation (a `NUMBER` type). Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_byte_property (property_name IN VARCHAR, property_value IN INT)
Checks whether property_name is null or exists. If it is not null, the procedure checks whether `property_value` is within -128 to 127 (8-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `byte` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_double_property (property_name IN VARCHAR, property_value IN DOUBLE PRECISION)
Checks whether `property_name` is null or exists. If it is not null, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_float_property (property_name IN VARCHAR, property_value IN FLOAT)
Checks whether `property_name` is null or exists. If it is not null, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_int_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If it is not null, the procedure checks whether `property_value` is within -2147483648 to 2147483647 (32-bits). This check is necessary because the `INT` datatype is 38 bits in PL/SQL and Oracle Database. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_long_property (property_name IN VARCHAR, property_value IN NUMBER)
Checks whether `property_name` is null or exists. If it is not null, the procedure stores `property_value`. In PL/SQL and Oracle Database, the `NUMBER` datatype is 38 bits. In Java, the long datatype is 64 bits. Therefore, no range check is needed. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_short_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If it is not null, the procedure checks whether `property_value` is within -32768 to 32767 (16-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `short` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_string_property (property_name IN VARCHAR, property_value IN VARCHAR)
Checks whether `property_name` is null or exists. If it is not null, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN
If the property with the corresponding property name passed in exists, and if it is a `BOOLEAN` property, then this function returns the value of the property. Otherwise it returns a null.

get_byte_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `BYTE` property, then this function returns the value of the property. Otherwise it returns a null.

get_double_property (property_name IN VARCHAR) RETURN DOUBLE PRECISION
If the property with the corresponding property name passed in exists, and if it is a `DOUBLE` property, then this function returns the value of the property. Otherwise it returns a null.

get_float_property (property_name IN VARCHAR) RETURN FLOAT
If the property with the corresponding property name passed in exists, and if it is a `FLOAT` property, then this function returns the value of the property. Otherwise it returns a null.

get_int_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `Integer` property, then this function returns the value of the property. Otherwise it returns a null.

get_long_property (property_name IN VARCHAR) RETURN NUMBER
If the property with the corresponding property name passed in exists, and if it is a `long` property, then this function returns the value of the property. Otherwise it returns a null.

get_short_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `short` property, then this function returns the value of the property. Otherwise it returns a null.

get_string_property (property_name IN VARCHAR) RETURN VARCHAR
If the property with the corresponding property name passed in exists, and if it is a `STRING` property, then this function returns the value of the property. Otherwise it returns a null.

Payload Methods

set_text (payload IN VARCHAR2)
Sets the payload, a `VARCHAR2` value, to an internal representation.

set_text (payload IN CLOB),
Sets the payload, a `CLOB` value, to an internal representation.

set_bytes (payload IN RAW)
Sets the payload, a `RAW` value, to an internal representation.

set_bytes (payload IN BLOB)
Sets the payload, a `BLOB` value, to an internal representation.

get_text (payload OUT VARCHAR2)
Puts the internal representation of the payload into a `VARCHAR2` variable payload.

get_text (payload OUT CLOB)
Puts the internal representation of the payload into a `CLOB` variable payload.

get_bytes (payload OUT RAW)
Puts the internal representation of the payload into a `RAW` variable payload.

get_bytes (payload OUT BLOB)
Puts the internal representation of the payload into a `BLOB` variable payload.

308.9.2 SYS.AQ$_JMS_TEXT_MESSAGE Type

This type is the ADT used to store a `TextMessage` in an Oracle Database Advanced Queuing queue.

This section contains these topics:

- CONSTRUCT aq$_jms_text_message Function
- JMS Header Methods
- System Properties Methods
- User Properties Methods
- Payload Methods

Syntax

```
TYPE AQ$_JMS_TEXT_MESSAGE AS OBJECT(
 header    aq$_jms_header,
 text_len  INT,
 text_vc   varchar2(4000),
 text_lob  clob,
 STATIC FUNCTION construct    RETURN aq$_jms_text_message,
 MEMBER PROCEDURE set_replyto  (replyto  IN  sys.aq$_agent),
 MEMBER PROCEDURE set_type     (type     IN  VARCHAR),
 MEMBER FUNCTION  get_replyto RETURN sys.aq$_agent,
 MEMBER FUNCTION  get_type    RETURN VARCHAR,
 MEMBER PROCEDURE set_userid   (userid   IN  VARCHAR),
 MEMBER PROCEDURE set_appid    (appid    IN  VARCHAR),
 MEMBER PROCEDURE set_groupid  (groupid  IN  VARCHAR),
 MEMBER PROCEDURE set_groupseq (groupseq IN  INT),
 MEMBER FUNCTION get_userid   RETURN VARCHAR,
 MEMBER FUNCTION get_appid    RETURN VARCHAR,
 MEMBER FUNCTION get_groupid  RETURN VARCHAR,
 MEMBER FUNCTION get_groupseq RETURN INT,
 MEMBER PROCEDURE clear_properties,
 MEMBER PROCEDURE set_boolean_property(property_name IN VARCHAR,
   property_value IN BOOLEAN),
 MEMBER PROCEDURE set_byte_property   (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_double_property (property_name IN VARCHAR,
   property_value IN DOUBLE PRECISION),
 MEMBER PROCEDURE set_float_property  (property_name IN VARCHAR,
   property_value IN FLOAT),
 MEMBER PROCEDURE set_int_property    (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_long_property   (property_name IN VARCHAR,
   property_value IN NUMBER),
 MEMBER PROCEDURE set_short_property  (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_string_property (property_name IN VARCHAR,
   property_value IN VARCHAR),
 MEMBER FUNCTION get_boolean_property (property_name IN VARCHAR)
   RETURN BOOLEAN,
```

```
MEMBER FUNCTION get_byte_property    (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_double_property  (property_name IN VARCHAR)
  RETURN DOUBLE PRECISION,
MEMBER FUNCTION get_float_property   (property_name IN VARCHAR) RETURN FLOAT,
MEMBER FUNCTION get_int_property     (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_long_property    (property_name IN VARCHAR) RETURN NUMBER,
MEMBER FUNCTION get_short_property   (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_string_property  (property_name IN VARCHAR)
  RETURN VARCHAR,
MEMBER PROCEDURE set_text            (payload  IN  VARCHAR2),
MEMBER PROCEDURE set_text            (payload  IN  CLOB),
MEMBER PROCEDURE get_text            (payload  OUT VARCHAR2),
MEMBER PROCEDURE get_text            (payload  OUT CLOB));
```

CONSTRUCT aq$_jms_text_message Function

STATIC FUNCTION construct RETURN aq$_jms_text_message
Creates an empty `aq$_jms_text_message`.

JMS Header Methods

set_replyto (replyto IN sys.aq$_agent)
Sets the `replyto` parameter, which corresponds to `JMSReplyTo` in JMS.

set_type (type IN VARCHAR)
Sets the JMS type, which can be any text, and which corresponds to `JMSType` in JMS.

get_replyto RETURN sys.aq$_agent
Returns `replyto`, which corresponds to `JMSReplyTo`.

get_type RETURN VARCHAR
Returns `type`, which corresponds to `JMSType`.

System Properties Methods

set_userid (userid IN VARCHAR)
Sets `userid`, which corresponds to `JMSXUserID` in JMS.

set_appid (appid IN VARCHAR)
Sets `appid`, which corresponds to `JMSXAppID` in JMS.

set_groupid (groupid IN VARCHAR)
Sets `groupid`, which corresponds to `JMSXGroupID` in JMS.

set_groupseq (groupseq IN INT)
Sets `groupseq`, which corresponds to `JMSXGroupSeq` in JMS.

get_userid RETURN VARCHAR
Returns `userid`, which corresponds to `JMSXUserID`.

get_appid RETURN VARCHAR
Returns `appid`, which corresponds to `JMSXAppID`.

get_groupid RETURN VARCHAR
Returns `groupid`, which corresponds to `JMSXGroupID`.

get_groupseq RETURN INT
Returns `groupseq`, which corresponds to `JMSXGroupSeq`.

User Properties Methods

clear_properties
Clears all user properties. This procedure does not affect system properties.

set_boolean_property (property_name IN VARCHAR, property_value IN BOOLEAN)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value` in an internal representation. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_byte_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -128 to 127 (8-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `BYTE` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_double_property (property_name IN VARCHAR, property_value IN DOUBLE PRECISION)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_float_property (property_name IN VARCHAR, property_value IN FLOAT)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_int_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -2147483648 to 2147483647 (32-bits). This check is necessary because in PL/SQL and Oracle Database, the `INT` datatype is 38 bits. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_long_property (property_name IN VARCHAR, property_value IN NUMBER)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. In PL/SQL and Oracle Database, the `NUMBER` datatype is 38 bits. In Java, the `long` datatype is 64 bits. Therefore, no range check is needed.Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_short_property property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -32768 to 32767 (16-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `short` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_string_property (property_name IN VARCHAR, property_value IN VARCHAR)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN
If the property with the corresponding property name passed in exists, and if it is a `BOOLEAN` property, then this function returns the value of the property. Otherwise it returns a null.

get_byte_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `BYTE` property, then this function returns the value of the property. Otherwise it returns a null.

get_double_property (property_name IN VARCHAR) RETURN DOUBLE PRECISION
If the property with the corresponding property name passed in exists, and if it is a `DOUBLE` property, then this function returns the value of the property. Otherwise it returns a null.

get_float_property (property_name IN VARCHAR) RETURN FLOAT
If the property with the corresponding property name passed in exists, and if it is a `FLOAT` property, then this function returns the value of the property. Otherwise it returns a null.

get_int_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `Integer` property, then this function returns the value of the property. Otherwise it returns a null.

get_long_property (property_name IN VARCHAR) RETURN NUMBER
If the property with the corresponding property name passed in exists, and if it is a `long` property, then this function returns the value of the property. Otherwise it returns a null.

get_short_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `short` property, then this function returns the value of the property. Otherwise it returns a null.

get_string_property (property_name IN VARCHAR) RETURN VARCHAR)
If the property with the corresponding property name passed in exists, and if it is a `STRING` property, then this function returns the value of the property. Otherwise it returns a null.

Payload Methods

set_text (payload IN VARCHAR2)
Sets the payload, a `VARCHAR2` value, to an internal representation.

set_text (payload IN CLOB)
Sets the payload, a `CLOB` value, to an internal representation.

get_text (payload OUT VARCHAR2)
Puts the internal representation of the payload into a `VARCHAR2` variable payload.

get_text (payload OUT CLOB)
Puts the internal representation of the payload into a `CLOB` variable payload.

308.9.3 SYS.AQ$_JMS_BYTES_MESSAGE Type

The `SYS.AQ$_JMS_BYTES_MESSAGE` type is the ADT used to store a `BytesMessage` in an Oracle Database Advanced Queuing queue.

This section contains these topics:

- CONSTRUCT aq$_jms_bytes_message Function
- JMS Header Methods
- System Properties Methods

- User Properties Methods
- Payload Methods

Syntax

```
TYPE AQ$_JMS_BYTES_MESSAGE AS OBJECT(
 header     aq$_jms_header,
 bytes_len  INT,
 bytes_raw  raw(2000),
 bytes_lob  blob,
 STATIC FUNCTION construct RETURN aq$_jms_bytes_message,
 MEMBER PROCEDURE set_replyto  (replyto IN sys.aq$_agent),
 MEMBER PROCEDURE set_type     (type    IN VARCHAR),
 MEMBER FUNCTION get_replyto RETURN sys.aq$_agent,
 MEMBER FUNCTION get_type    RETURN VARCHAR,
 MEMBER PROCEDURE set_userid   (userid   IN VARCHAR),
 MEMBER PROCEDURE set_appid    (appid    IN VARCHAR),
 MEMBER PROCEDURE set_groupid  (groupid  IN VARCHAR),
 MEMBER PROCEDURE set_groupseq (groupseq IN INT),
 MEMBER FUNCTION get_userid   RETURN VARCHAR,
 MEMBER FUNCTION get_appid    RETURN VARCHAR,
 MEMBER FUNCTION get_groupid  RETURN VARCHAR,
 MEMBER FUNCTION get_groupseq RETURN INT,
 MEMBER PROCEDURE clear_properties,
 MEMBER PROCEDURE set_boolean_property(property_name IN VARCHAR,
   property_value IN BOOLEAN),
 MEMBER PROCEDURE set_byte_property   (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_double_property (property_name IN VARCHAR,
   property_value IN DOUBLE PRECISION),
 MEMBER PROCEDURE set_float_property  (property_name IN VARCHAR,
   property_value IN FLOAT),
 MEMBER PROCEDURE set_int_property    (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_long_property   (property_name IN VARCHAR,
   property_value IN NUMBER),
 MEMBER PROCEDURE set_short_property  (property_name IN VARCHAR,
   property_valuE IN INT),
 MEMBER PROCEDURE set_string_property (property_name IN VARCHAR,
   property_value IN VARCHAR),
 MEMBER FUNCTION get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN,
 MEMBER FUNCTION get_byte_property    (property_name IN VARCHAR) RETURN INT,
 MEMBER FUNCTION get_double_property  (property_name IN VARCHAR)
   RETURN DOUBLE PRECISION,
 MEMBER FUNCTION get_float_property   (property_name IN VARCHAR) RETURN FLOAT,
 MEMBER FUNCTION get_int_property     (property_name IN VARCHAR) RETURN INT,
 MEMBER FUNCTION get_long_property    (property_name IN VARCHAR) RETURN NUMBER,
 MEMBER FUNCTION get_short_property   (property_name IN VARCHAR) RETURN INT,
 MEMBER FUNCTION get_string_property  (property_name IN VARCHAR) RETURN VARCHAR,
 MEMBER PROCEDURE set_bytes           (payload  IN RAW),
 MEMBER PROCEDURE set_bytes           (payload  IN BLOB),
 MEMBER PROCEDURE get_bytes           (payload  OUT RAW),
 MEMBER PROCEDURE get_bytes           (payload  OUT BLOB),
 MEMBER FUNCTION  prepare             (id IN PLS_INTEGER) RETURN PLS_INTEGER,
 MEMBER PROCEDURE reset               (id IN PLS_INTEGER),
 MEMBER PROCEDURE flush               (id IN PLS_INTEGER),
 MEMBER PROCEDURE clear_body          (id IN PLS_INTEGER),
 MEMBER PROCEDURE clean               (id IN PLS_INTEGER),
 STATIC PROCEDURE clean_all,
 MEMBER FUNCTION get_mode           (id IN PLS_INTEGER) RETURN PLS_INTEGER,
 MEMBER FUNCTION read_boolean       (id IN PLS_INTEGER) RETURN BOOLEAN,
 MEMBER FUNCTION read_byte          (id IN PLS_INTEGER) RETURN PLS_INTEGER,
```

```
MEMBER FUNCTION read_bytes         (id IN PLS_INTEGER,
  value OUT NOCOPY BLOB, length IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION read_char          (id IN PLS_INTEGER) RETURN CHAR,
MEMBER FUNCTION read_double        (id IN PLS_INTEGER) RETURN DOUBLE PRECISION,
MEMBER FUNCTION read_float         (id IN PLS_INTEGER) RETURN FLOAT,
MEMBER FUNCTION read_int           (id IN PLS_INTEGER) RETURN INT,
MEMBER FUNCTION read_long          (id IN PLS_INTEGER) RETURN NUMBER,
MEMBER FUNCTION read_short         (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION read_unsigned_byte  (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION read_unsigned_short (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER PROCEDURE read_utf          (id IN PLS_INTEGER, value OUT NOCOPY CLOB),
MEMBER PROCEDURE write_boolean     (id IN PLS_INTEGER, value IN BOOLEAN),
MEMBER PROCEDURE write_byte        (id IN PLS_INTEGER, value IN PLS_INTEGER),
MEMBER PROCEDURE write_bytes       (id IN PLS_INTEGER, value IN RAW),
MEMBER PROCEDURE write_bytes       (id IN PLS_INTEGER, value IN BLOB),
MEMBER PROCEDURE write_bytes       (id IN PLS_INTEGER, value IN RAW,
  offset IN PLS_INTEGER, length IN PLS_INTEGER),
MEMBER PROCEDURE write_bytes       (id IN PLS_INTEGER, value IN BLOB,
  offset IN INT, length IN INT),
MEMBER PROCEDURE write_char        (id IN PLS_INTEGER, value IN CHAR),
MEMBER PROCEDURE write_double      (id IN PLS_INTEGER,
  value IN DOUBLE PRECISION),
MEMBER PROCEDURE write_float       (id IN PLS_INTEGER, value IN FLOAT),
MEMBER PROCEDURE write_int         (id IN PLS_INTEGER, value IN PLS_INTEGER),
MEMBER PROCEDURE write_long        (id IN PLS_INTEGER, value IN NUMBER),
MEMBER PROCEDURE write_short       (id IN PLS_INTEGER, value IN PLS_INTEGER),
MEMBER PROCEDURE write_utf         (id IN PLS_INTEGER, value IN VARCHAR2),
MEMBER PROCEDURE write_utf         (id IN PLS_INTEGER, value IN CLOB));
```

CONSTRUCT aq$_jms_bytes_message Function

STATIC FUNCTION construct RETURN aq$_jms_bytes_message
Creates an empty `aq$_jms_bytes_message`.

JMS Header Methods

set_replyto (replyto IN sys.aq$_agent)
Sets the `replyto` parameter, which corresponds to `JMSReplyTo` in JMS.

set_type (type IN VARCHAR)
Sets the JMS type, which can be any text, and which corresponds to `JMSType` in JMS.

get_replyto RETURN sys.aq$_agent
Returns `replyto`, which corresponds to `JMSReplyTo`.

get_type RETURN VARCHAR
Returns `type`, which corresponds to `JMSType`.

System Properties Methods

set_userid (userid IN VARCHAR)
Sets `userid`, which corresponds to `JMSXUserID` in JMS.

set_appid (appid IN VARCHAR)
Sets `appid`, which corresponds to `JMSXAppID` in JMS.

set_groupid (groupid IN VARCHAR)
Sets `groupid`, which corresponds to `JMSXGroupID` in JMS.

set_groupseq (groupseq IN INT)
Sets `groupseq`, which corresponds to `JMSXGroupSeq` in JMS.

get_userid RETURN VARCHAR
Returns `userid`, which corresponds to `JMSXUserID`.

get_appid RETURN VARCHAR
Returns `appid`, which corresponds to `JMSXAppID`.

get_groupid RETURN VARCHAR
Returns `groupid`, which corresponds to `JMSXGroupID`.

get_groupseq RETURN NUMBER
Returns `groupseq`, which corresponds to `JMSXGroupSeq`.

User Properties Methods

clear_properties
Clears all user properties. This procedure does not affect system properties.

set_boolean_property (property_name IN VARCHAR, property_value IN BOOLEAN)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value` in an internal representation. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_byte_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -128 to 127 (8-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `BYTE` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_double_property (property_name IN VARCHAR, property_value IN DOUBLE PRECISION)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_float_property (property_name IN VARCHAR, property_value IN FLOAT)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_int_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -2147483648 to 2147483647 (32-bits). This check is necessary because in PL/SQL and Oracle Database, the `INT` datatype is 38 bits. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_long_property (property_name IN VARCHAR, property_value IN NUMBER)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. In PL/SQL and Oracle Database, the `NUMBER` datatype is 38 bits. In Java, the `long` datatype is 64 bits. Therefore, no range check is needed.Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_short_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -32768 to 32767 (16-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `short` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_string_property (property_name IN VARCHAR, property_value IN VARCHAR)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN
If the property with the corresponding property name passed in exists, and if it is a `BOOLEAN` property, then this function returns the value of the property. Otherwise it returns a null.

get_byte_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `BYTE` property, then this function returns the value of the property. Otherwise it returns a null.

get_double_property (property_name IN VARCHAR) RETURN DOUBLE PRECISION
If the property with the corresponding property name passed in exists, and if it is a `DOUBLE` property, then this function returns the value of the property. Otherwise it returns a null.

get_float_property (property_name IN VARCHAR) RETURN FLOAT
If the property with the corresponding property name passed in exists, and if it is a `FLOAT` property, then this function returns the value of the property. Otherwise it returns a null.

get_int_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `Integer` property, then this function returns the value of the property. Otherwise it returns a null.

get_long_property (property_name IN VARCHAR) RETURN NUMBER
If the property with the corresponding property name passed in exists, and if it is a `long` property, then this function returns the value of the property. Otherwise it returns a null.

get_short_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `short` property, then this function returns the value of the property. Otherwise it returns a null.

get_string_property (property_name IN VARCHAR) RETURN VARCHAR
If the property with the corresponding property name passed in exists, and if it is a `STRING` property, then this function returns the value of the property. Otherwise it returns a null.

Payload Methods

set_bytes (payload in RAW)
Sets the payload, a `RAW` value, to an internal representation.

set_bytes (payload in BLOB)
Sets the payload, a `BLOB` value, to an internal representation.

get_bytes (payload out RAW)
Puts the internal representation of the payload into a `RAW` variable payload. Raises exception ORA-24190 if the length of the internal payload is more than 32767 (the maximum length of `RAW` in PL/SQL).

get_bytes (payload out BLOB)
Puts the internal representation of the payload into a `BLOB` variable payload.

prepare (id IN PLS_INTEGER) RETURN PLS_INTEGER
Takes the byte array stored in aq$_jms_bytes_message and decodes it as a Java object in the Java stored procedure. The result of the decoding is stored as a static variable in Jserv session memory. Parameter `id` is used to identify the slot where the Java object is stored in the Oracle Database JVM session memory. If `id` is null, then a new slot is created for this PL/SQL object. Subsequent JMS operations on the payload need to provide this operation ID. This function also sets the message access mode to `MESSAGE_ACCESS_READONLY`. Subsequent calls of `write_XXX` procedure raise an ORA-24196 error. Users can call the `clear_body` procedure to set the message access mode to `MESSAGE_ACCESS_READONLY`.
This function raises ORA-24197 if the Java stored procedure throws an exception during execution, ORA-24198 if the operation ID is invalid, or ORA-24199 if the Java stored procedure message store overflows.

reset (id IN PLS_INTEGER)
Resets the starting position of the stream to the beginning and puts the bytes message in read-only mode. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

flush (id IN PLS_INTEGER)
Takes the static variable in Jserv and synchronizes the content back to the `aq$_jms_bytes_message`. This procedure will not affect the underlying access mode. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

clear_body (id IN PLS_INTEGER)
Sets the Java stored procedure static variable to empty payload. Parameter `id` is used to identify the slot where the Java object is stored in the Oracle Database JVM session memory. If `id` is null, a new slot is created for this PL/SQL object. Subsequent JMS operations on the payload need to provide this operation ID.
It also sets the message access mode to `MESSAGE_ACCESS_WRITEONLY`. Later calls of `read_XXX` procedure raise ORA-24196 error. Users can call the `reset` or `prepare` procedures to set the message access mode to `MESSAGE_ACCESS_READONLY`. Write-only and read-only modes affect only the payload functions of `AQ$_JMS_BYTES_MESSAGE`. They do not affect the header functions.
This function raises ORA-24197 if the Java stored procedure throws an exception during execution, ORA-24198 if the operation ID is invalid, or ORA-24199 if the Java stored procedure message store overflows.

clean (id IN PLS_INTEGER)
Closes and cleans up the `DataInputStream` or `DataOutputStream` at the Java stored procedure side corresponding to the operation ID. It is very important to call this procedure to avoid memory leaks. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

clean_all
Closes and cleans up all the messages in the corresponding type of message store at the Java stored procedure side. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution.

get_mode (id IN PLS_INTEGER) RETURN PLS_INTEGER
Returns the current mode of this message. The return value is either `SYS.dbms_jms.plsql.MESSAGE_ACCESS_READONLY` or

SYS.dbms_jms.plsql.MESSAGE_ACCESS_WRITEONLY. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

read_boolean (id IN PLS_INTEGER) RETURN BOOLEAN
Reads a Boolean value from the bytes message and returns the Boolean value read. Null is returned if the end of the message stream has been reached. Parameter id is the operation ID. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_byte (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads a BYTE value from the bytes message and returns the BYTE value read. Null is returned if the end of the stream has been reached. Because there is no BYTE type in PL/SQL, Oracle Database uses PLS_INTEGER to represent a BYTE. Although PL/SQL users get a PLS_INTEGER, they are guaranteed that the value is in the Java BYTE value range. If this value is issued with a write_byte function, then there will not be an out of range error. Parameter id is the operation ID. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_bytes (id IN PLS_INTEGER, value OUT NO COPY BLOB, length IN PLS_INTEGER) RETURN PLS_INTEGER
Reads length of the bytes from bytes message stream into value and returns the total number of bytes read. If there is no more data (because the end of the stream has been reached), then it returns -1. Raises exceptions ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_char (id IN PLS_INTEGER) RETURN CHAR
Reads a character value from the bytes message and returns the character value read. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_double (id IN PLS_INTEGER) RETURN DOUBLE PRECISION
Reads a double from the bytes message and returns the character value read. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_float (id IN PLS_INTEGER) RETURN FLOAT
Reads a float from the bytes message and returns the float read. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_int (id IN PLS_INTEGER) RETURN INT
Reads an INT from the bytes message and returns the INT read. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_long (id IN PLS_INTEGER) RETURN NUMBER
Reads a long from the bytes message and returns the long read. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-

only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_short (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads a short value from the bytes message and returns the short value read. Null is returned if the end of the stream has been reached. Because there is no short type in PL/SQL, `PLS_INTEGER` is used to represent a `SHORT`. Although PL/SQL users get an `PLS_INTEGER`, they are guaranteed that the value is in the Java short value range. If this value is issued with a `write_short` function, then there will not be an out of range error. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_unsigned_byte (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads an unsigned 8-bit number from the bytes message stream and returns the next byte from the bytes message stream, interpreted as an unsigned 8-bit number. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_unsigned_short (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads an unsigned 16-bit number from the bytes message stream and returns the next two bytes from the bytes message stream, interpreted as an unsigned 16-bit integer. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_utf (id IN PLS_INTEGER, value OUT NOCOPY CLOB)
Reads a string that has been encoded using a UTF-8 format from the bytes message. Null is returned if the end of the stream has been reached. Raises exception ORA-24196 if the bytes message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_boolean (id IN PLS_INTEGER, value IN BOOLEAN)
Writes a Boolean to the bytes message stream as a 1-byte value. The value `true` is written as the value (byte)1. The value `false` is written as the value (byte)0. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_byte (id IN PLS_INTEGER, value IN PLS_INTEGER)
Writes a byte to the bytes message. Because there is no `BYTE` type in PL/SQL, `PLS_INTEGER` is used to represent a `BYTE`. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN RAW)
Writes an array of bytes to the bytes message. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN BLOB)
Writes an array of bytes to the bytes message. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN RAW, offset IN PLS_INTEGER, length IN PLS_INTEGER)
Writes a portion of a byte array to the bytes message stream. Parameter `offset` is the initial offset within the byte array. If the range [offset, offset+length] exceeds the boundary of the byte array value, then a Java IndexOutOfBounds exception is thrown in the Java stored procedure and this procedure raises error ORA-24197. The index starts from 0. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN BLOB, offset IN INT, length IN INT)
Writes a portion of a byte array to the bytes message stream. Parameter `offset` is the initial offset within the byte array. If the range [offset, offset+length] exceeds the boundary of the byte array value, then a Java IndexOutOfBounds exception is thrown in the Java stored procedure and this procedure raises error ORA-24197. The index starts from 0. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_char (id IN PLS_INTEGER, value IN CHAR)
Writes a character value to the bytes message. If this value has multiple characters, it is the first character that is written. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_double (id IN PLS_INTEGER, value IN DOUBLE PRECISION)
Writes a double to the bytes message. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_float (id IN PLS_INTEGER, value IN FLOAT)
Writes a float to the bytes message. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_int (id IN PLS_INTEGER, value IN PLS_INTEGER)
Writes an `INT` to the bytes message. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_long (id IN PLS_INTEGER, value IN NUMBER)
Writes a long to the bytes message. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_short (id IN PLS_INTEGER, value IN PLS_INTEGER)
Writes a short to the bytes message as two bytes, high byte first. Because there is no short type in PL/SQL, `INT` is used to represent a short. Raises exception ORA-24193 if the parameter value exceeds the valid range, ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_utf (id IN PLS_INTEGER, value IN VARCHAR2)
Writes a string to the bytes message stream using UTF-8 encoding in a machine-independent manner. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197

if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_utf (id IN PLS_INTEGER, value IN CLOB)
Writes a string to the bytes message stream using UTF-8 encoding in a machine-independent manner. Raises exception ORA-24196 if the bytes message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

308.9.4 SYS.AQ$_JMS_MAP_MESSAGE Type

This type is the ADT used to store a `MapMessage` in an Oracle Database Advanced Queuing queue.

This section contains these topics:

- CONSTRUCT aq$_jms_map_message Function
- JMS Header Methods
- System Properties Methods
- User Properties Methods
- Payload Methods

Syntax

```
TYPE aq$_jms_map_message AS object(
 header     aq$_jms_header,
 bytes_len  int,
 bytes_raw  raw(2000),
 bytes_lob  blob,
 STATIC FUNCTION  construct   RETURN aq$_jms_map_message,
 MEMBER PROCEDURE set_replyto  (replyto IN sys.aq$_agent),
 MEMBER PROCEDURE set_type     (type    IN VARCHAR),
 MEMBER FUNCTION get_replyto  RETURN sys.aq$_agent,
 MEMBER FUNCTION get_type     RETURN VARCHAR,
 MEMBER PROCEDURE set_userid   (userid   IN VARCHAR),
 MEMBER PROCEDURE set_appid    (appid    IN VARCHAR),
 MEMBER PROCEDURE set_groupid  (groupid  IN VARCHAR),
 MEMBER PROCEDURE set_groupseq (groupseq IN INT),
 MEMBER FUNCTION get_userid   RETURN VARCHAR,
 MEMBER FUNCTION get_appid    RETURN VARCHAR,
 MEMBER FUNCTION get_groupid  RETURN VARCHAR,
 MEMBER FUNCTION get_groupseq RETURN INT,
 MEMBER PROCEDURE clear_properties,
 MEMBER PROCEDURE set_boolean_property(property_name IN VARCHAR,
   property_value IN BOOLEAN),
 MEMBER PROCEDURE set_byte_property   (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_double_property (property_name IN VARCHAR,
   property_value IN DOUBLE PRECISION),
 MEMBER PROCEDURE set_float_property  (property_name IN VARCHAR,
   property_value IN FLOAT),
 MEMBER PROCEDURE set_int_property    (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_long_property   (property_name IN VARCHAR,
   property_value IN NUMBER),
 MEMBER PROCEDURE set_short_property  (property_name IN VARCHAR,
   property_valuE IN INT),
 MEMBER PROCEDURE set_string_property (property_name IN VARCHAR,
```

```
  property_value IN VARCHAR),
MEMBER FUNCTION get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN,
MEMBER FUNCTION get_byte_property    (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_double_property  (property_name IN VARCHAR)
  RETURN DOUBLE PRECISION,
MEMBER FUNCTION get_float_property   (property_name IN VARCHAR) RETURN FLOAT,
MEMBER FUNCTION get_int_property     (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_long_property    (property_name IN VARCHAR) RETURN NUMBER,
MEMBER FUNCTION get_short_property   (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_string_property  (property_name IN VARCHAR) RETURN VARCHAR,
MEMBER PROCEDURE set_bytes   (payload IN RAW),
MEMBER PROCEDURE set_bytes   (payload IN BLOB),
MEMBER PROCEDURE get_bytes   (payload OUT RAW),
MEMBER PROCEDURE get_bytes   (payload OUT BLOB),
MEMBER FUNCTION  prepare     (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER PROCEDURE flush       (id  IN  PLS_INTEGER),
MEMBER PROCEDURE clear_body  (id IN PLS_INTEGER),
MEMBER PROCEDURE clean       (id  IN  PLS_INTEGER),
STATIC PROCEDURE clean_all,
MEMBER PROCEDURE set_boolean (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN BOOLEAN),
MEMBER PROCEDURE set_byte    (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  PLS_INTEGER),
MEMBER PROCEDURE set_bytes   (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  RAW),
MEMBER PROCEDURE set_bytes   (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN RAW, offset IN INT, length IN INT),
MEMBER PROCEDURE set_bytes   (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN BLOB),
MEMBER PROCEDURE set_bytes   (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN BLOB, offset IN INT, length IN INT),
MEMBER PROCEDURE set_char    (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  CHAR),
MEMBER PROCEDURE set_double  (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  DOUBLE PRECISION),
MEMBER PROCEDURE set_float   (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  FLOAT),
MEMBER PROCEDURE set_int     (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  PLS_INTEGER),
MEMBER PROCEDURE set_long    (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  NUMBER),
MEMBER PROCEDURE set_short   (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  PLS_INTEGER),
MEMBER PROCEDURE set_string  (id IN PLS_INTEGER, name IN VARCHAR2,
  value IN  VARCHAR2),
MEMBER PROCEDURE set_string  (id IN  PLS_INTEGER, name IN VARCHAR2,
  value IN  CLOB),
MEMBER FUNCTION get_boolean  (id IN PLS_INTEGER, name IN VARCHAR2)
  RETURN BOOLEAN,
MEMBER FUNCTION  get_byte    (id IN PLS_INTEGER, name IN VARCHAR2)
  RETURN PLS_INTEGER,
MEMBER PROCEDURE get_bytes   (id IN PLS_INTEGER, name IN VARCHAR2,
  value OUT NOCOPY BLOB),
MEMBER FUNCTION get_char     (id IN PLS_INTEGER, name IN VARCHAR2) RETURN CHAR,
MEMBER FUNCTION get_double   (id IN PLS_INTEGER, name IN VARCHAR2)
  RETURN DOUBLE PRECISION,
MEMBER FUNCTION get_float   (id IN PLS_INTEGER, name IN VARCHAR2) RETURN FLOAT,
MEMBER FUNCTION get_int     (id IN PLS_INTEGER, name IN VARCHAR2)
  RETURN PLS_INTEGER,
MEMBER FUNCTION get_long    (id IN PLS_INTEGER, name IN VARCHAR2)
  RETURN NUMBER,
MEMBER FUNCTION get_short   (id IN PLS_INTEGER, name IN VARCHAR2)
```

```
  RETURN PLS_INTEGER,
MEMBER PROCEDURE get_string (id IN PLS_INTEGER, name IN VARCHAR2,
  value OUT NOCOPY CLOB),
MEMBER FUNCTION get_names   (id IN PLS_INTEGER) RETURN aq$_jms_namearray,
MEMBER FUNCTION get_names   (id IN PLS_INTEGER, names OUT aq$_jms_namearray,
  offset IN  PLS_INTEGER, length IN  PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER PROCEDURE get_object (id IN PLS_INTEGER, name IN  VARCHAR2,
  value  OUT NOCOPY AQ$_JMS_VALUE),
MEMBER FUNCTION get_size    (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION item_exists (id IN PLS_INTEGER, name IN VARCHAR2)
  RETURN BOOLEAN);
```

CONSTRUCT aq$_jms_map_message Function

STATIC FUNCTION construct RETURN aq$_jms_map_message
Creates an empty `aq$_jms_map_message` object.

JMS Header Methods

set_replyto (replyto IN sys.aq$_agent)
Sets the `replyto` parameter, which corresponds to `JMSReplyTo` in JMS.

set_type (type IN VARCHAR)
Sets the JMS type, which can be any text, and which corresponds to `JMSType` in JMS.

get_replyto RETURN sys.aq$_agent
Returns `replyto`, which corresponds to `JMSReplyTo`.

get_type RETURN VARCHAR
Returns `type`, which corresponds to `JMSType`.

System Properties Methods

set_userid (userid IN VARCHAR)
Sets `userid`, which corresponds to `JMSXUserID` in JMS.

set_appid (appid IN VARCHAR)
Sets `appid`, which corresponds to `JMSXAppID` in JMS.

set_groupid (groupid IN VARCHAR)
Sets `groupid`, which corresponds to `JMSXGroupID` in JMS.

set_groupseq (groupseq IN INT)
Sets `groupseq`, which corresponds to `JMSXGroupSeq` in JMS.

get_userid RETURN VARCHAR
Returns `userid`, which corresponds to `JMSXUserID`.

get_appid RETURN VARCHAR
Returns `appid`, which corresponds to `JMSXAppID`.

get_groupid RETURN VARCHAR
Returns `groupid`, which corresponds to `JMSXGroupID`.

get_groupseq RETURN NUMBER
Returns `groupseq`, which corresponds to `JMSXGroupSeq`.

User Properties Methods

clear_properties
Clears all user properties. This procedure does not affect system properties.

set_boolean_property (property_name IN VARCHAR, property_value IN BOOLEAN)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value` in an internal representation. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_byte_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -128 to 127 (8-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `BYTE` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_double_property (property_name IN VARCHAR, property_value IN DOUBLE PRECISION)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_float_property (property_name IN VARCHAR, property_value IN FLOAT)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_int_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -2147483648 to 2147483647 (32-bits). This check is necessary because in PL/SQL and Oracle Database, the `INT` datatype is 38 bits. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_long_property (property_name IN VARCHAR, property_value IN NUMBER)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. In PL/SQL and Oracle Database, the `NUMBER` datatype is 38 bits. In Java, the `long` datatype is 64 bits. Therefore, no range check is needed.Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_short_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -32768 to 32767 (16-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `short` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_string_property (property_name IN VARCHAR, property_value IN VARCHAR)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN
If the property with the corresponding property name passed in exists, and if it is a `BOOLEAN` property, then this function returns the value of the property. Otherwise it returns a null.

get_byte_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `BYTE` property, then this function returns the value of the property. Otherwise it returns a null.

get_double_property (property_name IN VARCHAR) RETURN DOUBLE PRECISION
If the property with the corresponding property name passed in exists, and if it is a `DOUBLE` property, then this function returns the value of the property. Otherwise it returns a null.

get_float_property (property_name IN VARCHAR) RETURN FLOAT
If the property with the corresponding property name passed in exists, and if it is a `FLOAT` property, then this function returns the value of the property. Otherwise it returns a null.

get_int_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `Integer` property, then this function returns the value of the property. Otherwise it returns a null.

get_long_property (property_name IN VARCHAR) RETURN NUMBER
If the property with the corresponding property name passed in exists, and if it is a `long` property, then this function returns the value of the property. Otherwise it returns a null.

get_short_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `short` property, then this function returns the value of the property. Otherwise it returns a null.

get_string_property (property_name IN VARCHAR) RETURN VARCHAR
If the property with the corresponding property name passed in exists, and if it is a `STRING` property, then this function returns the value of the property. Otherwise it returns a null.

Payload Methods

set_bytes (payload IN RAW)
Sets the internal payload as a RAW variable without any interpretation. The payload of `aq$_jms_map_message` is stored as either `RAW` or `BLOB` in the database. This member function sets a payload as a RAW variable without interpreting it.

set_bytes (payload IN BLOB)
Sets the internal payload as a BLOB variable without any interpretation. The payload of `aq$_jms_map_message` is stored as either `RAW` or `BLOB` in the database. This member function sets a payload as a BLOB variable without interpreting it.

get_bytes (payload OUT RAW)
Puts the internal payload into a `RAW` variable without any interpretation. The payload of `aq$_jms_map_message` is stored as either `RAW` or `BLOB` in the database. This member function gets a payload as raw bytes without interpreting it. Raises exceptions ORA-24190 if the length of internal payload is more than 32767.

get_bytes (payload OUT BLOB)
Puts the internal payload into a `BLOB` variable without any interpretation. The payload of `aq$_jms_map_message` is stored as either `RAW` or `BLOB` in the database. This member function gets a payload as a `BLOB` without interpreting it.

prepare (id IN PLS_INTEGER) RETURN PLS_INTEGER
Takes the byte array stored in `aq$_jms_map_message` and decodes it as a Java object in the Java stored procedure. The result of the decoding is stored as a static variable in Jserv session memory. Parameter `id` is used to identify the slot where the Java object is stored in the Oracle Database JVM session memory. If `id` is null, then a new slot is created for this PL/SQL object. Subsequent JMS operations on the payload need to provide this operation ID. This function raises ORA-24197 if the Java stored procedure throws an exception during execution, ORA-24198 if the operation ID is invalid, or ORA-24199 if the Java stored procedure message store overflows.

flush (id IN PLS_INTEGER)
Takes the static variable in Jserv and synchronizes the content back to `aq$_jms_map_message`. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

clear_body (id IN PLS_INTEGER)
Sets the Java stored procedure static variable to empty payload. Parameter `id` is used to identify the slot where the Java object is stored in the Oracle Database JVM session memory. If `id` is null, a new slot is created for this PL/SQL object. Subsequent JMS operations on the payload need to provide this operation ID.
This function raises ORA-24197 if the Java stored procedure throws an exception during execution, ORA-24198 if the operation ID is invalid, or ORA-24199 if the Java stored procedure message store overflows.

clean (id IN PLS_INTEGER)
Closes and cleans up the `DataInputStream` or `DataOutputStream` at the Java stored procedure side corresponding to the operation ID. It is very important to call this procedure to avoid memory leaks. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

clean_all
Closes and cleans up all the messages in the corresponding type of message store at the Java stored procedure side. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution.

set_boolean (id IN PLS_INTEGER, name IN VARCHAR2, value IN BOOLEAN)
Sets the Boolean value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_byte (id IN PLS_INTEGER, name IN VARCHAR2, value IN PLS_INTEGER)
Sets the `BYTE` value with the specified name in the map. Because there is no `BYTE` type in PL/SQL, `PLS_INTEGER` is used to represent a byte. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_bytes (id IN PLS_INTEGER, name IN VARCHAR2, value IN RAW))
Sets the byte array value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_bytes (id IN PLS_INTEGER, name IN VARCHAR2, value IN RAW, offset IN INT, length IN INT)
Sets a portion of the byte array value with the specified name in the map. Parameter `offset` is the initial offset within the byte array, and parameter `length` is the number of bytes to use. If the range [offset … offset+length] exceeds the boundary of the byte array value, then a Java

`IndexOutOfBounds` exception is thrown in the Java stored procedure and this procedure raises an ORA-24197 error. The index starts from 0. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_bytes (id IN PLS_INTEGER, name IN VARCHAR2, value IN BLOB)
Sets the byte array value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_bytes (id IN PLS_INTEGER, name IN VARCHAR2, value IN BLOB, offset IN INT, length IN INT)
Sets a portion of the byte array value with the specified name in the map. Parameter `offset` is the initial offset within the byte array, and parameter `length` is the number of bytes to use. If the range [offset … offset+length] exceeds the boundary of the byte array value, then a Java `IndexOutOfBounds` exception is thrown in the Java stored procedure, and this procedure raises an ORA-24197 error. The index starts from 0. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_char (id IN PLS_INTEGER, name IN VARCHAR2, value IN CHAR)
Sets the character value with the specified name in the map. If this value has multiple characters, then it is the first character that is used. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_double (id IN PLS_INTEGER, name IN VARCHAR2, value IN DOUBLE PRECISION)
Sets the double value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_float (id IN PLS_INTEGER, name IN VARCHAR2, value IN FLOAT)
This procedure is to set the float value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_int (id IN PLS_INTEGER, name IN VARCHAR2, value IN PLS_INTEGER)
Sets the int value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_long (id IN PLS_INTEGER, name IN VARCHAR2, value IN NUMBER)
Sets the long value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_short (id IN PLS_INTEGER, name IN VARCHAR2, value IN PLS_INTEGER)
Sets the short value with the specified name in the map. Because there is no short type in PL/SQL, `PLS_INTEGER` is used to represent a short. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_string (id IN PLS_INTEGER, name IN VARCHAR2, value IN VARCHAR2)
Sets the string value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

set_string (id IN PLS_INTEGER, name IN VARCHAR2, value IN CLOB))
Sets the string value with the specified name in the map. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

get_boolean (id IN PLS_INTEGER, name IN VARCHAR2) RETURN BOOLEAN
Retrieves the Boolean value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_byte (id IN PLS_INTEGER, name IN VARCHAR2) RETURN PLS_INTEGER
Retrieves the `BYTE` value with the specified name. If there is no item by this name, then null is returned. Because there is no `BYTE` type in PL/SQL, `PLS_INTEGER` is used to represent a byte. Although the PL/SQL users get an `PLS_INTEGER`, they are guaranteed that the value is in the Java `BYTE` value range. If this value is issued with a `set_byte` function, then there will not be an out of range error. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_bytes (id IN PLS_INTEGER, name IN VARCHAR2, value OUT NOCOPY BLOB)
Retrieves the byte array value with the specified name. If there is no item by this name, then null is returned. Because the size of the array might be larger than the limit of PL/SQL `RAW` type, a `BLOB` is always returned here. The `BLOB` returned is a copy, which means it can be modified without affecting the message payload. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_char (id IN PLS_INTEGER, name IN VARCHAR2) RETURN CHAR
Retrieves and returns the character value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid.

get_double (id IN PLS_INTEGER, name IN VARCHAR2) RETURN DOUBLE PRECISION
Retrieves and returns the double value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid.

get_float (id IN PLS_INTEGER, name IN VARCHAR2) RETURN FLOAT
Retrieves the float value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_int (id IN PLS_INTEGER, name IN VARCHAR2) RETURN PLS_INTEGER
Retrieves the `INT` value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_long (id IN PLS_INTEGER, name IN VARCHAR2) RETURN NUMBER
Retrieves the long value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_short (id IN PLS_INTEGER, name IN VARCHAR2) RETURN PLS_INTEGER
Retrieves the short value with the specified name. If there is no item by this name, then null is returned. Because there is no `short` type in PL/SQL, `INT` is used to represent a `short`. Although the PL/SQL users get an `PLS_INTEGER`, they are guaranteed that the value is in the Java short value range. If this value is issued with a `set_short` function, then there will not be an out of range error. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_string (id IN PLS_INTEGER, name IN VARCHAR2, value OUT NOCOPY CLOB)
Retrieves the string value with the specified name. If there is no item by this name, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

get_names (id IN PLS_INTEGER) RETURN aq$_jms_namearray
Retrieves all the names within the map message and returns them in a varray. Because `aq$_jms_namearray` has a size as 1024 and each element is a `VARCHAR(200)`, this function will return an error if the size of the name array of the payload exceeds the limit. Raises exception ORA-24195 if the size of the name array or the size of a name exceeds the limit.

get_names (id IN PLS_INTEGER, names OUT aq$_jms_namearray, offset IN PLS_INTEGER, length IN PLS_INTEGER) RETURN PLS_INTEGER
Retrieves a portion of the names within the map message. Because `aq$_jms_namearray` has a size as 1024 and each element is a `VARCHAR(200)`, this function will return an error if either limits are exceeded during the retrieval. (This means there is no sense to put a `length` parameter greater than 1024.) The index of the names of a map messages begins from 0. Parameter `offset` is the offset from which to start retrieving.
The function returns the number of names that have been retrieved. The names retrieved is the intersection of the interval [offset, offset+length-1] and interval [0, size-1] where size is the size of this map message. If the intersection is an empty set, then names will be returned as null and the function returns 0 as the number of names retrieved. If users iterate the names by retrieving in small steps, then this can be used to test that there are no more names to read from map message.
Raises exception ORA-24195 if the size of the name array or the size of a name exceed the limit, ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

get_object (id IN PLS_INTEGER, name IN VARCHAR2, value OUT NOCOPY AQ$_JMS_VALUE)
Returns a general value ADT `AQ$_JMS_VALUE`. If there is no item by this name, then null is returned.Users can use the `type` attribute of this ADT to interpret the data. See the map in the `AQ$_JMS_VALUE` ADT for the correspondence among `dbms_jms_plsql` package constants, Java datatype and `AQ$_JMS_VALUE` attribute. Note this member procedure might bring additional overhead compared to other `get` member procedures or functions. It is used only if the user does not know the datatype of the fields within a message before hand. Otherwise it is a good idea to use a specific `get` member procedure or function. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

get_size (id IN PLS_INTEGER) RETURN PLS_INTEGER
Retrieves the size of the map message. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

item_exists (id IN PLS_INTEGER, name IN VARCHAR2) RETURN BOOLEAN
Indicates that an item exists in this map message by returning `TRUE`. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

308.9.5 SYS.AQ$_JMS_STREAM_MESSAGE Type

This type is the ADT used to store a `StreamMessage` in an Oracle Database Advanced Queuing queue.

This section contains these topics:

- CONSTRUCT aq$_jms_stream_message Function
- JMS Header Methods
- System Properties Methods
- User Properties Methods
- Payload Methods

Syntax

```
TYPE aq$_jms_stream_message AS object(
 header     aq$_jms_header,
 bytes_len  int,
 bytes_raw  raw(2000),
 bytes_lob  blob,
 STATIC FUNCTION  construct RETURN aq$_jms_stream_message,
 MEMBER PROCEDURE set_replyto  (replyto IN sys.aq$_agent),
 MEMBER PROCEDURE set_type     (type    IN VARCHAR),
 MEMBER FUNCTION get_replyto  RETURN sys.aq$_agent,
 MEMBER FUNCTION get_type     RETURN VARCHAR,
 MEMBER PROCEDURE set_userid   (userid   IN VARCHAR),
 MEMBER PROCEDURE set_appid    (appid    IN VARCHAR),
 MEMBER PROCEDURE set_groupid  (groupid  IN VARCHAR),
 MEMBER PROCEDURE set_groupseq (groupseq IN INT),
 MEMBER FUNCTION get_userid   RETURN VARCHAR,
 MEMBER FUNCTION get_appid    RETURN VARCHAR,
 MEMBER FUNCTION get_groupid  RETURN VARCHAR,
 MEMBER FUNCTION get_groupseq RETURN INT,
 MEMBER PROCEDURE clear_properties,
 MEMBER PROCEDURE set_boolean_property(property_name IN VARCHAR,
   property_value IN BOOLEAN),
 MEMBER PROCEDURE set_byte_property   (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_double_property (property_name IN VARCHAR,
   property_value IN DOUBLE PRECISION),
 MEMBER PROCEDURE set_float_property  (property_name IN VARCHAR,
   property_value IN FLOAT),
 MEMBER PROCEDURE set_int_property    (property_name IN VARCHAR,
   property_value IN INT),
 MEMBER PROCEDURE set_long_property   (property_name IN VARCHAR,
   property_value IN NUMBER),
 MEMBER PROCEDURE set_short_property  (property_name IN VARCHAR,
   property_valuE IN INT),
 MEMBER PROCEDURE set_string_property (property_name IN VARCHAR,
   property_value IN VARCHAR),
 MEMBER FUNCTION get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN,
 MEMBER FUNCTION get_byte_property    (property_name IN VARCHAR) RETURN INT,
 MEMBER FUNCTION get_double_property  (property_name IN VARCHAR)
```

```
  RETURN DOUBLE PRECISION,
MEMBER FUNCTION get_float_property   (property_name IN VARCHAR) RETURN FLOAT,
MEMBER FUNCTION get_int_property     (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_long_property    (property_name IN VARCHAR) RETURN NUMBER,
MEMBER FUNCTION get_short_property   (property_name IN VARCHAR) RETURN INT,
MEMBER FUNCTION get_string_property  (property_name IN VARCHAR) RETURN VARCHAR,
MEMBER PROCEDURE set_bytes           (payload IN RAW),
MEMBER PROCEDURE set_bytes           (payload IN BLOB),
MEMBER PROCEDURE get_bytes           (payload OUT RAW),
MEMBER PROCEDURE get_bytes           (payload OUT BLOB),
MEMBER FUNCTION  prepare             (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER PROCEDURE reset               (id IN PLS_INTEGER),
MEMBER PROCEDURE flush               (id IN PLS_INTEGER),
MEMBER PROCEDURE clear_body          (id IN PLS_INTEGER),
MEMBER PROCEDURE clean               (id IN PLS_INTEGER),
STATIC PROCEDURE clean_all,
MEMBER FUNCTION  get_mode       (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION  read_boolean   (id IN PLS_INTEGER) RETURN BOOLEAN,
MEMBER FUNCTION  read_byte      (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION  read_bytes     (id IN PLS_INTEGER) RETURN BLOB,
MEMBER PROCEDURE read_bytes     (id IN PLS_INTEGER, value OUT NOCOPY BLOB),
MEMBER FUNCTION  read_char      (id IN PLS_INTEGER) RETURN CHAR,
MEMBER FUNCTION  read_double    (id IN PLS_INTEGER) RETURN DOUBLE PRECISION,
MEMBER FUNCTION  read_float     (id IN PLS_INTEGER) RETURN FLOAT,
MEMBER FUNCTION  read_int       (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION  read_long      (id IN PLS_INTEGER) RETURN NUMBER,
MEMBER FUNCTION  read_short     (id IN PLS_INTEGER) RETURN PLS_INTEGER,
MEMBER FUNCTION  read_string RETURN CLOB,
MEMBER PROCEDURE read_string    (id IN PLS_INTEGER, value OUT NOCOPY CLOB),
MEMBER PROCEDURE read_object    (id IN PLS_INTEGER,
  value OUT NOCOPY AQ$_JMS_VALUE),
MEMBER PROCEDURE write_boolean  (id IN PLS_INTEGER, value IN BOOLEAN),
MEMBER PROCEDURE write_byte     (id IN PLS_INTEGER, value IN INT),
MEMBER PROCEDURE write_bytes    (id IN PLS_INTEGER, value IN RAW),
MEMBER PROCEDURE write_bytes    (id IN PLS_INTEGER, value IN RAW,
  offset IN INT, length IN INT),
MEMBER PROCEDURE write_bytes    (id IN PLS_INTEGER, value IN BLOB),
MEMBER PROCEDURE write_bytes    (id IN PLS_INTEGER, value IN BLOB,
  offset IN INT, length IN INT),
MEMBER PROCEDURE write_char     (id IN PLS_INTEGER, value IN CHAR),
MEMBER PROCEDURE write_double   (id IN PLS_INTEGER, value IN DOUBLE PRECISION),
MEMBER PROCEDURE write_float    (id IN PLS_INTEGER, value IN FLOAT),
MEMBER PROCEDURE write_int      (id IN PLS_INTEGER, value IN PLS_INTEGER),
MEMBER PROCEDURE write_long     (id IN PLS_INTEGER, value IN NUMBER),
MEMBER PROCEDURE write_short    (id IN PLS_INTEGER, value IN PLS_INTEGER),
MEMBER PROCEDURE write_string   (id IN PLS_INTEGER, value IN VARCHAR2),
MEMBER PROCEDURE write_string   (id IN PLS_INTEGER, value IN CLOB));
```

CONSTRUCT aq$_jms_stream_message Function

STATIC FUNCTION construct RETURN aq$_jms_stream_message
Creates an empty `aq$_jms_stream_message` object.

JMS Header Methods

set_replyto (replyto IN sys.aq$_agent)
Sets the `replyto` parameter, which corresponds to `JMSReplyTo` in JMS.

set_type (type IN VARCHAR)
Sets the JMS type, which can be any text, and which corresponds to `JMSType` in JMS.

get_replyto RETURN sys.aq$_agent
Returns `replyto`, which corresponds to `JMSReplyTo`.

get_type RETURN VARCHAR
Returns `type`, which corresponds to `JMSType`.

System Properties Methods

set_userid (userid IN VARCHAR)
Sets `userid`, which corresponds to `JMSXUserID` in JMS.

set_appid (appid IN VARCHAR)
Sets `appid`, which corresponds to `JMSXAppID` in JMS.

set_groupid (groupid IN VARCHAR)
Sets `groupid`, which corresponds to `JMSXGroupID` in JMS.

set_groupseq (groupseq IN INT)
Sets `groupseq`, which corresponds to `JMSXGroupSeq` in JMS.

get_userid RETURN VARCHAR
Returns `userid`, which corresponds to `JMSXUserID`.

get_appid RETURN VARCHAR
Returns `appid`, which corresponds to `JMSXAppID`.

get_groupid RETURN VARCHAR
Returns `groupid`, which corresponds to `JMSXGroupID`.

get_groupseq RETURN NUMBER
Returns `groupseq`, which corresponds to `JMSXGroupSeq`.

User Properties Methods

clear_properties
Clears all user properties. This procedure does not affect system properties.

set_boolean_property (property_name IN VARCHAR, property_value IN BOOLEAN)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value` in an internal representation. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_byte_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -128 to 127 (8-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `BYTE` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_double_property (property_name IN VARCHAR, property_value IN DOUBLE PRECISION)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_float_property (property_name IN VARCHAR, property_value IN FLOAT)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_int_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -2147483648 to 2147483647 (32-bits). This check is necessary because in PL/SQL and Oracle Database, the `INT` datatype is 38 bits. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_long_property (property_name IN VARCHAR, property_value IN NUMBER)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. In PL/SQL and Oracle Database, the `NUMBER` datatype is 38 bits. In Java, the `long` datatype is 64 bits. Therefore, no range check is needed.Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

set_short_property (property_name IN VARCHAR, property_value IN INT)
Checks whether `property_name` is null or exists. If not, the procedure checks whether `property_value` is within -32768 to 32767 (16-bits). This check is necessary because neither PL/SQL nor RDBMS defines the `short` datatype. Raises exception ORA-24191 if the property name exists, ORA-24192 if the property name is null, or ORA-24193 if the property value exceeds the valid range.

set_string_property (property_name IN VARCHAR, property_value IN VARCHAR)
Checks whether `property_name` is null or exists. If not, the procedure stores `property_value`. Raises exception ORA-24191 if the property name exists or ORA-24192 if the property name is null.

get_boolean_property (property_name IN VARCHAR) RETURN BOOLEAN
If the property with the corresponding property name passed in exists, and if it is a `BOOLEAN` property, then this function returns the value of the property. Otherwise it returns a null.

get_byte_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `BYTE` property, then this function returns the value of the property. Otherwise it returns a null.

get_double_property (property_name IN VARCHAR) RETURN DOUBLE PRECISION
If the property with the corresponding property name passed in exists, and if it is a `DOUBLE` property, then this function returns the value of the property. Otherwise it returns a null.

get_float_property (property_name IN VARCHAR) RETURN FLOAT
If the property with the corresponding property name passed in exists, and if it is a `FLOAT` property, then this function returns the value of the property. Otherwise it returns a null.

get_int_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `Integer` property, then this function returns the value of the property. Otherwise it returns a null.

get_long_property (property_name IN VARCHAR) RETURN NUMBER
If the property with the corresponding property name passed in exists, and if it is a `long` property, then this function returns the value of the property. Otherwise it returns a null.

get_short_property (property_name IN VARCHAR) RETURN INT
If the property with the corresponding property name passed in exists, and if it is a `short` property, then this function returns the value of the property. Otherwise it returns a null.

get_string_property (property_name IN VARCHAR) RETURN VARCHAR
If the property with the corresponding property name passed in exists, and if it is a `STRING` property, then this function returns the value of the property. Otherwise it returns a null.

Payload Methods

get_bytes (payload OUT RAW)
Puts the internal payload into a RAW variable without any interpretation. The payload of type `aq$_jms_stream_message` is stored as either `RAW` or `BLOB` in the database. This member function gets a payload as raw bytes without interpreting it. Raises exception ORA-24190 if the length of internal payload is more than 32767.

get_bytes (payload OUT BLOB)
Puts the internal payload into a BLOB variable without any interpretation. The payload of type `aq$_jms_stream_message` is stored as either `RAW` or `BLOB` in the database. This member function gets a payload as a `BLOB` variable without interpreting it.

set_bytes (payload IN RAW)
Sets the internal payload as the RAW variable without any interpretation. The payload of type `aq$_jms_stream_message` is stored as either `RAW` or `BLOB` in the database. This member function sets a payload as raw bytes without interpreting it.

set_bytes (payload IN BLOB)
Sets the internal payload as the `BLOB` variable without any interpretation. The payload of type `aq$_jms_stream_message` is stored as either `RAW` or `BLOB` in the database. This member function sets a payload as a `BLOB` variable without interpreting it.

prepare (id IN PLS_INTEGER) RETURN PLS_INTEGER
Takes the byte array stored in `aq$_jms_stream_message` and decodes it as a Java object in the Java stored procedure. The result of the decoding is stored as a static variable in Jserv session memory. Parameter `id` is used to identify the slot where the Java object is stored in the Oracle Database JVM session memory. If `id` is null, then a new slot is created for this PL/SQL object. Subsequent JMS operations on the payload need to provide this operation ID. This function also sets the message access mode to `MESSAGE_ACCESS_READONLY`. Subsequent calls of `write_XXX` procedure raise an ORA-24196 error. Users can call the `clear_body` procedure to set the message access mode to `MESSAGE_ACCESS_READONLY`.
This function raises ORA-24197 if the Java stored procedure throws an exception during execution, ORA-24198 if the operation ID is invalid, or ORA-24199 if the Java stored procedure message store overflows.

reset (id IN PLS_INTEGER)
Resets the starting position of the stream to the beginning and puts the stream message in `MESSAGE_ACCESS_READONLY` mode.

flush (id IN PLS_INTEGER)
Takes the static variable in Jserv and synchronizes the content back to `aq$_jms_stream_message`. This procedure will not affect the underlying access mode. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

clear_body (id IN PLS_INTEGER)
Sets the Java stored procedure static variable to empty payload. Parameter `id` is used to identify the slot where the Java object is stored in the Oracle Database JVM session memory. If `id` is null, a new slot is created for this PL/SQL object. Subsequent JMS operations on the payload need to provide this operation ID.
It also sets the message access mode to `MESSAGE_ACCESS_WRITEONLY`. Later calls of `read_XXX` procedure raise ORA-24196 error. Users can call the `reset` or `prepare` procedures to set the message access mode to `MESSAGE_ACCESS_READONLY`. Write-only and read-only modes affect only the payload functions of `AQ$_JMS_BYTES_MESSAGE`. They do not affect the header functions.
This function raises ORA-24197 if the Java stored procedure throws an exception during execution, ORA-24198 if the operation ID is invalid, or ORA-24199 if the Java stored procedure message store overflows.

clean (id IN PLS_INTEGER)
Closes and cleans up the `DataInputStream` or `DataOutputStream` at the Java stored procedure side corresponding to the operation ID. It is very important to call this procedure to avoid memory leaks. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

clean_all
Closes and cleans up all the messages in the corresponding type of message store at the Java stored procedure side. This procedure raises ORA-24197 if the Java stored procedure throws an exception during execution.

get_mode (id IN PLS_INTEGER) RETURN PLS_INTEGER
Returns the current mode of this message. The return value is either `SYS.dbms_aqjms.READ_ONLY` or SYS.`dbms_aqjms.WRITE_ONLY`. Raises exception ORA-24197 if the Java stored procedure throws an exception during execution or ORA-24198 if the operation ID is invalid.

read_boolean (id IN PLS_INTEGER) RETURN BOOLEAN
Reads and returns a Boolean value from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_byte (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads and returns a byte value from the stream message. If the end of the message stream has been reached, then null is returned. Because there is no `BYTE` type in PL/SQL, `INT` is used to represent a byte. Although PL/SQL users get an `INT`, they are guaranteed that the value is in the Java `BYTE` value range. If this value is issued with a `write_byte` function, then there will not be an out of range error. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_bytes (id IN PLS_INTEGER) RETURN BLOB
Reads and returns a byte array from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid or ORA-24196 if the stream message is in write-only mode.

read_bytes (id IN PLS_INTEGER, value OUT NOCOPY BLOB)
Reads a byte array from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_char (id IN PLS_INTEGER) RETURN CHAR
Reads and returns a character value from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_double (id IN PLS_INTEGER) RETURN DOUBLE PRECISION
Reads and returns a double from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_float (id IN PLS_INTEGER) RETURN FLOAT
Reads and returns a float from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_int (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads and returns an `INT` from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_long (id IN PLS_INTEGER) RETURN NUMBER
Reads and returns a long from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_short (id IN PLS_INTEGER) RETURN PLS_INTEGER
Reads and returns a short value from the stream message. If the end of the message stream has been reached, then null is returned. Because there is no short type in PL/SQL, `PLS_INTEGER` is used to represent a `SHORT`. Although PL/SQL users get an `INT`, they are guaranteed that the value is in the Java short value range. If this value is issued with a `write_short` function, then there will not be an out of range error. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_string RETURN CLOB
Reads and returns a string from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion

between the type of real value and the expected type is invalid or ORA-24196 if the stream message is in write-only mode.

read_string (id IN PLS_INTEGER, value OUT NOCOPY CLOB)
Reads a string from the stream message. If the end of the message stream has been reached, then null is returned. Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

read_object (id IN PLS_INTEGER, value OUT NOCOPY AQ$_JMS_VALUE)
Returns a general value ADT `AQ$_JMS_VALUE`. Users can use the type attribute of this ADT to interpret the data. See Table 308-2 for the correspondence among `dbms_jms_plsql` package constants, Java datatype and `AQ$_JMS_VALUE` attribute. This member procedure might bring additional overhead compared to other read member procedures or functions. It is used only if the user does not know the datatype of the fields within a message beforehand. Otherwise it is a good idea to use a specific read member procedure or function.
Raises exception ORA-24194 if the type conversion between the type of real value and the expected type is invalid, ORA-24196 if the stream message is in write-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_boolean (id IN PLS_INTEGER, value IN BOOLEAN)
Writes a Boolean to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_byte (id IN PLS_INTEGER, value IN INT)
Writes a byte to the stream message. Because there is no `BYTE` type in PL/SQL, `INT` is used to represent a byte. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN RAW)
Writes a byte array field to the stream message. Consecutively written byte array fields are treated as two distinct fields when the fields are read. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN RAW, offset IN INT, length IN INT)
Writes a portion of a byte array as a byte array field to the stream message. Consecutively written byte array fields are treated as two distinct fields when the fields are read. Parameter `offset` is the initial offset within the byte array, and parameter `length` is the number of bytes to use. If the range [offset, offset+length] exceeds the boundary of the byte array value, then a Java `IndexOutOfBounds` exception is thrown in the Java stored procedure. The index starts from 0.
Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN BLOB)
Writes a byte array field to the stream message. Consecutively written byte array fields are treated as two distinct fields when the fields are read. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_bytes (id IN PLS_INTEGER, value IN BLOB, offset IN INT, length IN INT)
Writes a portion of a byte array as a byte array field to the stream message. Consecutively written byte array fields are treated as two distinct fields when the fields are read. Parameter `offset` is the initial offset within the byte array, and parameter `length` is the number of bytes to use. If the range [offset, offset+length] exceeds the boundary of the byte array value, then a Java `IndexOutOfBounds` exception is thrown in the Java stored procedure. The index starts from 0.
Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_char (id IN PLS_INTEGER, value IN CHAR)
Writes a character value to the stream message. If this value has multiple characters, then it is the first character that is written. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_double (id IN PLS_INTEGER, value IN DOUBLE PRECISION)
Writes a double to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_float (id IN PLS_INTEGER, value IN FLOAT)
Writes a float to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_int (id IN PLS_INTEGER, value IN PLS_INTEGER)
Writes an `INT` to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_long (id IN PLS_INTEGER, value IN NUMBER)
Writes a long to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_short (id IN PLS_INTEGER, value IN PLS_INTEGER)
Writes a short to the stream message. Because there is no short type in PL/SQL, `INT` is used to represent a short. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_string (id IN PLS_INTEGER, value IN VARCHAR2)
Writes a string to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

write_string (id IN PLS_INTEGER, value IN CLOB)
Writes a string to the stream message. Raises exceptions ORA-24196 if the stream message is in read-only mode, ORA-24197 if the Java stored procedure throws an exception during execution, or ORA-24198 if the operation ID is invalid.

308.9.6 SYS.AQ$_JMS_OBJECT_MESSAGE Type

This type is the ADT used to store an `ObjectMessage` in an Oracle Database Advanced Queuing queue.

Syntax

```
TYPE aq$_jms_object_message AS object(
  header     aq$_jms_header,
  bytes_len  int,
  bytes_raw  raw(2000),
  bytes_lob  blob);
```

308.9.7 SYS.AQ$_JMS_NAMEARRAY Type

This type represents the name array returned by the `get_names` procedure of `aq$_jms_map_message`.

The maximum number of names this type can hold is 1024. The maximum length of each name is 200 characters.

Syntax

```
CREATE OR REPLACE TYPE AQ$_JMS_NAMEARRAY AS VARRAY(1024) OF VARCHAR(100);
```

Usage Notes

If the names array in the message payload is greater than 1024, then use the following function to retrieve the names in multiple portions:

```
MEMBER FUNCTION get_names(id IN PLS_INTEGER, names OUT aq$_jms_namearray,
   offset IN PLS_INTEGER, length IN PLS_INTEGER) RETURN PLS_INTEGER;
```

308.9.8 SYS.AQ$_JMS_VALUE Type

This type represents the general data returned by the `get_object` procedure of `aq$_jms_map_message` and the `read_object` procedure of `aq$_jms_stream_message`.

The `type` field in this ADT is used to decide which type of data this object is really holding. The following table lists the mapping between the `sys.dbms_jms_plsql` type constants, the corresponding Java type, and the data field of ADT `aq$_jms_value` which effectively holds the data.

Syntax

```
CREATE OR REPLACE TYPE AQ$_JMS_VALUE AS object(
 type      number(2),
 num_val   number,
 char_val  char(1),
 text_val  clob,
 bytes_val blob);
```

Table 308-2 AQ$_JMS_VALUE Type Fields and Java Fields

Type	Java Type	aq$_jms_value Data Field
DBMS_JMS_PLSQL.DATA_TYPE_BYTE	byte	num_val

Table 308-2 (Cont.) AQ$_JMS_VALUE Type Fields and Java Fields

Type	Java Type	aq$_jms_value Data Field
DBMS_JMS_PLSQL.DATA_TYPE_SHORT	short	num_val
DBMS_JMS_PLSQL.DATA_TYPE_INTEGER	int	num_val
DBMS_JMS_PLSQL.DATA_TYPE_LONG	long	num_val
DBMS_JMS_PLSQL.DATA_TYPE_FLOAT	float	num_val
DBMS_JMS_PLSQL.DATA_TYPE_DOUBLE	double	num_val
DBMS_JMS_PLSQL.DATA_TYPE_BOOLEAN	boolean	num_val: 0 FALSE, 1 TRUE
DBMS_JMS_PLSQL.DATA_TYPE_CHARACTER	char	char_val
DBMS_JMS_PLSQL.DATA_TYPE_STRING	java.lang.String	text_val
DBMS_JMS_PLSQL.DATA_TYPE_BYTES	byte[]	bytes_val

308.9.9 SYS.AQ$_JMS_EXCEPTION Type

This type represents a Java exception thrown on the Java stored procedure side.

The `id` field is reserved for future use. The `exp_name` stores the Java exception name, the `err_msg` field stores the Java exception error message, and the `stack` field stores the stack trace of the Java exception.

Syntax

```
CREATE OR REPLACE TYPE AQ$_JMS_EXCEPTION AS OBJECT (
   id          number, -- Reserved and not used. Right now always return 0.
   exp_name    varchar(200),
   err_msg     varchar(500),
   stack       varchar(4000));
```

309
JSON Data Structures

To work with JSON data in PL/SQL, you can use these data structures.

- JSON_ELEMENT_T Object Type
- JSON_OBJECT_T Object Type
- JSON_ARRAY_T Object Type
- JSON_SCALAR_T Object Type
- JSON_KEY_LIST Type

> **Note:**
>
> The time component in the `DATE` value of the `get_Date` function is truncated. If you want to preserve the time component, use `get_Timestamp` and then use `CAST AS DATE` to convert to a SQL date with the time component.

Related Topics

- *Oracle Database JSON Developer's Guide*
- The JSON standard

309.1 JSON_ELEMENT_T Object Type

`JSON_ELEMENT_T` is the supertype for the `JSON_OBJECT_T`, `JSON_SCALAR_T`, and `JSON_ARRAY_T` object types.

Description

Note the following:

- To create an instance of JSON_ELEMENT_T, use the `parse` or `load` function. See the "Constructors" section below for details.
- You cannot create an empty JSON_ELEMENT_T instance. To create an empty JSON container, create it based on one of the subtypes.
- To cast a JSON_ELEMENT_T into a subtype (for example, JSON_OBJECT_T), you need to perform an explicit cast using TREAT AS. For example:

```
TREAT (elem AS JSON_OBJECT_T)
```

Constructors

You create a JSON_ELEMENT_T instance using the `parse` or `load` function. The `parse` function takes as input VARCHAR2, CLOB, or BLOB data, and returns a JSON_ELEMENT_T instance. Similarly, the `load` function takes JSON type as an input.

```
STATIC FUNCTION parse(jsn VARCHAR2) RETURN JSON_ELEMENT_T
STATIC FUNCTION parse(jsn CLOB)     RETURN JSON_ELEMENT_T
STATIC FUNCTION parse(jsn BLOB)     RETURN JSON_ELEMENT_T
STATIC FUNCTION parse(jsn BLOB, FORMAT IN VARCHAR2) RETURN JSON_ELEMENT_T
STATIC FUNCTION load(jsn JSON)      RETURN JSON_Element_T,
```

Only UTF8-encoded JSON is passed as a BLOB.

The `parse` function takes a JSON string as input and sets up an internal representation of the JSON data. If the provided input is not valid JSON, then an error message is raised. Valid JSON has to pass the lax check of the “IS JSON” SQL condition. The input for the `load` function is a JSON type, therefore, no JSON syntax check is needed.

Serialization and Conversions

A `JSON_ELEMENT_T` instance (and all subtypes) can be serialized to a JSON string, converted to a JSON type or (if it is a scalar value) converted to a SQL value like Date or Number. Serialization is the inverse of the parse function: a string representation of the in-memory representation of the JSON data is being generated and returned as an appropriate SQL type.

The serialization and conversion functions are:

```
MEMBER FUNCTION to_String    RETURN VARCHAR2
MEMBER FUNCTION stringify    RETURN VARCHAR2

MEMBER FUNCTION to_Clob      RETURN CLOB
MEMBER FUNCTION to_Blob      RETURN BLOB
MEMBER PROCEDURE to_Clob(c IN OUT CLOB)
MEMBER PROCEDURE to_Blob(c IN OUT BLOB)

MEMBER FUNCTION to_Json      RETURN JSON

MEMBER FUNCTION to_Number    RETURN NUMBER
MEMBER FUNCTION to_Date      RETURN DATE
MEMBER FUNCTION to_Timestamp RETURN TIMESTAMP
MEMBER FUNCTION to_Boolean   RETURN BOOLEAN
```

The FUNCTION stringify is synonym of to_String. It has the same functionality.

The `to_Clob` and `to_Blob` *procedures* accept a CLOB or BLOB input and enable you to provide a LOB to be used as the serialization destination. For instance, `EMPTY_LOB` can be provided. The input LOB cannot be `NULL`.

If the `to_Clob` *function* is used, then a new CLOB is created. If you do not want to create a CLOB first, then you can use the `to_Clob` or `to_Blob` *functions*, which take no parameter and generate a temp LOB.

`to_Blob` serializes to UTF8 format only.

Introspection

Introspection enables you to discover properties of JSON objects without modifying them. The introspection functions are:

```
MEMBER FUNCTION has(key VARCHAR2) RETURN BOOLEAN,
MEMBER FUNCTION is_Object         RETURN BOOLEAN
MEMBER FUNCTION is_Array          RETURN BOOLEAN
MEMBER FUNCTION is_Scalar         RETURN BOOLEAN
```

```
MEMBER FUNCTION is_String         RETURN BOOLEAN
MEMBER FUNCTION is_Number         RETURN BOOLEAN
MEMBER FUNCTION is_Boolean        RETURN BOOLEAN
MEMBER FUNCTION is_True           RETURN BOOLEAN
MEMBER FUNCTION is_False          RETURN BOOLEAN
MEMBER FUNCTION is_Null           RETURN BOOLEAN
MEMBER FUNCTION is_Date           RETURN BOOLEAN
MEMBER FUNCTION is_Timestamp      RETURN BOOLEAN
MEMBER FUNCTION get_Size          RETURN NUMBER
```

The `has` function checks if a given name exists in the JSON_OBJECT_T object.

The return value of the `get_size` function depends on the JSON type:

- For a scalar, it returns `1`.
- For an object, it returns the number of keys.
- For an array, it returns the number of items.

Note that textual JSON does not support dates and timestamps natively. Instead they are typically modeled as strings. You can use the JSON type to preserve dates and timestamps natively. The Document Object Model (DOM) interface enables you to add dates and timestamps as scalar values and preserve them until serialization to JSON, where they are printed as string following the ISO 8601 format. If a SQL value of type date or timestamp has been added, then the `is_Date` and `is_Timestamp` functions return true. If a date has been added as a string (e.g. as ISO 8601), then the `is_Date` and `is_Timestamp` functions return false. You can use Oracle conversion function `to_Date` and `to_Timestamp` to convert a string representation of a date, and timestamp to the Oracle representation.

Comparing

These functions enables the users to compare two JSON elements or objects or arrays or scalars . The comparing functions are:

```
MEMBER FUNCTION equals(cmp <JSON_ELEMENT_T or JSON>) RETURN BOOLEAN
MEMBER FUNCTION equals(cmp JSON_ELEMENT_T,
order_by_compare BOOLEAN DEFAULT FALSE)              RETURN NUMBER
MEMBER FUNCTION equals(cmp JSON,
order_by_compare BOOLEAN DEFAULT FALSE)              RETURN NUMBER
MEMBER FUNCTION compare_to(JSON_SCALAR_T)            RETURN NUMBER
```

The comparing functions returns a numeric values corresponding to the outcome of the comparison function. 0 indicates that the compared values are same. -1 indicates the value is smaller than the compared value. +1 indicates that the value is greater than the compared value. NULL value is returned if the values can not be compared due to type incompatibility. If the Boolean argument `order_by_compare` is set to true, incompatible types could also be compared. In this case, a comparison value of NULL would never be returned.

Error Handling

You can set the level of error handling for JSON processing. You might not want an error to be raised for every mismatch. The `on_Error` procedure enables you to specify when errors should be raised:

```
MEMBER PROCEDURE on_Error(value NUMBER)
```

The `on_Error` procedure defines what happens if an error is encountered during a PL/SQL operation, for instance, a get call.

The default is to not raise an error but to return NULL instead.

You invoke `On_error` on a JSON_ELEMENT_T instance and it sets the error behavior for all subsequent calls. To reset the behavior to the default, you can call `on_Error(0)`.

Values for the `value` parameter are:

Table 309-1 Values for the value Parameter in the ON_ERROR Procedure

Value	Description
0	Reset to the default behavior, which is to return NULL instead of raising an error.
1	Raise all errors.
2	Raise an error if no value is detected.
3	Raise an error if the data types do not match, for example, if you call GET_NUMBER on a string value.
4	Raise an error if the input is invalid, for example, if the array is out of bounds.

You can combine values. For example, you can specify 7 to indicate a combination of 3 and 4.

In the following example, an error is raised because the value of “a” is “xyz”, which cannot be converted to a number. If the `on_Error` procedure had not been called, then NULL would be returned, and no error would be raised.

```
declare
     jo JSON_OBJECT_T;
begin
     jo := JSON_OBJECT_T.parse('{a:"xyz"}');
     jo.On_error(1);
     dbms_output.put_line(jo.get_Number('a'));
end;
/
```

309.2 JSON_OBJECT_T Object Type

`JSON_OBJECT_T` is a subtype of the JSON_ELEMENT_T object type. It corresponds to the JSON object structure.

Constructors

You can create an empty JSON_OBJECT_T instance using the following constructor.

```
   CONSTRUCTOR FUNCTION JSON_OBJECT_T RETURN SELF AS RESULT
```

You can create a JSON_OBJECT_T instance using one of the following `parse` functions:

```
STATIC FUNCTION parse(jsn VARCHAR2) RETURN JSON_OBJECT_T
STATIC FUNCTION parse(jsn CLOB)     RETURN JSON_OBJECT_T
STATIC FUNCTION parse(jsn BLOB)     RETURN JSON_OBJECT_T
STATIC FUNCTION parse(jsn BLOB, FORMAT IN VARCHAR2) RETURN JSON_OBJECT_T
```

You can also create a JSON_OBJECT_T instance using one of the following constructors:

```
CONSTRUCTOR FUNCTION JSON_OBJECT_T(jsn VARCHAR2) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_OBJECT_T(jsn CLOB) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_OBJECT_T(jsn BLOB) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_Object_T(jsn BLOB, FORMAT IN VARCHAR2) RETURN SELF AS RESULT
```

```
CONSTRUCTOR FUNCTION JSON_OBJECT_T(jsn JSON) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_OBJECT_T(e JSON_ELEMENT_T) RETURN SELF AS RESULT
```

Only UTF8-encoded JSON is passed as a BLOB.

The `parse` function takes a JSON string as input and sets up an internal representation of the JSON data. If the provided input is not a valid JSON object, then an error message is raised. The input has to specify a JSON object, not an array.

Get Functions and Procedures

The following functions and procedures enable you to retrieve the value of the JSON object:

```
MEMBER FUNCTION get(key VARCHAR2)              RETURN JSON_ELEMENT_T
MEMBER FUNCTION get_String(key VARCHAR2)       RETURN VARCHAR2
MEMBER FUNCTION get_Number(key VARCHAR2)       RETURN NUMBER
MEMBER FUNCTION get_Date(key VARCHAR2)         RETURN DATE
MEMBER FUNCTION get_Timestamp(key VARCHAR2)    RETURN TIMESTAMP
MEMBER FUNCTION get_TimestampTZ(key VARCHAR2) RETURN TIMESTAMP WITH TIME ZONE
MEMBER FUNCTION get_Boolean(key VARCHAR2)      RETURN BOOLEAN
MEMBER FUNCTION get_Clob(key VARCHAR2)         RETURN CLOB
MEMBER FUNCTION get_Blob(key VARCHAR2)         RETURN BLOB
MEMBER FUNCTION get_Object(key VARCHAR2)       RETURN JSON_OBJECT_T
MEMBER FUNCTION get_Array(key VARCHAR2)        RETURN JSON_ARRAY_T
MEMBER FUNCTION get_Json(key NUMBER)           RETURN JSON

MEMBER PROCEDURE get_Json(key NUMBER, j OUT NOCOPY JSON)
MEMBER PROCEDURE get_Clob(key NUMBER, c IN OUT CLOB)
MEMBER PROCEDURE get_Blob(key NUMBER, c IN OUT BLOB)
```

Note:

- The `get` function has reference semantics. This means that if the returned JSON_ELEMENT_T is modified, then the containing JSON_ELEMENT_T is also changed. See the Reference Semantics section below for details.
- The `GET_STRING` function converts the value to a string if it is not already a string. Thus, the `GET_STRING` function returns a non-null value even if IS_STRING returns false.
- All the “get” functions perform a conversion if possible. If no conversion is possible, then an error might be raised, depending on what `ON_ERROR` is set to.

The `GET_CLOB` and `GET_BLOB` *procedures*, which accept a CLOB or BLOB as input, enable you to provide a LOB to be used as serialization destination. For instance, EMPTY_LOB can be provided. If you use the `GET_CLOB` *function* instead, then a new CLOB will be created implicitly. The input LOB cannot be NULL. If you do not want to create a BLOB or CLOB first, you can use the GET_CLOB or GET_BLOB functions, which take no parameter and generates a temp LOB.

`GET_BLOB` serializes to UTF8 format only.

Set Procedures

The following procedures enable you to set the value of a JSON object. Existing values, if present, will be overwritten.

```
MEMBER PROCEDURE put(key VARCHAR2, val JSON_ELEMENT_T)
MEMBER PROCEDURE put(key VARCHAR2, val VARCHAR2)
MEMBER PROCEDURE put(key VARCHAR2, val NUMBER)
MEMBER PROCEDURE put(key VARCHAR2, val BOOLEAN)
MEMBER PROCEDURE put(key VARCHAR2, val DATE)
MEMBER PROCEDURE put(key VARCHAR2, val TIMESTAMP)
```

```
MEMBER PROCEDURE put_Null(key VARCHAR2)
MEMBER PROCEDURE put(key VARCHAR2, val JSON)
```

Introspection Functions

Introspection enables you to discover properties of JSON objects without modifying them. The introspection functions are:

```
MEMBER FUNCTION has(key VARCHAR2)      RETURN BOOLEAN
MEMBER FUNCTION get_Type(key VARCHAR2) RETURN VARCHAR2
MEMBER FUNCTION get_Keys               RETURN JSON_KEY_LIST
```

The get_Keys function returns an object type of JSON_KEY_LIST, which is a varray of VARCHAR2(4000). The varray contains the names of keys in the JSON object. The get_Keys function returns at most 32767 field names for a given JSON object. An error is raised if it is applied to an object with more than 32767 fields.

The example below walks through the items of the varray to build a JSON_ARRAY_T object that contains all key names.

```
declare
  jo JSON_OBJECT_T;
  ja JSON_ARRAY_T;
  keys JSON_KEY_LIST;
  keys_string VARCHAR2(100);

begin
  ja := new JSON_ARRAY_T;
  jo := JSON_OBJECT_T.parse('{"name":"fred",
                             "jobTitle":"codemonkey",
                             "projects":["json", "xml"]}');
  keys := jo.get_keys;
  for i in 1..keys.count loop
     ja.append(keys(i));
  end loop;

  keys_string := ja.to_string;
  dbms_output.put_line(keys_string);
end;
/
```

The output is:

```
["name","jobTitle","projects"]
```

Modification Procedures

The following procedures enable you to remove or rename keys in a JSON object.

```
MEMBER PROCEDURE remove(key VARCHAR2)
MEMBER PROCEDURE rename_Key(keyOld VARCHAR2, keyNew VARCHAR2)
```

Duplicate key names are not supported and will raise an error.

Clone Function

This function makes a copy of the JSON object. Reference semantics is changed to value semantics.

```
MEMBER FUNCTION clone RETURN JSON_OBJECT_T
```

Reference Semantics

Calling the `get` function that returns a JSON_ELEMENT_T object always returns a reference to the complex values instead of a copy. This means that changing the returned value affects its container. See the example below.

```
declare
  data    JSON_OBJECT_T;
  address JSON_OBJECT_T;
  zip     number;

begin

  data := new JSON_OBJECT_T('{
    "first": "John",
    "last": "Doe",
    "address": {
        "country": "USA",
        "zip": "94065"
    }
  }');

  address := data.get_object('address');
  dbms_output.put_line(address.to_string);

  -- 1) VALUE SEMANTICS for scalar values
  -- (changing the value has no effect on container)
  zip := address.get_number('zip');
  dbms_output.put_line(zip);
  zip := 12345;
  dbms_output.put_line(zip);
  -- address is still the same
  dbms_output.put_line(address.to_string);


  -- 2) REFERENCE SEMANTICS for complex values
  --    'address' is a reference to the complex address values inside 'data'
  address.put('zip', 12345);
  address.put('street', 'Detour Road');
  dbms_output.put_line(data.to_string);
end;
/
```

In cases where you do not want the reference semantics, you can use the `clone` function to create a copy of the returned object. This decouples the value from its container. In the example above, you can create a copy of the “address” object by replacing this line:

```
address := data.get_object('address');
```

with this line:

```
address := data.get_object('address').clone;
```

After this, changing the address with have no effect on the value of the ‘data’ containing object.

Update Example

The following example updates the price of an item by 10%.

```
WITH
  FUNCTION updatePrice(jsonTxt in VARCHAR2 ) RETURN VARCHAR2 IS
```

```
  jo JSON_OBJECT_T;
  oldPrice NUMBER;

  BEGIN
    jo := new JSON_OBJECT_T(jsonTxt);
    oldPrice := jo.get_number('price');
    jo.put('price', oldPrice * 1.1);
    RETURN jo.to_string();
  END;
SELECT updatePrice(col)
FROM   t1;
```

309.3 JSON_ARRAY_T Object Type

`JSON_ARRAY_T` is a subtype of the `JSON_ELEMENT_T` object type. `JSON_ARRAY_T` corresponds to the JSON array structure.

Constructors

You can create an empty JSON_ARRAY_T instance using the following constructor.

```
CONSTRUCTOR FUNCTION JSON_ARRAY_T RETURN SELF AS RESULT
```

You can create a JSON_ARRAY_T instance using one the following `parse` functions:

```
STATIC FUNCTION parse(jsn VARCHAR2)                   RETURN JSON_ARRAY_T
STATIC FUNCTION parse(jsn CLOB)                       RETURN JSON_ARRAY_T
STATIC FUNCTION parse(jsn BLOB)                       RETURN JSON_ARRAY_T
STATIC FUNCTION parse(jsn BLOB, FORMAT IN VARCHAR2)   RETURN JSON_ARRAY_T
```

You can also create a JSON_ARRAY_T instance using one the following constructors:

```
CONSTRUCTOR FUNCTION JSON_Array_T(jsn VARCHAR2) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_Array_T(jsn CLOB) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_Array_T(jsn BLOB) RETURN SELF AS RESULT
CONSTRUCTOR FUNCTION JSON_Array_T(e JSON_ELEMENT_T) RETURN SELF AS RESULT
```

Only UTF8-encoded JSON is passed as a BLOB.

The `parse` function takes a JSON string as input and sets up an internal representation of the JSON data. If the provided input is not valid JSON, then an error message is raised. The input has to specify a JSON array, not an object.

Get Functions and Procedures

The following functions and procedures enable you to retrieve the value of the JSON array:

```
MEMBER FUNCTION get(pos NUMBER)           RETURN JSON_ELEMENT_T
MEMBER FUNCTION get_String(pos NUMBER)    RETURN VARCHAR2
MEMBER FUNCTION get_Number(pos NUMBER)    RETURN NUMBER
MEMBER FUNCTION get_Date(pos NUMBER)      RETURN DATE
MEMBER FUNCTION get_Timestamp(pos NUMBER) RETURN TIMESTAMP
MEMBER FUNCTION get_Boolean(pos NUMBER)   RETURN BOOLEAN
MEMBER FUNCTION get_Clob(pos NUMBER)      RETURN CLOB
MEMBER FUNCTION get_Blob(pos NUMBER)      RETURN BLOB
MEMBER FUNCTION get_Json(pos NUMBER)      RETURN JSON

MEMBER PROCEDURE get_Json(pos NUMBER, j IN OUT NOCOPY JSON)
```

```
MEMBER PROCEDURE get_Clob(pos NUMBER, c IN OUT CLOB)
MEMBER PROCEDURE get_Blob(pos NUMBER, c IN OUT BLOB)
```

Note:

- The `get` function has reference semantics. This means that if the returned JSON_ELEMENT_T is modified, then the containing JSON_ELEMENT_T is changed too.
- The `GET_STRING` function converts the value to a string if it is not already a string. The function returns a non-null value even if `IS_STRING` returns false.
- All the “get” functions perform a conversion if possible. If no conversion is possible, then an error might be raised, depending on what `ON_ERROR` is set to.

The `GET_CLOB` and `GET_BLOB` *procedures*, which accept a `CLOB` or `BLOB` as input, enable you to provide a LOB to be used as serialization destination. For instance, `EMPTY_LOB` can be provided. If you use the `GET_CLOB` *function* instead, then a new `CLOB` will be created implicitly. The input `LOB` cannot be NULL. If you do not want to create a `BLOB` or `CLOB` first, you can use the `GET_CLOB` or `GET_BLOB` functions, which take no parameter and generates a temp LOB.

`GET_BLOB` serializes to UTF8 format only.

Set Procedures

The following procedures enable you to set the value at the specified position in the JSON array. These procedures insert (not overwrite) at the specified position unless overwrite is requested.

```
MEMBER PROCEDURE put(pos NUMBER, val VARCHAR2, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put(pos NUMBER, val NUMBER, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put(pos NUMBER, val BOOLEAN, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put(pos NUMBER, val DATE, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put(pos NUMBER, val JSON_ELEMENT_T, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put(pos NUMBER, val JSON, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put(pos NUMBER, val TIMESTAMP, overwrite BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE put_Null(pos NUMBER, overwrite BOOLEAN DEFAULT FALSE)
```

The following procedures append the specified value to the end of the JSON array:

```
MEMBER PROCEDURE append(val JSON_ELEMENT_T)
MEMBER PROCEDURE append(val VARCHAR2)
MEMBER PROCEDURE append(val NUMBER)
MEMBER PROCEDURE append(val BOOLEAN)
MEMBER PROCEDURE append(val DATE)
MEMBER PROCEDURE append(val JSON)
MEMBER PROCEDURE append(val TIMESTAMP)
MEMBER PROCEDURE append_all(arr JSON_ELEMENT_T, excl_exis BOOLEAN DEFAULT FALSE)
MEMBER PROCEDURE append_Null
```

> **Note:**
>
> No sorting happens in append_all procedure. Deduplication happens only in the arr2, the new array that is being appended. If the original array has duplicates, they are left unaltered.

Introspection Function

Introspection enables you to discover properties of the JSON array without modifying them.

```
MEMBER FUNCTION get_Type(pos NUMBER) RETURN VARCHAR2
MEMBER FUNCTION has_value(val <data_type>) RETURN BOOLEAN
```

The `has_value` function can be called on a JSON_ARRAY_T type instance. It checks if the array contains the `val` argument and returns TRUE; otherwise it returns FALSE. The function takes an argument `val` of data type VARCHAR2, NUMBER, BOOLEAN, DATE, TIMESTAMP, JSON, and JSON_ELEMENT_T.

Modification Procedure

The following procedures enables you to remove the value at the specified position, removes the duplicate values, or sort the elements of the array targeted by the specified path in the JSON array.

```
MEMBER PROCEDURE remove(pos NUMBER)
MEMBER PROCEDURE sort(ascending BOOLEAN DEFAULT TRUE)
MEMBER PROCEDURE deduplicate
```

The `deduplicate` procedure does not reorder the items.

Clone Function

This function makes a copy of the JSON array. Reference semantics is changed to value semantics.

```
MEMBER FUNCTION clone RETURN JSON_ARRAY_T
```

309.4 JSON_SCALAR_T Object Type

`JSON_SCALAR_T` is a subtype of the `JSON_ELEMENT_T` object type.

Description

A `JSON_SCALAR_T` instance captures one scalar values, for example, the string “fred” or the number 1. This type does not have any functions or procedures other than the ones inherited from `JSON_ELEMENT_T`. You cannot create an instance of this type directly.

309.5 JSON_KEY_LIST Type

`JSON_KEY_LIST` is a varray of `VARCHAR2(4000)`.

Description

This type is used by the `get_Keys` function in the JSON_OBJECT_T Object Type.

310

Logical Change Record TYPEs

This chapter describes the logical change record (LCR) types.

This chapter contains these topics:

- Overview
- Security Model
- Summary of Logical Change Record Types
- Common Subprograms for LCR$_DDL_RECORD and LCR$_ROW_RECORD

310.1 Logical Change Record TYPEs Overview

In Replication (Oracle GoldenGate and XStream), logical change records (LCRs) are message payloads that contain information about changes to a database. These changes can include changes to the data, which are data manipulation language (DML) changes, and changes to database objects, which are data definition language (DDL) changes.

When you use Replication (Oracle GoldenGate and XStream), the capture process captures changes in the form of LCRs and enqueues them into a queue. Finally, the apply process can apply LCRs at a destination database. You also have the option of creating, enqueuing, and dequeuing LCRs manually.

310.2 Logical Change Record Types Security Model

`PUBLIC` is granted `EXECUTE` privilege on the types described in this chapter.

310.3 Summary of Logical Change Record Types

This table lists the Logical Change Record TYPEs and briefly describes them.

Table 310-1 Logical Change Record (LCR) Types

Type	Description
LCR$_DDL_RECORD Type	Represents a data definition language (DDL) change to a database object
LCR$_ROW_RECORD Type	Represents a data manipulation language (DML) change to a database object
LCR$_ROW_LIST Type	Identifies a list of column values for a row in a table
LCR$_ROW_UNIT Type	Identifies the value for a column in a row

These logical change record (LCR) types can be used with the following Oracle-supplied PL/SQL packages:

- `DBMS_APPLY_ADM`
- `DBMS_AQ`

- DBMS_AQADM
- DBMS_CAPTURE_ADM
- DBMS_PROPAGATION_ADM
- DBMS_RULE
- DBMS_RULE_ADM
- DBMS_STREAMS
- DBMS_STREAMS_ADM
- DBMS_TRANSFORM

310.3.1 LCR$_DDL_RECORD Type

This type represents a data definition language (DDL) change to a database object.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

If you create or modify a DDL logical change record (DDL LCR), then make sure the `ddl_text` is consistent with the `base_table_name`, `base_table_owner`, `object_type`, `object_owner`, `object_name`, and `command_type` attributes.

This topic contains information about the constructor for row LCRs and information about the member subprograms for this type.

> **Note:**
>
> - When passing a name as a parameter to an LCR constructor, you can enclose the name in double quotes to handle names that use mixed case or lower case for database objects. For example, if a name contains any lower case characters, then you must enclose it in double quotes.
> - The application does not need to specify a transaction identifier or SCN when it creates an LCR because the apply process generates these values and stores them in memory. If a transaction identifier or SCN is specified in the LCR, then the apply process ignores it and assigns a new value.

LCR$_DDL_RECORD Constructor

Creates a `SYS.LCR$_DDL_RECORD` object with the specified information.

```
STATIC FUNCTION CONSTRUCT(
   source_database_name  IN  VARCHAR2,
   command_type          IN  VARCHAR2,
```

```
   object_owner          IN  VARCHAR2,
   object_name           IN  VARCHAR2,
   object_type           IN  VARCHAR2,
   ddl_text              IN  CLOB,
   logon_user            IN  VARCHAR2,
   current_schema        IN  VARCHAR2,
   base_table_owner      IN  VARCHAR2,
   base_table_name       IN  VARCHAR2,
   tag                   IN  RAW       DEFAULT NULL,
   transaction_id        IN  VARCHAR2  DEFAULT NULL,
   scn                   IN  NUMBER    DEFAULT NULL,
   position              IN  RAW       DEFAULT NULL,
   edition_name          IN  VARCHAR2  DEFAULT NULL,
   root_name             IN  VARCHAR2  DEFAULT NULL)
RETURN SYS.LCR$_DDL_RECORD;
```

LCR$_DDL_RECORD Constructor Function Parameters

Table 310-2 Constructor Function Parameters for LCR$_DDL_RECORD

Parameter	Description
source_database_name	The database where the DDL statement occurred If the LCRs originated in a multitenant container database (CDB), then this field specifies the global name of the container where the DDL change occurred. If you do not include the domain name, then the function appends the local domain to the database name automatically. For example, if you specify DBS1 and the local domain is EXAMPLE.COM, then the function specifies DBS1.EXAMPLE.COM automatically. Set this parameter to a non-NULL value.
command_type	The type of command executed in the DDL statement Set this parameter to a non-NULL value. **See Also:** The "SQL Command Codes" table in the *Oracle Call Interface Programmer's Guide* for a complete list of command types The following command types *are not supported* in DDL LCRs: ALTER MATERIALIZED VIEW ALTER MATERIALIZED VIEW LOG ALTER SUMMARY CREATE SCHEMA CREATE MATERIALIZED VIEW CREATE MATERIALIZED VIEW LOG CREATE SUMMARY DROP MATERIALIZED VIEW DROP MATERIALIZED VIEW LOG DROP SUMMARY RENAME The snapshot equivalents of the materialized view command types are also not supported.
object_owner	The user who owns the object on which the DDL statement was executed
object_name	The database object on which the DDL statement was executed

Table 310-2 (Cont.) Constructor Function Parameters for LCR$_DDL_RECORD

Parameter	Description
object_type	The type of object on which the DDL statement was executed The following are valid object types: CLUSTER FUNCTION INDEX LINK OUTLINE PACKAGE PACKAGE BODY PROCEDURE SEQUENCE SYNONYM TABLE TRIGGER TYPE USER VIEW LINK represents a database link. NULL is also a valid object type. Specify NULL for all object types not listed. The GET_OBJECT_TYPE member procedure returns NULL for object types not listed.
ddl_text	The text of the DDL statement Set this parameter to a non-NULL value.
logon_user	The user whose session executed the DDL statement
current_schema	The schema that is used if no schema is specified explicitly for the modified database objects in ddl_text If a schema is specified in ddl_text that differs from the one specified for current_schema, then the function uses the schema specified in ddl_text. Set this parameter to a non-NULL value.
base_table_owner	If the DDL statement is a table-related DDL (such as CREATE TABLE and ALTER TABLE), or if the DDL statement involves a table (such as creating a trigger on a table), then base_table_owner specifies the owner of the table involved. Otherwise, base_table_owner is NULL.
base_table_name	If the DDL statement is a table-related DDL (such as CREATE TABLE and ALTER TABLE), or if the DDL statement involves a table (such as creating a trigger on a table), then base_table_name specifies the name of the table involved. Otherwise, base_table_name is NULL.
tag	A binary tag that enables tracking of the LCR For example, this tag can be used to determine the original source database of the DDL statement if apply forwarding is used.
transaction_id	The identifier of the transaction
scn	The SCN at the time when the change record for a captured LCR was written to the redo log The SCN value is meaningless for a user-created LCR.

Table 310-2 (Cont.) Constructor Function Parameters for LCR$_DDL_RECORD

Parameter	Description
position	The position of the LCR LCR position is commonly used in XStream configurations. Using XStream requires purchasing a license for the Oracle GoldenGate product. **See Also:** *Oracle Database XStream Guide*
edition_name	The name of the edition in which the DDL statement was executed
root_name	If the LCRs is associated with a CDB, then this field specifies the global name of the root in the CDB. If the LCR is associated with a non-CDB, then this field is NULL.

Summary of LCR$_DDL_RECORD Subprograms

Table 310-3 LCR$_DDL_RECORD Type Subprograms

Subprogram	Description
EXECUTE Member Procedure	Executes the LCR under the security domain of the current user
GET_BASE_TABLE_NAME Member Function	Gets the base (dependent) table name
GET_BASE_TABLE_OWNER Member Function	Gets the base (dependent) table owner
GET_CURRENT_SCHEMA Member Function	Gets the default schema (user) name
GET_DDL_TEXT Member Procedure	Gets the DDL text in a CLOB
GET_EDITION_NAME Member Function	Gets the name of the edition in which the DDL statement was executed
GET_LOGON_USER Member Function	Gets the logon user name
GET_OBJECT_TYPE Member Function	Gets the type of the object involved for the DDL
SET_BASE_TABLE_NAME Member Procedure	Sets the base (dependent) table name
SET_BASE_TABLE_OWNER Member Procedure	Sets the base (dependent) table owner
SET_CURRENT_SCHEMA Member Procedure	Sets the default schema (user) name
SET_DDL_TEXT Member Procedure	Sets the DDL text
SET_EDITION_NAME Member Procedure	Sets the name of the edition in which the DDL statement was executed
SET_LOGON_USER Member Procedure	Sets the logon user name
SET_OBJECT_TYPE Member Procedure	Sets the object type
Common Subprograms	See "Common Subprograms for LCR$_DDL_RECORD and LCR$_ROW_RECORD" for a list of subprograms common to the SYS.LCR$_ROW_RECORD and SYS.LCR$_DDL_RECORD types

EXECUTE Member Procedure

Executes the DDL LCR under the security domain of the current user. Apply handlers are not run when the LCR is applied using this procedure.

Syntax

```
MEMBER PROCEDURE EXECUTE;
```

GET_BASE_TABLE_NAME Member Function

Gets the base (dependent) table name.

Syntax

```
MEMBER FUNCTION GET_BASE_TABLE_NAME()
RETURN VARCHAR2;
```

GET_BASE_TABLE_OWNER Member Function

Gets the base (dependent) table owner.

Syntax

```
MEMBER FUNCTION GET_BASE_TABLE_OWNER()
RETURN VARCHAR2;
```

GET_CURRENT_SCHEMA Member Function

Gets the current schema name.

Syntax

```
MEMBER FUNCTION GET_CURRENT_SCHEMA()
RETURN VARCHAR2;
```

GET_DDL_TEXT Member Procedure

Gets the DDL text in a `CLOB`.

For example, the following PL/SQL code uses this procedure to get the DDL text in a DDL LCR:

```
CREATE OR REPLACE PROCEDURE ddl_in_lcr (ddl_lcr in SYS.LCR$_DDL_RECORD)
IS
  ddl_text   CLOB;
BEGIN
  DBMS_OUTPUT.PUT_LINE( '  ----------------------------------------' );
  DBMS_OUTPUT.PUT_LINE( '  Displaying DDL text in a DDL LCR: ' );
  DBMS_OUTPUT.PUT_LINE( '  ----------------------------------------' );
  DBMS_LOB.CREATETEMPORARY(ddl_text, true);
  ddl_lcr.GET_DDL_TEXT(ddl_text);
  DBMS_OUTPUT.PUT_LINE('DDL text:' || ddl_text);
  DBMS_LOB.FREETEMPORARY(ddl_text);
END;
/
```

> **Note:**
>
> `GET_DDL_TEXT` is a member procedure and not a member function to make it easier for you to manage the space used by the `CLOB`. Notice that the previous example creates temporary space for the `CLOB` and then frees the temporary space when it is no longer needed.

Syntax

```
MEMBER FUNCTION GET_DDL_TEXT(
  ddl_text  IN/OUT  CLOB);
```

Parameter

Table 310-4 GET_DDL_TEXT Procedure Parameter

Parameter	Description
`ddl_text`	The DDL text in the DDL LCR

GET_EDITION_NAME Member Function

Gets the name of the edition in which the DDL statement was executed.

> **See Also:**
>
> *Oracle Database Development Guide*

Syntax

```
MEMBER FUNCTION GET_EDITION_NAME()
RETURN VARCHAR2;
```

GET_LOGON_USER Member Function

Gets the logon user name.

Syntax

```
MEMBER FUNCTION GET_LOGON_USER()
RETURN VARCHAR2;
```

GET_OBJECT_TYPE Member Function

Gets the type of the object involved for the DDL.

Syntax

```
MEMBER FUNCTION GET_OBJECT_TYPE()
RETURN VARCHAR2;
```

SET_BASE_TABLE_NAME Member Procedure

Sets the base (dependent) table name.

Syntax

```
MEMBER PROCEDURE SET_BASE_TABLE_NAME(
   base_table_name  IN  VARCHAR2);
```

Parameter

Table 310-5 SET_BASE_TABLE_NAME Procedure Parameter

Parameter	Description
base_table_name	The name of the base table

SET_BASE_TABLE_OWNER Member Procedure

Sets the base (dependent) table owner.

Syntax

```
MEMBER PROCEDURE SET_BASE_TABLE_OWNER(
   base_table_owner  IN  VARCHAR2);
```

Parameter

Table 310-6 SET_BASE_TABLE_OWNER Procedure Parameter

Parameter	Description
base_table_owner	The name of the base table owner

SET_CURRENT_SCHEMA Member Procedure

Sets the default schema (user) name.

Syntax

```
MEMBER PROCEDURE SET_CURRENT_SCHEMA(
   current_schema  IN  VARCHAR2);
```

Parameter

Table 310-7 SET_CURRENT_SCHEMA Procedure Parameter

Parameter	Description
current_schema	The name of the schema to set as the current schema Set this parameter to a non-NULL value.

SET_DDL_TEXT Member Procedure

Sets the DDL text.

Syntax

```
MEMBER PROCEDURE SET_DDL_TEXT(
   ddl_text  IN  CLOB);
```

Parameter

Table 310-8 SET_DDL_TEXT Procedure Parameter

Parameter	Description
ddl_text	The DDL text Set this parameter to a non-NULL value.

SET_EDITION_NAME Member Procedure

Sets the name of the edition in which the DDL statement was executed.

> **See Also:**
>
> *Oracle Database Development Guide*

Syntax

```
MEMBER PROCEDURE SET_EDITION_NAME(
   edition_name  IN  VARCHAR2);
```

Parameter

Table 310-9 SET_EDITION_NAME Procedure Parameter

Parameter	Description
edition_name	Name of the edition

SET_LOGON_USER Member Procedure

Sets the logon user name.

Syntax

```
MEMBER PROCEDURE SET_LOGON_USER(
   logon_user  IN  VARCHAR2);
```

Parameter

Table 310-10 SET_LOGON_USER Procedure Parameter

Parameter	Description
logon_user	The name of the schema to set as the logon user

SET_OBJECT_TYPE Member Procedure

Sets the object type.

Syntax

```
MEMBER PROCEDURE SET_OBJECT_TYPE(
   object_type  IN  VARCHAR2);
```

Parameter

Table 310-11 SET_OBJECT_TYPE Procedure Parameter

Parameter	Description

`object_type`	The object type
	The following are valid object types:
	`CLUSTER`
	`FUNCTION`
	`INDEX`
	`LINK`
	`OUTLINE`
	`PACKAGE`
	`PACKAGE BODY`
	`PROCEDURE`
	`SEQUENCE`
	`SYNONYM`
	`TABLE`
	`TRIGGER`
	`TYPE`
	`USER`
	`VIEW`
	`LINK` represents a database link.
	`NULL` is also a valid object type. Specify `NULL` for all object types not listed. The `GET_OBJECT_TYPE` member procedure returns `NULL` for object types not listed.

310.3.2 LCR$_ROW_RECORD Type

This type represents a data manipulation language (DML) change to a row in a table. This type uses the `LCR$_ROW_LIST` type.

> **Note:**
>
> A multitenant container database is the only supported architecture in Oracle Database 21c and later releases. While the documentation is being revised, legacy terminology may persist. In most cases, "database" and "non-CDB" refer to a CDB or PDB, depending on context. In some contexts, such as upgrades, "non-CDB" refers to a non-CDB from a previous release.

If you create or modify a row logical change record (row LCR), then make sure the `command_type` attribute is consistent with the presence or absence of old column values and the presence or absence of new column values.

This topic contains information about the constructor for DDL LCRs and information about the member subprograms for this type.

Note:

- When passing a name as a parameter to an LCR constructor, you can enclose the name in double quotes to handle names that use mixed case or lower case for database objects. For example, if a name contains any lower case characters, then you must enclose it in double quotes.
- The application does not need to specify a transaction identifier or SCN when it creates an LCR because the apply process generates these values and stores them in memory. If a transaction identifier or SCN is specified in the LCR, then the apply process ignores it and assigns a new value.

See Also:

LCR$_ROW_LIST Type

LCR$_ROW_RECORD Constructor

Creates a `SYS.LCR$_ROW_RECORD` object with the specified information.

```
STATIC FUNCTION CONSTRUCT(
   source_database_name  IN  VARCHAR2,
   command_type          IN  VARCHAR2,
   object_owner          IN  VARCHAR2,
   object_name           IN  VARCHAR2,
   tag                   IN  RAW                DEFAULT NULL,
   transaction_id        IN  VARCHAR2           DEFAULT NULL,
   scn                   IN  NUMBER             DEFAULT NULL,
   old_values            IN  SYS.LCR$_ROW_LIST  DEFAULT NULL,
   new_values            IN  SYS.LCR$_ROW_LIST  DEFAULT NULL,
   position              IN  RAW                DEFAULT NULL,

   statement             IN  VARCHAR2           DEFAULT NULL,

   bind_variables        IN  SYS.LCR$_ROW_LIST  DEFAULT NULL,

   bind_by_position      IN  VARCHAR2           DEFAULT 'N',

   root_name             IN  VARCHAR2           DEFAULT NULL)
RETURN SYS.LCR$_ROW_RECORD;
```

LCR$_ROW_RECORD Constructor Function Parameters

Table 310-12 Constructor Function Parameters for LCR$_ROW_RECORD

Parameter	Description
source_database_name	The database where the row change occurred If the LCRs originated in a CDB, then this field specifies the global name of the container where the row change occurred. If you do not include the domain name, then the function appends the local domain to the database name automatically. For example, if you specify DBS1 and the local domain is EXAMPLE.COM, then the function specifies DBS1.EXAMPLE.COM automatically. Set this parameter to a non-NULL value.
command_type	The type of command executed in the DML statement Set this parameter to a non-NULL value. Valid values are the following: INSERT UPDATE DELETE LOB ERASE LOB WRITE LOB TRIM If INSERT, then ensure that the LCR has a new_values collection that is not empty and an empty or NULL old_values collection. If UPDATE, then ensure that the LCR has a new_values collection that is not empty and an old_values collection that is not empty. If DELETE, then ensure that the LCR has a NULL or empty new_values collection and an old_values collection that is not empty. If LOB ERASE, LOB WRITE, or LOB TRIM, then ensure that the LCR has a new_values collection that is not empty and an empty or NULL old_values collection.
object_owner	The user who owns the table on which the row change occurred Set this parameter to a non-NULL value.
object_name	The table on which the DML statement was executed Set this parameter to a non-NULL value.
tag	A binary tag that enables tracking of the LCR For example, this tag can be used to determine the original source database of the DML change when apply forwarding is used.
transaction_id	The identifier of the transaction
scn	The SCN at the time when the change record was written to the redo log The SCN value is meaningless for a user-created LCR.
old_values	The column values for the row before the DML change If the DML statement is an UPDATE or a DELETE statement, then this parameter contains the values of columns in the row before the DML statement. If the DML statement is an INSERT statement, then there are no old values.

Table 310-12 (Cont.) Constructor Function Parameters for LCR$_ROW_RECORD

Parameter	Description
new_values	The column values for the row after the DML change If the DML statement is an UPDATE or an INSERT statement, then this parameter contains the values of columns in the row after the DML statement. If the DML statement is a DELETE statement, then there are no new values. If the LCR reflects a LOB operation, then this parameter contains the supplementally logged columns and any relevant LOB information.
position	The position of the LCR LCR position is commonly used in XStream configurations. Using XStream requires purchasing a license for the Oracle GoldenGate product. **See Also:** *Oracle Database XStream Guide*
statement	This parameter is reserved for internal use only.
bind_variables	This parameter is reserved for internal use only.
bind_by_position	This parameter is reserved for internal use only.
root_name	If the LCRs is associated with a CDB, then this field specifies the global name of the root in the CDB. If the LCR is associated with a non-CDB, then this field is NULL.

Summary of LCR$_ROW_RECORD Subprograms

Table 310-13 LCR$_ROW_RECORD Type Subprograms

Subprogram	Description
ADD_COLUMN Member Procedure	Adds the value as old or new, depending on the value type specified, for the column
CONVERT_LONG_TO_LOB_CHUNK Member Procedure	Converts LONG data in a row LCR into fixed width CLOB, or converts LONG RAW data in a row LCR into a BLOB
DELETE_COLUMN Member Procedure	Deletes the old value, the new value, or both, for the specified column, depending on the value type specified
EXECUTE Member Procedure	Executes the LCR under the security domain of the current user
GET_JSON_INFORMATION Member Function	Gets JSON information.
GET_LOB_INFORMATION Member Function	Gets the LOB information for the column
GET_LOB_OFFSET Member Function	Gets the LOB offset for the specified column
GET_LOB_OPERATION_SIZE Member Function	Gets the operation size for the LOB column
GET_LONG_INFORMATION Member Function	Gets the LONG information for the column
GET_ROW_TEXT Member Procedure	Gets the SQL statement for the change that is encapsulated in the LCR
GET_VALUE Member Function	Gets the old or new value for the specified column, depending on the value type specified

Table 310-13 (Cont.) LCR$_ROW_RECORD Type Subprograms

Subprogram	Description
GET_VALUES Member Function	Gets a list of old or new values, depending on the value type specified
GET_WHERE_CLAUSE Member Procedure	Gets a `WHERE` clause for the change that is encapsulated in the row LCR
GET_XML_INFORMATION Member Function	Gets the XML information for the specified column
IS_STATEMENT_LCR Member Function	Reserved for internal use only
KEEP_COLUMNS Member Procedure	Keeps a list of columns a row LCR
RENAME_COLUMN Member Procedure	Renames a column in an LCR
SET_JSON_INFORMATION Member Procedure	Sets the JSON information for a column.
SET_LOB_INFORMATION Member Procedure	Sets LOB information for the column
SET_LOB_OFFSET Member Procedure	Sets the LOB offset for the specified column
SET_LOB_OPERATION_SIZE Member Procedure	Sets the operation size for the LOB column
SET_ROW_TEXT Member Procedure	Reserved for internal use only
SET_VALUE Member Procedure	Overwrites the value of the specified column
SET_VALUES Member Procedure	Replaces the existing old or new values for the LCR, depending on the value type specified
SET_XML_INFORMATION Member Procedure	Sets the XML information for the column
Common Subprograms	See Common Subprograms for LCR$_DDL_RECORD and LCR$_ROW_RECORD for a list of subprograms common to the `SYS.LCR$_ROW_RECORD` and `SYS.LCR$_DDL_RECORD` types

ADD_COLUMN Member Procedure

Adds the value as old or new, depending on the value type specified, for the column. An error is raised if a value of the same type already exists for the column.

> **Note:**
>
> To set a column value that already exists, run `SET_VALUE`.

> **See Also:**
>
> SET_VALUE Member Procedure

Considerations for LOB Columns

When processing a row LCR with LOB columns with a procedure DML handler or error handler and the handler is using LOB assembly (the `assemble_lobs` parameter is set to `TRUE` for the

handler), you use this member procedure in the handler procedure to add a LOB column to a row LCR. If `assemble_lobs` is set to `FALSE` for the handler, then you cannot use this member procedure to add a LOB column to a row LCR.

To use a DML or error handler to add a LOB column, specify the LOB locator for the `column_value` parameter in the member procedure. The `ADD_COLUMN` member procedure verifies that an `ANYDATA` encapsulated LOB locator is processed with a DML or error handler that is using LOB assembly. An error is raised under the following conditions:

- The handler attempts to enqueue a row LCR with an `ANYDATA` encapsulated LOB locator.
- An attempt is made to add an LOB column that is set incorrectly.

If an error is raised because of one of these conditions, then the transaction that includes the row LCR is moved to the error queue, and the LOB is represented by the original (nonassembled) row LCRs.

> **Note:**
>
> - Database compatibility must be `10.2.0` or higher to use LOB assembly.
> - When you are processing a row LCR with a rule-based transformation, you cannot use this member procedure to add a LOB column.
> - When you are processing a row LCR with a rule-based transformation, procedure DML handler, or error handler, you cannot use this member procedure to add a `LONG` or `LONG RAW` column.

Syntax

```
MEMBER PROCEDURE ADD_COLUMN(
   value_type    IN  VARCHAR2,
   column_name   IN  VARCHAR2,
   column_value  IN  ANYDATA);
```

Parameters

Table 310-14 ADD_COLUMN Procedure Parameters

Parameter	Description
`value_type`	The type of value to add for the column Specify `old` to add the old value of the column. Specify `new` to add the new value of the column.
`column_name`	The column name This name is not validated. An error can be raised during application of the LCRs if an invalid name is specified.
`column_value`	The value of the column If `NULL`, then this procedure raises an error. If the member procedure is used in a procedure DML handler or error handler that uses LOB assembly, then a LOB locator can be specified. A `NULL` column value can be specified by encapsulating the `NULL` value in an `ANYDATA` wrapper.

CONVERT_LONG_TO_LOB_CHUNK Member Procedure

Converts `LONG` data in a row LCR into a `CLOB`, or converts `LONG RAW` data in a row LCR into a `BLOB`.

This procedure can change the operation code from `LONG WRITE` to `LOB WRITE` for the row LCR.

This member procedure can be used in rule-based transformations.

The following restrictions apply to this member procedure:

- This member procedure cannot be used in apply handlers.
- `LONG` data can be sent as a part of a row LCR with one of the following operation codes: `INSERT`, `UPDATE`, or `LONG_WRITE`. Because `LONG` data can be sent in multiple pieces, make sure that this method is invoked on either none or all `LONG` pieces.
- LOB to `LONG` conversion is not supported.
- A row LCR on which this procedure is executed must have been created by a capture process. That is, this procedure does not support persistent row LCRs.

Syntax

```
MEMBER PROCEDURE CONVERT_LONG_TO_LOB_CHUNK;
```

DELETE_COLUMN Member Procedure

Deletes the old value, the new value, or both, for the specified column, depending on the value type specified.

Syntax

```
MEMBER PROCEDURE DELETE_COLUMN(
   column_name  IN  VARCHAR2,
   value_type   IN  VARCHAR2  DEFAULT '*');
```

Parameters

Table 310-15 DELETE_COLUMN Procedure Parameters

Parameter	Description
`column_name`	The column name An error is raised if the column does not exist in the LCR.
`value_type`	The type of value to delete for the column Specify `old` to delete the old value of the column. Specify `new` to delete the new value of the column. If `*` is specified, then the procedure deletes both the old and new values.

EXECUTE Member Procedure

Executes the row LCR under the security domain of the current user. Any apply handlers that would be run for an LCR are not run when the LCR is applied using this procedure.

This member procedure can be run on a row LCR under any of the following conditions:

- The LCR is being processed by an apply handler.

- The LCR has been constructed using the `LCR$_ROW_RECORD` constructor function but has not been enqueued.
- The LCR is in the error queue.

> **Note:**
>
> Do not run this member procedure in a custom rule-based transformation on a row LCR. Doing so could execute the row LCR outside of its transactional context.

Considerations for LOB Columns

When processing a row LCR with LOB columns with a procedure DML handler or error handler, and the handler is using LOB assembly (the `assemble_lobs` parameter is set to `TRUE` for the handler), this member procedure executes the assembled row LCR. An assembled row LCR represents a LOB value with a LOB locator or `NULL`.

If `assemble_lobs` is set to `FALSE` for the handler, then this member procedure executes the nonassembled row LCRs. Nonassembled row LCRs represent LOB values with `VARCHAR2` and `RAW` data types. These nonassembled row LCRs might have been modified by the handler.

An error is raised under the following conditions:

- A DML or error handler configured with `assemble_lobs` set to `FALSE` attempts to execute a row LCR that contains a LOB locator.
- A DML or error handler configured with `assemble_lobs` set to `TRUE` attempts to execute a row LCR that contains one or more LOB values represented with `VARCHAR2` or `RAW` data types.

If an error is raised because of one of these conditions, then the transaction that includes the row LCR is moved to the error queue, and the LOB is represented by the original (nonassembled) row LCRs.

Syntax

```
MEMBER PROCEDURE EXECUTE(
   conflict_resolution  IN  BOOLEAN);
```

Parameters

Table 310-16 EXECUTE Procedure Parameters

Parameter	Description
`conflict_resolution`	If `TRUE`, then any conflict resolution defined for the table using the `SET_UPDATE_CONFLICT_HANDLER` procedure in the `DBMS_APPLY_ADM` package is used to resolve conflicts resulting from the execution of the LCR. If `FALSE`, then conflict resolution is not used. An error is raised if this parameter is not specified or is set to `NULL`.

GET_JSON_INFORMATION Member Function

Gets JSON information.

The return value can be one of the following:

```
DBMS_LCR.NOT_OSON CONSTANT NUMBER := 1;
DBMS_LCR.OSON_DOC CONSTANT NUMBER := 2;
```

> **Note:**
>
> OSON is Oracle binary format for JSON column.

They represent either a column not containing OSON data or full OSON document.

Syntax

```
MEMBER FUNCTION GET_JSON_INFORMATION(
   column_name  IN  VARCHAR2)
RETURN NUMBER;
```

Parameters

Table 310-17 GET_JSON_INFORMATION Member Function Parameter

Parameter	Description
column_name	Name of column to obtain JSON information.

GET_LOB_INFORMATION Member Function

Gets the LOB information for the column.

The return value can be one of the following:

```
DBMS_LCR.NOT_A_LOB        CONSTANT NUMBER := 1;
DBMS_LCR.NULL_LOB         CONSTANT NUMBER := 2;
DBMS_LCR.INLINE_LOB       CONSTANT NUMBER := 3;
DBMS_LCR.EMPTY_LOB        CONSTANT NUMBER := 4;
DBMS_LCR.LOB_CHUNK        CONSTANT NUMBER := 5;
DBMS_LCR.LAST_LOB_CHUNK   CONSTANT NUMBER := 6;
```

Returns NULL if the specified column does not exist.

If the command type of the row LCR is UPDATE, then specifying 'Y' for the use_old parameter is a convenient way to get the value of the columns.

Syntax

```
MEMBER FUNCTION GET_LOB_INFORMATION(
  value_type   IN  VARCHAR2,
  column_name  IN  VARCHAR2,
  use_old      IN  VARCHAR2  DEFAULT 'Y')
RETURN NUMBER;
```

Parameters

Table 310-18 GET_LOB_INFORMATION Function Parameters

Parameter	Description
value_type	The type of value to return for the column, either old or new
column_name	The name of the column
use_old	If Y and value_type is new, and no new value exists, then the function returns the corresponding old value. If N and value_type is new, then the function does not return the old value if no new value exists. If value_type is old or if the command_type of the row LCR is not UPDATE, then the function ignores the value of the use_old parameter. NULL is not a valid specification for the use_old parameter.

GET_LOB_OFFSET Member Function

Gets the LOB offset for the specified column in the number of characters for CLOB columns and the number of bytes for BLOB columns. Returns a non-NULL value only if all of the following conditions are met:

- The value exists for the column
- The column value is an out-of-line LOB. That is, the information is DBMS_LCR.LAST_LOB_CHUNK or DBMS_LCR.LOB_CHUNK
- The command type is LOB ERASE or LOB WRITE

Otherwise, returns NULL.

Syntax

```
GET_LOB_OFFSET(
   value_type   IN  VARCHAR2,
   column_name  IN  VARCHAR2)
RETURN NUMBER;
```

Parameters

Table 310-19 GET_LOB_OFFSET Function Parameters

Parameter	Description
value_type	The type of value to return for the column Currently, only new can be specified.
column_name	The name of the LOB column

GET_LOB_OPERATION_SIZE Member Function

Gets the operation size for the LOB column in the number of characters for CLOB columns and the number of bytes for BLOB columns. Returns a non-NULL value only if all of the following conditions are met:

- The value exists for the column
- The column value is an out-of-line LOB

- The command type is `LOB ERASE` or `LOB TRIM`
- The information is `DBMS_LCR.LAST_LOB_CHUNK`

Otherwise, returns `NULL`.

Syntax

```
MEMBER FUNCTION GET_LOB_OPERATION_SIZE(
  value_type   IN  VARCHAR2,
  column_name  IN  VARCHAR2)
RETURN NUMBER,
```

Parameters

Table 310-20 GET_LOB_OPERATION_SIZE Function Parameters

Parameter	Description
`value_type`	The type of value to return for the column Currently, only `new` can be specified.
`column_name`	The name of the LOB column

GET_LONG_INFORMATION Member Function

Gets the `LONG` information for the column.

The return value can be one of the following:

```
DBMS_LCR.NOT_A_LONG        CONSTANT NUMBER := 1;
DBMS_LCR.NULL_LONG         CONSTANT NUMBER := 2;
DBMS_LCR.INLINE_LONG       CONSTANT NUMBER := 3;
DBMS_LCR.LONG_CHUNK        CONSTANT NUMBER := 4;
DBMS_LCR.LAST_LONG_CHUNK   CONSTANT NUMBER := 5;
```

Returns `NULL` if the specified column does not exist.

If the command type of the row LCR is `UPDATE`, then specifying `'Y'` for the `use_old` parameter is a convenient way to get the value of the columns.

Syntax

```
MEMBER FUNCTION GET_LONG_INFORMATION(
  value_type   IN  VARCHAR2,
  column_name  IN  VARCHAR2,
  use_old      IN  VARCHAR2  DEFAULT 'Y')
RETURN NUMBER;
```

Parameters

Table 310-21 GET_LONG_INFORMATION Function Parameters

Parameter	Description
`value_type`	The type of value to return for the column, either `old` or `new`
`column_name`	The name of the column

Table 310-21 (Cont.) GET_LONG_INFORMATION Function Parameters

Parameter	Description
use_old	If Y and value_type is new, and no new value exists, then the function returns the corresponding old value. If N and value_type is new, then the function does not return the old value if no new value exists. If value_type is old or if the command_type of the row LCR is not UPDATE, then the function ignores the value of the use_old parameter. NULL is not a valid specification for the use_old parameter.

GET_ROW_TEXT Member Procedure

Gets the SQL statement for the change that is encapsulated in the row LCR. This method performs SQL generation in PL/SQL.

This method is overloaded. The different functionality of each form of syntax is presented along with the definitions.

Syntax

The following procedure returns the SQL statement in a CLOB datatype.

```
MEMBER PROCEDURE GET_ROW_TEXT(
   row_text  IN/OUT  CLOB);
```

The following procedure returns the SQL statement with bind variables in a CLOB datatype.

```
MEMBER PROCEDURE GET_ROW_TEXT(
   row_text         IN/OUT  CLOB,
   variable_list    IN/OUT  LCR$_ROW_LIST,
   bind_var_syntax  IN      VARCHAR2  DEFAULT ':');
```

> **See Also:**
>
> "LCR$_ROW_LIST Type"

Parameters

Table 310-22 GET_ROW_TEXT Procedure Parameters

Parameter	Description
row_text	The SQL statement for the change that is encapsulated in the LCR
variable_list	The values for the bind variables in the order of the bind variables
bind_var_syntax	The syntax for the bind variables One of the following values is valid: • Specify :, the default, for bind values to be in the form :1, :2, and so on. • Specify ? for bind values to be in the form ?.

GET_VALUE Member Function

Gets the old or new value for the specified column, depending on the value type specified.

If the command type of the row LCR is `UPDATE`, then specifying `'Y'` for the `use_old` parameter is a convenient way to get the value of a column.

Syntax

```
MEMBER FUNCTION GET_VALUE(
   value_type   IN  VARCHAR2,
   column_name  IN  VARCHAR2,
   use_old      IN  VARCHAR2  DEFAULT 'Y')
RETURN ANYDATA;
```

Parameters

Table 310-23 GET_VALUE Function Parameters

Parameter	Description
`value_type`	The type of value to return for the column Specify `old` to get the old value for the column. Specify `new` to get the new value for the column.
`column_name`	The column name If the column is present and has a `NULL` value, then the function returns an `ANYDATA` instance containing a `NULL` value. If the column value is absent, then the function returns a `NULL`.
`use_old`	If `Y` and `value_type` is `new`, and no new value exists, then the function returns the corresponding old value. If `N` and `value_type` is `new`, then the function returns `NULL` if no new value exists. If `value_type` is `old` or if the `command_type` of the row LCR is not `UPDATE`, then the function ignores the value of the `use_old` parameter. `NULL` is not a valid specification for the `use_old` parameter.

GET_VALUES Member Function

Gets a list of old or new values, depending on the value type specified.

If the command type of the row LCR is `UPDATE`, then specifying `'Y'` for the `use_old` parameter is a convenient way to get the values of all columns.

Syntax

```
MEMBER FUNCTION GET_VALUES(
   value_type  IN  VARCHAR2,
   use_old     IN  VARCHAR2  DEFAULT 'Y')
RETURN SYS.LCR$_ROW_LIST;
```

Parameters

Table 310-24 GET_VALUES Function Parameters

Parameter	Description
value_type	The type of values to return Specify old to return a list of old values. Specify new to return a list of new values.
use_old	If Y and value_type is new, then the function returns a list of all new values in the LCR. If a new value does not exist in the list, then the function returns the corresponding old value. Therefore, the returned list contains all existing new values and the old values where there are no new values. If N and value_type is new, then the function returns a list of all new values in the LCR without returning any old values. If value_type is old or if the command_type of the row LCR is not UPDATE, then the function ignores the value of the use_old parameter. NULL is not a valid specification for the use_old parameter.

GET_WHERE_CLAUSE Member Procedure

Gets a WHERE clause for the change that is encapsulated in the row LCR.

Use the WHERE clause returned by GET_WHERE_CLAUSE instead of using the ROWID, because the ROWID is not ANSI compatible. The generated WHERE clause might not match the WHERE clause in the original DML operation.

The ROWID of an INSERT statement is the ROWID of the new row created by the INSERT. The WHERE clause generated for an INSERT operation identifies the new row. Therefore, the generated WHERE clause includes all of the new values inserted.

For example, consider the following insert into the hr.departments table:

```
INSERT INTO hr.departments (
   department_id, department_name, manager_id, location_id)
   VALUES (10, 'HR', 20, 40);
```

The generated WHERE clause represents the row with the values 10, 'HR', 20, and 40. Hence, the generated WHERE clause is the following:

```
WHERE "DEPARTMENT_ID" = 10 AND "DEPARTMENT_NAME" = 'HR' AND
      "MANAGER_ID" = 20 AND "LOCATION_ID" = 40
```

The ROWID of an UPDATE statement is the ROWID of the row that was updated. The WHERE clause generated for an UPDATE operation identifies the row after the UPDATE executes. The generated WHERE clause is based on the old and new values of the UPDATE.

For example, consider the following update to the hr.departments table:

```
UPDATE hr.departments SET department_name='Management'
 WHERE department_name='Administration' AND location_id = 20 AND
       manager_id = 30 AND department_id = 10;
```

The values of the row after the UPDATE are 10, 'Management', 30, and 20. Hence, the generated WHERE clause to identify the row is the following:

```
WHERE "DEPARTMENT_ID" = 10 AND "DEPARTMENT_NAME" = 'MANAGEMENT' AND
      "MANAGER_ID" = 30 AND "LOCATION_ID" = 20
```

Notice that the new value is used for "DEPARTMENT_NAME", because the new value is the value of the column after the UPDATE. For the rest of the columns, the old values are used.

The ROWID of a DELETE operation is the row that existed before it was deleted. The generated WHERE clause consists of all the old column values present in the DELETE operation.

LOB columns do not appear in generated WHERE clauses. The generated WHERE clause is not affected by the presence of LOB columns in the LCR.

This method is overloaded. The different functionality of each form of syntax is presented along with the definitions.

Syntax

The following procedure returns the WHERE clause of a SQL statement in a CLOB datatype.

```
MEMBER PROCEDURE GET_WHERE_CLAUSE(
   where_clause  IN/OUT  CLOB);
```

The following procedure returns the WHERE clause of a SQL statement with bind variables in a CLOB datatype.

```
MEMBER PROCEDURE GET_WHERE_CLAUSE(
   where_clause     IN/OUT  CLOB,
   variable_list    IN/OUT  LCR$_ROW_LIST,
   bind_var_syntax  IN      VARCHAR2  DEFAULT ':');
```

> **See Also:**
>
> LCR$_ROW_LIST Type

Parameters

Table 310-25 GET_WHERE_CLAUSE Procedure Parameters

Parameter	Description
where_clause	The WHERE clause of the SQL statement for the change that is encapsulated in the LCR
variable_list	The values for the bind variables in the order of the bind variables
bind_var_syntax	The syntax for the bind variables One of the following values is valid: • Specify :, the default, for bind values to be in the form :1, :2, and so on. • Specify ? for bind values to be in the form ?.

GET_XML_INFORMATION Member Function

Gets the XML information for the specified column.

The return value can be one of the following:

```
DBMS_LCR.NOT_XML   CONSTANT NUMBER := 1;
DBMS_LCR.XML_DOC   CONSTANT NUMBER := 2;
DBMS_LCR.XML_DIFF  CONSTANT NUMBER := 3;

DBMS_LCR.XML_DIFF  CONSTANT NUMBER := 3;
```

`DBMS_LCR.NOT_XML` indicates that the column is not an `XMLType` column.

`DBMS_LCR.XML_DOC` indicates that the column contains an XML document.

`DBMS_LCR.XML_DIFF` indicates that the column contains an XML document that describes a change made by an update operation. This XML document describes the differences in the column's XML document. The entire XML document is not replaced.

`DBMS_LCR.XML_DIFF` indicates that the column contains differences between old and new XML documents for an update operation.

Returns `NULL` if the specified column does not exist.

Syntax

```
MEMBER FUNCTION GET_XML_INFORMATION(
   column_name  IN  VARCHAR2)
RETURN NUMBER;
```

Parameter

Table 310-26 GET_XML_INFORMATION Function Parameter

Parameter	Description
`column_name`	The column name

IS_STATEMENT_LCR Member Function

This function is reserved for internal use only.

KEEP_COLUMNS Member Procedure

This procedure keeps a list of columns in a row LCR. The procedure deletes columns that are not in the list from the row LCR.

Syntax

```
MEMBER PROCEDURE KEEP_COLUMNS(
   column_list  IN  VARCHAR2,
   value_type   IN  VARCHAR2  DEFAULT '*');
```

Parameters

Table 310-27 KEEP_COLUMNS Procedure Parameters

Parameter	Description
`column_list`	The names of the columns kept for the row LCR Specify a comma-delimited list of type `VARCHAR2`. This procedure removes columns that are not in the list from the current row LCR.

Table 310-27 (Cont.) KEEP_COLUMNS Procedure Parameters

Parameter	Description
value_type	The type of value for which to keep the columns Specify old to keep the old values of the columns. An error is raised if the old values do not exist in the LCR. Specify new to keep the new values of the columns. An error is raised if the new values do not exist in the LCR. If * is specified, then the procedure keeps both the old and the new columns.

RENAME_COLUMN Member Procedure

Renames a column in a row LCR.

Syntax

```
MEMBER PROCEDURE RENAME_COLUMN(
   from_column_name  IN  VARCHAR2,
   to_column_name    IN  VARCHAR2,
   value_type        IN  VARCHAR2  DEFAULT '*');
```

Parameters

Table 310-28 RENAME_COLUMN Procedure Parameters

Parameter	Description
from_column_name	The existing column name
to_column_name	The new column name An error is raised if a column with the specified name already exists.
value_type	The type of value for which to rename the column Specify old to rename the old value of the column. An error is raised if the old value does not exist in the LCR. Specify new to rename the new value of the column. An error is raised if the new value does not exist in the LCR. If * is specified, then the procedure renames the column names for both old and new value. The procedure raises an error if either column value does not exist in the LCR.

SET_JSON_INFORMATION Member Procedure

Sets the JSON information for a column.

Syntax

```
MEMBER PROCEDURE SET_JSON_INFORMATION(
  self             IN OUT NOCOPY LCR$_ROW_RECORD,
  value_type       IN VARCHAR2,
  column_name      IN  VARCHAR2,
  json_information IN NUMBER);
```

Parameters

Table 310-29 SET_JSON_INFORMATION Procedure Parameters

Parameter	Description
self	
value_type	The type of value to set for the column. The supported values are: • NEW • OLD
column_name	The name of the column. An exception is raised if the column does not exists in the LCR.
json_information	The supported values are: • DBMS_LCR.NOT_OSON • DBMS_LCR.OSON_DOC

SET_LOB_INFORMATION Member Procedure

Sets LOB information for the column.

> **Note:**
>
> When you are processing a row LCR with a rule-based transformation, procedure DML handler, or error handler, you cannot use this member procedure.

Syntax

```
MEMBER PROCEDURE SET_LOB_INFORMATION(
  value_type       IN  VARCHAR2,
  column_name      IN  VARCHAR2,
  lob_information  IN  NUMBER);
```

Parameters

Table 310-30 SET_LOB_INFORMATION Procedure Parameters

Parameter	Description
value_type	The type of value to set for the column, either old or new Specify old only if lob_information is set to DBMS_LCR.NOT_A_LOB.
column_name	The name of the column. An exception is raised if the column value does not exist. You might need to set this parameter for non-LOB columns.

Table 310-30 (Cont.) SET_LOB_INFORMATION Procedure Parameters

Parameter	Description
lob_information	Specify one of the following values: DBMS_LCR.NOT_A_LOB CONSTANT NUMBER := 1; DBMS_LCR.NULL_LOB CONSTANT NUMBER := 2; DBMS_LCR.INLINE_LOB CONSTANT NUMBER := 3; DBMS_LCR.EMPTY_LOB CONSTANT NUMBER := 4; DBMS_LCR.LOB_CHUNK CONSTANT NUMBER := 5; DBMS_LCR.LAST_LOB_CHUNK CONSTANT NUMBER := 6;

SET_LOB_OFFSET Member Procedure

Sets the LOB offset for the specified column in the number of characters for `CLOB` columns and the number of bytes for `BLOB` columns.

> **Note:**
>
> When you are processing a row LCR with a rule-based transformation, procedure DML handler, or error handler, you cannot use this member procedure.

Syntax

```
MEMBER PROCEDURE SET_LOB_OFFSET(
   value_type    IN  VARCHAR2,
   column_name   IN  VARCHAR2,
   lob_offset    IN  NUMBER);
```

Parameters

Table 310-31 SET_LOB_OFFSET Procedure Parameters

Parameter	Description
value_type	The type of value to set for the column Currently, only `new` can be specified.
column_name	The column name An error is raised if the column value does not exist in the LCR.
lob_offset	The LOB offset number Valid values are `NULL` or a positive integer less than or equal to `DBMS_LOB.LOBMAXSIZE`.

SET_LOB_OPERATION_SIZE Member Procedure

Sets the operation size for the LOB column in the number of characters for `CLOB` columns and bytes for `BLOB` columns.

> **Note:**
>
> When you are processing a row LCR with a rule-based transformation, procedure DML handler, or error handler, you cannot use this member procedure.

Syntax

```
MEMBER PROCEDURE SET_LOB_OPERATION_SIZE(
  value_type          IN  VARCHAR2,
  column_name         IN  VARCHAR2,
  lob_operation_size  IN  NUMBER);
```

Parameters

Table 310-32 SET_LOB_OPERATION_SIZE Procedure Parameters

Parameter	Description
`value_type`	The type of value to set for the column Currently, only `new` can be specified.
`column_name`	The name of the LOB column An exception is raised if the column value does not exist in the LCR.
`lob_operation_size`	If `lob_information` for the LOB is or will be `DBMS_LCR.LAST_LOB_CHUNK`, then this parameter can be set to either a valid `LOB ERASE` value or a valid `LOB TRIM` value. A `LOB ERASE` value must be a positive integer less than or equal to `DBMS_LOB.LOBMAXSIZE`. A `LOB TRIM` value must be a nonnegative integer less than or equal to `DBMS_LOB.LOBMAXSIZE`. Otherwise, set to `NULL`.

SET_ROW_TEXT Member Procedure

This procedure is reserved for internal use only.

SET_VALUE Member Procedure

Overwrites the old or new value of the specified column.

One reason to overwrite an old value for a column is to resolve an error that resulted from a conflict.

> **Note:**
>
> To add a column to a row LCR, run `ADD_COLUMN`.

> **See Also:**
>
> ADD_COLUMN Member Procedure

Considerations for LOB Columns

When processing a row LCR with LOB columns with a procedure DML handler or error handler, and the handler is using LOB assembly (the `assemble_lobs` parameter is set to `TRUE` for the handler), you can use this member procedure in the handler procedure on a LOB column in a row LCR. If `assemble_lobs` is set to `FALSE` for the handler, then you cannot use this member procedure on a LOB column.

To use a DML or error handler to set the value of a LOB column, specify the LOB locator for the `column_value` parameter in the member procedure. The `SET_VALUE` member procedure verifies that an `ANYDATA` encapsulated LOB locator is processed with a DML or error handler that is using LOB assembly. An error is raised under the following conditions:

- The handler attempts to enqueue a row LCR with an `ANYDATA` encapsulated LOB locator.
- An attempt is made to set a LOB column incorrectly.

If an error is raised because of one of these conditions, then the transaction that includes the row LCR is moved to the error queue, and the LOB is represented by the original (nonassembled) row LCRs.

Note:

- Database compatibility must be `10.2.0` or higher to use LOB assembly.
- When you are processing a row LCR with a rule-based transformation, you cannot use this member procedure on a LOB column.
- When you are processing a row LCR with a rule-based transformation, procedure DML handler, or error handler, you cannot use this member procedure on a `LONG` or `LONG RAW` column.

Considerations for XMLType Columns

When processing a row LCR with `XMLType` columns with a procedure DML handler or error handler, any `XMLType` columns and LOB columns in the LCR are always assembled using LOB assembly. You can use this member procedure in the handler procedure on a row LCR that contains one or more `XMLType` columns.

To use a DML or error handler to set the value an `XMLType` column, specify the `XMLType` for the `column_value` parameter. The `SET_VALUE` member procedure verifies that an `ANYDATA` encapsulated `XMLType` is processed with a DML or error handler. An error is raised under the following conditions:

- The handler attempts to enqueue a row LCR with an `ANYDATA` encapsulated `XMLType`.
- An attempt is made to set a `XMLType` column incorrectly.

If an error is raised because of one of these conditions, then the transaction that includes the row LCR is moved to the error queue, and the `XMLType` column is represented by the original (nonassembled) row LCRs.

> **Note:**
>
> - Database compatibility must be `11.1.0` or higher to process row LCRs with `XMLType` columns.
> - When you are processing a row LCR with a rule-based transformation, you cannot use this member procedure on `XMLType` columns.

Syntax

```
MEMBER PROCEDURE SET_VALUE(
   value_type    IN  VARCHAR2,
   column_name   IN  VARCHAR2,
   column_value  IN  ANYDATA);
```

Parameters

Table 310-33 SET_VALUE Procedure Parameters

Parameter	Description
`value_type`	The type of value to set Specify `old` to set the old value of the column. Specify `new` to set the new value of the column.
`column_name`	The column name An error is raised if the specified `column_value` does not exist in the LCR for the specified `column_type`.
`column_value`	The new value of the column If `NULL` is specified, then this procedure raises an error. To set the value to `NULL`, encapsulate the `NULL` in an `ANYDATA` instance. If the member procedure is used in a procedure DML handler or error handler that uses LOB assembly, then specify a LOB locator for LOB columns.

SET_VALUES Member Procedure

Replaces all old values or all new values for the LCR, depending on the value type specified.

Considerations for LOB Columns

You can use this procedure when processing a row LCR with LOB columns with a procedure DML handler or error handler. If the handler is using LOB assembly (the `assemble_lobs` parameter is set to `TRUE` for the handler), then you can use this member procedure in the handler procedure. If `assemble_lobs` is set to `FALSE` for the handler, then you cannot use this member procedure on a row LCR.

To use a DML or error handler to set the value of one or more LOB columns in a row LCR, specify a LOB locator for each LOB column in the `value_list` parameter. The `SET_VALUES` member procedure verifies that an `ANYDATA` encapsulated LOB locator is processed with a DML or error handler that is using LOB assembly. An error is raised under the following conditions:

- The handler attempts to enqueue a row LCR with an `ANYDATA` encapsulated LOB locator.
- An attempt is made to set a LOB column incorrectly.

If an error is raised because of one of these conditions, then the transaction that includes the row LCR is moved to the error queue, and the LOB columns are represented by the original (nonassembled) row LCRs.

> **Note:**
>
> - Database compatibility must be `10.2.0` or higher to use LOB assembly.
> - When you are processing a row LCR with a rule-based transformation, you cannot use this member procedure on LOB columns.
> - When you are processing a row LCR with a rule-based transformation, procedure DML handler, or error handler, you cannot use this member procedure on `LONG` or `LONG RAW` columns.

Considerations for XMLType Columns

When processing a row LCR with `XMLType` columns with a procedure DML handler or error handler, any `XMLType` and LOB columns in the LCR are always assembled using LOB assembly. You can use this member procedure in the handler procedure on a row LCR that contains one or more `XMLType` columns.

To use a DML or error handler to set the value of one or more `XMLType` columns in a row LCR, specify an `XMLType` for each `XMLType` column in the `value_list` parameter. The `SET_VALUES` member procedure verifies that an `ANYDATA` encapsulated `XMLType` is processed with a DML or error handler. An error is raised under the following conditions:

- The handler attempts to enqueue a row LCR with an `ANYDATA` encapsulated `XMLType`.
- An attempt is made to set a `XMLType` incorrectly.

If an error is raised because of one of these conditions, then the transaction that includes the row LCR is moved to the error queue, and the `XMLType` columns are represented by the original (nonassembled) row LCRs.

> **Note:**
>
> - Database compatibility must be `11.1.0` or higher to process row LCRs with `XMLType` columns.
> - When you are processing a row LCR with a rule-based transformation, you cannot use this member procedure on `XMLType` columns.

Syntax

```
MEMBER PROCEDURE SET_VALUES(
   value_type  IN  VARCHAR2,
   value_list  IN  SYS.LCR$_ROW_LIST);
```

Parameters

Table 310-34 SET_VALUES Procedure Parameters

Parameter	Description
value_type	The type of values to replace Specify old to replace the old values. Specify new to replace the new values.
value_list	List of values to replace the existing list Use a NULL or an empty list to remove all values. If the member procedure is used in a procedure DML handler or error handler that uses LOB assembly, then specify one or more LOB locators for LOB columns.

SET_XML_INFORMATION Member Procedure

Sets the XML information for the column.

Syntax

```
MEMBER PROCEDURE SET_XML_INFORMATION(
   column_name      IN  VARCHAR2,
   xml_information  IN  NUMBER);
```

Parameters

Table 310-35 SET_XML_INFORMATION Procedure Parameters

Parameter	Description
column_name	The name of the column An exception is raised if the column value does not exist in the LCR.

xml_information	Specify one of the following values:
	`DBMS_LCR.NOT_XML CONSTANT NUMBER := 1;`
	`DBMS_LCR.XML_DOC CONSTANT NUMBER := 2;`
	`DBMS_LCR.XML_DIFF CONSTANT NUMBER := 3;`
	`DBMS_LCR.XML_DIFF CONSTANT NUMBER := 3;`
	DBMS_LCR.NOT_XML indicates that the column is not an XMLType column.
	DBMS_LCR.XML_DOC indicates that the column contains an XML document.
	DBMS_LCR.XML_DIFF indicates that the column contains differences between old and new XML documents for an update operation.
	DBMS_LCR.XML_DIFF indicates that the column contains an XML document that describes a change made by an update operation. This XML document describes the differences in the column's XML document. The entire XML document is not replaced.

310.3.3 Common Subprograms for LCR$_DDL_RECORD and LCR$_ROW_RECORD

These functions and procedures are common to both the `LCR$_DDL_RECORD` and `LCR$_ROW_RECORD` type.

> **See Also:**
>
> For descriptions of the subprograms for these types that are exclusive to each type:
>
> - "LCR$_DDL_RECORD Type"
> - "LCR$_ROW_RECORD Type"

Table 310-36 Summary of Common Subprograms for DDL and Row LCR Types

Subprogram	Description
GET_COMMAND_TYPE Member Function	Gets the command type of the logical change record (LCR)
GET_COMMIT_SCN Member Function	Gets the commit system change number (SCN) of the transaction to which the current LCR belongs
GET_COMMIT_SCN_FROM_POSITION Static Function	Gets the commit SCN of a transaction from the input position, which is generated by an XStream outbound server
GET_COMMIT_TIME	Gets the commit time of the transaction to which the current LCR belongs
GET_COMPATIBLE Member Function	Gets the minimal database compatibility required to support the LCR
GET_EXTRA_ATTRIBUTE Member Function	Gets the value for the specified extra attribute in the LCR
GET_OBJECT_NAME Member Function	Gets the name of the object that is changed by the LCR
GET_OBJECT_OWNER Member Function	Gets the owner of the object that is changed by the LCR
GET_POSITION Member Function	Gets the position of the current LCR
GET_ROOT_NAME Member Function	Gets the global name of the root for a CDB.
GET_SCN Member Function	Gets the SCN of the LCR
GET_SCN_FROM_POSITION Static Function	Gets the SCN from the input position, which is generated by an XStream outbound server
GET_SOURCE_DATABASE_NAME Member Function	Gets the source database name.
GET_SOURCE_TIME Member Function	Gets the time when the change in an LCR captured by a capture process was generated in the redo log of the source database, or the time when a persistent LCR was created
GET_TAG Member Function	Gets the tag for the LCR
GET_THREAD_NUMBER Member Function	Gets the thread number of the database instance that made the change that is encapsulated in the LCR

Table 310-36 (Cont.) Summary of Common Subprograms for DDL and Row LCR Types

Subprogram	Description
GET_TRANSACTION_ID Member Function	Gets the transaction identifier of the LCR
IS_NULL_TAG Member Function	Returns `Y` if the tag for the LCR is `NULL`, or returns `N` if the tag for the LCR is not `NULL`
SET_COMMAND_TYPE Member Procedure	Sets the command type in the LCR
SET_EXTRA_ATTRIBUTE Member Procedure	Sets the value for the specified extra attribute in the LCR
SET_OBJECT_NAME Member Procedure	Sets the name of the object that is changed by the LCR
SET_OBJECT_OWNER Member Procedure	Sets the owner of the object that is changed by the LCR
SET_ROOT_NAME Member Procedure	Sets the global name of the root in a CDB.
SET_SOURCE_DATABASE_NAME Member Procedure	Sets the source database name of the object that is changed by the LCR
SET_TAG Member Procedure	Sets the tag for the LCR

GET_COMMAND_TYPE Member Function

Gets the command type of the LCR.

> **See Also:**
>
> The "SQL Command Codes" table in the *Oracle Call Interface Programmer's Guide* for a complete list of command types

Syntax

```
MEMBER FUNCTION GET_COMMAND_TYPE()
RETURN VARCHAR2;
```

GET_COMMIT_SCN Member Function

Gets the commit system change number (SCN) of the transaction to which the current LCR belongs.

The commit SCN for a transaction is available only during apply or during error transaction execution. This function can be used only in a procedure DML handler, DDL handler, or error handler.

The commit SCN might not be available for an LCR that is part of an incomplete transaction. For example, persistent LCRs might not have a commit SCN. If the commit SCN is not available for an LCR, then this function returns `NULL`.

Syntax

```
MEMBER FUNCTION GET_COMMIT_SCN()
RETURN NUMBER;
```

GET_COMMIT_SCN_FROM_POSITION Static Function

Gets the commit system change number (SCN) of a transaction from the input position, which is generated by an XStream outbound server.

Syntax

```
STATIC FUNCTION GET_COMMIT_SCN_FROM_POSITION(
   position  IN  RAW)
RETURN NUMBER;
```

Parameters

Table 310-37 GET_COMMIT_SCN_FROM_POSITION Function Parameter

Parameter	Description
position	The position You can obtain the position by using the GET_POSITION member function or by querying the DBA_XSTREAM_OUTBOUND_PROGRESS data dictionary view.

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

GET_COMMIT_TIME

Gets the commit time of the transaction to which the current LCR belongs.

The commit time for a transaction is available only during apply or during error transaction execution. This function can be used only in a procedure DML handler, DDL handler, or error handler.

The commit time might not be available for an LCR that is part of an incomplete transaction. For example, persistent LCRs might not have a commit time. If the commit time is not available for an LCR, then this function returns NULL.

Syntax

```
MEMBER FUNCTION GET_COMMIT_TIME()
RETURN DATE;
```

GET_COMPATIBLE Member Function

Gets the minimal database compatibility required to support the LCR. You control the compatibility of an Oracle database using the COMPATIBLE initialization parameter.

The return value for this function can be one of the following:

Return Value	COMPATIBLE Initialization Parameter Equivalent
DBMS_STREAMS.COMPATIBLE_9_2	9.2.0
DBMS_STREAMS.COMPATIBLE_10_1	10.1.0

Return Value	COMPATIBLE Initialization Parameter Equivalent
DBMS_STREAMS.COMPATIBLE_10_2	10.2.0
DBMS_STREAMS.COMPATIBLE_11_1	11.1.0
DBMS_STREAMS.COMPATIBLE_11_2	11.2.0

DDL LCRs always return DBMS_STREAMS.COMPATIBLE_9_2.

You can use the following functions in the DBMS_STREAMS package for constant compatibility return values:

- The COMPATIBLE_9_2 function returns the DBMS_STREAMS.COMPATIBLE_9_2 constant.
- The COMPATIBLE_10_1 function returns DBMS_STREAMS.COMPATIBLE_10_1 constant.
- The COMPATIBLE_10_2 function returns DBMS_STREAMS.COMPATIBLE_10_2 constant.
- The COMPATIBLE_11_1 function returns DBMS_STREAMS.COMPATIBLE_11_1 constant.
- The COMPATIBLE_11_2 function returns DBMS_STREAMS.COMPATIBLE_11_2 constant.
- The MAX_COMPATIBLE function returns an integer that is greater than the highest possible compatibility constant for the current release of Oracle Database.

You can use these functions with the GET_COMPATIBLE member function for an LCR in rule conditions and apply handlers.

See Also:

- *Oracle Database Reference* and *Oracle Database Upgrade Guide* for more information about the COMPATIBLE initialization parameter

Syntax

```
MEMBER FUNCTION GET_COMPATIBLE()
RETURN NUMBER;
```

GET_EXTRA_ATTRIBUTE Member Function

Gets the value for the specified extra attribute in the LCR. The returned extra attribute is contained within an ANYDATA instance. You can use the INCLUDE_EXTRA_ATTRIBUTE procedure in the DBMS_CAPTURE_ADM package to instruct a capture process to capture one or more extra attributes.

See Also:

INCLUDE_EXTRA_ATTRIBUTE Procedure

Syntax

```
MEMBER FUNCTION GET_EXTRA_ATTRIBUTE(
   attribute_name  IN  VARCHAR2)
RETURN ANYDATA;
```

Parameters

Table 310-38 GET_EXTRA_ATTRIBUTE Function Parameter

Parameter	Description
attribute_name	The name of the extra attribute to return Valid names are: • row_id The rowid of the row changed in a row LCR. This attribute is not included in DDL LCRs, nor in row LCRs for index-organized tables. The type is UROWID. • serial# The serial number of the session that performed the change captured in the LCR. The type is NUMBER. • session# The identifier of the session that performed the change captured in the LCR. The type is NUMBER. • thread# The thread number of the instance in which the change captured in the LCR was performed. Typically, the thread number is relevant only in an Oracle Real Application Clusters (Oracle RAC) environment. The type is NUMBER. • tx_name The name of the transaction that includes the LCR. The type is VARCHAR2. • username The name of the current user who performed the change captured in the LCR. The type is VARCHAR2. An error is raised if the specified attribute_name is not valid. If no value exists for the specified extra attribute, then the function returns a NULL. **See Also:** *Oracle Database PL/SQL Language Reference* for more information about the current user

GET_OBJECT_NAME Member Function

Gets the name of the object that is changed by the LCR.

Syntax

```
MEMBER FUNCTION GET_OBJECT_NAME()
RETURN VARCHAR2;
```

GET_OBJECT_OWNER Member Function

Gets the owner of the object that is changed by the LCR.

Syntax

```
MEMBER FUNCTION GET_OBJECT_OWNER()
RETURN VARCHAR2;
```

GET_POSITION Member Function

Gets the position of the current LCR. The position uniquely identifies each LCR. The position strictly increases within each transaction and across transactions.

LCR position is commonly used in XStream configurations.

Syntax

```
MEMBER FUNCTION GET_POSITION()
RETURN RAW;
```

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

GET_ROOT_NAME Member Function

Gets the global name of the root in a CDB, which is the root name for the LCR.

Syntax

```
MEMBER FUNCTION GET_ROOT_NAME()
RETURN VARCHAR2;
```

GET_SCN Member Function

Gets the system change number (SCN) of the LCR.

Syntax

```
MEMBER FUNCTION GET_SCN()
RETURN NUMBER;
```

GET_SCN_FROM_POSITION Static Function

Gets the system change number (SCN) from the input position, which is generated by an XStream outbound server.

Syntax

```
STATIC FUNCTION GET_SCN_FROM_POSITION(
   position  IN  RAW)
RETURN NUMBER;
```

Parameters

Table 310-39 GET_SCN_FROM_POSITION Function Parameter

Parameter	Description
`position`	The position You can obtain the position by using the `GET_POSITION` member function or by querying the `DBA_XSTREAM_OUTBOUND_PROGRESS` data dictionary view.

> **Note:**
>
> Using XStream requires purchasing a license for the Oracle GoldenGate product. See *Oracle Database XStream Guide*.

GET_SOURCE_DATABASE_NAME Member Function

Gets the global name of the source database. The source database is the database where the change occurred.

Syntax

```
MEMBER FUNCTION GET_SOURCE_DATABASE_NAME()
RETURN VARCHAR2;
```

GET_SOURCE_TIME Member Function

Gets the time when the change in an LCR captured by a capture process was generated in the redo log of the source database, or the time when a persistent LCR was created.

Syntax

```
MEMBER FUNCTION GET_SOURCE_TIME()
RETURN DATE;
```

GET_TAG Member Function

Gets the tag for the LCR. An LCR tag is a binary tag that enables tracking of the LCR. For example, this tag can be used to determine the original source database of the DML or DDL change when apply forwarding is used.

Syntax

```
MEMBER FUNCTION GET_TAG()
RETURN RAW;
```

GET_THREAD_NUMBER Member Function

Gets the thread number of the database instance that made the change that is encapsulated in the LCR. Typically, the thread number is relevant in an Oracle Real Application Clusters configuration.

> **See Also:**
>
> *Oracle Real Application Clusters Administration and Deployment Guide*

Syntax

```
MEMBER FUNCTION GET_THREAD_NUMBER()
RETURN NUMBER;
```

GET_TRANSACTION_ID Member Function

Gets the transaction identifier of the LCR.

Syntax

```
MEMBER FUNCTION GET_TRANSACTION_ID()
RETURN VARCHAR2;
```

IS_NULL_TAG Member Function

Returns `Y` if the tag for the LCR is `NULL`, or returns `N` if the tag for the LCR is not `NULL`.

Syntax

```
MEMBER FUNCTION IS_NULL_TAG()
RETURN VARCHAR2;
```

SET_COMMAND_TYPE Member Procedure

Sets the command type in the LCR. If the command type specified cannot be interpreted, then this procedure raises an error. For example, changing `INSERT` to `GRANT` would raise an error.

> **See Also:**
>
> - The description of the `command_type` parameter in LCR$_DDL_RECORD Constructor Function Parameters
> - The description of the command_type parameter in LCR$_ROW_RECORD Type
> - The "SQL Command Codes" table in the *Oracle Call Interface Programmer's Guide* for a complete list of command types

Syntax

```
MEMBER PROCEDURE SET_COMMAND_TYPE(
   command_type  IN  VARCHAR2);
```

Parameter

Table 310-40 SET_COMMAND_TYPE Procedure Parameter

Parameter	Description
`command_type`	The command type Set this parameter to a non-`NULL` value.

SET_EXTRA_ATTRIBUTE Member Procedure

Sets the value for the specified extra attribute in the LCR. You can use the `INCLUDE_EXTRA_ATTRIBUTE` procedure in the `DBMS_CAPTURE_ADM` package to instruct a capture process to capture one or more extra attributes.

> **See Also:**
>
> INCLUDE_EXTRA_ATTRIBUTE Procedure

Syntax

```
MEMBER PROCEDURE SET_EXTRA_ATTRIBUTE(
   attribute_name   IN  VARCHAR2,
   attribute_value  IN  ANYDATA);
```

Parameters

Table 310-41 SET_EXTRA_ATTRIBUTE Procedure Parameter

Parameter	Description
attribute_name	The name of the extra attribute to set Valid names are: • row_id The rowid of the row changed in a row LCR. This attribute is not included in DDL LCRs, nor in row LCRs for index-organized tables. The type is VARCHAR2. • serial# The serial number of the session that performed the change captured in the LCR. The type is NUMBER. • session# The identifier of the session that performed the change captured in the LCR. The type is NUMBER. • thread# The thread number of the instance in which the change captured in the LCR was performed. Typically, the thread number is relevant only in an Oracle Real Application Clusters (Oracle RAC) environment. The type is NUMBER. • tx_name The name of the transaction that includes the LCR. The type is VARCHAR2. • username The name of the current user who performed the change captured in the LCR. The type is VARCHAR2. An error is raised if the specified attribute_name is not valid. **See Also:** *Oracle Database PL/SQL Language Reference* for more information about the current user
attribute_value	The value to which the specified extra attribute is set If set to NULL, then this procedure removes the specified extra attribute from the LCR. To set to NULL, encapsulate the NULL in an ANYDATA instance.

SET_OBJECT_NAME Member Procedure

Sets the name of the object that is changed by the LCR.

Syntax

```
MEMBER PROCEDURE SET_OBJECT_NAME(
   object_name  IN  VARCHAR2);
```

Parameter

Table 310-42 SET_OBJECT_NAME Procedure Parameter

Parameter	Description
object_name	The name of the object

SET_OBJECT_OWNER Member Procedure

Sets the owner of the object that is changed by the LCR.

Syntax

```
MEMBER PROCEDURE SET_OBJECT_OWNER(
   object_owner  IN  VARCHAR2);
```

Parameter

Table 310-43 SET_OBJECT_OWNER Procedure Parameter

Parameter	Description
object_owner	The schema that contains the object

SET_ROOT_NAME Member Procedure

Sets the global name of the root in a CDB. The setting is the root name for the LCR.

Syntax

```
MEMBER PROCEDURE SET_ROOT_NAME(
   root_name  IN  VARCHAR2);
```

Parameter

Table 310-44 SET_ROOT_NAME Procedure Parameter

Parameter	Description
root_name	The global name of the root.

SET_SOURCE_DATABASE_NAME Member Procedure

Sets the source database name of the object that is changed by the LCR.

Syntax

```
MEMBER PROCEDURE SET_SOURCE_DATABASE_NAME(
   source_database_name  IN  VARCHAR2);
```

Parameter

Table 310-45 SET_SOURCE_DATABASE_NAME Procedure Parameter

Parameter	Description
source_database_name	The source database of the change If you do not include the domain name, then the procedure appends the local domain to the database name automatically. For example, if you specify DBS1 and the local domain is EXAMPLE.COM, then the procedure specifies DBS1.EXAMPLE.COM automatically. Set this parameter to a non-NULL value.

SET_TAG Member Procedure

Sets the tag for the LCR. An LCR tag is a binary tag that enables tracking of the LCR. For example, this tag can be used to determine the original source database of the change when apply forwarding is used.

Syntax

```
MEMBER PROCEDURE SET_TAG(
   tag  IN  RAW);
```

Parameter

Table 310-46 SET_TAG Procedure Parameter

Parameter	Description
tag	The binary tag for the LCR The size limit for a tag value is two kilobytes.

310.3.4 LCR$_ROW_LIST Type

This type identifies a list of column values for a row in a table.

It uses the LCR$_ROW_UNIT type and is used in the LCR$_ROW_RECORD type.

> **See Also:**
>
> - LCR$_ROW_UNIT Type
> - LCR$_ROW_RECORD Type

Syntax

```
CREATE TYPE SYS.LCR$_ROW_LIST AS TABLE OF SYS.LCR$_ROW_UNIT
/
```

310.3.5 LCR$_ROW_UNIT Type

This type identifies the value for a column in a row.

It is used in the `LCR$_ROW_LIST` type.

> **See Also:**
>
> LCR$_ROW_LIST Type

Syntax

```
CREATE TYPE LCR$_ROW_UNIT AS OBJECT (
  column_name         VARCHAR2(4000),
  data                ANYDATA,
  lob_information     NUMBER,
  lob_offset          NUMBER,
  lob_operation_size  NUMBER
  long_information    NUMBER);
/
```

Attributes

Table 310-47 LCR$_ROW_UNIT Attributes

Attribute	Description
`column_name`	The name of the column
`data`	The data contained in the column

`lob_information`	Contains the LOB information for the column and contains one of the following values:
	`DBMS_LCR.NOT_A_LOB CONSTANT NUMBER := 1;`
	`DBMS_LCR.NULL_LOB CONSTANT NUMBER := 2;`
	`DBMS_LCR.INLINE_LOB CONSTANT NUMBER := 3;`
	`DBMS_LCR.EMPTY_LOB CONSTANT NUMBER := 4;`
	`DBMS_LCR.LOB_CHUNK CONSTANT NUMBER := 5;`
	`DBMS_LCR.LAST_LOB_CHUNK CONSTANT NUMBER := 6;`

`lob_offset`	The LOB offset specified in the number of characters for `CLOB` columns and the number of bytes for `BLOB` columns Valid values are `NULL` or a positive integer less than or equal to `DBMS_LOB.LOBMAXSIZE`.
`lob_operation_size`	If `lob_information` for the LOB is `DBMS_LCR.LAST_LOB_CHUNK`, then this parameter can be set to either a valid `LOB ERASE` value or a valid `LOB TRIM` value. A `LOB ERASE` value must be a positive integer less than or equal to `DBMS_LOB.LOBMAXSIZE`. A `LOB TRIM` value must be a nonnegative integer less than or equal to `DBMS_LOB.LOBMAXSIZE`. If `lob_information` is not `DBMS_LCR.LAST_LOB_CHUNK` and for all other operations, is `NULL`.

Table 310-47 (Cont.) LCR$_ROW_UNIT Attributes

Attribute	Description
long_information	Contains the LONG information for the column and contains one of the following values: DBMS_LCR.not_a_long CONSTANT NUMBER := 1; DBMS_LCR.null_long CONSTANT NUMBER := 2; DBMS_LCR.inline_long CONSTANT NUMBER := 3; DBMS_LCR.long_chunk CONSTANT NUMBER := 4; DBMS_LCR.last_long_chunk CONSTANT NUMBER := 5;

311
MGD_ID Package Types

The MGD_ID package provides an extensible framework that supports current radio-frequency ID (RFID) tags with the standard family of EPC bit encodings for the supported encoding types. The MGD_ID Package also supports new and evolving tag encodings that are not included in the current EPC standard (EPC v1.1 specification).

The MGD_ID package contains several predefined types.

> **See Also:**
>
> *Oracle Database Development Guide* for more information.

This chapter contains the following topics:

- Security Model
- Summary of Types
- Summary of MGD_ID Subprograms

The method described in this reference chapter show examples based on the examples shown in the constructor functions.

The examples in this chapter assume that the you have run the following set of commands before running the contents of each script:

```
SQL> connect / as sysdba;
Connected.
SQL> create user mgduser identified by mgduser;
SQL> grant connect, resource to mgduser;
SQL> connect mgduser
Enter password: mgduserpassword
Connected.
SQL> set serveroutput on;
```

311.1 MGD_ID Package Types Security Model

You must run the catmgd.sql script to load the DBMS_MGD_ID_UTL package and create the required Identity Code Package schema objects in the MGDSYS schema.

MGD_ID is a MGDSYS-owned object type. Any MGD_ID subprogram called from an anonymous PL/SQL block is run using the privileges of the current user.

A user must be granted connect and resource roles to use the MGD_ID object type and its subprograms.

EXECUTE privilege is granted to PUBLIC for this ADT: MGD_ID.

A public synonym, by the same name, is created for this ADT: MGD_ID.

311.2 Summary of Types

This table lists and briefly describes the `MGD_ID` Package object types.

Table 311-1 MGD_ID Package Object Types

Object Type Name	Description
MGD_ID_COMPONENT Object Type	Datatype that specifies the name and value pair attributes that define a component
MGD_ID_COMPONENT_VARRAY Object Type	Datatype that specifies a list of up to 128 components as name-value attribute pairs used in two constructor functions for creating an identity code type object
MGD_ID Object Type	Represents an `MGD_ID` object that specifies the category identifier for the code category for this identity code and its list of components

311.2.1 MGD_ID_COMPONENT Object Type

The `MGD_ID_COMPONENT` type is a datatype that specifies the name and value pair attributes that define a component.

Syntax

```
MGD_ID_COMPONENT as object (name  VARCHAR2(256),
                            value VARCHAR2(1024));
```

Attributes

Table 311-2 MGD_ID_COMPONENT Attributes

Attribute	Description
`name`	Name of component
`value`	Value of the component as a character

Examples

See the MGD_ID Constructor Function for an example.

311.2.2 MGD_ID_COMPONENT_VARRAY Object Type

The `MGD_ID_COMPONENT_VARRAY` type is a datatype that specifies a list of up to 128 components as name-value attribute pairs for use in two constructor functions for creating a product code type object with its list of components.

Syntax

```
MGD_ID_COMPONENT_VARRAY is VARRAY (128) of MGD_ID_COMPONENT;
```

Examples

See the MGD_ID Constructor Function for an example.

311.2.3 MGD_ID Object Type

The `MGD_ID` type represents an identity code in an RFID application. This type represents RFID tags with standard EPC bit encoding as well as tag encodings that are not included in the EPC standard.

Syntax

```
MGD_ID as object (category_id  VARCHAR2(256),
                  components   MGD_ID_COMPONENT_VARRAY);
```

Attributes

Table 311-3 MGD_ID Object Type Attributes

Attribute	Description
`category_id`	Category identifier for the code category of this code
`components`	List of components as name-value attributes

Methods

Table 311-5 describes the methods of the `MGD_ID` object type.

Table 311-4 MGD_ID Methods

Method	Description
`MGD_ID` constructor function	Creates an `MGD_ID` object based on the parameters passed in and returns self as a result
`FORMAT` function	Returns the string representation of the `MGD_ID` in the specified format
`GET_COMPONENT` function	Returns the string value of the specified `MGD_ID` component
`TO_STRING` function	Returns the string value of semicolon (;) separated component name value pairs of the `MGD_ID` object
`TRANSLATE` function	Returns the result of the conversion of the identifier from one format to the specified format

Examples

See the Summary of MGD_ID Subprograms section and the section about using the Identity Code package in Using the Identity Code Package in *Oracle Database Development Guide* for examples.

311.3 Summary of MGD_ID Subprograms

This table describes the subprograms in the `MGD_ID` object type.

All the values and names passed to the procedures defined in the `MGD_ID` object type are case insensitive unless otherwise mentioned. To preserve the case, enclose the values with double quotation marks.

Table 311-5 MGD_ID Object Type Subprograms

Subprogram	Description
MGD_ID Constructor Function	Creates an MGD_ID object based on the parameters passed in and returns self as a result
FORMAT Function	Returns the string representation of the MGD_ID object in the specified format
GET_COMPONENT Function	Returns the string value of the specified MGD_ID component
TO_STRING Function	Returns the string value of semicolon (;) separated component name value pairs of the MGD_ID object
TRANSLATE Function	Returns the result of the conversion of the identifier from one format to the specified format

311.3.1 MGD_ID Constructor Function

This constructor function constructs an identity code type object, MGD_ID. The constructor function is overloaded. The different functionality of each form of syntax is presented along with the definitions.

Syntax

Constructs an MGD_ID object type based on the category ID and a list of components.

```
MGD_ID (
   category_id      IN VARCHAR2,
   components       IN MGD_ID_COMPONENT_VARRAY)
 RETURN SELF AS RESULT DETERMINISTIC;
```

Constructs an MGD_ID object type based on the category ID, the identifier string, and the list of additional parameters required to create it.

```
MGD_ID (
   category_id        VARCHAR2,
   identifier         VARCHAR2,
   parameter_list      VARCHAR2)
 RETURN SELF AS RESULT DETERMINISTIC;
```

Constructs an MGD_ID object type based on the category name, category version, and a list of components.

```
MGD_ID (
   category_name      VARCHAR2,
   category_version   VARCHAR2,
   components         MGD_ID_COMPONENT_VARRAY)
 RETURN SELF AS RESULT DETERMINISTIC;
```

Constructs an MGD_ID object type based on the category name, category version, the identifier string, and the list of additional parameters required to create it.

```
MGD_ID (
   category_name      VARCHAR2,
   category_version   VARCHAR2,
   identifier         VARCHAR2,
   parameter_list      VARCHAR2)
 RETURN SELF AS RESULT DETERMINISTIC;
```

Parameters

Table 311-6 MGD_ID Constructor Function Parameters

Parameter	Description
category_id	Category identifier
components	List of component name value pairs
category_name	Category name, such as EPC
category_version	Category version. If NULL, the latest version for the specified category name will be used.
identifier	Identifier string in any format of an encoding scheme in the specified category. For example, for SGTIN-96 encoding, the identifier can be in the format of BINARY, PURE_IDENTITY, TAG_ENCODING, or LEGACY. Express this identifier as a string according to the appropriate grammar or pattern in the tag data translation (TDT) markup file. For example, a binary string consisting of characters 0 and 1, a URI (either tag-encoding or pure-identity formats), or a serialized legacy code expressed as a string format for input, such as gtin=00037000302414;serial=10419703 for a SGTIN coding scheme.
parameter_list	List of additional parameters required to create the object in the representation. The list is expressed as a parameter string containing key-value pairs, separated by the semicolon (;) as a delimiter between key-value pairs. For example, for a GTIN code, the parameter string would look as follows: filter=3;companyprefixlength=7;taglength=96

Usage Notes

- Use MGD_ID_UTL.EPC_ENCODING_CATEGORY_ID as category_id.
- If the category is not already registered, an error is raised.
- If the bit_length parameter is NULL, the bit_length is 8* the length of bit_encoding.
- If the component list does not contain all required components, an exception MGD_ID_UTL.e_LackComponent will be thrown.

Examples

The following examples construct identity code type objects.

Construct an MGD_ID object (SGTIN-64) passing in the category ID and a list of components.

```
--Contents of constructor11.sql
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
call DBMS_MGD_ID_UTL.refresh_category('1');
select MGD_ID('1',
              MGD_ID_COMPONENT_VARRAY(
               MGD_ID_COMPONENT('companyprefix','0037000'),
               MGD_ID_COMPONENT('itemref','030241'),
               MGD_ID_COMPONENT('serial','1041970'),
               MGD_ID_COMPONENT('schemes','SGTIN-64')
              )
             ) from dual;
call DBMS_MGD_ID_UTL.remove_proxy();
```

```
SQL> @constructor11.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY(MGD_ID_COMPONENT('companyprefix', '0037000'),
                                    MGD_ID_COMPONENT('itemref', '030241'),
                                    MGD_ID_COMPONENT('serial', '1041970'),
                                    MGD_ID_COMPONENT('schemes', 'SGTIN-64')))
.
.
.
```

Constructs an `MGD_ID` object (SGTIN-64) passing in the category ID, the tag identifier, and the list of additional parameters that may be required to create it.

```
--Contents of constructor22.sql
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
call DBMS_MGD_ID_UTL.refresh_category('1');
select MGD_ID('1',
              'urn:epc:id:sgtin:0037000.030241.1041970',
              'filter=3;scheme=SGTIN-64') from dual;
call DBMS_MGD_ID_UTL.remove_proxy();

SQL> @constructor22.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY(MGD_ID_COMPONENT('filter', '3'),
                                    MGD_ID_COMPONENT('schemes', 'SGTIN-64'),
                                    MGD_ID_COMPONENT('companyprefixlength', '7'),
                                    MGD_ID_COMPONENT('companyprefix', '0037000'),
                                    MGD_ID_COMPONENT('scheme', 'SGTIN-64'),
                                    MGD_ID_COMPONENT('serial', '1041970'),
                                    MGD_ID_COMPONENT('itemref', '030241')))
.
.
.
```

Construct an `MGD_ID` object (SGTIN-64) passing in the category name, category version (if `NULL`, then the latest version will be used), and a list of components.

```
--Contents of constructor33.sql
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
call DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
select MGD_ID('EPC', NULL,
              MGD_ID_COMPONENT_VARRAY(
               MGD_ID_COMPONENT('companyprefix','0037000'),
               MGD_ID_COMPONENT('itemref','030241'),
               MGD_ID_COMPONENT('serial','1041970'),
               MGD_ID_COMPONENT('schemes','SGTIN-64')
              )
             ) from dual;
call DBMS_MGD_ID_UTL.remove_proxy();

SQL> @constructor33.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY(MGD_ID_COMPONENT('companyprefix', '0037000'),
                                    MGD_ID_COMPONENT('itemref', '030241'),
                                    MGD_ID_COMPONENT('serial', '1041970'),
                                    MGD_ID_COMPONENT('schemes', 'SGTIN-64')))
```

```
.
.
.
```

Constructs an `MGD_ID` object (SGTIN-64) passing in the category name and category version, the tag identifier, and the list of additional parameters that may be required to create it.

```
--Contents of constructor44.sql
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
call DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
select MGD_ID('EPC', NULL,
              'urn:epc:id:sgtin:0037000.030241.1041970',
              'filter=3;scheme=SGTIN-64') from dual;
call DBMS_MGD_ID_UTL.remove_proxy();

SQL> @constructor4.sql
.
.
.
MGD_ID('1', MGD_ID_COMPONENT_VARRAY(MGD_ID_COMPONENT('filter', '3'),
                                    MGD_ID_COMPONENT('schemes', 'SGTIN-64'),
                                    MGD_ID_COMPONENT('companyprefixlength', '7'),
                                    MGD_ID_COMPONENT('companyprefix', '0037000'),
                                    MGD_ID_COMPONENT('scheme', 'SGTIN-64'),
                                    MGD_ID_COMPONENT('serial', '1041970'),
                                    MGD_ID_COMPONENT('itemref', '030241')))
.
.
.
```

311.3.2 FORMAT Function

This function returns the string representation of the `MGD_ID` object in the specified format.

Syntax

```
FORMAT (parameter_list  IN VARCHAR2,
        output_format   IN VARCHAR2)
 RETURN VARCHAR2 DETERMINISTIC;
```

Parameters

Table 311-7 FORMAT Function Parameters

Parameter	Description
`parameter_list`	List of additional parameters required to create the object in the representation. The list is expressed as a parameter string containing key-value pairs, separated by the semicolon (;) as a delimiter between key-value pairs. For example, for a GTIN code, the parameter string would look as follows: `filter=3;companyprefixlength=7;taglength=96`

Table 311-7 (Cont.) FORMAT Function Parameters

Parameter	Description
output_format	One of the supported output formats into which an MGD_ID component is formatted: • BINARY • LEGACY • TAG_ENCODING • PURE_IDENTITY • ONS_HOSTNAME

Examples

See the example for the GET_COMPONENT Function.

311.3.3 GET_COMPONENT Function

This function returns the value of the specified MGD_ID component.

Syntax

```
GET_COMPONENT (
   component_name  IN  VARCHAR2)
 RETURN VARCHAR2 DETERMINISTIC;
```

Parameters

Table 311-8 GET_COMPONENT Function Parameter

Parameter	Description
component_name	Name of component

Usage Notes

- If the code is an invalid code, meaning its structure is not defined in the metadata table, an error is raised.
- If the code is valid, but it does not contain the required component, NULL is returned.

Examples

The following example returns the general manager, object class, and serial number components for this GID-96 identity component:

```
--Contents of get_components.sql file
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
DECLARE
id  MGD_ID;
BEGIN
 DBMS_MGD_ID_UTL.set_java_logging_level(DBMS_MGD_ID_UTL.LOGGING_LEVEL_OFF);
 DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
    ---------------
    --PURE_IDENTIT
    ---------------
    dbms_output.put_line('..Testing constructor with pure identity');
```

```
    ------------------------------------------------------------------------
    -- PURE_IDENTITY representation can be translated to BINARY and
    -- TAG_ENCODING ONLY when BOTH scheme and filer are provided.
    ------------------------------------------------------------------------
    id := MGD_ID('EPC', NULL, 'urn:epc:id:sgtin:0037000.030241.1041970', 'scheme=SGTIN-64;filter=3');
    dbms_output.put_line(id.to_string);
    dbms_output.put_line('filter         = ' || id.get_component('filter'));
    dbms_output.put_line('company prefix = ' || id.get_component('companyprefix'));
    dbms_output.put_line('itemref        = ' || id.get_component('itemref'));
    dbms_output.put_line('serial         = ' || id.get_component('serial'));
    dbms_output.put_line('BINARY format  = ' || id.format(NULL, 'BINARY'));
    dbms_output.put_line('PURE_IDENTITY format  = ' || id.format(NULL, 'PURE_IDENTITY'));
    dbms_output.put_line('TAG_ENCODING format   = ' || id.format(NULL, 'TAG_ENCODING'));
END;
/
SHOW ERRORS;
call DBMS_MGD_ID_UTL.remove_proxy();
SQL> @get_component.sql
.
.
.
..Testing constructor with pure identity
category_id =1;filter = 3;schemes = SGTIN-64;companyprefixlength =
7;companyprefix = 0037000;scheme = SGTIN-64;serial = 1041970;itemref = 030241
filter         = 3
company prefix = 0037000
itemref        = 030241
serial         = 1041970
BINARY format  =10011000000000000010000011101100010000100000111111100110001100010
PURE_IDENTITY format  = urn:epc:id:sgtin:0037000.030241.1041970
TAG_ENCODING format   = urn:epc:tag:sgtin-64:3.0037000.030241.1041970
PL/SQL procedure successfully completed.
.
.
.
```

311.3.4 TO_STRING Function

This function returns the semicolon (;) separated component name value pairs of the `MGD_ID` object.

Syntax

```
TO_STRING
 RETURN VARCHAR2;
```

Examples

The following example converts the `MGD_ID` object into a string value:

```
-- Contents of tostring3.sql file
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
DECLARE
id          MGD_ID;
BEGIN
 DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
 dbms_output.put_line('..Testing to_string');
   id := mgd_id('EPC', NULL, 'urn:epc:id:gid:0037000.30241.1041970', 'scheme=GID-96');
      DBMS_OUTPUT.PUT_LINE('mgd_id object as a string');
      DBMS_OUTPUT.PUT_LINE(id.to_string);
END;
```

```
/
SHOW ERRORS;
call DBMS_MGD_ID_UTL.remove_proxy();
connect / as sysdba;
drop user mgduser cascade;

SQL> @tostring3.sql
.
.
.
..Testing to_string
mgd_id object as a string
category_id =1;schemes = GID-96;objectclass = 30241;generalmanager =
0037000;scheme = GID-96;1 = 1;serial = 1041970
PL/SQL procedure successfully completed.
.
.
```

311.3.5 TRANSLATE Function

This static function translates between different representations directly without first constructing an `MGD_ID` object.

This method is overloaded. The different functionality of each form of syntax is presented along with the definitions.

Syntax

Converts the identifier in one format to another given the category name, the tag identifier, the parameter list, and the output format.

```
TRANSLATE (
   category_name    IN VARCHAR2,
   identifier       IN VARCHAR2,
   parameter_list   IN VARCHAR2,
   output_format    IN VARCHAR2)
 RETURN VARCHAR2 DETERMINISTIC;
```

Converts the identifier in one format to another given the category name, category version, the tag identifier, the parameter list, and the output format.

```
TRANSLATE (
   category_name    IN VARCHAR2,
   category_version IN VARCHAR2,
   identifier       IN VARCHAR2,
   parameter_list   IN VARCHAR2,
   output_format    IN VARCHAR2)
 RETURN VARCHAR2 DETERMINISTIC;
```

Parameters

Table 311-9 TRANSLATE Function Parameters

Parameter	Description
`category_name`	Name of category
`category_version`	Category version. If `NULL`, the latest version of the specified category name will be used.

Table 311-9 (Cont.) TRANSLATE Function Parameters

Parameter	Description
identifier	EPC identifier, expressed as a string in accordance with one of the grammars or patterns in the TDT markup file. For example, a binary string consisting of characters 0 and 1, a URI (either tag-encoding or pure-identity formats), or a serialized legacy code expressed as a string format for input, such as gtin=00037000302414;serial=10419703 for a SGTIN coding scheme.
parameter_list	List of additional parameters required to create the object in the representation. The list is expressed as a parameter string containing key-value pairs, separated by the semicolon (;) as a delimiter between key-value pairs. For example, for a GTIN code, the parameter string would look as follows: filter=3;companyprefixlength=7;taglength=96
output_format	One of the supported output formats into which an MGD_ID component shall be converted: • BINARY • LEGACY • TAG_ENCODING • PURE_IDENTITY • ONS_HOSTNAME

Usage Notes

When converting from a pure identity representation to a binary representation, the filter value must be supplied as a value using the parameter_list parameter.

Examples

The following examples translates one GID-96 representation into another:

```
-- Contents of translate1.sql file
call DBMS_MGD_ID_UTL.set_proxy('www-proxy.example.com', '80');
DECLARE
id          MGD_ID;
BEGIN
 DBMS_MGD_ID_UTL.refresh_category(DBMS_MGD_ID_UTL.get_category_id('EPC', NULL));
  dbms_output.put_line('Category ID is EPC, Identifier is BINARY, Output format is BINARY');
  dbms_output.put_line(
    mgd_id.translate('EPC',
NULL,'00110101000000000000100100001000100000000000011101100010000100000000000000000011111110011000110010'
, NULL, 'BINARY'));
  dbms_output.put_line('Category ID is EPC, Identifier is BINARY, Output format is PURE_IDENTITY');
  dbms_output.put_line(
    mgd_id.translate('EPC',
NULL,'00110101000000000000100100001000100000000000011101100010000100000000000000000011111110011000110010'
, NULL, 'PURE_IDENTITY'));
dbms_output.put_line('Category ID is EPC, Identifier is BINARY, Output format is TAG_ENCODING');
  dbms_output.put_line(
    mgd_id.translate('EPC',
NULL,'00110101000000000000100100001000100000000000011101100010000100000000000000000011111110011000110010'
, NULL, 'TAG_ENCODING'));
  dbms_output.put_line('Category ID is EPC, Identifier is TAG_ENCODING, Output format is BINARY');
  dbms_output.put_line(
    mgd_id.translate('EPC', NULL,
                     'urn:epc:tag:gid-96:0037000.30241.1041970',
```

```
                    NULL, 'BINARY'));
  dbms_output.put_line('Category ID is EPC, Identifier is TAG_ENCODING, Output format is
PURE_IDENTITY');
  dbms_output.put_line(
    mgd_id.translate('EPC', NULL,
                    'urn:epc:tag:gid-96:0037000.30241.1041970',
                    NULL, 'PURE_IDENTITY'));
dbms_output.put_line('Category ID is EPC, Identifier is TAG_ENCODING, Output format is TAG_ENCODING');
  dbms_output.put_line(
    mgd_id.translate('EPC', NULL,
                    'urn:epc:tag:gid-96:0037000.30241.1041970',
                    NULL, 'TAG_ENCODING'));
  dbms_output.put_line('Category ID is EPC, Identifier is PURE_IDENTITY, Output format is BINARY');
  dbms_output.put_line(
    mgd_id.translate('EPC', NULL,
                    'urn:epc:id:gid:0037000.30241.1041970',
                    NULL, 'BINARY'));
  dbms_output.put_line('Category ID is EPC, Identifier is PURE_IDENTITY, Output format is
PURE_IDENTITY');
  dbms_output.put_line(
    mgd_id.translate('EPC', NULL,
                    'urn:epc:id:gid:0037000.30241.1041970',
                    NULL, 'PURE_IDENTITY'));
  dbms_output.put_line('Category ID is EPC, Identifier is PURE_IDENTITY, Output format is TAG_ENCODING');
  dbms_output.put_line(
    mgd_id.translate('EPC', NULL,
                    'urn:epc:id:gid:0037000.30241.1041970',
                    NULL, 'TAG_ENCODING'));
END;
/
SHOW ERRORS;
call DBMS_MGD_ID_UTL.remove_proxy();

SQL> @translate1.sql
.
.
.
Category ID is EPC, Identifier is BINARY, Output format is BINARY
00110101000000000000100100001000100000000000011101100010000100000000000000000011111110011000110010
Category ID is EPC, Identifier is BINARY, Output format is PURE_IDENTITY
urn:epc:id:gid:37000.30241.1041970
Category ID is EPC, Identifier is BINARY, Output format is TAG_ENCODING
urn:epc:tag:gid-96:37000.30241.1041970
Category ID is EPC, Identifier is TAG_ENCODING, Output format is BINARY
00110101000000000000100100001000100000000000011101100010000100000000000000000011111110011000110010
Category ID is EPC, Identifier is TAG_ENCODING, Output format is PURE_IDENTITY
urn:epc:id:gid:0037000.30241.1041970
Category ID is EPC, Identifier is TAG_ENCODING, Output format is TAG_ENCODING
urn:epc:tag:gid-96:0037000.30241.1041970
Category ID is EPC, Identifier is PURE_IDENTITY, Output format is BINARY
00110101000000000000100100001000100000000000011101100010000100000000000000000011111110011000110010
Category ID is EPC, Identifier is PURE_IDENTITY, Output format is PURE_IDENTITY
urn:epc:id:gid:0037000.30241.1041970
Category ID is EPC, Identifier is PURE_IDENTITY, Output format is TAG_ENCODING
urn:epc:tag:gid-96:0037000.30241.1041970
PL/SQL procedure successfully completed.
.
.
.
```

312
Rule TYPEs

This chapter describes the types used with rules, rule sets, and evaluation contexts.

This chapter contains the following topics:

- Overview
- Security Model
- Summary of Rule Types

312.1 Rule TYPEs Overview

These types are used in rules and enable clients to evaluate rules with the rules engine.

> **See Also:**
>
> - DBMS_RULE
> - DBMS_RULE_ADM

312.2 Rule TYPEs Security Model

`PUBLIC` is granted `EXECUTE` privilege on these RULE types.

> **See Also:**
>
> *Oracle Database Security Guide* for more information about user group `PUBLIC`

312.3 Summary of Rule Types

This table lists the `Rule Types` and briefly describes them.

Table 312-1 Rule Types

Type	Description
RE$ATTRIBUTE_VALUE Type	Specifies the value of a variable attribute
RE$ATTRIBUTE_VALUE_LIST Type	Identifies a list of attribute values
RE$COLUMN_VALUE Type	Specifies the value of a table column
RE$COLUMN_VALUE_LIST Type	Identifies a list of column values
RE$NAME_ARRAY Type	Identifies a list of names

Table 312-1 (Cont.) Rule Types

Type	Description
RE$NV_ARRAY Type	Identifies a list of name-value pairs
RE$NV_LIST Type	Identifies an object containing a list of name-value pairs and methods that operate on this list. This object type is used to represent the event context and the action context for a rule
RE$NV_NODE Type	Identifies a name-value pair
RE$RULE_HIT Type	Specifies a rule found because of evaluation
RE$RULE_HIT_LIST Type	Identifies a list of rules found because of evaluation
RE$RULE_NAME Type	Allows users to pass a list of rule as argument.
RE$RULE_NAME_LIST Type	Identifies a list of Rule Name values
RE$TABLE_ALIAS Type	Provides the table corresponding to an alias used in a rule evaluation context
RE$TABLE_ALIAS_LIST Type	Identifies a list of table aliases used in a rule evaluation context
RE$TABLE_VALUE Type	Specifies the value of a table row using a `ROWID`
RE$TABLE_VALUE_LIST Type	Identifies a list of table values
RE$VALUE Type	Represents one datapoint value element
RE$VALUE_LIST Type	Defines a list of `RE$VALUE` type
RE$VARIABLE_TYPE Type	Provides the type of a variable used in a rule evaluation context
RE$VARIABLE_TYPE_LIST Type	Identifies a list of variables and their types used in a rule evaluation context
RE$VARIABLE_VALUE Type	Specifies the value of a variable
RE$VARIABLE_VALUE_LIST Type	Identifies a list of variable values

Rule types are used with the following Oracle-supplied PL/SQL packages:

- `DBMS_RULE`
- `DBMS_RULE_ADM`

You can use the `DBMS_RULE_ADM` package to create and administer rules, rule sets, and evaluation contexts, and you can use the `DBMS_RULE` package to evaluate rules.

When you use Replication (Oracle GoldenGate and XStream), rules determine which changes are captured by a capture process, which messages are propagated by a propagation, which messages are applied by an apply process, and which messages are dequeued by a messaging client. The following Replication (Oracle GoldenGate and XStream) packages use rules:

- `DBMS_APPLY_ADM`
- `DBMS_CAPTURE_ADM`
- `DBMS_PROPAGATION_ADM`

312.3.1 RE$ATTRIBUTE_VALUE Type

`RE$ATTRIBUTE_VALUE` type specifies the value of a variable attribute.

> **Note:**
>
> Enclose the variable name and attribute name in double quotation marks (") if the name contains special characters.

Syntax

```
TYPE SYS.RE$ATTRIBUTE_VALUE (
   variable_name    VARCHAR2(32),
   attribute_name   VARCHAR2(4000),
   attribute_value  ANYDATA);
```

Attributes

Table 312-2 RE$ATTRIBUTE_VALUE Attributes

Attribute	Description
`variable_name`	Specifies the variable used in a rule
`attribute_name`	Specifies the attribute name. The attribute name can be a multi-component name, such as `a1.b2.c3`.
`attribute_value`	Specifies the attribute value

312.3.2 RE$ATTRIBUTE_VALUE_LIST Type

The `RE$ATTRIBUTE_VALUE_LIST` type identifies a list of attribute values.

Syntax

```
TYPE SYS.RE$ATTRIBUTE_VALUE_LIST AS VARRAY(1024) OF SYS.RE$ATTRIBUTE_VALUE;
```

312.3.3 RE$COLUMN_VALUE Type

`RE$COLUMN_VALUE` type specifies the value of a table column.

> **Note:**
>
> Enclose the column name in double quotation marks (") if the name contains special characters.

Syntax

```
TYPE SYS.RE$COLUMN_VALUE (
   table_alias   VARCHAR2(32),
```

```
  column_name   VARCHAR2(4000),
  column_value  ANYDATA);
```

Attributes

Table 312-3 RE$COLUMN_VALUE Attributes

Attribute	Description
table_alias	Specifies the alias used for the table in a rule
column_name	Specifies the column name
column_value	Specifies the column value

312.3.4 RE$COLUMN_VALUE_LIST Type

RE$COLUMN_VALUE_LIST type identifies a list of column values.

Syntax

```
TYPE SYS.RE$COLUMN_VALUE_LIST AS VARRAY(1024) OF SYS.RE$COLUMN_VALUE;
```

312.3.5 RE$NAME_ARRAY Type

RE$NAME_ARRAY type identifies a list of names.

Syntax

```
TYPE SYS.RE$NAME_ARRAY AS VARRAY(1024) OF VARCHAR2(30);
```

312.3.6 RE$NV_ARRAY Type

RE$NV_ARRAY type identifies a list of name-value pairs.

Syntax

```
TYPE SYS.RE$NV_ARRAY AS VARRAY(1024) OF SYS.RE$NV_NODE;
```

312.3.7 RE$NV_LIST Type

identifies an object containing a list of name-value pairs and methods that operate on this list.

This object type is used to represent the event context for rule set evaluation and the action context for a rule.

Syntax

```
TYPE SYS.RE$NV_LIST AS OBJECT(
  actx_list  SYS.RE$NV_ARRAY);
```

Attributes

Table 312-4 RE$NV_LIST Attributes

Attribute	Description
actx_list	The list of name-value pairs

RE$NV_LIST Subprograms

This section describes the following member procedures and member functions of the SYS.RE$NV_LIST type:

- ADD_PAIR Member Procedure
- GET_ALL_NAMES Member Function
- GET_VALUE Member Function
- REMOVE_PAIR Member Procedure

ADD_PAIR Member Procedure

Adds a name-value pair to the list of name-value pairs.

> **Note:**
>
> Enclose the name in double quotation marks (") if the name contains special characters.

Syntax

```
MEMBER PROCEDURE ADD_PAIR(
   name   IN  VARCHAR2,
   value  IN  ANYDATA);
```

Parameters

Table 312-5 ADD_PAIR Procedure Parameters

Parameter	Description
name	The name in the name-value pair being added to the list. If the name exists in the list, then this procedure raises an error.
value	The value in the name-value pair being added to the list

GET_ALL_NAMES Member Function

Returns a list of all the names in the name-value pair list.

Syntax

```
MEMBER FUNCTION GET_ALL_NAMES()
RETURN SYS.RE$NAME_ARRAY;
```

GET_VALUE Member Function

Returns the value for the specified name in a name-value pair list.

> **Note:**
>
> Enclose the name in double quotation marks (") if the name contains special characters.

Syntax

```
MEMBER FUNCTION GET_VALUE(
   name  IN  VARCHAR2)
RETURN ANYDATA;
```

Parameters

Table 312-6 GET_VALUE Function Parameters

Parameter	Description
name	The name whose value to return

REMOVE_PAIR Member Procedure

Removes the name-value pair with the specified name from the name-value pair list.

> **Note:**
>
> Enclose the name in double quotation marks (") if the name contains special characters.

Syntax

```
MEMBER PROCEDURE REMOVE_PAIR(
   name  IN  VARCHAR2);
```

Parameters

Table 312-7 REMOVE_PAIR Procedure Parameters

Parameter	Description
name	The name of the pair to remove

312.3.8 RE$NV_NODE Type

`RE$COLUMN_VALUE` type identifies a name-value pair.

> **Note:**
>
> Enclose the name in double quotation marks (") if the name contains special characters.

Syntax

```
TYPE SYS.RE$NV_NODE (
   nvn_name   VARCHAR2(30),
   nvn_value  ANYDATA);
```

Attributes

Table 312-8 RE$NV_NODE Attributes

Attribute	Description
`nvn_name`	Specifies the name in the name-value pair
`nvn_value`	Specifies the value in the name-value pair

312.3.9 RE$RULE_HIT Type

The `RE$RULE_HIT` type specifies a rule found because of an evaluation.

> **See Also:**
>
> - CREATE_RULE Procedure
> - ALTER_RULE Procedure

Syntax

```
TYPE SYS.RE$RULE_HIT (
   rule_name           VARCHAR2(65),
   rule_action_context  RE$NV_LIST);
```

Attributes

Table 312-9 RE$RULE_HIT Attributes

Attribute	Description
`rule_name`	The rule name in the form *`schema_name.rule_name`*. For example, a rule named `employee_rule` in the `hr` schema is returned in the form `"hr"."employee_rule"`.

Table 312-9 (Cont.) RE$RULE_HIT Attributes

Attribute	Description
rule_action_context	The rule action context as specified in the CREATE_RULE or ALTER_RULE procedure of the DBMS_RULE_ADM package

312.3.10 RE$RULE_HIT_LIST Type

The RE$RULE_HIT_LIST type identifies a list of rules found because of an evaluation.

Syntax

```
TYPE SYS.RE$RULE_HIT_LIST AS VARRAY(1024) OF SYS.RE$RULE_HIT;
```

312.3.11 RE$RULE_NAME Type

The RE$RULE_HIT type is create to allow use to pass a list of rule as argument.

Syntax

```
CREATE OR REPLACE TYPE SYS.RE$RULE_NAME AS OBJECT(
  rule_name            IN       varchar2,
  evaluation_context   IN       varchar2);
```

Attributes

Table 312-10 RE$RULE_NAME Attributes

Attribute	Description
rule_name	The rule name in the form *SCHEMA_NAME.RULE_NAME*. For example, a rule named employee_rule in the hr schema is returned in the form "hr"."employee_rule".
Evaluation_context	The evaluation context as specified in the CREATE_RULE or ALTER_RULE procedure of the DBMS_RULE_ADM package

312.3.12 RE$RULE_NAME_LIST Type

RE$RULE_HIT_LIST type Identifies a list of Rule Name values.

Syntax

```
CREATE OR REPLACE TYPE SYS.RE$RULE_NAME_LIST AS VARRAY(1024) OF SYS.RE$RULE_NAME;
```

312.3.13 RE$TABLE_ALIAS Type

RE$TABLE_ALIAS type provides the table corresponding to an alias used in a rule evaluation context.

A specified table name must satisfy the schema object naming rules.

> **Note:**
>
> Enclose the table name in double quotation marks (") if the name contains special characters.

> **See Also:**
>
> *Oracle Database SQL Language Reference* for information about schema object naming rules

Syntax

```
TYPE SYS.RE$TABLE_ALIAS IS OBJECT(
   table_alias  VARCHAR2(32),
   table_name   VARCHAR2(194));
```

Attributes

Table 312-11 RE$TABLE_ALIAS Attributes

Attribute	Description
`table_alias`	The alias used for the table in a rule
`table_name`	The table name referred to by the alias. A synonym can be specified. The table name is resolved in the evaluation context schema. The format is one of the following: *`schema_name.table_name`* *`table_name`* For example, if the *`schema_name`* is `hr` and the *`table_name`* is `employees`, then enter the following: `hr.employees`

312.3.14 RE$TABLE_ALIAS_LIST Type

The `RE$TABLE_ALIAS_LIST` type identifies a list of table aliases used in a rule evaluation context.

Syntax

```
TYPE SYS.RE$TABLE_ALIAS_LIST AS VARRAY(1024) OF SYS.RE$TABLE_ALIAS;
```

312.3.15 RE$TABLE_VALUE Type

The `RE$TABLE_VALUE` type specifies the value of a table row using a `ROWID`.

Syntax

```
TYPE SYS.RE$TABLE_VALUE(
   table_alias  VARCHAR2(32),
   table_rowid  VARCHAR2(18));
```

Attributes

Table 312-12 RE$TABLE_VALUE Attributes

Attribute	Description
`table_alias`	Specifies the alias used for the table in a rule
`table_rowid`	Specifies the rowid for the table row

312.3.16 RE$TABLE_VALUE_LIST Type

The `RE$TABLE_VALUE_LIST` type identifies a list of table values.

> **Note:**
>
> Each table alias in the list in the list must be unique.

Syntax

```
TYPE SYS.RE$TABLE_VALUE_LIST AS VARRAY(1024) OF SYS.RE$TABLE_VALUE;
```

312.3.17 RE$VALUE Type

The `RE$VALUE` type represents one datapoint value element.

Syntax

```
TYPE SYS.RE$VALUE  (
     table_values          IN       SYS.RE$TABLE_VALUE_LIST:= NULL,
     column_values         IN       SYS.RE$COLUMN_VALUE_LIST:=NULL,
     variable_values       IN       SYS.RE$VARIABLE_VALUE_LIST:=NULL,
     attribute_values      IN       SYS.RE$ATTRIBUTE_VALUE_LIST:=NULL)
```

Attributes

Table 312-13 RE$VALUE Attributes

Attribute	Description
`table_values`	ROWID of table row for expression evaluation
`column_ values`	Values of columns referred in the expression

Table 312-13 (Cont.) RE$VALUE Attributes

Attribute	Description
variable_values	Values of variables referred in expressio
attribute_values	Values of attributes referred in expression

312.3.18 RE$VALUE_LIST Type

RE$TABLE_VALUE_LIST type is a varray for defining a list of RE$VALUE type. It represents a list of datapoint element.

Syntax

```
CREATE OR REPLACE TYPE SYS.RE$VALUE _LIST AS VARRAY(1024) OF SYS.RE$VALUE;
```

312.3.19 RE$VARIABLE_TYPE Type

The RE$VARIABLE_TYPE type provides the type of a variable used in a rule evaluation context. A specified variable name must satisfy the schema object naming rules.

> **Note:**
>
> Enclose the variable name in double quotation marks (") if the name contains special characters.

> **See Also:**
>
> *Oracle Database SQL Language Reference* for information about schema object naming rules

Syntax

```
TYPE SYS.RE$VARIABLE_TYPE (
   variable_name             VARCHAR2(32),
   variable_type             VARCHAR2(4000),
   variable_value_function   VARCHAR2(228),
   variable_method_function  VARCHAR2(228));
```

Attributes

Table 312-14 RE$VARIABLE_TYPE Attributes

Attribute	Description
variable_name	The variable name used in a rule

Table 312-14 (Cont.) RE$VARIABLE_TYPE Attributes

Attribute	Description
variable_type	The type that is resolved in the evaluation context schema. Any valid Oracle built-in datatype, user-defined type, or Oracle-supplied type can be specified. See the *Oracle Database SQL Language Reference* for more information about these types.
variable_value_function	A value function that can be specified for implicit variables. A synonym can be specified. The function name is resolved in the evaluation context schema. It is executed on behalf of the owner of a rule set using the evaluation context or containing a rule that uses the evaluation context. See the "Usage Notes" for more information.
variable_method_function	Specifies a value function, which can return the result of a method invocation. Specifying such a function can speed up evaluation, if there are many simple rules that invoke the method on the variable. The function can be a synonym or a remote function. The function name is resolved in the evaluation context schema. It is executed on behalf of the owner of a rule set using the evaluation context or containing a rule that uses the evaluation context. See the "Usage Notes" for more information.

Usage Notes

The functions for both the for the `variable_value_function` parameter and `variable_method_function` parameter have the following format:

```
schema_name.package_name.function_name@dblink
```

Any of the following parts of the format can be omitted: *schema_name*, *package_name*, and *@dblink*.

For example, if the *schema_name* is `hr`, the *package_name* is `var_pac`, the *function_name* is `func_value`, and the *dblink* is `dbs1.net`, then enter the following:

```
hr.var_pac.func_value@dbs1.net
```

The following sections describe the signature of the functions.

Signature for variable_value_function

The function must have the following signature:

```
FUNCTION variable_value_function_name(
  evaluation_context_schema  IN VARCHAR2,
  evaluation_context_name    IN VARCHAR2,
  variable_name              IN VARCHAR2,
  event_context              IN SYS.RE$NV_LIST )
RETURN SYS.RE$VARIABLE_VALUE;
```

Signature for variable_method_function

This function must have the following signature:

```
FUNCTION variable_method_function_name(
  evaluation_context_schema  IN VARCHAR2,
```

```
  evaluation_context_name    IN VARCHAR2,
  variable_value             IN SYS.RE$VARIABLE_VALUE,
  method_name                IN VARCHAR2,
  event_context              IN SYS.RE$NV_LIST)
RETURN SYS.RE$ATTRIBUTE_VALUE;
```

312.3.20 RE$VARIABLE_TYPE_LIST Type

The `RE$VARIABLE_TYPE_LIST` type identifies a list of variables and their types used in a rule evaluation context.

Syntax

```
TYPE SYS.RE$VARIABLE_TYPE_LIST AS VARRAY(1024) OF SYS.RE$VARIABLE_TYPE;
```

312.3.21 RE$VARIABLE_VALUE Type

The `RE$VARIABLE_VALUE` type specifies the value of a variable.

> **Note:**
>
> Enclose the variable name in double quotation marks (") if the name contains special characters.

Syntax

```
TYPE SYS.RE$VARIABLE_VALUE (
   variable_name  VARCHAR2(32),
   variable_data  ANYDATA);
```

Attributes

Table 312-15 RE$VARIABLE_VALUE Attributes

Attribute	Description
`variable_name`	Specifies the variable name used in a rule
`variable_data`	Specifies the data for the variable value

312.3.22 RE$VARIABLE_VALUE_LIST Type

The `RE$VARIABLE_VALUE_LIST` type identifies a list of variable values.

Syntax

```
TYPE SYS.RE$VARIABLE_VALUE_LIST AS VARRAY(1024) OF SYS.RE$VARIABLE_VALUE;
```

313
SODA Types

There are several SODA types: SODA_DOCUMENT_T, SODA_COLLECTION_T, SODA_OPERATION_T, and SODA_CURSOR_T. SODA_DOCUMENT_T and SODA_COLLECTION_T represent two primary abstractions provided by SODA: document and collections. SODA_OPERATION_T is used for specifying condition of operations on the collection. SODA_CURSOR_T is a cursor over results of read operations on the collection.

This chapter contains the following topics:

- SODA Types Overview
- SODA Types Security Model
- Summary of SODA Types

313.1 SODA Types Overview

There are several SODA types: SODA_DOCUMENT_T, SODA_COLLECTION_T, SODA_OPERATION_T, and SODA_CURSOR_T. SODA_DOCUMENT_T and SODA_COLLECTION_T represent two primary abstractions provided by SODA: document and collections. SODA_OPERATION_T is used for specifying condition of operations on the collection. SODA_CURSOR_T is a cursor over results of read operations on the collection.

> **See Also:**
>
> - *Oracle Database SODA for PL/SQL Developer's Guide*
> - DBMS_SODA

313.2 SODA Types Security Model

The SODA Types are available to users with the SODA_APP role.

All SODA types are SYS types. PUBLIC is granted EXECUTE privilege on the SODA types described in this chapter:

- TYPE SODA_Collection_T
- TYPE SODA_Document_T
- TYPE SODA_Operation_T
- TYPE SODA_Cursor_T

313.3 Summary of SODA Types

This chapter lists the SODA types and describes them.

Table 313-1 SODA Types

Type	Description
SODA_Collection_T Type	This `SODA` type represents a `SODA` collection. This type is not persistable.
SODA_Document_T Type	This `SODA` type represents a document with content, usually in `JSON` format. This type is not persistable.
SODA_OPERATION_T Type	This `SODA` type performs read/write operations, such as document finds with filtering and pagination, removes, and replaces on a `SODA` collection. This type is not persistable.
SODA_CURSOR_T Type	This `SODA` type represents the result set of documents. This type is not persistable.

313.3.1 SODA_Collection_T Type

This `SODA` type represents a `SODA` collection. A reference of `SODA` collection can only be obtained by either calling `DBMS_SODA.CREATE_COLLECTION()` or `DBMS_SODA.OPEN_COLLECTION()`.

Table 313-2 SODA_Collection_T Type Subprograms

Subprogram	Description
CREATE_INDEX Function	Creates an index using an index specification expressed in JSON. Three types of specifications are supported. Each specifying a different type of index: for B-tree, JSON search with Data Guide, and Spatial.
CREATE_VIEW_FROM_DG Function	Creates a view with relational columns, using scalar `JSON` fields as specified in the data guide.
DROP_INDEX Function	Drops the named index.
FIND Function	Returns the `SODA_OPERATION_T` object. This is the only way to get the reference of `SODA_Operation_T` as there is no constructor.
FIND_ONE Function	Fetches the document matching the key.
GET_DATA_GUIDE Function	Returns the JSON data guide as a CLOB.
GET_INDEX Function	This function returns the specification for the supplied index created on the collection.
GET_METADATA Function	Returns the metadata of the collection in `JSON` format.
GET_NAME Function	Returns the name of the collection.
INSERT_ONE Function	Inserts a document into the collection.

Table 313-2 (Cont.) SODA_Collection_T Type Subprograms

Subprogram	Description
INSERT_ONE_AND_GET Function	Inserts a document into the collection and returns a result document with all components except for content.
LIST_INDEXES Function	This function returns the specifications for all the indexes created on the collection.
REMOVE_ONE Function	Removes the document matching the key.
REPLACE_ONE Function	Replaces the content and (optionally) the media type of the document matching the key.
REPLACE_ONE_AND_GET Function	Replaces the content and (optionally) the media type of the document matching the key and returns a result document with all components (except content).
SAVE Function	Saves a document into the collection.
SAVE_AND_GET Function	Saves a document into the collection.
TRUNCATE Function	Deletes all documents in the collection.

313.3.1.1 CREATE_INDEX Function

This function creates an index using an index specification expressed in JSON. Three types of specifications are supported. Each specifying a different type of index: for B-tree, JSON search with Data Guide, and Spatial.

Syntax

```
CREATE_INDEX (
     specification VARCHAR2)
  RETURN NUMBER;
```

Parameters

Table 313-3 CREATE_INDEX Function Parameters

Parameter	Description
`specification`	The index specification.

Example 313-1 Return Values

The function returns:

- `1`—if the index was successfully created
- `0`—if the index was not created

Exceptions

`Error`—If an error occurs creating the index.

> **See Also:**
>
> For more information about SODA Index specifications, see:
>
> - Overview of SODA Indexing
> - SODA Index Specifications (Reference)

313.3.1.2 CREATE_VIEW_FROM_DG Function

This function creates a view with relational columns, using scalar JSON fields as specified in the data guide. A data guide enabled JSON search index is not required for this function; the data guide is passed to the function. An error is thrown if the data guide passed to the function is invalid.

This procedure is available only for Autonomous Database starting 19c release.

Syntax

```
CREATE_VIEW_FROM_DG (
   data_guide               IN   CLOB,
   view_name                IN   VARCHAR2,
   materialize              IN   BOOLEAN DEFAULT FALSE,
   mv_refresh_mode          IN   NUMBER DEFAULT 1,
   path                     IN   VARCHAR2 DEFAULT '$',
   resolve_name_conflicts   IN   BOOLEAN   DEFAULT FALSE,
   col_name_prefix          IN   VARCHAR2 DEFAULT NULL,
   mixed_case_columns       IN   BOOLEAN DEFAULT FALSE)
  RETURN NUMBER;
```

Parameters

Table 313-4 CREATE_VIEW_FROM_DG Function Parameters

Parameter	Description
data_guide	Data guide of the collection.
view_name	Name of the view to be created.
materialize	A boolean value to indicate if the view should be materialized or not. The default value is FALSE.
mv_refresh_mode	The materialized view refresh mode. Possible values are: • DBMS_SODA.MV_REFRESH_ON_STATEMENT (default) • DBMS_SODA.MV_REFRESH_ON_COMMIT • DBMS_SODA.MV_REFRESH_ON_DEMAND

Table 313-4 (Cont.) CREATE_VIEW_FROM_DG Function Parameters

Parameter	Description
`path`	The path of the JSON field to be expanded. It uses JSON path expression syntax. For example: • `$` will create a view starting from the JSON document root • `$.purchaseOrder` will create a view starting from purchaseOrder. It expands the children or descendants under `purchaseOrder`, and create view columns for every scalar value. The default value is `$`.
`resolve_name_conflicts`	By default, if there are conflicts among `o:preferred_column_name`, an error is raised. If you set this parameter to `TRUE`, the procedure automatically resolves the view column name conflicts by appending a sequence number. The default value is `FALSE`.
`col_name_prefix`	By default, the view column name is the same as the JSON field name. This parameter allows you to provide a prefix to prepend to the view column names. The default value is `NULL`.
`mixed_case_columns`	By default, the view column names are case sensitive. This parameter allows you to change the behavior to case insensitive. The default value is `FALSE`.

Return Values

The function returns:

- `1`—if the procedure is successfully completed
- `0`—if the procedure could not be successfully completed

Exceptions

`Error`—If an error occurs if the function was unable to create a view.

> **See Also:**
>
> For more info on the JSON data guide, see JSON Data Guide

313.3.1.3 DROP_INDEX Function

This function drops the named index.

Syntax

```
DROP_INDEX (
     index_Name IN VARCHAR2,
```

```
    force      IN BOOLEAN DEFAULT FALSE)
  RETURN NUMBER;
```

Example 313-2 Parameters

Table 313-5 DROP_INDEX Function Parameters

Parameter	Description
index_Name	The name of the index.
force	The force parameter can be TRUE or FALSE. Should only be set to TRUE for dropping a JSON search index or spatial index (not B-tree index). For more information, see DROP INDEX

Return Values

The function returns:

- 1—If the index was successfully dropped
- 0—If the index was not dropped. For example, if there was no existing index with the specified name.

Exceptions

Error—if an error occurs while dropping the index.

313.3.1.4 FIND Function

This function returns the operation type for the collection. The operation type allows building and executing various read/write operations. This is the only way to get the reference of SODA_Operation_T as there is no constructor.

Syntax

```
FIND ()
 RETURN SODA_Operation_T;
```

Return Values

This function returns SODA_OPERATION_T object.

Exceptions

This function does not throw any exception.

313.3.1.5 FIND_ONE Function

This function fetches the document matching the given key.

Syntax

```
FIND_ONE (
     key          IN VARCHAR2)
 RETURN SODA_Document_T;
```

Parameters

Table 313-6 FIND_ONE Function Parameters

Parameter	Description
`key`	The key of the document to be fetched.

Return Values

This function returns the document that matches the key. Returns `NULL` if no match is found.

Exceptions

`Error`—If an error occurs while finding the document.

313.3.1.6 GET_DATA_GUIDE Function

This function fetches the JSON data guide as a CLOB. The JSON data guide is essentially inferred schema for the JSON documents in the collection. In order to be able to return the JSON data guide, a collection must have a JSON Search Index defined on it, with the data guide enabled.

Syntax

```
GET_DATA_GUIDE ()
 RETURN CLOB;
```

Return Values

The function returns the JSON data guide as a CLOB.

> **See Also:**
>
> For more info on the JSON data guide, see JSON Data Guide

Exceptions

`Error`—If an error occurs while fetching the data guide.

313.3.1.7 GET_INDEX Function

This function returns the specification for the supplied index created on the collection.

Syntax

```
GET_INDEX (
     index_name           IN VARCHAR2,
     schema_name          IN VARCHAR2 DEFAULT NULL)
  RETURN VARCHAR2;
```

Parameters

Table 313-7 GET_INDEX Function Parameters

Parameter	Description
index_name	The name of the index to be described.
schema_name	Name of the schema containing the index. This parameter is optional. If this parameter is not set, the method will look for the specified index in the schema from which this method is called.

Return Values

The function returns the index specification in JSON format.

Exceptions

Error—If an error occurs while returning the index specification.

313.3.1.8 GET_METADATA Function

This function returns the metadata of the collection in JSON format.

Syntax

```
GET_METADATA ()
  RETURN VARCHAR2;
```

Return Values

This function returns the metadata of the collection in JSON format.

313.3.1.9 GET_NAME Function

This function returns the name of the collection.

Syntax

```
GET_NAME ()
 RETURN NVARCHAR2;
```

Return Values

This function returns the name of the collection.

313.3.1.10 INSERT_ONE Function

This function inserts a document into the collection.

Syntax

```
INSERT_ONE (
     document      IN SODA_Document_T)
 RETURN NUMBER;
```

Parameters

Table 313-8 INSERT_ONE Function Parameters

Parameter	Description
`document`	The input document.

Return Values

The function returns a number– `1` if the doc was inserted successfully, `0` otherwise.

Exceptions

`Error`—If an error occurs while inserting the document into the collection.

313.3.1.11 INSERT_ONE_AND_GET Function

This function inserts a document into the collection.

Syntax

```
INSERT_ONE_AND_GET (
     document      IN SODA_Document_T,
     hint          IN VARCHAR2 DEFAULT NULL)
 RETURN SODA_Document_T;
```

Parameters

Table 313-9 INSERT_ONE_AND_GET Function Parameters

Parameter	Description
`document`	The input document.
`hint`	A hint string in Oracle SQL format, without the enclosing `/*+` and `*/`. This parameter is optional.

Return Values

The function returns the result document containing all document components supported by the given collection, with the exception of content.

Exceptions

`Error`—If an error occurs while inserting the document into the collection.

313.3.1.12 LIST_INDEXES Function

This function returns the specifications for all the indexes created on the collection.

Syntax

```
LIST_INDEXES ( )
  RETURN SODA_Index_List_T;
```

Return Values

The function returns index specifications in `JSON` format as an instance of `SODA_Index_List_T`.

Exceptions

`Error`—If an error occurs while returning the index specification.

313.3.1.13 REMOVE_ONE Function

This function removes the document matching the given key.

Syntax

```
REMOVE_ONE (
     key        IN VARCHAR2)
 RETURN NUMBER;
```

Parameters

Table 313-10 REMOVE_ONE Function Parameters

Parameter	Description
`key`	The key of the document.

Return Values

This function returns the following values:

- `1`–If the document was successfully removed.
- `0`–If the document with the specified key was not found.

Exceptions

`Error`—If an error occurs while deleting the document from the collection.

313.3.1.14 REPLACE_ONE Function

This function updates the existing document with a new content and media type using the key. Any components set in `document` with the exception of content and media type are not used during the replace. They are ignored.

Syntax

```
REPLACE_ONE (
     key            IN VARCHAR2,
     document       IN SODA_Document_T)
 RETURN NUMBER;
```

Parameters

Table 313-11 REPLACE_ONE Parameters

Parameter	Description
`key`	The key of the document.
`document`	The document with the new content and media type to replace the old one.

Return Values

This function returns a number—`1` if the document was replaced, `0` otherwise.

Exceptions

`Error`—If an error occurs while replacing the document in the collection.

313.3.1.15 REPLACE_ONE_AND_GET Function

This function updates the existing document with a new content and media type using the key. Any components set in `document` with the exception of content and media type are not used during the replace. They are ignored.

Syntax

```
REPLACE_ONE_AND_GET (
     key            IN VARCHAR2,
     document       IN SODA_Document_T)
 RETURN SODA_Document_T;
```

Parameters

Table 313-12 REPLACE_ONE_AND_GET Function Parameters

Parameter	Description
`key`	The key of the document.
`document`	The document with the new content and media type to replace the old one.

Return Values

The function returns the result document containing all document components supported by the given collection, with the exception of content. Last-modified and version components, if supported by the given collection, will be updated with new values. If no document in the collection had the supplied key, `NULL` is returned instead of the result document.

Exceptions

`Error`—If an error occurs while replacing the document in the collection.

313.3.1.16 SAVE Function

This function saves a document into the collection. This function is equivalent to the `INSERT_ONE(document)` function except that if client-assigned keys are used, and the document with the specified key already exists in the collection, it will be replaced with the input document. The key is automatically created, unless this collection is configured with client-assigned keys and the key is provided in the input document.

Syntax

```
SAVE (
     document IN  SODA_Document_T)
  RETURN NUMBER;
```

Parameters

Table 313-13 SAVE Parameters

Parameter	Description
`document`	The input document. This cannot be null.

Return Values

The function returns a number- `1` if the function successfully completed, `0` otherwise.

Exceptions

`Error`—If an error occurs while saving the document.

313.3.1.17 SAVE_AND_GET Function

This function saves a document into the collection. This method is equivalent to `INSERT_ONE_AND_GET(document)` except that if client-assigned keys are used, and the document with the specified key already exists in the collection, it will be replaced with the input document. The key will be automatically created, unless this collection is configured with client-assigned keys and the key is provided in the input document.

Syntax

```
SAVE_AND_GET (
     document      IN SODA_Document_T,
     hint          IN VARCHAR2 DEFAULT NULL)
  RETURN SODA_Document_T;
```

Parameters

Table 313-14 SAVE_AND_GET Function Parameters

Parameter	Description
`document`	The input document. This cannot be null.
`hint`	A hint string in Oracle SQL format, without the enclosing `/*+` and `*/`. This parameter is optional.

Return Values

The function returns the result document containing all document components supported by the given collection, with the exception of content.

Exceptions

`Error`—If an error occurs while saving or getting the document.

313.3.1.18 TRUNCATE Function

This function deletes all documents in the collection.

Syntax

```
TRUNCATE ( )
  RETURN Number;
```

Return Values

The function returns:

- `1`—if the function is successfully completed
- `0`—if the function could not be successfully completed

Exceptions

`Error`—if an error occurs while deleting the documents in the collection.

313.3.2 SODA_Document_T Type

This `SODA` type represents a document with content, that is usually in `JSON` format.

This type is not persistable `pl/sql` type. However, `SODA` is a system that basically provides persistence — it has read and write operations. So you do not persist `SODA_DOCUMENT_T` directly, but you pass it to a write operation (like `insert` or `replace`), which is defined on `SODA_COLLECTION_T`, in order to write the document content and other components to the database.

A document has the following components:

- key
- content
- created-on timestamp
- last-modified timestamp
- version
- media type

Table 313-15 SODA_Document_T Type Subprograms

Subprogram	Description
GET_BLOB Function	Fetches the BLOB content of a `BLOB`-based document.

Table 313-15 (Cont.) SODA_Document_T Type Subprograms

Subprogram	Description
GET_CLOB Function	Fetches the CLOB content of a CLOB-based document.
GET_CREATED_ON Function	Fetches the created-on timestamp in VARCHAR2.
GET_DATA_TYPE Function	Fetches the SQL datatype of the document content with which it was created.
GET_JSON Function	Fetches the JSON content of a JSON-based document.
GET_KEY Function	Fetches the document key in VARCHAR2.
GET_LAST_MODIFIED Function	Fetches the last modified timestamp in VARCHAR2.
GET_MEDIA_TYPE Function	Fetches the media type of the document content in VARCHAR2.
GET_VARCHAR2 Function	Fetches the VARCHAR2 content of a VARCHAR2-based document.
GET_VERSION Function	Fetches the version of the document in VARCHAR2.
SODA_Document_T Function	There are three different SODA_DOCUMENT_T constructor functions. Each constructor function instantiates a document object using key, content, and media type.

313.3.2.1 GET_BLOB Function

This functions fetches the BLOB content of the document. It assumes that the document was constructed with BLOB content, or was returned from a collection with BLOB content. Otherwise, an error is returned.

Syntax

```
GET_BLOB ()
 RETURN BLOB;
```

Return Values

This function returns the BLOB content of a document.

Exceptions

SODA Error: If the document was initially not created with BLOB content.

313.3.2.2 GET_CLOB Function

The function fetches CLOB content of the document. It assumes that the document was constructed with CLOB content, or was returned from a collection with CLOB content. Otherwise, an error is returned.

Syntax

```
GET_CLOB ()
 RETURN CLOB;
```

Return Values

This function returns the `CLOB` content of a document.

Exceptions

`SODA Error`: If the document was initially not created with `CLOB` content.

313.3.2.3 GET_CREATED_ON Function

This function fetches the created-on timestamp. The timestamp string is in `ISO-8601` format, in particular this form: `YYYY-MM-DDThh:mm:ss.ssssssZ` format. As indicated by the `Z` at the end, timestamps are returned in UTC (`Z` indicates zero UTC offset).

Syntax

```
GET_CREATED_ON ()
 RETURN VARCHAR2;
```

Return Values

This function returns the created-on timestamp.

313.3.2.4 GET_DATA_TYPE Function

This function fetches the SQL datatype of the document content with which it was created.

Syntax

```
GET_DATA_TYPE ()
 RETURN PLS_INTEGER;
```

Return Values

Table 313-16 GET_DATA_TYPE Return Values

Constant	Value	Description
DOC_VARCHAR2 CONSTANT PLS_INTEGER	1	VARCHAR2
DOC_BLOB CONSTANT PLS_INTEGER	2	BLOB
DOC_CLOB CONSTANT PLS_INTEGER	3	CLOB
DOC_JSON CONSTANT PLS_INTEGER	4	JSON

313.3.2.5 GET_JSON Function

This functions fetches the `JSON` content of the document. It assumes that the document was constructed with `JSON` type content, or was returned from a collection with `JSON` type content. Otherwise, an error is returned.

Syntax

```
GET_JSON ()
 RETURN JSON;
```

Return Values

This function returns the `JSON` content of a document.

Exceptions

`SODA Error`: If the document was initially not created with `JSON` content.

313.3.2.6 GET_KEY Function

This function fetches the document key.

Syntax

```
GET_KEY ()
 RETURN VARCHAR2;
```

Return Values

This function returns the document key.

313.3.2.7 GET_LAST_MODIFIED Function

This function fetches the last modified timestamp. The timestamp string is in `ISO-8601` format, in particular this form: `YYYY-MM-DDThh:mm:ss.ssssssZ` format. As indicated by the `Z` at the end, timestamps are returned in UTC (`Z` indicates zero UTC offset).

Syntax

```
GET_LAST_MODIFIED ()
 RETURN VARCHAR2;
```

Return Values

This function returns the last modified timestamp.

313.3.2.8 GET_MEDIA_TYPE Function

This function fetches the media type of the document content.

Syntax

```
GET_MEDIA_TYPE ()
 RETURN VARCHAR2;
```

Return Values

This function returns the media type of the document content. `application/JSON` is the media type for `JSON` documents (default).

313.3.2.9 GET_VARCHAR2 Function

This function fetches the VARCHAR2 content of the document. It assumes that the document was constructed with VARCHAR2 content, or was returned from a collection with VARCHAR2 content. Otherwise, an error is returned.

Syntax

```
GET_VARCHAR2 ()
 RETURN VARCHAR2;
```

Return Values

This function returns the VARCHAR2 content of a document.

Exceptions

SODA Error: If the document was initially not created with VARCHAR2 content.

313.3.2.10 GET_VERSION Function

This function fetches the version of the document.

Syntax

```
GET_VERSION ()
 RETURN VARCHAR2;
```

Return Values

This function returns the version of the document.

313.3.2.11 SODA_Document_T Function

This function instantiates a document object using key, content, and media type. There are three different SODA_DOCUMENT_T constructor functions. The second parameter (<v|b|c>_Content) is different in each constructor. It is VARCHAR2 in the first variant, BLOB in the second, and CLOB in the third.

Key and media type are optional parameters (will be defaulted to NULL). All three parameters can be set to NULL. If media_Type is set to NULL, it will be defaulted to application/json.

Use key and j_Content with the constructor to instantiate a document object using key and content. Media type parameter is not present in this constructor as the data is of JSON type. Therefore, media type is understood to be application/json.

Syntax

```
SODA_DOCUMENT_T (
     key           IN VARCHAR2 DEFAULT NULL,
     v_Content     IN VARCHAR2,
     media_Type    IN VARCHAR2 DEFAULT NULL)
 RETURN SODA_Document_T;

SODA_DOCUMENT_T (
     key           IN VARCHAR2 DEFAULT NULL,
     b_Content     IN BLOB,
     media_Type    IN VARCHAR2 DEFAULT NULL)
```

```
 RETURN SODA_Document_T;

SODA_DOCUMENT_T (
     key           IN VARCHAR2 DEFAULT NULL,
     c_Content     IN CLOB,
     media_Type    IN VARCHAR2 DEFAULT NULL)
 RETURN SODA_Document_T;

SODA_DOCUMENT_T (
     key           IN VARCHAR2 DEFAULT NULL,
     j_Content     IN JSON)
 RETURN SODA_Document_T;
```

Parameters

Table 313-17 SODA_Document_T Parameters

Parameter	Description
`key`	The key of the document.
`v_Content`	The content of the document in `VARCHAR2` format.
`b_Content`	The content of the document in `BLOB` format.
`c_Content`	The content of the document in `CLOB` format.
`j_Content`	The content of the document as `JSON` type instance.
`media_Type`	The media type of the document. The media type could be `application/json` for `JSON` documents.

> **Note:**
>
> `v_Content`, `b_Content`, and `c_Content` are not all parameters of a single `SODA_DOCUMENT_T` constructor function. Each one corresponds to a particular variant of the constructor function as shown in the `Syntax` section.

Return Values

This function returns a document of type `SODA_Document_T`.

313.3.3 SODA_OPERATION_T Type

This `SODA` type is used to perform read/write operations, such as document finds with filtering and pagination, removes, and replaces on a `SODA` collection.

Table 313-18 SODA_OPERATION_T Type Subprograms

Subprogram	Description
ACQUIRE_LOCK Function	This function ensures that the document(s) affected by a `read` operation are locked for update (which is equivalent to SQL '`select for update`').

Table 313-18 (Cont.) SODA_OPERATION_T Type Subprograms

Subprogram	Description
AS_OF_SCN Function	This function sets the SCN value for the operation.
AS_OF_TIMESTAMP Function	This function sets the the timestamp value for the operation.
COUNT Function	Returns a count of the number of documents in the collection that match the criteria. If skip(...) or limit(...) were chained together with this count(), an exception is raised.
FILTER Function	Sets the filter (also known as QBE or query-by-example) criteria on the operation. Returns the same `SODA_OPERATION_T` object so that further criteria can be chained together if required.
GET_CURSOR Function	Returns a `SODA_CURSOR_T` object that can be used to iterate over the documents that match the criteria.
GET_DATA_GUIDE Function	This function gets the data guide.
GET_ONE Function	Returns a single `SODA_DOCUMENT_T` object that matches the criteria. Note that, if multiple documents match the criteria, only the first document is returned.
HINT Function	This function sets the hint attribute of the operation.
KEY Function	Specifies that the document with the specified key should be returned. This causes any previous calls made to this function and `KEYS`(...), when they appear in the same chain, to be ignored. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.
KEYS Function	Specifies that documents that match the keys supplied to this function should be returned. This causes any previous calls made to this function and `KEY`(...), when they appear in the same chain, to be ignored. Returns the same `SODA_OPERATION_T` object, so that further operation criteria can be chained together, if needed.
LIMIT Function	Sets a limit on the specified number of documents the operation should return. This setting is only usable for read operations such as `GET_CURSOR`. For write operations, any value set using this method is ignored. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.
REMOVE Function	Removes all of the documents in the collection that match the criteria. Returns the number of documents that was removed.

Table 313-18 (Cont.) SODA_OPERATION_T Type Subprograms

Subprogram	Description
REPLACE_ONE Function	Replaces a single document in the collection with the specified document. Returns a number that indicates if the document was replaced or not. Currently, before calling this function, you must call the function `KEY(...)` to uniquely identify the document being replaced. Any components set in the input document with the exception of content and media type are not used during the replace. They are ignored.
REPLACE_ONE_AND_GET Function	Replaces a single document in the collection with the specified document. Returns a result document if the document was replaced, `NULL` otherwise. Currently, before calling this function, you must call the function `KEY(...)` to uniquely identify the document being replaced. This function is similar to `REPLACE_ONE`. The only difference is that `REPLACE_ONE_AND_GET` also returns the result document with updated components, such as version and last-modified timestamp. The result document does not contain the content component. Any components set in the input document with the exception of content and media type are not used during the replace. They are ignored.
SAMPLE Function	Sets the sampling parameters to be used for the operation.
SKIP Function	Sets the number of documents that match the operation criteria that will be skipped from the operation result. This setting is only usable for read operations such as `GET_CURSOR`. For write operations, any value set using this method is ignored. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.
VERSION Function	Specifies that only documents with the supplied version should be returned. Typically, this is chained together with `KEY(...)` to implement optimistic locking for write operations such as `REMOVE` and `REPLACE`. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.

313.3.3.1 ACQUIRE_LOCK Function

This function ensures that the the document(s) affected by a read operation are locked for update (which is equivalent to SQL '`select for update`'). An operation involving `LOCK()` would be followed by another operation that updates the document in the collection. For example, via

`REPLACE` or `REMOVE` functions. The lock will prevent other transactions from modifying the document in the meantime.

The next commit or rollback, performed after the operation involving this `ACQUIRE_LOCK()` function, will unlock the document, i.e. commit or rollback the transaction which is holding the lock on this document.

This function should only be used in conjunction with `read` operations, other than `COUNT()` and `GET_DATA_GUIDE()` functions.

Specifying it in conjunction with `SKIP` and `LIMIT` functions is also not allowed (an error is thrown).

If specified in conjunction with a `write` operation, such as, `REPLACE` or `REMOVE`, it's simply ignored and has no effect.

Syntax

```
ACQUIRE_LOCK ( )
 RETURN SODA_Operation_T;
```

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.2 AS_OF_SCN Function

This function sets the SCN value for the operation.

Syntax

```
AS_OF_SCN (
     scn          IN NUMBER)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-19 AS_OF_SCN Function Parameters

Parameter	Description
`scn`	The input value for SCN. This cannot be `NULL`.

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.3 AS_OF_TIMESTAMP Function

This function sets the the timestamp value for the operation.

Syntax

```
AS_OF_TIMESTAMP (
     timestamp          IN VARCHAR2)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-20 AS_OF_TIMESTAMP Function Parameters

Parameter	Description
timestamp	The input timestamp. This cannot be NULL.

Return Values

This function returns the same SODA_OPERATION_T object it was invoked on.

313.3.3.4 COUNT Function

This function returns a count of the number of documents in the collection that match the criteria. If skip(...) or limit(...) were chained together with this count(), an exception is raised.

Syntax

```
COUNT ()
 RETURN NUMBER;
```

Return Values

This function returns the number of documents matching the criteria specified in the operation.

Exceptions

Error—If an error occurs while finding the count.

313.3.3.5 FILTER Function

Sets the filter (also known as QBE or query-by-example) criteria on the operation. Returns the same SODA_OPERATION_T object so that further criteria can be attached if needed.

Syntax

```
FILTER (
     qbe          IN VARCHAR2)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-21 FILTER Function Parameters

Parameter	Description
qbe	The string representing the query by example.

Return Values

This function returns the same SODA_OPERATION_T object it was invoked on.

> **See Also:**
>
> - Overview of SODA Filter Specifications (QBEs)
> - SODA Filter Specifications (Reference)

313.3.3.6 GET_CURSOR Function

Returns a `SODA_CURSOR_T` object that can be used to iterate over the documents that match the criteria.

Syntax

(Optional) Enter syntax information here.

```
GET_CURSOR ()
 RETURN SODA_Cursor_T;
```

Return Values

This function returns a `SODA_CURSOR_T` object that can be used to iterate over the documents that match the read operation criteria.

Exceptions

`SODA Error`:If an error occurs while fetching the cursor.

313.3.3.7 GET_DATA_GUIDE Function

This function gets the data guide.

Syntax

```
GET_DATA_GUIDE (
  format    IN PLS_INTEGER DEFAULT 1,
  flag      IN PLS_INTEGER DEFAULT 0)
 RETURN CLOB;
```

Parameters

Table 313-22 GET_DATA_GUIDE Function Parameters

Parameter	Description
`format`	The format of the data guide. This parameter can have one of the following values: • `DBMS_SODA.DATAGUIDE_FORMAT_HIERARCHICAL CONSTANT PLS_INTEGER := 1;` • `DBMS_SODA.DATAGUIDE_FORMAT_FLAT CONSTANT PLS_INTEGER := 2;` The default value is `1`.

Table 313-22 (Cont.) GET_DATA_GUIDE Function Parameters

Parameter	Description
`flag`	This parameter can have one of the following values: • `DBMS_SODA.DATAGUIDE_PRETTY CONSTANT PLS_INTEGER := 1;` • `DBMS_SODA.DATAGUIDE_GEOJSON CONSTANT PLS_INTEGER := 2;` • `DBMS_SODA.DATAGUIDE_GATHER_STATS CONSTANT PLS_INTEGER := 4;` The default value is `0`.

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.8 GET_ONE Function

Returns a single `SODA_DOCUMENT_T` object that matches the criteria. Note that, if multiple documents match the criteria, only the first document is returned.

Syntax

```
GET_ONE ()
 RETURN SODA_Document_T;
```

Return Values

The first matching document.

Exceptions

`Error`—If an error occurs while fetching the document.

313.3.3.9 HINT Function

This function sets the hint attribute of the operation.

Syntax

```
HINT (
     hint        IN VARCHAR2)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-23 HINT Function Parameters

Parameter	Description
`hint`	A hint string in Oracle SQL format, with out the enclosing `/*+` and `*/`.

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.10 KEY Function

Specifies that the document with the specified key should be returned. This causes any previous calls made to this function and `KEYS`(...), when they appear in the same chain, to be ignored. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.

Syntax

```
KEY (
     key          IN VARCHAR2)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-24 KEY Function Parameters

Parameter	Description
`key`	The key to be used for the operations.

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.11 KEYS Function

Specifies that documents that match the keys supplied to this function should be returned. This causes any previous calls made to this function and `key`(...), when they appear in the same chain, to be ignored. Returns the same `SODA_OPERATION_T` object, so that further operation criteria can be chained together, if needed.

Syntax

```
KEYS (
     key_List  IN SODA_Key_List_T)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-25 KEYS Function Parameters

Parameter	Description
`key_List`	The parameter is a `SODA_Key_List_T` which is a list of `VARCHAR2` values representing keys. Assuming `key_list` is a variable of type `SODA_Key_List_T`, it can be initialized as follows: `key_list := SODA_Key_List_T('key1', 'key2', 'key3', etc);`

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.12 LIMIT Function

This function sets a limit on the specified number of documents the operation should return. This setting is only usable for read operations such as `GET_CURSOR`. For write operations, any value set using this method is ignored. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.

Syntax

```
LIMIT (
     limit          IN NUMBER)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-26 LIMIT Function Parameters

Parameter	Description
`limit`	A limit on the number of results returned by read operations.

Return Values

This function returns the same `SODA_OPERATION_T` object it was invoked on.

313.3.3.13 REMOVE Function

This function removes all of the documents in the collection that match the criteria. Returns the number of documents that was removed.

Syntax

```
REMOVE ()
 RETURN NUMBER;
```

Return Values

This function returns the number of matching documents that were removed in the operation.

Exceptions

`Error`—If an error occurs while removing the documents.

313.3.3.14 REPLACE_ONE Function

This function replaces a single document in the collection with the specified document. Returns a number that indicates if the document was replaced or not. Currently, before calling this function, you must call the function `KEY`(...) to uniquely identify the document being replaced. Any components set in the input document with the exception of content and media type are not used during the replace. They are ignored.

Syntax

```
REPLACE_ONE (
      document IN SODA_Document_T)
 RETURN NUMBER;
```

Parameters

Table 313-27 REPLACE_ONE Function Parameters

Parameter	Description
`document`	The document object with the new content and media type to be used for replacement.

Return Values

This function returns a number—`1` if the document was replaced, `0` otherwise.

Exceptions

`Error`—If an error occurs while updating the collection.

313.3.3.15 REPLACE_ONE_AND_GET Function

Replaces a single document in the collection with the specified document. Returns a result document if the document was replaced, `NULL` otherwise. Currently, before calling this function, you must call the function `KEY`(...) to uniquely identify the document being replaced. This function is similar to `REPLACE_ONE`. The only difference is that `REPLACE_ONE_AND_GET` also returns the result document with updated components, such as version and last-modified timestamp. The result document does not contain the content component. Any components set in the input document with the exception of content and media type are not used during the replace. They are ignored.

Syntax

```
REPLACE_ONE_AND_GET (
     document         IN SODA_Document_T)
 RETURN SODA_Document_T;
```

Parameters

Table 313-28 REPLACE_ONE_AND_GET Function Parameters

Parameter	Description
`document`	The document object with the new content and media type to be used for replacement.

Return Values

The function returns the result document containing all document components supported by the given collection, with the exception of content. Last-modified and version components, if supported by the given collection, will be updated with new values. If no document in the collection had the supplied key, `NULL` is returned instead of the result document.

Exceptions

`Error`—If an error occurs while updating the collection

313.3.3.16 SAMPLE Function

This function sets the sampling parameters to be used for the operation.

Syntax

```
SAMPLE (
     pct     IN   DOUBLE,
     seed    IN   NUMBER      DEFAULT NULL,
     method  IN   PLS_INTEFER DEFAULT DBMS_SODA.SAMPLE_ROW)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-29 SAMPLE Function Parameters

Parameter	Description
`pct`	The percentage of the total documents or block count to be included in the sample. The value must be in the range `.000001` to, but not including, `100`. This percentage indicates the probability of each row or each cluster of rows in the case of block sampling, being selected as part of the sample. It does not mean that the database will retrieve exactly the percentage of documents in the collection.
`seed`	Specify this attribute to instruct the database to attempt to return the same sample from one execution to the next. The seed value must be an integer between `0` and `4294967295`. If you omit this attribute, then the resulting sample will change from one execution to the next.
`method`	The type of sampling method to be used. Valid values are: • `DBMS_SODA.SAMPLE_ROW` • `DBMS_SODA.SAMPLE_BLOCK`

Return Values

This function returns the `SELF` operation object.

313.3.3.17 SKIP Function

This function sets the number of documents that match the operation criteria that will be skipped from the operation result. This setting is only usable for read operations such as `GET_CURSOR`. For write operations, any value set using this method is ignored. Returns the same `SODA_OPERATION_T` object so that further operation criteria can be chained together, if needed.

Syntax

```
SKIP (
     offset          IN NUMBER)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-30 SKIP Function Parameters

Parameter	Description
offset	The number of documents to skip.

Return Values

This function returns the same SODA_OPERATION_T object it was invoked on

313.3.3.18 VERSION Function

This function specifies that only documents with the supplied version should be returned. Typically, this is chained together with KEY(...) to implement optimistic locking for write operations such as REMOVE and REPLACE. Returns the same SODA_OPERATION_T object so that further operation criteria can be chained together, if needed.

Syntax

```
VERSION (
     version          IN VARCHAR2)
 RETURN SODA_Operation_T;
```

Parameters

Table 313-31 VERSION Function Parameters

Parameter	Description
version	Document version to be used for the operation.

Return Values

This function returns the same SODA_OPERATION_T object it was invoked on

313.3.4 SODA_CURSOR_T Type

This SODA type is used to represent a result set of documents.

Table 313-32 SODA_CURSOR_T Type Subprograms

Subprogram	Description
CLOSE Function	Closes the cursor.
HAS_NEXT Function	Returns TRUE, if the next document is available for the cursor. Otherwise, returns FALSE.

Table 313-32 (Cont.) SODA_CURSOR_T Type Subprograms

Subprogram	Description
NEXT Function	Returns the next SODA documented pointed by the cursor.

313.3.4.1 CLOSE Function

This function closes the cursor.

Syntax

```
CLOSE ()
 RETURN BOOLEAN;
```

Example 313-3 Return Values

This function returns a boolean value.

313.3.4.2 HAS_NEXT Function

This function returns `TRUE`, if the next document is available for the cursor. Otherwise, returns `FALSE`.

Syntax

```
HAS_NEXT ()
 RETURN BOOLEAN;
```

Return Values

This function returns a boolean value. `TRUE`, if the next document is available for the cursor. Otherwise, returns `FALSE`.

Exceptions

`Error`—If an error occurs while checking if the next document is available.

313.3.4.3 NEXT Function

This function returns the next `SODA` documented pointed by the cursor.

Syntax

```
NEXT ()
 RETURN SODA_Document_T;
```

Return Values

This function returns the next `SODA` documented pointed by the cursor. Returns `NULL` when the `HAS_NEXT function` returns `FALSE`.

Exceptions

`Error`—If an error occurs while getting the next document.

314
UTL Streams Types

`UTL Streams Types` describe abstract types used with Oracle XML functionality.

Four abstract PL/SQL streams are introduced and defined within the '`SYS`' schema. The streams may be referenced by `PUBLIC` and are described in the following sections.

This chapter contains the following topics:

- Security Model
- Summary of UTL Binary Streams Types

> **See Also:**
>
> For more information, see Oracle XML DB Developer's Guide

314.1 UTL Streams Types Security Model

`EXECUTE` on UTL Streams Types is granted to `PUBLIC`.

314.2 Summary of UTL Binary Streams Types

This table lists the `UTL Binary Streams Types` and briefly describes them.

Table 314-1 UTL Streams Types

Type	Description
UTL_BINARYINPUTSTREAM Type	Reads bytes and closes a stream.
UTL_BINARYOUTPUTSTREAM Type	Writes bytes and closes a stream.
UTL_CHARACTERINPUTSTREAM Type	Reads chars and closes a stream.
UTL_CHARACTEROUTPUTSTREAM Type	Writes chars and closes a stream.

314.2.1 UTL_BINARYINPUTSTREAM Type

This type is similar to `java.io.InputStream` in that it can only read and close a stream.

Syntax

```
CREATE OR REPLACE TYPE Utl_BinaryInputStream AS OBJECT (

   MEMBER FUNCTION available (
```

```
      self    IN OUT NOCOPY  Utl_BinaryInputStream)
   RETURN INTEGER,

  MEMBER FUNCTION read (                                  -- #1
     self      IN OUT NOCOPY  Utl_BinaryInputStream,
     numBytes  IN             INTEGER DEFAULT 1)
   RETURN RAW,

  MEMBER PROCEDURE read (                                 -- #2
     self       IN OUT NOCOPY   Utl_BinaryInputStream,
     bytes      IN OUT NOCOPY  RAW,
     numBytes   IN OUT         INTEGER),

  MEMBER PROCEDURE read (                                 -- #3
     self       IN OUT NOCOPY   Utl_BinaryInputStream,
     bytes      IN OUT NOCOPY   RAW,
     offset     IN INTEGER,
     numBytes   IN OUT           INTEGER),

  member function close (
     self       In Out Nocopy   Utl_BinaryInputStream)

) NOT FINAL;
```

Attributes

Table 314-2 UTL_BINARYINPUTSTREAM Type Member Subprograms

Member Subprogram	Description
AVAILABLE	Returns the number of bytes available to be read
READ	• #1 - Reads the number of bytes specified by numBytes (default is 1) and returns the bytes as a RAW. If there are no remaining bytes a value of NULL is returned. • #2 - Reads the number of bytes specified in numBytes into the parameter bytes. Additionally, the actual number of bytes read is returned in parameter numBytes. If this parameter is set to 0 then there are no more bytes to be read. • #3 - Reads the number of bytes specified in numBytes into the parameter bytes, beginning at the offset specified by parameter offset. The actual number of bytes read is returned in parameter numBytes. If this value is 0, then there are no additional bytes to be read.
CLOSE	Releases all resources held on the node to support the stream

314.2.2 UTL_BINARYOUTPUTSTREAM Type

This type is similar to java.io.OutputStream in that it can only write and close a stream.

Syntax

```
CREATE OR REPLACE TYPE Utl_BinaryOutputStream AS OBJECT (

  MEMBER FUNCTION write (                                 -- #1
       self       IN OUT NOCOPY   sys.utl_BinaryOutputStream,
       bytes      IN              RAW,
       numBytes   IN              INTEGER DEFAULT 1)
  RETURN INTEGER,

  MEMBER PROCEDURE write (                                 -- #2
       self       IN OUT NOCOPY  sys.utl_BinaryOutputStream,
```

```
      bytes      IN NOCOPY      RAW,
      numBytes   IN OUT         INTEGER),

  MEMBER PROCEDURE write (                                  -- #3
      self       IN OUT NOCOPY  utl_BinaryOutputStream,        bytes      IN
NOCOPY      RAW,
      offset     IN             INTEGER,
      numBytes   IN OUT         INTEGER),

  MEMBER PROCEDURE flush (
      self       IN OUT NOCOPY  utl_BinaryOutputStream),

  MEMBER PROCEDURE close (
      self       IN OUT NOCOPY utl_BinaryOutputStream)

 ) NOT FINAL;
```

Attributes

Table 314-3 UTL_BINARYOUTPUTSTREAM Type Member Subprograms

Member Subprogram	Description
WRITE	• #1 - Writes the number of bytes specified by `numBytes` (default is 1) from `RAW` into the stream. The actual number of bytes written is returned. • #2 - Writes the number of bytes specified in parameter `numBytes` from parameter bytes to the stream. The actual number of bytes written is returned in parameter `numBytes`. • #3 - Writes the number of bytes specified by `numBytes` to the stream, beginning at the offset specified by parameter offset. The actual number of bytes written is returned in parameter `numBytes`.
FLUSH	Insures that any buffered bytes are copied to the node destination
CLOSE	Frees all resources associated with the stream

314.2.3 UTL_CHARACTERINPUTSTREAM Type

This type is similar to `java.io.Reader` in that it can only read characters (chars) and close a stream.

Syntax

```
CREATE OR REPLACE TYPE Utl_CharacterInputStream AS OBJECT (

   MEMBER FUNCTION available (
       self       IN OUT NOCOPY   utl_CharacterInputStream)
   RETURN INTEGER,

   MEMBER FUNCTION read (                                   -- #1
       self       IN OUT NOCOPY   utl_CharacterInputStream,
       numChars   IN              INTEGER DEFAULT 1,
       lineFeed   IN              BOOLEAN DEFAULT FALSE)
   RETURN VARCHAR2,

   MEMBER PROCEDURE read (                                  -- #2
       self       IN OUT NOCOPY    utl_CharacterInputStream,
       chars      IN OUT NOCOPY    VARCHAR2,
       numChars   IN OUT           INTEGER,
```

```
      lineFeed   IN BOOLEAN DEFAULT FALSE),

   MEMBER PROCEDURE read (                                  -- #3
      self       IN OUT NOCOPY   utl_CharacterInputStream,
      chars      IN OUT NOCOPY    VARCHAR2,
      offset     IN               INTEGER,
      numChars   IN OUT           INTEGER,
      lineFeed   IN               BOOLEAN DEFAULT FALSE),

   MEMBER PROCEDURE close (
      self       IN OUT NOCOPY    utl_CharacterInputStream)
) NOT FINAL;
```

Attributes

Table 314-4 UTL_CHARACTERINPUTSTREAM Type Member Subprograms

Member Subprogram	Description
AVAILABLE	Returns the number of bytes available to be read
READ	• #1 - Returns the number of characters remaining to be read • #2 - Reads the number of characters specified by numChars (default value is 1) and returns the characters as a VARCHAR2. If the value of lineFeed is true (default value is FALSE) then the reading stops if a linefeed character is found. If there are no remaining characters a value of NULL is returned. • #3 - Reads reads the number of characters specified by parameter numChars into the parameter chars. Additionally, the actual number of characters read is returned in parameter numChars. If this value is 0, then there are no more characters to be read. If the value of lineFeed is TRUE (default is FALSE), then reading stops if a linefeed character is encountered.
CLOSE	Releases all resources held by the stream

314.2.4 UTL_CHARACTEROUTPUTSTREAM Type

This type is similar to java.io.Reader in that it can only read characters (chars) and close a stream.

Syntax

```
CREATE OR REPLACE TYPE utl_CharacterOutputStream  AS OBJECT (

 MEMBER FUNCTION write (                                  -- #1
     self      IN OUT NOCOPY   utl_CharacterOutputStream,
     chars     IN              VARCHAR2,
     numChars  IN              INTEGER DEFAULT 1,
     lineFeed  IN              BOOLEAN DEFAULT FALSE)
  RETURN INTEGER,

 MEMBER PROCEDURE write (                                  -- #2
     self      IN OUT NOCOPY   utl_CharacterOutputStream,
     chars     IN OUT NOCOPY   VARCHAR2,
     numChars  IN OUT          INTEGER,
     lineFeed  IN              BOOLEAN DEFAULT FALSE),

 member procedure write (                                  -- #3
     self     IN OUT NOCOPY    utl_CharacterOutputStream,
     chars    IN  NOCOPY       varchar2,
```

```
    offset   IN              integer,
    numChars IN OUT          integer,
   lineFeed  IN              boolean default false),

 MEMBER PROCEDURE flush (
    self     IN OUT NOCOPY    utl_CharacterOutputStream),

MEMBER PROCEDURE close (
    self     IN OUT NOCOPY utl_CharacterOutputStream)

 ) NOT FINAL;
```

Attributes

Table 314-5 UTL_CHARACTEROUTPUTSTREAM Type Member Subprograms

Member Subprogram	Description
WRITE	• #1 - Writes the number of characters specified by numChars (default is 1) from parameter chars into the stream and returns the actual number of characters written. If the value of lineFeed is TRUE (default is FALSE) a lineFeed character is inserted after the last character. • #2 - writes the number of characters specified by parameter numChars, from parameter chars into the stream. The actual number of characters written is returned in parameter numChars. If the value of lineFeed is true (default is FALSE) a lineFeed character is inserted after the last character. • #3 - Writes the number of characters specified by parameter numChars, from parameter chars, beginning at offset specified by parameter offset. The actual number of characters written is returned in parameter numChars. If the value of lineFeed is true (default is FALSE) a lineFeed character is inserted after the last character .
FLUSH	Copies all characters that may be contained within buffers to the node value
CLOSE	Releases all resources held by the stream

315
XMLTYPE

`XMLType` is a system-defined opaque type for handling XML data. It as predefined member functions on it to extract XML nodes and fragments.

You can create columns of `XMLType` and insert XML documents into it. You can also generate XML documents as `XMLType` instances dynamically using the `SYS_XMLAGG SQL` function.

This chapter contains the following topics:

- Summary of XMLType Subprograms

> **See Also:**
>
> – *Oracle XML DB Developer's Guide*

315.1 Summary of XMLType Subprograms

This table summarizes functions and procedures of `XMLType`.

Table 315-1 XMLTYPE Subprograms

Method	Description
CREATENONSCHEMABASEDXML	Creates a non schema based XML from the input schema based instance.
CREATESCHEMABASEDXML	Creates a schema based `XMLType` instance from the non-schema based instance using the input schema URL.
CREATEXML	Static function for creating and returning an `XMLType` instance.
EXISTSNODE	Takes a `XMLType` instance and a `XPath` and returns 1 or 0 indicating if applying the `XPath` returns a non-empty set of nodes.
EXTRACT	Takes a `XMLType` instance and an `XPath`, applies the XPath expression and returns the results as an `XMLType`.
GETBLOBVAL	Returns the value of the `XMLType` instance as a `BLOB`
GETCLOBVAL	Returns the value of the `XMLType` instance as a `CLOB`.
GETNAMESPACE	Returns the namespace for the top level element in a schema based document.
GETNUMBERVAL	Returns the value of the `XMLType` instance as a `NUMBER`. This is only valid if the input `XMLType` instance contains a simple text node and is convertible to a number.
GETROOTELEMENT	Returns the root element of the input instance. Returns `NULL` if the instance is a fragment
GETSCHEMAURL	Returns the XML schema URL if the input is an XML Schema based.

Table 315-1 (Cont.) XMLTYPE Subprograms

Method	Description
GETSTRINGVAL	Returns the value of the `XMLType` instance as a string.
ISFRAGMENT	Checks if the input `XMLType` instance is a fragment or not. A fragment is a XML instance, which has more than one root element.
ISSCHEMABASED	Returns 1 or 0 indicating if the input `XMLType` instance is a schema based one or not.
ISSCHEMAVALID	Checks if the input instance is schema valid according to the given schema URL.
ISSCHEMAVALIDATED	Checks if the instance has been validated against the schema.
SCHEMAVALIDATE	Validates the input instance according to the XML Schema. Raises error if the input instance is non-schema based.
SETSCHEMAVALIDATED	Sets the schema valid flag to avoid costly schema validation.
TOOBJECT	Converts the `XMLType` instance to an object type.
TRANSFORM	Takes an `XMLType` instance and an associated stylesheet (which is also an `XMLType` instance), applies the stylesheet and returns the result as XML.
XMLTYPE	Constructs an instance of the `XMLType` datatype. The constructor can take in the XML as a `CLOB`, `VARCHAR2` or take in a object type.

315.1.1 CREATENONSCHEMABASEDXML

This member function creates a non-schema based XML document from a schema based instance.

Syntax

```
MEMBER FUNCTION CREATENONSCHEMABASEDXML
return XMLType deterministic;
```

315.1.2 CREATESCHEMABASEDXML

This member function creates a schema based `XMLType` instance from a non-schema based `XMLType` value.

It uses either the supplied `SCHEMA` URL, or the `SCHEMALOCATION` attribute of the instance.

Syntax

```
MEMBER FUNCTION createSchemaBasedXML(
schema IN varchar2 := NULL)
return XMLType deterministic;
```

Table 315-2 CREATESCHEMABASEDXML Subprogram Parameters

Parameter	Description
`schema`	Optional XMLSchema URL used to convert the value to the specified schema.

315.1.3 CREATEXML

This static function creates and returns an `XMLType` instance. The string and `CLOB` parameters used to pass in the date must contain well-formed and valid XML documents.

The options are described in the following table.

Table 315-3 CREATEXML Subprograms

Syntax	Description

`STATIC FUNCTION createXML(`
`  xmlData IN varchar2)`
`RETURN XMLType deterministic;`
	Creates the `XMLType` instance from a string.

`STATIC FUNCTION createXML(`
`  xmlData IN clob)`
`RETURN XMLType deterministic;`
	Creates the `XMLType` instance from a `CLOB`.

`STATIC FUNCTION createXML (`
`  xmlData IN clob,`
`  schema IN varchar2,`
`  validated IN number := 0,`
`  wellformed IN number := 0 )`
`RETURN XMLType deterministic;`
	This static function creates a schema-based `XMLType` instance using the specified schema and xml data parameters.

`STATIC FUNCTION createXML (`
`  xmlData IN varchar2,`
`  schema IN varchar2,`
`  validated IN number := 0,`
`  wellformed IN number := 0)`
`RETURN XMLType deterministic;`
	This static function creates a schema-based `XMLType` instance using the specified schema and xml data parameters.

`STATIC FUNCTION createXML (`
`  xmlData IN "<ADT_1>",`
`  schema IN varchar2 := NULL,`
`  element IN varchar2 := NULL,`
`  validated IN NUMBER := 0)`
`RETURN XMLType deterministic;`
	Creates an XML instance from an instance of an user-defined type.

`STATIC FUNCTION createXML (`
`  xmlData IN SYS_REFCURSOR,`
`  schema in varchar2 := NULL,`
`  element in varchar2 := NULL,`
`  validated in number := 0)`
`RETURN XMLType deterministic;`
	Creates an XML instance from a cursor reference. You can pass in any arbitrary SQL query as a `CURSOR`.

Table 315-3 (Cont.) CREATEXML Subprograms

Syntax	Description

`STATIC FUNCTION createXML (`
`xmlData IN AnyData,`
`schema in varchar2 := NULL,`
`element in varchar2 := NULL,`
`validated in number := 0)`
`RETURN sys.XMLType deterministic`
`parallel_enable`
	Creates an XML instance from `ANYDATA`.If the `ANYDATA`instance contains an ADT, the `XMLType`returned is the same as would be returned for a call directly on the ADT. If the `ANYDATA`contains a scalar, the `XMLType` contains a leaf node with the scalar value. The element name for this node is taken from the optional element string if present, and is "`ANYDATA`" if it is not.

`STATIC FUNCTION createXML (`
`xmlData IN blob,`
`csid IN number,`
`schema IN varchar2,`
`validated IN number := 0,`
`wellformed IN number := 0)`
`return sys.XMLType deterministic`
	Creates an XML instance from a `BLOB`.

`STATIC FUNCTION createXML (`
`xmlData IN bfile,`
`csid IN number,`
`Schema IN varchar2,`
`validated IN number := 0,`
`wellformed IN number := 0)`
`return sys.XMLType deterministic`
	Creates an XML instance from a `BFILE`.

Table 315-4 CREATEXML Parameters

Parameter	Description
`xmlData`	The actual data in the form of a `BFILE`, `BLOB`, `CLOB`, `REF` cursor, `VARCHAR2` or object type.
`schema`	Optional Schema URL to be used to make the input conform to the given schema. **Caution**: Oracle does not support use of types generated by Schema Registration (see *Oracle XML DB Developer's Guide*).
`validated`	Flag to indicate that the instance is valid according to the given XML Schema. (Default is `0`)
`wellformed`	Flag to indicate that the input is well formed. If set, then the database would not do well formed check on the input instance. (Default is `0`)
`element`	Optional element name in the case of the `ADT_1` or `REF CURSOR` constructors. (Default is `NULL`). **Caution**: Oracle does not support use of types generated by Schema Registration (see *Oracle XML DB Developer's Guide*).
`csid`	The character set id of input XML data.

315.1.4 EXISTSNODE

This member function checks if the node exists.

If the XPath string is `NULL` or the document is empty, then a value of 0 is returned, otherwise returns 1.

The options are described in the following table.

Syntax	Description

`MEMBER FUNCTION existsNode(`
`  xpath IN varchar2)`
`RETURN number deterministic;`
	Given an XPath expression, checks if the XPath applied over the document can return any valid nodes.

`MEMBER FUNCTION existsNode(`
`  xpath in varchar2,`
`  nsmap in varchar2)`
`RETURN number deterministic;`
	This member function uses the XPath expression with the namespace information and checks if applying the XPath returns any nodes or not.

Table 315-5 EXISTSNODE Subprogram Parameters

Parameter	Description
`xpath`	The XPath expression to test.
`nsmap`	Optional namespace mapping.

315.1.5 EXTRACT

This member function extracts an `XMLType` fragment and returns an `XMLType` instance containing the result nodes. If the XPath does not result in any nodes, it then returns `NULL`.

The options are described in the following table.

Syntax	Description

`MEMBER FUNCTION extract(`
`  xpath IN varchar2)`
`RETURN XMLType deterministic;`
	Given an XPath expression, applies the XPath to the document and returns the fragment as an `XMLType`.

`MEMBER FUNCTION extract(`
`  xpath IN varchar2,`
`  nsmap IN varchar2)`
`RETURN XMLType deterministic;`
	This member function applies the XPath expression and namespace mapping, over the XML data to return a `XMLType` instance containing the resultant fragment.

Table 315-6 EXTRACT Subprogram Parameters

Parameter	Description
`xpath`	The XPath expression to apply.
`nsmap`	Optional prefix to namespace mapping information.

315.1.6 GETBLOBVAL

This member function returns a BLOB containing the serialized XML representation. If the BLOB returned is temporary, it must be freed after use.

Syntax

```
MEMBER FUNCTION getBlobVal(
   csid    IN    NUMBER)
 RETURN BLOB DETERMINISTIC;
```

Table 315-7 GETBLOBVAL Subprogram Parameters

Parameter	Description
csid	The desired character set ID of output BLOB

315.1.7 GETCLOBVAL

This member function returns a CLOB containing the serialized XML representation. If the CLOB returned is temporary, it must be freed after use. The CLOBs returned by this function are read-only.

Syntax

```
MEMBER FUNCTION getClobVal()
RETURN clob deterministic;
```

315.1.8 GETNAMESPACE

GETNAMESPACE is a member function. It returns the namespace of the top level element in the instance. It returns NULL if the input is a fragment or is a non-schema based instance.

Syntax

```
MEMBER FUNCTION getNamespace
return varchar2 deterministic;
```

315.1.9 GETNUMBERVAL

This is a member function. It returns a numeric value, formatted from the text value pointed to by the XMLType instance. The XMLType must point to a valid text node that contains a numerical value.

The options are described in the following table.

Syntax

```
MEMBER FUNCTION getNumberVal()
RETURN number deterministic;
```

315.1.10 GETROOTELEMENT

this member function gets the root element of the `XMLType` instance. It returns `NULL` if the instance is a fragment.

Syntax

```
MEMBER FUNCTION getRootElement
return varchar2 deterministic;
```

315.1.11 GETSCHEMAURL

This member function returns the XML Schema URL corresponding to the `XMLType` instance, if the `XMLType` instance is a schema-based document. Otherwise, it returns `NULL`.

Syntax

```
MEMBER FUNCTION getSchemaURL
return varchar2 deterministic;
```

315.1.12 GETSTRINGVAL

This member function returns the document as a string. It returns a string containing the seralized XML representation, or in the case of text nodes, the text itself.

If the XML document exceeds the `VARCHAR2` maximum size (4000), then an error is raised at run time.

Syntax

```
MEMBER FUNCTION getStringVal()
RETURN varchar2 deterministic;
```

315.1.13 ISFRAGMENT

`ISFRAGMENT` determines if the `XMLType` instance corresponds to a well-formed document, or a fragment. It returns 1 or 0 indicating if the `XMLType` instance contains a fragment or a well-formed document.

Syntax

```
MEMBER FUNCTION isFragment()
RETURN number deterministic;
```

315.1.14 ISSCHEMABASED

This member function determines whether the `XMLType` instance is schema-based or not. It returns 1 or 0 depending on whether the `XMLType` instance is schema-based.

Syntax

```
MEMBER FUNCTION isSchemaBased
return number deterministic;
```

315.1.15 ISSCHEMAVALID

This member function checks if the input instance conforms to a specified schema. I does not change the validation status of the XML instance.

If an XML Schema URL is not specified and the xml document is schema based, the conformance is checked against the `XMLType` instance's own schema.

Syntax

```
member function isSchemaValid(
schurl IN VARCHAR2 := NULL,
elem IN VARCHAR2 := NULL)
return NUMBER deterministic;
```

Table 315-8 ISSCHEMAVALID Subprogram Parameters

Parameter	IN / OUT	Description
`schurl`	`(IN)`	The URL of the XML Schema against which to check conformance.
`elem`	`(IN)`	Element of a specified schema, against which to validate. This is useful when we have a XML Schema which defines more than one top level element, and we want to check conformance against a specific one of these elements.

315.1.16 ISSCHEMAVALIDATED

This member function returns the validation status of the `XMLType` instance to tell if a schema-based instance has been actually validated against its schema. It returns 1 if the instance has been validated against the schema, 0 otherwise.

Syntax

```
MEMBER FUNCTION isSchemaValidated
return NUMBER deterministic;
```

315.1.17 SCHEMAVALIDATE

This member procedure validates the XML instance against its schema, if it has not already been done.

For non-schema based documents an error is raised. If validation fails an error is raised; else, the document's status is changed to validated.

Syntax

```
MEMBER PROCEDURE schemaValidate(
   self IF OUT NOCOPY XMLType);
```

Table 315-9 SCHEMAVALIDATE Subprogram Parameters

Parameter	IN / OUT	Description
`self`	`(OUT)`	XML instance being validated against the schema.

315.1.18 SETSCHEMAVALIDATED

This member function sets the `VALIDATION` state of the input XML instance.

Syntax

```
MEMBER PROCEDURE setSchemaValidated(
self IF OUT NOCOPY XMLType,
   flag IN BINARY_INTEGER := 1);
```

Table 315-10 SERTSSCHEMAVALIDATED Subprogram Parameters

Parameter	IN / OUT	Description
`self`	`(OUT)`	XML instance.
`flag`	`(IN)`	0 - `NOT VALIDATED`; 1 - `VALIDATED` (Default)

315.1.19 TOOBJECT

This member procedure converts the XML value to an object type using the `XMLSCHEMA` mapping, if available. If a `SCHEMA` is not supplied or the input is a non-schema based XML, the procedure uses cannonical mapping between elements and object type attributes.

Syntax

```
MEMBER PROCEDURE toObject(
SELF in XMLType,
object OUT "<ADT_1>",
schema in varchar2 := NULL,
element in varchar2 := NULL);
```

Table 315-11 TOOBJECT Subprogram Parameters

Parameter	IN / OUT	Description
`SELF`	`(IN)`	Instance to be converted. Implicit if used as a member procedure.
`object`	`(IN)`	Converted object. An object instance of the required type may be passed in to this function
`schema`	`(IN)`	Schema URL. The mapping of the `XMLType` instance to the converted object instance may be specified using a schema. **Caution**: Oracle does not support use of types generated by Schema Registration (see *Oracle XML DB Developer's Guide*).
`element`	`(IN)`	Top-level element name. An XML Schema document does not specify the top-level element for a conforming XML instance document without this parameter. **Caution**: Oracle does not support use of types generated by Schema Registration (see *Oracle XML DB Developer's Guide*).

315.1.20 TRANSFORM

This member function transforms the XML data using the XSL stylesheet argument and the top-level parameters passed as a string of name=value pairs

If any of the arguments other than the parammap is `NULL`, then a `NULL` is returned.

Syntax

```
MEMBER FUNCTION transform(
xsl IN XMLType,
parammap in varchar2 := NULL)
RETURN XMLType deterministic;
```

Table 315-12 TRANSFORM Subprogram Parameters

Parameter	IN / OUT	Description
xsl	(IN)	The XSL stylesheet describing the transformation
parammap	(IN)	Top level parameters to the XSL - string of name=value pairs

315.1.21 XMLTYPE

This is an `XMLType` constructor.

The options are described in the following table.

Table 315-13 XMLTYPE Member Subprogram Parameters

Syntax	Description

`constructor function XMLType(`
`xmlData IN clob,`
`schema IN varchar2 := NULL,`
`validated IN number := 0,`
`wellformed IN Number := 0)`
`return self as result deterministic;`
	This constructor function creates an optionally schema-based `XMLType` instance using the specified schema and xml data parameters.

`constructor function XMLType(`
`xmlData IN varchar2,`
`schema IN varchar2 := NULL,`
`validated IN number := 0,`
`wellformed IN number := 0)`
`return self as result deterministic;`
	This constructor function creates an optionally schema-based `XMLType` instance using the specified schema and xml data parameters.

`constructor function XMLType (`
`xmlData IN "w<ADT_1>",`
`schema IN varchar2 := NULL,`
`element IN varchar2 := NULL,`
`validated IN number := 0)`
`return self as result deterministic;`
	This constructor function creates an optionally schema-based `XMLType` instance from the specified object type parameter.

`constructor function XMLType(`
`xmlData IN SYS_REFCURSOR,`
`schema in varchar2 := NULL,`
`element in varchar2 := NULL,`
`validated in number := 0)`
`return self as result deterministic;`
	This constructor function creates an optionally schema-based `XMLType` instance from the specified `REF CURSOR` parameter.

Table 315-13 (Cont.) XMLTYPE Member Subprogram Parameters

Syntax	Description

`constructor function XMLType(`
`  xmlData IN AnyData,`
`  schema IN varchar2 := NULL,`
`  element IN varchar2 := NULL,`
`  validated IN number := 0)`
`return self as result deterministic`
`  parallel_enable`
	This constructor function creates an optionally schema-based `XMLType` instance from the specified `ANYDATA` parameter.If the `ANYDATA` instance contains an ADT, the `XMLType` returned is the same as would be returned for a call directly on the ADT. If the `ANYDATA` contains a scalar, the `XMLType` contains a leaf node with the scalar value. The element name for this node is taken from the optional element string if present, and is "`ANYDATA`" if it is not.

`constructor function XMLType(`
`  xmlData IN blob, csid IN number,`
`  schema IN varchar2 := NULL,`
`  validated IN number := 0,`
`  wellformed IN number := 0)`
`return self as result deterministic`
	This constructor function creates an optionally schema-based `XMLType` instance from the specified `BLOB` parameter.

`constructor function XMLType(`
`  xmlData IN bfile,`
`  csid IN number,`
`  schema IN varchar2 := NULL,`
`  validated IN number := 0,`
`  wellformed IN number := 0)`
`return self as result deterministic`
	This constructor function creates an optionally schema-based `XMLType` instance from the specified `BFILE` parameter.

Table 315-14 XMLTYPE Constructor Subprogram Parameters

Parameter	Description
`xmlData`	The data in the form of a `BFILE`, `BLOB`, `CLOB`, `REF`s, `VARCHAR2` or object type.
`schema`	Optional Schema URL to be used to make the input conform to the given schema. **Caution**: Oracle does not support use of types generated by Schema Registration (see *Oracle XML DB Developer's Guide*).
`validated`	Indicates that the instance is valid to the given XML Schema.
`wellformed`	Indicates that the input is well formed. If set, then the database would not do well formed check on the input instance.
`element`	Optional element name in the case of the `ADT_1` or `REF CURSOR` constructors. (Default is `NULL`). **Caution**: Oracle does not support use of types generated by Schema Registration (see *Oracle XML DB Developer's Guide*).
`csid`	The character set id of input XML data.

Index

A

ABORT procedure, *161-2*
ABORT_GLOBAL_INSTANTIATION procedure, *44-3*
ABORT_REDEF_COLLECTION procedure, *177-7*
ABORT_REDEF_TABLE procedure, *154-5*
ABORT_ROLLBACK Procedure, *154-5*
ABORT_SCHEMA_INSTANTIATION procedure, *44-4*
ABORT_TABLE_INSTANTIATION procedure, *44-5*
ABORT_UPDATE Procedure, *154-6*
ABORTED_REQUEST_THRESHOLD procedure, *175-2*
ACCEPT_SQL_PATCH Procedure, *186-5*
ACCEPT_SQL_PLAN_BASELINE Procedure, *181-4*
ACCEPT_SQL_PROFILE procedure, *190-15*
ACCHK_REPORT Procedure, *22-1*
ACCHK_SET Procedure, *21-2*, *21-3*, *21-5*, *21-6*
ACLCHECKPRIVILEGES Function, *223-4*
ACTIVE_INSTANCES Procedure, *210-6*
ACTIVITY_REPORT Function, *39-1*
ADD procedure, *155-2*
ADD_ALLOWED_CONTEXT procedure (of DBMS_SQL_FIREWALL), *183-4*
ADD_AUTO_CDR procedure, *87-3*
ADD_AUTO_CDR_COLUMN_GROUP procedure, *87-5*
ADD_AUTO_CDR_DELTA_RES procedure, *87-6*
ADD_CAPTURE Procedure, *216-4*
ADD_COLORED_SQL Procedure, *217-5*
ADD_COLUMN member procedure, *310-10*
ADD_COLUMN procedure, *239-5*
ADD_COOKIES procedure, *286-29*
ADD_FILTER Procedure, *215-2*, *216-5*
ADD_GLOBAL_PROPAGATION_RULES procedure, *239-8*
ADD_GLOBAL_RULES procedure, *239-12*
ADD_GROUPED_POLICY Procedure, *162-4*
ADD_OUTBOUND procedure, *239-16*
ADD_PAIR member procedure, *312-4*
ADD_POLICY Procedure, *80-2*, *153-7*, *162-5*, *203-3*
ADD_POLICY_CONTEXT Procedure, *162-10*
ADD_POOL Procedure, *49-2*
ADD_PRIVILEGE Procedure, *131-6*
ADD_REFERENCE function, *189-5*
ADD_RULE procedure, *166-2*
ADD_SCHEDULE_ORDERING Function, *216-5*
ADD_SCHEMA_PROPAGATION_RULES procedure, *239-20*
ADD_SCHEMA_RULES procedure, *239-24*
ADD_SCHEME procedure, *124-4*
ADD_SENSITIVE_COLUMN Procedure, *202-2*
ADD_SENSITIVE_TYPE Procedure, *202-4*
ADD_SQL_CONNECTION_TEST Procedure, *20-2*, *21-6*
ADD_SQLSET_REFERENCE function, *190-19*
ADD_SQLWKLD_REF Procedure, *17-5*
ADD_SQLWKLD_STATEMENT Procedure, *17-6*
ADD_SUBSCRIBER Procedure, *125-12*
ADD_SUBSET_OUTBOUND_RULES procedure, *239-28*
ADD_SUBSET_PROPAGATION_RULES procedure, *239-31*
ADD_SUBSET_RULES procedure, *239-34*
ADD_TABLE_PROPAGATION_RULES procedure, *239-39*
ADD_TABLE_RULES procedure, *239-43*
ADD_TABLE_TO_APPLICATION Procedure, *83-4*
ADD_TO_ILM Procedure, *97-3*
ADD_VIRTUAL_COLUMNS procedure, *107-4*
ADD_WARNING_SETTING_CAT procedure, *213-2*
ADD_WARNING_SETTING_NUM procedure, *213-2*
ADD2MULTI procedure, *251-2*
ADDATTR member procedure
of ANYTYPE TYPE, *304-3*
ADDHTTPEXPIREMAPPING Procedure, *221-4*
ADDINSTANCE member procedure
of ANYDATASET TYPE, *303-2*
ADDMIMEMAPPING Procedure, *221-4*
ADDREPOSITORYRESCONFIG Procedure, *157-2*
ADDRESCONFIG Function, *157-3*
ADDRESS function
of HTF package, *242-11*
ADDRESS procedure
of HTP package, *243-11*

ADDSCHEMALOCMAPPING Procedure, *221-5*
ADDSERVLET Procedure, *221-5*
ADDSERVLETMAPPING Procedure, *221-6*
ADDSERVLETSECROLE Procedure, *221-7*
ADDXDBNAMESPACE, *233-3*
ADDXMLEXTENSION Procedure, *221-7*
ADM_DROP_CHUNKS Procedure, *137-7*
ADM_DROP_TASK Procedure, *137-7*
ADM_GET_TASK_STATUS Procedure, *137-8*
ADM_STOP_TASK Procedure, *137-8*
ADMIN_TABLES procedure, *156-4*
ADVISE_COMMIT procedure, *200-2*
ADVISE_NOTHING procedure, *200-3*
ADVISE_ROLLBACK procedure, *200-3*
ADVISOR privilege, *54-2*
aggregation management, *54-2*
aggregation operators (OLAP), *54-19*
ALLOCATE_UNIQUE procedure, *115-3*
ALLOCATE_UNIQUE_AUTONOMOUS procedure, *115-5*
ALTER_AGENT Procedure, *125-15*
ALTER_APPLY procedure, *24-5*
ALTER_AUTO_CDR procedure, *87-7*
ALTER_AUTO_CDR_COLUMN_GROUP procedure, *87-10*
ALTER_CAPTURE procedure, *44-6*
ALTER_COMPILE procedure, *66-2*
ALTER_EVALUATION_CONTEXT procedure, *166-4*
ALTER_GROUPED_POLICY Procedure, *162-12*
ALTER_INBOUND procedure, *239-47*
ALTER_JOB Procedure, *125-16*
ALTER_MSGSYSTEM_LINK Procedure for TIB/Rendezvous, *125-18*
ALTER_MSGSYSTEM_LINK Procedure for WebSphere MQ, *125-19*
ALTER_OUTBOUND procedure, *239-48*
ALTER_PARAM Procedure, *49-4*
ALTER_PARTITION_INTERVAL procedure, *31-8*
ALTER_POLICY Procedure, *153-13*, *162-11*, *203-4*
ALTER_PROPAGATION procedure, *150-2*
ALTER_PROPAGATION_SCHEDULE Procedure, *125-20*
ALTER_REWRITE_EQUIVALENCE Procedure, *16-2*
ALTER_RULE procedure, *166-6*
ALTER_SENSITIVE_COLUMN Procedure, *202-3*
ALTER_SQL_PATCH Procedure, *186-7*
ALTER_SQL_PLAN_BASELINE Function, *181-8*
ALTER_SQL_PLAN_DIRECTIVE Procedure, *180-2*
ALTER_SQL_PROFILE procedure, *190-20*
ALTER_STATS_HISTORY_RETENTION procedure, *192-18*
ALTER_SUBSCRIBER Procedure, *125-21*
ALTER_SYNC_CAPTURE procedure, *44-11*
ALTER_TABLE_NOT_REFERENCEABLE procedure, *66-3*
ALTER_TABLE_REFERENCEABLE procedure, *66-4*
AMATCH function, *249-3*
analytic views, *91-1*
  hierarchy validation, *91-10*
  validation, *91-4*, *91-5*, *91-7*
  validation log table, *91-3*, *91-7*
  verify validation, *91-9*
ANALYZE Function, *93-2*
ANALYZE Procedure, *41-1*
ANALYZE_DATABASE Procedure, *210-6*
ANALYZE_DB Procedure, *15-3*
ANALYZE_INST Procedure, *15-4*
ANALYZE_PART_OBJECT Procedure, *210-7*
ANALYZE_PARTIAL Procedure, *15-5*
ANALYZE_SCHEMA Procedure, *210-8*
ANCHOR function
  of HTF package, *242-12*
ANCHOR procedure
  of HTP package, *243-11*
ANCHOR2 function
  of HTF package, *242-12*
ANCHOR2 procedure
  of HTP package, *243-12*
anonymous PL/SQL blocks
  dynamic SQL and, *182-2*
AnyData datatype
  queues
    creating, *239-91*
    removing, *239-77*
APEX packages, *2-1*
APPEND Procedures, *114-13*
APPEND_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-5*
APPEND_ALLOW_LIST_SINGLE_SQL procedure (of DBMS_SQL_FIREWALL), *183-6*
APPEND_HOST_ACE Procedure, *131-7*
APPEND_HOST_ACL Procedure, *131-8*
APPEND_WALLET_ACE Procedure, *131-9*
APPEND_WALLET_ACL Procedure, *131-10*
APPENDCHILD function, *228-20*
APPENDDATA procedure, *228-21*
APPENDRESCONFIG Procedure, *157-3*
APPENDRESOURCEMETADATA Procedure, *223-4*
APPLETCLOSE function
  of HTF package, *242-13*
APPLETCLOSE procedure
  of HTP package, *243-13*
APPLETOPEN function
  of HTF package, *242-13*

APPLETOPEN procedure
of HTP package, 243-13
apply process
altering, 24-5
clear key columns
removing, 24-10
conflict handlers
setting, 24-28, 24-75
creating, 24-13, 239-12, 239-24, 239-34, 239-43
DBMS_APPLY_ADM package, 24-1
DDL handler
setting, 24-5, 24-13
dropping, 24-21
enqueuing events, 24-38
error handlers
setting, 24-34
error queue
deleting errors, 24-20
executing errors, 24-23, 24-24
getting error messages, 24-26
instantiation
global SCN, 24-41
schema SCN, 24-70
table SCN, 24-73
message handler
setting, 24-5, 24-13
parameters
setting, 24-48
precommit handler
setting, 24-5, 24-13
procedure DML handlers
setting, 24-34
rules
defining global, 239-12
defining schema, 239-24
defining subset, 239-34
defining table, 239-43
removing, 239-78
specifying execution, 24-40
starting, 24-80
stopping, 24-81
substitute key columns
setting, 24-45
APPLY_POLICY_EXPR_TO_COL Procedure, 153-20
APPLY$_ENQUEUE, 24-38
APPLY$_EXECUTE, 24-40
AQ$_AGENT Type, 305-2
AQ$_AGENT_LIST_T Type, 305-2
AQ$_DESCRIPTOR Type, 305-3
AQ$_NTFN_MSGID_ARRAY Type, 305-4
AQ$_POST_INFO Type, 305-4
AQ$_POST_INFO_LIST Type, 305-4
AQ$_PURGE_OPTIONS_T Type, 305-5
AQ$_RECIPIENT_LIST_T Type, 305-5
AQ$_REG_INFO Type, 305-6
AQ$_REG_INFO_LIST Type, 305-8
AQ$_SUBSCRIBER_LIST_T Type, 305-8
ARCHIVESTATENAME Function, 97-4
AREA function
of HTF package, 242-14
AREA procedure
of HTP package, 243-14
arrays
BIND_ARRAY procedure, 182-7
bulk DML using DBMS_SQL, 182-28
ASA_RECO_ROW Record Type, 178-2
ASA_RECO_ROW_TB Table Type, 178-3
ASA_RECOMMENDATIONS Function, 178-3
ASH_GLOBAL_REPORT_HTML Function, 217-5
ASH_GLOBAL_REPORT_TEXT Function, 217-7
ASH_REPORT_ANALYTICS function, 217-9
ASH_REPORT_HTML Function, 217-10
ASH_REPORT_TEXT Function, 217-12
ASSIGN_ACL Procedure, 131-10
ASSIGN_GROUP_TO_INSTANCE Procedure, 216-6
ASSIGN_WALLET_ACL Procedure, 131-12
ASSM_SEGMENT_VERIFY Procedure, 179-4
ASSM_TABLESPACE_VERIFY Procedure, 179-5
ASSOCIATE_POLICY Procedure, 203-5
ATTACH_SESSION procedure, 67-15
AUTH Function and Procedure, 298-8
AUTHORIZE function, 246-1
AUTHORIZE_DAD Procedure, 78-4
AUTO_CACHE_CREATE Procedure, 40-2
AUTO_CACHE_DISABLE Procedure, 40-3
AUTO_CACHE_ENABLE Procedure, 40-3
AUTO_CACHE_MODIFY Procedure, 40-4
AUTO_CACHE_REFRESH Procedure, 40-4
AUTO_CACHE_REMOVE Procedure, 40-5
AUTO_CACHE_STAR_DISABLE Procedure, 40-6
AUTO_CACHE_STAR_ENABLE Procedure, 40-6
AUTO_CACHE_STAR_MODIFY Procedure, 40-7
AUTO_CACHE_STAR_REFRESH Procedure, 40-7
AUTO_CACHE_TUNE Procedure, 40-5
AVAILABLE function, 299-5
AWR_BASELINE_METRIC_TYPE Object Type, 217-2
AWR_BASELINE_METRIC_TYPE_TABLE Table Type, 217-2
AWR_DIFF_REPORT_HTML Function, 217-15
AWR_DIFF_REPORT_TEXT Function, 217-15
AWR_EXP Procedure, 217-16
AWR_GLOBAL_DIFF_REPORT_HTML Functions, 217-17
AWR_GLOBAL_DIFF_REPORT_TEXT Functions, 217-18
AWR_GLOBAL_REPORT_HTML Functions, 217-19

AWR_GLOBAL_REPORT_TEXT Functions, *217-20*
AWR_IMP Procedure, *217-21*
AWR_REPORT_HTML function, *217-21*
AWR_REPORT_TEXT function, *217-22*
AWR_SET_REPORT_THRESHOLDS Procedure, *217-23*
AWR_SQL_REPORT_HTML Function, *217-24*
AWR_SQL_REPORT_TEXT Function, *217-25*

B

BACKTRACE_DEPTH Function, *281-4*
BACKTRACE_LINE Function, *281-5*
BACKTRACE_UNIT Function, *281-5*
BASE function
  of HTF package, *242-15*
BASE procedure
  of HTP package, *243-15*
BASE64_DECODE function, *284-2*
BASE64_ENCODE function, *284-2*
BASEFONT function
  of HTF package, *242-15*
BASEFONT procedure
  of HTP package, *243-15*
BEGIN_OPERATION Function, *184-2*
BEGIN_PREPARE Procedure, *76-3*
BEGIN_REPLAY_SCHEDULE Procedure, *216-7*
BEGIN_REQUEST function, *286-30*
BEGIN_SQL_BLOCK Procedure, *158-5*
BEGIN_TRANSACTION Procedure, *106-8*
BEGIN_UPGRADE Procedure, *76-3*
BEGINCREATE static procedure
  of ANYDATA TYPE, *302-3*
  of ANYDATASET TYPE, *303-3*
  of ANYTYPE TYPE, *304-2*
BEST_POSSIBLE_RETENTION function, *207-2*
BFILE_TABLE Table Type, *182-20*
BGSOUND function
  of HTF package, *242-16*
BGSOUND procedure
  of HTP package, *243-16*
BIG function
  of HTF package, *242-16*
BIG procedure
  of HTP package, *243-16*
BINARY_DOUBLE_TABLE Table Type, *182-20*
BINARY_FLOAT_TABLE Table Type, *182-20*
BIND_ARRAY procedures, *182-27*
BIND_INOUT_VARIABLE Procedure, *95-2*
BIND_INOUT_VARIABLE_RAW Procedure, *95-3*
BIND_OUT_VARIABLE Procedure, *95-4*
BIND_OUT_VARIABLE_RAW Procedure, *95-5*
BIND_VARIABLE Procedure, *95-6*, *182-7*
BIND_VARIABLE procedures, *182-32*
BIND_VARIABLE_PKG Procedure, *182-7*, *182-34*
BIND_VARIABLE_RAW Procedure, *95-7*
BIND_VARIABLES function, *252-3*
BIT_AND function, *294-3*
BIT_COMPLEMENT function, *294-4*
BIT_OR function, *294-4*
BIT_XOR function, *294-5*
BLACK_LIST Function, *160-3*
BLACK_LIST_ADD Procedure, *160-3*
BLACK_LIST_REMOVE, *160-3*
BLAS Level 1 (Vector-Vector Operations) Subprograms, *293-2*
BLAS Level 2 (Matrix-Vector Operations) Subprograms, *293-3*
BLAS Level 3 (Matrix-Matrix Operations) Subprograms, *293-4*
BLAS_ASUM Functions, *293-13*
BLAS_AXPY Procedures, *293-14*
BLAS_COPY Procedures, *293-15*
BLAS_DOT Functions, *293-16*
BLAS_GBMV Procedures, *293-17*
BLAS_GEMM Procedures, *293-19*
BLAS_GEMV Procedures, *293-21*
BLAS_GER Procedures, *293-23*
BLAS_IAMAX Functions, *293-24*
BLAS_NRM2 Functions, *293-25*
BLAS_ROT Procedures, *293-26*
BLAS_ROTG Procedures, *293-27*
BLAS_SBMV Procedures, *293-33*
BLAS_SCAL Procedure, *293-27*
BLAS_SPMV Procedures, *293-28*
BLAS_SPR Procedures, *293-30*
BLAS_SPR2 Procedures, *293-31*
BLAS_SWAP Procedure, *293-34*
BLAS_SYMM Procedures, *293-35*
BLAS_SYMV Procedures, *293-37*
BLAS_SYR Procedures, *293-39*
BLAS_SYR2 Procedures, *293-40*
BLAS_SYRK Procedures, *293-44*
BLAS_TBMV Procedures, *293-46*
BLAS_TBSV Procedures, *293-48*
BLAS_TPMV Procedures, *293-50*
BLAS_TPSV Procedures, *293-51*
BLAS_TRMM Procedures, *293-53*
BLAS_TRMV Procedures, *293-55*
BLAS_TRSM Procedures, *293-56*
BLAS_TRSV Procedures, *293-58*
BLOB_TABLE Table Type, *182-21*
BLOCK_HEAT_MAP Function, *90-2*
BLOCK_LIST Function, *160-9*
BLOCK_LIST_ADD Procedure, *160-9*
BLOCK_LIST_CLEAR Procedure, *160-6*
BLOCK_LIST_REMOVE, *160-10*
BLOCK_LIST_REMOVE Procedure, *160-6*
BLOCKQUOTECLOSE function
  of HTF package, *242-17*

BLOCKQUOTECLOSE procedure
of HTP package, 243-17
BLOCKQUOTEOPEN function
of HTF package, 242-17
BLOCKQUOTEOPEN procedure
of HTP package, 243-17
BODYCLOSE function
of HTF package, 242-18
BODYCLOSE procedure
of HTP package, 243-17
BODYOPEN function
of HTF package, 242-18
BODYOPEN procedure
of HTP package, 243-18
BOLD function
of HTF package, 242-19
BOLD procedure
of HTP package, 243-18
BOOLEAN_TABLE Table Type, 182-21
BR function
of HTF package, 242-19
BR procedure
of HTP package, 243-19
BREAKPOINT_INFO Record Type, 67-11
BROKEN procedure, 105-3
BUILD procedure, 44-12, 54-13
BUILD_CHAIN_ROWS_TABLE procedure, 103-1
BUILD_EXCEPTIONS_TABLE procedure, 103-2
BUILD_PART_INDEX procedure, 139-3
BUILD_SAFE_REWRITE_EQUIVALENCE
Procedure, 16-3
BYPASS Procedure, 160-3

C

CALENDARPRINT procedures, 252-3
CALIBRATE Function, 216-7
CALIBRATE_IO Procedure, 158-5
CAN_REDEF_COLLECTION procedure, 177-7
CAN_REDEF_TABLE procedure, 154-7
CANCEL_ADVISOR_TASK procedure, 192-18
CANCEL_ANALYSIS_TASK Procedure, 187-2
CANCEL_DIAGNOSIS_TASK Procedure, 186-7
CANCEL_EVOLVE_TASK Procedure, 181-10
CANCEL_MIGRATION Procedure, 29-3
CANCEL_REPLAY Procedure, 216-8
CANCEL_TASK Procedure, 17-9
CANCEL_TUNING_TASK procedure, 190-21
CANONICALIZE Procedure, 210-9
capture process
altering, 44-6
building a Streams data dictionary, 44-12
creating, 239-12, 239-24, 239-34, 239-43
DBMS_CAPTURE_ADM package, 44-1
instantiation
aborting database preparation, 44-3
capture process (continued)
instantiation (continued)
aborting schema preparation, 44-4
aborting table preparation, 44-5
preparing a database for, 44-27
preparing a schema for, 44-28
preparing a table for, 44-30
parameters
setting, 44-31
rules
defining global, 239-12
defining schema, 239-24
defining subset, 239-34
defining table, 239-43
removing, 239-78
starting, 44-46
stopping, 44-46
CAPTURE_CURSOR_CACHE procedure, 189-6
CAPTURE_CURSOR_CACHE_SQLSET
Procedure, 190-22
CAPTURE_DEPENDENCY_PRIVS Procedure,
147-3
CAST_FROM_BINARY_DOUBLE function, 294-6
CAST_FROM_BINARY_FLOAT function, 294-7
CAST_FROM_BINARY_INTEGER function,
294-8
CAST_FROM_NUMBER function, 294-8
CAST_TO_BINARY_DOUBLE function, 294-9
CAST_TO_BINARY_FLOAT function, 294-10
CAST_TO_BINARY_INTEGER function, 294-11
CAST_TO_NUMBER function, 294-11
CAST_TO_NVARCHAR2 function, 294-12
CAST_TO_RAW function, 294-13
CAST_TO_VARCHAR2 function, 294-13
CB$ table prefix, 54-2
CELLSPRINT procedures, 252-4
CENTER function
of HTF package, 242-20
CENTER procedure
of HTP package, 243-19
CENTERCLOSE function
of HTF package, 242-20
CENTERCLOSE procedure
of HTP package, 243-20
CENTEROPEN function
of HTF package, 242-20
CENTEROPEN procedure
of HTP package, 243-20
CFG_GET Function, 221-8
CFG_REFRESH Procedure, 221-8
CFG_UPDATE Procedure, 221-8
CHANGE functions and procedures, 249-5
CHANGE procedure, 105-4, 155-2
CHANGEOWNER Procedure, 223-5
CHANGEPRIVILEGES Function, 223-5
CHARARR Table Type, 136-5

CHECK_MIGRATED_MESSAGES Procedure, *29-4*
CHECK_MIGRATION_TO_TXEVENTQ Procedure, *29-4*
CHECK_OBJECT procedure, *156-5*
CHECK_PATCH_INSTALLED Function, *151-3*
CHECK_PLUG_COMPATIBILITY Function, *140-3*
CHECK_PRIVILEGE Function, *131-13*
CHECK_PRIVILEGE_ACLID Function, *131-14*
CHECK_STATUS Procedure, *29-5*
CHECKACCESS Function, *62-24*, *63-4*
CHECKIN function, *224-2*
CHECKOUT procedure, *224-3*
CHECKPRIVILEGES Function, *223-6*
CHECKSPI Functions and Procedures, *62-24*
CHECKSUM functions, *248-2*
CHOOSE_DATE procedure, *252-6*
CITE function
  of HTF package, *242-21*
CITE procedure
  of HTP package, *243-20*
CLEAN_AUDIT_TRAIL procedure, *31-8*
CLEANUP_GATEWAY Procedure, *125-23*
CLEANUP_GIDX Procedure, *138-2*
CLEANUP_GIDX_JOB Procedure, *138-2*
CLEANUP_ONLINE_OP Procedure, *138-3*
CLEANUPUNUSEDBACKUPFILES Procedure, *64-7*
CLEAR Procedure, *227-7*
CLEAR_ALL_CONTEXT Procedure, *173-3*
CLEAR_ALL_LOCAL_CONTEXTS Procedure, *173-3*
CLEAR_AUDIT_TRAIL_PROPERTY procedure, *31-11*
CLEAR_CONTEXT Procedure, *173-3*
CLEAR_COOKIES procedure, *286-31*
CLEAR_HEAT_MAP_ALL Procedur, *98-5*
CLEAR_HEAT_MAP_TABLE Procedure, *98-5*
CLEAR_IDENTIFIER Procedure, *173-4*
CLEAR_KEY_COLUMNS procedure, *24-10*
CLEAR_LAST_ARCHIVE_TIMESTAMP procedure, *31-12*
CLEAR_PENDING_AREA procedure, *158-7*
CLEAR_PLSQL_TRACE procedure, *199-5*
CLEAR_UNSUPPORTED_FEATURE_TABLE Procedure, *29-7*
CLEAR_USED procedure, *135-2*
CLEARKEYCOLUMNLIST procedure, *236-2*
CLEARUPDATECOLUMNLIST procedure, *236-2*
CLIENT_ID_STAT_DISABLE procedure, *128-2*
CLIENT_ID_STAT_ENABLE procedure, *128-2*
CLIENT_ID_TRACE_DISABLE procedure, *128-3*
CLIENT_ID_TRACE_ENABLE procedure, *128-3*
CLOB_TABLE Table Type, *182-21*
CLOB2FILE procedure, *114-14*
CLONEDB_RENAMEFILE Procedure, *75-1*
CLONENODE function, *228-21*
CLOSE Procedure, *114-15*
CLOSE_ALL_CONNECTIONS procedure, *299-6*
CLOSE_CONNECTION procedure, *299-7*
CLOSE_CONNECTION Procedure, *298-9*
CLOSE_CURSOR procedure, *182-35*
CLOSE_CURSOR Procedure, *95-8*
CLOSE_DATA Function and Procedure, *298-10*
CLOSE_DATABASE_LINK Procedure, *173-4*
CLOSE_ITERATOR procedure, *165-2*
CLOSE_PERSISTENT_CONN procedure, *286-32*
CLOSE_PERSISTENT_CONNS procedure, *286-32*
CLOSECONTEXT procedure, *229-2*, *236-3*
CODE function
  of HTF package, *242-21*
CODE procedure
  of HTP package, *243-21*
COLLECT_STATS procedure, *297-12*
collections
  table items, *182-27*
COLUMN_T Type, *197-23*
COLUMN_VALUE procedure, *182-36*
COLUMN_VALUE_LONG procedure, *182-39*
COLUMNS_NEW_T Type, *197-25*
COLUMNS_T Type, *197-25*
COLUMNS_WITH_TYPE_T Type, *197-25*
COMMA_TO_TABLE Procedures, *210-10*
COMMAND function and procedure, *298-10*
COMMAND_REPLIES function, *298-11*
COMMENT function
  of HTF package, *242-22*
COMMENT procedure
  of HTP package, *243-21*
COMMIT procedure, *200-3*
COMMIT_ACL procedure, *174-3*
COMMIT_COMMENT procedure, *200-3*
COMMIT_FORCE procedure, *200-3*
COMMIT_MIGRATION Procedure, *29-8*
COMMIT_TRANSACTION Procedure, *106-9*
COMPARE function, *294-14*
COMPARE Function, *47-11*
COMPARE Functions, *114-16*
COMPARE_CAPTURE_REPLAY_REPORT Function, *15-7*
COMPARE_DATABASES Function, *15-8*
COMPARE_INSTANCES Function, *15-9*
COMPARE_OLD_VALUES procedure, *24-11*
COMPARE_PERIOD_REPORT Procedure, *216-9*
COMPARE_PLANS Function, *237-7*
COMPARE_REPLAY_REPLAY_REPORT Function, *15-10*
COMPARE_SQLSET_REPORT Function, *216-9*
COMPARISON_TYPE type, *47-9*, *47-10*
COMPATIBLE_10_1 function, *310-34*
COMPATIBLE_10_2 function, *310-34*

COMPATIBLE_9_2 function, *310-34*
COMPILE_FROM_REMOTE procedure, *113-2*
COMPILE_SCHEMA Procedure, *210-10*
COMPILESCHEMA procedure, *232-5*
CONCAT function, *294-15*
CONCATENATE_SUBPROGRAM Function, *281-6*
CONFIGURE Procedure, *14-2*, *36-2*, *39-4*, *181-10*
CONFIGURE_ADVISOR_FILTER function, *192-20*
CONFIGURE_ADVISOR_OBJ_FILTER function, *192-25*
CONFIGURE_ADVISOR_OPR_FILTER function, *192-26*
CONFIGURE_ADVISOR_RULE_FILTER function, *192-29*
CONFIGURE_NODE procedure, *206-2*
CONFIGURE_POOL Procedure, *49-5*
conflicts
  detection
    stopping, *24-11*
CONFORMING Function, *234-1*
CONNECT_TCP Procedure, *68-2*
constants
  DBMS_HIERARCHY package, *91-1*
  DBMS_ILM_ADMIN package, *98-1*
  DBMS_MGWMSG package, *126-1*
CONTAINS_HOST Function, *132-3*
CONTINUE function, *67-16*
CONTROL_RESTRICTED_SNAPSHOT procedure, *217-26*
CONVERGE Procedure, *47-12*
CONVERT function, *115-6*, *229-3*, *294-16*
CONVERT_LONG_TO_LOB_CHUNK member procedure, *310-10*
CONVERT_RAW_VALUE procedures, *192-30*
CONVERT_RAW_VALUE_NVARCHAR procedure, *192-31*
CONVERT_RAW_VALUE_ROWID procedure, *192-32*
CONVERTTOBLOB procedure, *114-18*
CONVERTTOCLOB Procedure, *114-20*
COPIES function, *294-17*
COPY Procedures, *114-23*
COPY_COLLECTION_DEPENDENTS procedure, *177-7*
COPY_DBFS_LINK Procedures, *114-24*
COPY_FILE procedure, *81-2*
COPY_FROM_DBFS_LINK, *114-25*
COPY_TABLE_DEPENDENTS procedure, *154-7*
COPY_TABLE_STATS Procedure, *192-32*
COPYEVOLVE procedure, *232-5*
CQ_NOTIFICATION_QUERYID Function, *50-17*
CQ_NOTIFICATION$_DESCRIPTOR Object Type, *50-11*
CQ_NOTIFICATION$_QUERY Object Type, *50-11*
CQ_NOTIFICATION$_QUERY_ARRAY Object (Array) Type, *50-12*
CQ_NOTIFICATION$_REG_INFO Object Type, *50-14*
CQ_NOTIFICATION$_ROW Object Type, *50-13*
CQ_NOTIFICATION$_TABLE Object Type, *50-12*
CQ_NOTIFICATION$_TABLE_ARRAY Object (Array) Type, *50-13*
CREATE ANY DIMENSION privilege, *54-2*
CREATE ANY MATERIALIZED VIEW privilege, *54-2*
CREATE DIMENSION privilege, *54-2*
CREATE MATERIALIZED VIEW privilege, *54-2*
CREATE PACKAGE BODY command, *1-2*
CREATE PACKAGE command, *1-2*
CREATE SESSION privilege, *54-2*
CREATE_ACL Procedure, *131-15*
CREATE_ADVISOR_TASK function, *192-34*
CREATE_AFFECTED_TABLE Procedure, *76-4*
CREATE_AGENT Procedure, *125-25*
CREATE_ALTER_TYPE_ERROR_TABLE Procedure, *210-11*
CREATE_ANALYSIS_TASK Functions, *187-3*
CREATE_APPLY procedure, *24-13*
CREATE_BASELINE Function and Procedure, *217-26*
CREATE_BASELINE_TEMPLATE Procedures, *217-28*
CREATE_BIN_CAT procedure, *58-17*
CREATE_BIN_NUM procedure, *58-18*
CREATE_CAPTURE procedure
  capture process
    creating, *44-13*
CREATE_CAPTURE Procedure, *147-3*
CREATE_CAPTURE procedure (of DBMS_SQL_FIREWALL), *183-8*
CREATE_CATEGORY function, *124-7*
CREATE_CATEGORY Procedure, *158-7*
CREATE_CDB_PLAN Procedure, *158-8*
CREATE_CDB_PLAN_DIRECTIVE Procedure, *158-8*
CREATE_CDB_PROFILE_DIRECTIVE Procedure, *158-10*
CREATE_CHUNKS_BY_NUMBER_COL Procedure, *137-9*
CREATE_CHUNKS_BY_ROWID Procedure, *137-9*
CREATE_CHUNKS_BY_SQL Procedure, *137-10*
CREATE_CLIP procedure, *58-20*
CREATE_COL_REM procedure, *58-21*
CREATE_COMPARISON Procedure, *47-13*
CREATE_CONSUMER_GROUP procedure, *158-11*

CREATE_COVERAGE_TABLES Procedure, *144-3*
CREATE_CREDENTIAL Procedure, *51-3*
CREATE_DAD Procedure, *78-5*
CREATE_DIAGNOSIS_TASK Functions, *186-8*
CREATE_ERROR_LOG Procedure, *79-1*
CREATE_ERROR_TABLE Procedure, *76-4*
CREATE_EVALUATION_CONTEXT procedure, *166-8*
CREATE_EVOLVE_TASK Function, *181-13*
CREATE_EXTDDL_FOR_HIVE Procedure, *89-2*
CREATE_EXTENDED_STATS Function, *192-35*
CREATE_FILE Procedure, *17-11*
CREATE_FILTER_SET Procedure, *216-10*
CREATE_HYBRID_PARTNED_TABLE Procedure, *89-3*
CREATE_INBOUND procedure, *239-53*
CREATE_INDEX_COST procedure, *178-5*
CREATE_INTERIM_COLLECTION procedure, *177-8*
CREATE_JOB Procedure, *125-27*
CREATE_LINK procedure, *206-2*
CREATE_MISS_CAT procedure, *58-23*
CREATE_MISS_NUM procedure, *58-24*
CREATE_MSGSYSTEM_LINK Procedure for TIB/Rendezvous, *125-30*
CREATE_MSGSYSTEM_LINK Procedure for WebSphere MQ, *125-31*
CREATE_MVIEW function, *54-5*, *54-29*
CREATE_NORM_LIN procedure, *58-25*
CREATE_OBJECT Procedure, *17-12*
CREATE_OBJECT_DEPENDENCY procedure, *24-19*
CREATE_OR_REPLACE_VIEW procedure
  of DBMS_HS_PARALLEL package, *94-1*
CREATE_OUTBOUND procedure, *239-55*
CREATE_OUTLINE procedure, *135-3*
CREATE_PENDING_AREA procedure, *158-11*
CREATE_PIPE function, *143-15*
CREATE_PLAN procedure, *158-13*
CREATE_PLAN_DIRECTIVE procedure, *158-14*
CREATE_POLICY_EXPRESSION Procedure, *153-21*
CREATE_POLICY_GROUP Procedure, *162-13*
CREATE_PROFILE Procedure, *185-5*
CREATE_PROPAGATION procedure, *150-3*
CREATE_PURGE_JOB procedure, *31-14*
CREATE_REMOTE_SNAPSHOT function and procedure, *217-29*
CREATE_REQUEST_CONTEXT Function, *286-33*
CREATE_RULE procedure, *166-9*
CREATE_RULE_SET procedure, *166-11*
CREATE_SERVICE procedure, *172-7*
CREATE_SIMPLE_PLAN procedure, *158-18*
CREATE_SNAPSHOT function and procedure, *217-30*
CREATE_SNAPSHOT Procedure, *14-3*
CREATE_SQL_PATCH Function, *186-9*
CREATE_SQL_PLAN_BASELINE Procedure, *190-24*
CREATE_SQLSET procedure, *189-8*, *190-25*
CREATE_SQLWKLD Procedure, *17-14*
CREATE_STAT_TABLE procedure, *192-37*
CREATE_STGTAB procedure, *189-9*
CREATE_STGTAB_BASELINE Procedure, *181-14*
CREATE_STGTAB_DIRECTIVE Procedure, *180-3*
CREATE_STGTAB_SQLPATCH Procedure, *186-10*
CREATE_STGTAB_SQLPROF Procedure, *190-25*
CREATE_STGTAB_SQLSET Procedure, *190-26*
CREATE_SYNC_CAPTURE procedure
  synchronous capture
    creating, *44-23*
CREATE_TABLE_COST procedures, *178-5*
CREATE_TABLE_COST_COLINFO Object Type, *178-1*
CREATE_TABLE_TEMPLAGE procedure
  of DBMS_HS_PARALLEL package, *94-3*
CREATE_TABLES Function, *93-4*
CREATE_TASK Procedure, *137-8*
CREATE_TASK Procedures, *17-15*
CREATE_TEMP_HISTORY_TABLE Procedure, *83-4*
CREATE_TOPOLOGY procedure, *206-3*
CREATE_TRANSFORMATION procedure, *201-1*
CREATE_TRIGGER_TABLE Procedure, *76-5*
CREATE_TUNING_TASK functions, *190-28*
CREATE_VALIDATE_LOG_TABLE procedure
  of DBMS_HIERARCHY package, *91-3*
CREATE_VIEW procedure, *107-6*
CREATE_VIEW_FOR_FACT_ROWS Procedure
  of DBMS_HIERARCHY package, *91-4*
CREATE_VIEW_FOR_STAR_ROWS Procedure
  of DBMS_HIERARCHY package, *91-5*
CREATE_VIEW_ON_PATH procedure, *107-7*
CREATE_WRAPPED Procedure, *66-4*
CREATEATTRIBUTE function, *228-22*
CREATEBUCKET Procedure, *64-7*
CREATECDATASECTION function, *228-23*
CREATECOMMENT function, *228-23*
CREATEDATEINDEX Procedure, *230-2*
CREATEDIRECTORY Procedures, *62-25*, *63-5*
CREATEDOCUMENT function, *228-24*
CREATEDOCUMENTFRAGMENT function, *228-24*
CREATEELEMENT function, *228-25*
CREATEENTITYREFERENCE function, *228-26*

CREATEFILE Procedures, *62-26*, *63-6*
CREATEFILESYSTEM Procedure, *65-3*
CREATEFOLDER Function, *223-6*
CREATELINK Procedures, *62-27*, *63-6*
CREATENONCEKEY Procedure, *220-1*, *226-2*
CREATENONSCHEMABASEDXML function, *315-2*
CREATENUMBERINDEX Procedure, *230-3*
CREATEOIDPATH Function, *223-7*
CREATEPROCESSINGINSTRUCTION function, *228-26*
CREATEREFERENCE Procedures, *62-28*, *63-7*
CREATERESOURCE Functions, *223-7*
CREATESCHEMABASEDXML function, *315-2*
CREATESTORE Procedure, *64-8*, *65-4*
CREATETEMPORARY Procedures, *114-26*
CREATETEXTNODE function, *228-27*
CREATEURI function, *307-5*, *307-9*, *307-12*
CREATEXML function, *315-3*
creating
  packages, *1-2*
CSTORE_EXISTS function, *197-29*
CSTORE_GET procedure, *197-29*
CTX_ADM package, *3-1*
CTX_ANL package, *4-1*
CTX_CLS package, *6-1*
CTX_DDL package, *5-1*
CTX_DOC package, *7-1*
CTX_ENTITY package, *8-1*
CTX_OUTPUT package, *9-1*
CTX_QUERY package, *10-1*
CTX_REPORT package, *11-1*
CTX_THES package, *12-1*
CTX_ULEXER package, *13-1*
cube aggregation, *54-13*
Cube Build log, *56-4*
Cube Dimension Compile log, *56-5*
cube log verbosity levels, *56-3*
cube logging targets, *56-2*
cube logging types, *56-2*
cube maintenance, *54-1*
cube materialized views, *54-2*
  optimizing, *55-1*
Cube Operations log, *56-6*
cube refresh, *54-13*
Cube Rejected Records Log, *56-7*
CUBE_BUILD_LATEST view, *54-21*
CUBE_UPGRADE_INFO table, *54-43*
cubes
  optimizer statistics, *41-1*
CURRENT_EDITION Function, *281-5*
CURRENT_INSTANCE Function, *210-12*
CURRENT_IS_ROLE_ENABLED Function, *173-5*
CURRENT_SESSION_ID Function, *68-4*
cursors
  DBMS_SQL package, *182-5*
CUSTOMIZE_ILM Procedure, *98-6*

D

data dictionary
  removing XStream information, *239-74*
DATA function and procedure, *298-12*
data mining, *57-113*
  data transformation, *57-156*, *57-159*
  PMML, *57-155*
DATA_BLOCK_ADDRESS_BLOCK Function, *210-12*
DATA_BLOCK_ADDRESS_FILE Function, *210-13*
database
  locking
    OWA_OPT_LOCK package, *248-1*
database dictionary tables, *72-1*
database tables
  creating for DBMS_TRACE, *199-3*
DATABASE_TRACE_DISABLE Procedure, *128-4*
DATABASE_TRACE_ENABLE Procedure, *128-4*
datatypes
  DBMS_DESCRIBE, *70-7*
  PL/SQL
    numeric codes for, *70-4*
DATE_TABLE Table Type, *182-21*
DB_CONNECT_INFO Procedure, *125-32*
DB_VERSION Procedure, *210-13*
DBFS_LINK_GENERATE_PATH Functions, *114-27*
DBLINK_ARRAY Table Type, *210-3*
DBMS, *228-20*
DBMS_ALERT package, *19-1*
DBMS_APP_CONT_ADMIN package, *21-1*
DBMS_APP_CONT_REPORT package, *22-1*
DBMS_APPLY_ADM package, *24-1*
DBMS_AQADM Constants, *26-2*
DBMS_AQELM package, *27-1*
DBMS_AQIN package, *28-1*
DBMS_AUDIT_MGMT package
  ALTER_PARTITION_INTERVAL procedure, *31-8*
  CLEAN_AUDIT_TRAIL procedure, *31-8*
  CLEAR_AUDIT_TRAIL_PROPERTY procedure, *31-11*
  CLEAR_LAST_ARCHIVE_TIMESTAMP procedure, *31-12*
  CREATE_PURGE_JOB procedure, *31-14*
  DEINIT_CLEANUP procedure, *31-15*
  DROP_PURGE_JOB procedure, *31-17*
  GET_AUDIT_COMMIT_DELAY function, *31-17*
  INIT_CLEANUP procedure, *31-20*
  IS_CLEANUP_INITIALIZED function, *31-22*

DBMS_AUDIT_MGMT package *(continued)*
SET_AUDIT_TRAIL_LOCATION procedure, *31-24*
SET_AUDIT_TRAIL_PROPERTY procedure, *31-26*
SET_LAST_ARCHIVE_TIMESTAMP procedure, *31-28*
SET_PURGE_JOB_INTERVAL procedure, *31-30*
SET_PURGE_JOB_STATUS procedure, *31-31*
DBMS_AUDIT_UTIL package
DECODE_RLS_INFO_ATRAIL_FGA function, *32-2*
DECODE_RLS_INFO_ATRAIL_STD function, *32-3*
DECODE_RLS_INFO_ATRAIL_UNI function, *32-4*
DECODE_RLS_INFO_ATRAIL_XML function, *32-5*
overview, *32-1*
security model, *32-1*
views used by, *32-1*
DBMS_AUTO_CLUSTERING package
APPLY_RECOMMENDATION procedure, *33-6*
CONFIGURE procedure, *33-2*
GET_RECOMMENDATION function, *33-7*
RECOMMEND_CLUSTERING_METHOD Function, *33-4*
REPORT ACTIVITY function, *33-7*
REPORT_LAST_ACTIVITY function, *33-8*
VERIFY_RECOMMENDATION procedure, *33-5*
DBMS_AUTO_INDEX package, *35-1*
CONFIGURE procedure, *35-1*
DROP_AUTO_INDEXES Procedure, *35-6*
DROP_SECONDARY_INDEXES function, *35-7*
REPORT_ACTIVITY function, *35-8*
REPORT_LAST_ACTIVITY function, *35-10*
DBMS_AUTO_MV package
RECOMMEND function, *36-5*
REPORT_ACTIVITY function, *36-7*
REPORT_LAST_ACTIVITY function, *36-8*
DBMS_AUTO_TASK_ADMIN package, *38-1*
DBMS_AUTO_ZONEMAP package, *39-1*
DBMS_AVTUNE package, *40-1*
DBMS_BLOCKCHAIN_TABLE package, *42-1*
ADD_INTERVAL_PARTITIONING Procedure(), *42-4*
COUNTERSIGN_ROW Procedure(), *42-5*
COUNTERSIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure(), *42-6*
DELETE_EXPIRED_ROWS Procedure(), *42-7*
GET_BLOCKCHAIN_DIGEST Function(), *42-8*
GET_BLOCKCHAIN_DIGEST_FOR_SELECTED_ROWS Function(), *42-9*
GET_BYTES_FOR_ROW_HASH Procedure(), *42-10*
DBMS_BLOCKCHAIN_TABLE package *(continued)*
GET_BYTES_FOR_ROW_HASH_SPECIFIED_BY_KEY_C
Procedure(), *42-12*
GET_BYTES_FOR_ROW_SIGNATURE Procedure(), *42-13*
GET_BYTES_FOR_ROW_SIGNATURE_SPECIFIED_BY_K
Procedure(), *42-14*
GET_SIGNED_BLOCKCHAIN_DIGEST Function(), *42-15*
GET_SIGNED_BLOCKCHAIN_DIGEST_FOR_SELECTED_
42-16
overview, *42-1*
Security Model, *42-2*
SIGN_ROW Procedure(), *42-18*
SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS Procedure(
SIGN_ROW_SPECIFIED_BY_KEY_COLUMNS_WITH_CO
Procedure(), *42-21*
SIGN_ROW_WITH_COUNTERSIGNATURE Procedure(), *4*
VERIFY_ROWS Procedure(), *42-24*
VERIFY_TABLE_BLOCKCHAIN Procedure(), *42-25*
VERIFY_USER_BLOCKCHAIN_ROWS Procedure(), *42-26*
DBMS_CAPTURE package, *305-1*
DBMS_CAPTURE_ADM package, *44-1*
DBMS_CHANGE_NOTIFICATION package, *50-1*
DBMS_CLOUD, *46-1*
DBMS_COMPARISON package, *47-1*
constants, *47-3*
data structures, *47-9*
index columns, *47-5*
lead index column, *47-5*
requirements, *47-5*
subprograms, *47-10*
views, *47-5*
DBMS_CONNECTION_POOL package, *49-1*
DBMS_CONTENT_CONTEXT_T Object Type, *306-2*
DBMS_CSX_ADMIN package, *53-1*
DBMS_CUBE package, *54-1*
DBMS_CUBE_ADVISE package, *55-1*
DBMS_CUBE_LOG package, *56-1*
DBMS_DATA_MINING
ESA, *57-25*
ESM Settings, *57-27*
Exponential Smoothing, *57-27*
DBMS_DATA_MINING package, *57-1*, *57-113*
ADD_COST_MATRIX procedure, *57-67*
algorithm specific settings
machine learning models, *57-18*
algorithms, *57-1*
ALTER_REVERSE_EXPRESSION procedure, *57-71*
APPLY procedure, *57-74*
Automatic Data Preparation, *57-6*
COMPUTE_CONFUSION_MATRIX procedure, *57-77*
COMPUTE_CONFUSION_MATRIX_PART procedure, *57-83*
COMPUTE_LIFT procedure, *57-89*

DBMS_DATA_MINING package *(continued)*
COMPUTE_LIFT_PART procedure, *57-94*
COMPUTE_ROC procedure, *57-99*
COMPUTE_ROC_PART procedure, *57-103*
Create Model Using Registration Information, *57-113*
CREATE_MODEL procedure, *57-108*
CREATE_MODEL2 procedure, *57-112*
data transformation, *57-108*, *57-159*
data types, *57-59*
deprecated types, *57-60*
DROP_ALGORITHM procedure, *57-114*
DROP_MODEL procedure, *57-115*
DROP_PARTITION procedure, *57-115*
EXPORT_MODEL procedure, *57-116*
EXPORT_SERMODEL procedure, *57-119*
FETCH_JSON_SCHEMA procedure, *57-120*
GET_ASSOCIATION_RULES function, *57-120*
GET_FREQUENT_ITEMSETS function, *57-125*
GET_MODEL_COST_MATRIX function, *57-127*
GET_MODEL_DETAILS_AI function, *57-128*
GET_MODEL_DETAILS_EM function, *57-130*
GET_MODEL_DETAILS_EM_COMP function, *57-131*
GET_MODEL_DETAILS_EM_PROJ function, *57-133*
GET_MODEL_DETAILS_GLM function, *57-135*
GET_MODEL_DETAILS_GLOBAL function, *57-138*
GET_MODEL_DETAILS_KM function, *57-140*
GET_MODEL_DETAILS_NB function, *57-142*
GET_MODEL_DETAILS_NMF function, *57-144*
GET_MODEL_DETAILS_OC function, *57-145*
GET_MODEL_DETAILS_SVM function, *57-150*, *57-152*
GET_MODEL_DETAILS_XML function, *57-154*
GET_MODEL_SETTINGS function, *57-147*
GET_MODEL_SIGNATURE function, *57-148*
GET_MODEL_TRANSFORMATIONS function, *57-156*
GET_TRANSFORM_LIST procedure, *57-159*
IMPORT_MODEL procedure, *57-162*
IMPORT_ONNX_MODEL procedure, *57-166*
IMPORT_SERMODEL procedure, *57-167*
JSON schema for R extensible algorithm, *57-168*
overview, *57-1*
PMML, *57-155*
RANK_APPLY procedure, *57-173*
REGISTER_ALGORITHM procedure, *57-172*
DBMS_DATA_MINING package *(continued)*
REMOVE_COST_MATRIX procedure, *57-176*
RENAME_MODEL procedure, *57-177*
scoring, *57-74*, *57-82*, *57-88*, *57-92*, *57-97*, *57-172*, *57-176*
settings
ADMM, *57-58*
algorithm names, *57-5*
CUR Matrix Decomposition, *57-20*
Decision Tree, *57-20*
GLM, *57-32*
global, *57-13*
k-Means, *57-35*
LBFGS, *57-59*
machine learning functions, *57-7*
machine learning models, *57-5*
MSET-SPRT, *57-37*
Naive Bayes, *57-38*
neural network, *57-39*
NMF, *57-42*
O-Cluster, *57-42*
R model extensibility, *57-18*
Random Forest, *57-43*
solver, *57-57*
SVM, *57-45*
XGBoost, *57-47*
subprograms, *57-65*
transactional data, *57-15*
DBMS_DATA_MINING Package
ADD_PARTITION procedure, *57-70*
DBMS_DATA_MINING_TRANSFORM
datatypes, *58-13*
package, *58-1*
subprograms, *58-15*
DBMS_DATAPUMP package, *59-1*
ADD_FILE procedure, *59-13*
ATTACH function, *59-21*
DATA_FILTER procedure, *59-22*
DETACH procedure, *59-26*
GET_DUMPFILE_INFO procedure, *59-27*
GET_STATUS procedure, *59-30*
LOG_ENTRY procedure, *59-32*
METADATA_FILTER procedure, *59-33*
METADATA_REMAP procedure, *59-37*
METADATA_TRANSFORM procedure, *59-39*
OPEN function, *59-44*
roles used by, *59-1*
SET_PARALLEL procedure, *59-47*
SET_PARAMETER procedure, *59-48*
START_JOB procedure, *59-57*
STOP_JOB procedure, *59-58*
types used by, *59-3*
WAIT_FOR_JOB procedure, *59-60*
DBMS_DB_VERSION package, *60-1*
DBMS_DBFS_CONTENT package, *62-1*

DBMS_DBFS_CONTENT_LIST_ITEM_T Object Type, *306-2*
DBMS_DBFS_CONTENT_LIST_ITEMS_T Table Type, *306-3*
DBMS_DBFS_CONTENT_PROPERTIES_T Table Type, *306-4*
DBMS_DBFS_CONTENT_PROPERTY_T Object Type, *306-3*
DBMS_DBFS_CONTENT_RAW_T Table Type, *306-4*
DBMS_DBFS_CONTENT_SPI package, *63-1*
DBMS_DBFS_HS package, *64-1*
DBMS_DBFS_SFS package, *65-1*
DBMS_DDL package, *66-1*
DBMS_DEBUG_JDWP package, *68-1*
DBMS_DESCRIBE package, *70-1*
DBMS_DG package, *71-1*
  using, *71-1*
DBMS_DG Package
  INITIATE_FS_FAILOVER procedure, *71-2*
dbms_dictionary_check
  CboHiLo, *72-8*
  CheckDual, *72-8*
  ChkIotTs, *72-8*
  Critical, *72-8*
  DbmsSchemaCopy, *72-8*
  DictOwnNonDefaultSYSTEM, *72-8*
  DroppedROTS, *72-8*
  DuplicateBlockUse, *72-8*
  DuplicateDataobj, *72-8*
  ExtentlessSeg, *72-8*
  FetUet, *72-8*
  FilBlkZero, *72-8*
  Full, *72-8*
  IcolDepCnt, *72-8*
  IdnseqObj, *72-8*
  IdnseqSeq, *72-8*
  IndexPartitionSeg, *72-8*
  IndexSeg, *72-8*
  IndexSubPartitionSeg, *72-8*
  IndIndparMismatch, *72-8*
  IndPartObj, *72-8*
  InvalidTSMaxSCN, *72-8*
  InvCorrAudit, *72-8*
  LobNotInObj, *72-8*
  MaxControlfSeq, *72-8*
  Mview, *72-8*
  NextObject, *72-8*
  NoSegmentIndex, *72-8*
  ObjectNames, *72-8*
  ObjError, *72-8*
  ObjIndDobj, *72-8*
  ObjNotLob, *72-8*
  ObjNotTrigger, *72-8*
  OBJRecycleBin, *72-8*
  ObjSeq, *72-8*
dbms_dictionary_check *(continued)*
  ObjSyn, *72-8*
  ObjType0, *72-8*
  OIDOnObjCol, *72-8*
  OversizedFiles, *72-8*
  PartCol, *72-8*
  PartSubPartMismatch, *72-8*
  PoorDefaultStorage, *72-8*
  PoorStorage, *72-8*
  PublicObjects, *72-8*
  SegFreelist, *72-8*
  SeglessUET, *72-8*
  SegNotInDeferredStg, *72-8*
  SourceNotInObj, *72-8*
  StmtAuditOnCommit, *72-8*
  SystemNotRfile1, *72-8*
  TabComPartObj, *72-8*
  TablePartitionSeg, *72-8*
  TableSeg, *72-8*
  TableSubPartitionSeg, *72-8*
  TabPartCountMismatch, *72-8*
  TrgAfterUpgrade, *72-8*
  Uet0Check, *72-8*
  UndoSeg, *72-8*
  ValidateSeg, *72-8*
  ValidateTrigger, *72-8*
  ValidDepends, *72-8*
  ValidDir, *72-8*
  ValidInd, *72-8*
  ValidOwner, *72-8*
  ValidTab, *72-8*
  WarnIcolDep, *72-8*
  WarningTSMaxSCN, *72-8*
DBMS_DIMENSION package, *73-1*
DBMS_DISTRIBUTED_TRUST_ADMIN package, *74-1*
DBMS_EDITIONS_UTILITIES package, *77-1*
  CLEAN_UNUSABLE_EDITIONS Procedure, *77-2*
DBMS_EPG package, *78-1*
DBMS_FILE_TRANSFER package, *81-1*
DBMS_FLASHBACK package, *82-1*
  types, *82-2*
DBMS_FLASHBACK_ARCHIVE_MIGRATE Examples, *84-3*
DBMS_FLASHBACK_ARCHIVE_MIGRATE Operational Notes, *84-2*
DBMS_FLASHBACK_ARCHIVE_MIGRATE Overview, *84-1*
DBMS_FLASHBACK_ARCHIVE_MIGRATE package, *84-1*
DBMS_FLASHBACK_ARCHIVE_MIGRATE Security Model, *84-1*
DBMS_FS
  UNMOUNT_ORACLE_FS, *86-9*

DBMS_FS pacakge
overview, 86-1
DBMS_FS package, 86-1
DESTROY_ORACLE_FS procedure, 86-2
FS_EXISTS procedure, 86-3
LIST_FILES procedure, 86-3
MAKE_ORACLE_FS procedure, 86-4
MOUNT_ORACLE_FS procedure, 86-5
security model, 86-1
DBMS_GOLDENGATE_ADM package, 87-1
DBMS_GOLDENGATE_AUTH package, 88-1
DBMS_HADOOP package, 89-1
dbms_hcheck
critical, 72-1
full, 72-1
DBMS_HIERARCHY package, 91-1
overview, 91-1
security model, 91-1
DBMS_HPROF package, 93-1
DBMS_HS_PARALLEL package, 94-1
CREATE_OR_REPLACE_VIEW procedure, 94-1
CREATE_TABLE_TEMPLATE procedure, 94-3
DROP_VIEW procedure, 94-3
LOAD_TABLE procedure, 94-4
DBMS_HS_PASSTHROUGH package, 95-1
DBMS_ILM_ADMIN package
constants, 98-1
DBMS_IMMUTABLE_TABLE package, 99-1
ADD_INTERVAL_PARTITIONING Procedure(), 99-2
DELETE_EXPIRED_ROWS Procedure(), 99-2
overview, 99-1
Security Model, 99-1
DBMS_INMEMORY package
POPULATE procedure, 100-3
REPOPULATE procedure, 100-4
security model, 100-2
SEGMENT_DEALLOCATE_VERSIONS procedure, 100-4
DBMS_INMEMORY_ADMIN package, 101-1
FASTSTART_ENABLE procedure, 101-7
FASTSTART_MIGRATE_STORAGE procedure, 101-9
GET_FASTSTART_TABLESPACE function, 101-10
IME_CAPTURE_EXPRESSIONS procedure, 101-11
IME_CLOSE_CAPTURE_WINDOW procedure, 101-13
IME_DROP_ALL_EXPRESSIONS procedure, 101-14
IME_OPEN_CAPTURE_WINDOW procedure, 101-15
DBMS_INMEMORY_ADMIN package *(continued)*
IME_POPULATE_EXPRESSIONS procedure, 101-16
operational notes, 101-3
overview, 101-1
POPULATE_WAIT function, 101-16
security, 101-3
DBMS_IOT package, 103-1
DBMS_JAVA package, 104-1
DBMS_JSON constants, 107-1
DBMS_JSON Package, 107-1, 107-4
DBMS_JSON_DUALITY, 106-2
DBMS_JSON_DUALITY Package, 106-2
DBMS_JSON_SCHEMA constants, 108-1
DBMS_JSON_SCHEMA Package, 108-1, 108-2
DBMS_LOB package, 114-1
CLOB2FILE procedure, 114-14
DBMS_LOCK package, 115-1
DBMS_LOGMNR package
ADD_LOGFILE procedure, 116-5
COLUMN_PRESENT function, 116-7
END_LOGMNR procedure, 116-8
MINE_VALUE function, 116-9
REMOVE_LOGFILE procedure, 116-10
START_LOGMNR procedure, 116-11
DBMS_LOGMNR_D package
BUILD procedure, 118-2
SET_TABLESPACE procedure, 118-5
DBMS_LOGSTDBY package, 117-1
APPLY_SET procedure, 117-4
APPLY_UNSET procedure, 117-7
BUILD procedure, 117-7
INSTANTIATE_TABLE procedure, 117-8
MAP_PRIMARY_SCN function, 117-10
overview of managing SQL Apply, 117-1
PREPARE_FOR_NEW_PRIMARY procedure, 117-11
PURGE_SESSION procedure, 117-12
SET_TABLESPACE procedure, 117-13
SKIP procedure, 117-14
SKIP_ERROR procedure, 117-23
SKIP_TRANSACTION procedure, 117-26
UNSKIP procedure, 117-28
UNSKIP_ERROR procedure, 117-29
UNSKIP_TRANSACTION procedure, 117-30
DBMS_MEMOPTIMIZE package, 120-1
DBMS_MEMOPTIMIZE_ADMIN package, 121-1
DBMS_METADATA package
ADD_TRANSFORM function, 122-6
CLOSE procedure, 122-10
CONVERT functions and procedures, 122-10
GET_DDL function, 122-15
GET_QUERY function, 122-19
GET_XML function, 122-15
OPEN function, 122-20
OPENW function, 122-27

DBMS_METADATA package *(continued)*
PUT function, *122-28*
security, *122-2*
SET_COUNT procedure, *122-29*
SET_FILTER procedure, *122-30*
SET_PARSE_ITEM procedure, *122-41*
SET_REMAP_PARAM procedure, *122-44*
SET_TRANSFORM_PARAM procedure, *122-44*
DBMS_METADATA_DIFF package
ADD_DOCUMENT procedure, *123-4*
FETCH_CLOB functions and procedures, *123-5*
OPENC function, *123-4*
DBMS_MGD_ID_UTL package, *124-1*
ADD_SCHEME procedure, *124-4*
CREATE_CATEGORY function, *124-7*
EPC_TO_ORACLE_SCHEME function, *124-8*
GET_CATEGORY_ID function, *124-11*
GET_COMPONENTS function, *124-12*
GET_ENCODINGS function, *124-13*
GET_JAVA_LOGGING_LEVEL function, *124-13*
GET_PLSQL_LOGGING_LEVEL function, *124-14*
GET_SCHEMENAMES function, *124-15*
GET_TDT_XML function, *124-15*
GET_VALIDATOR function, *124-17*
REFRESH_CATEGORY function, *124-20*
REMOVE_CATEGORY procedure, *124-21*
REMOVE_PROXY procedure, *124-22*
REMOVE_SCHEME procedure, *124-22*
SET_JAVA_LOGGING_LEVEL procedure, *124-23*
SET_PLSQL_LOGGING_LEVEL procedure, *124-23*
SET_PROXY procedure, *124-24*
VALIDATE_SCHEME function, *124-25*
DBMS_MGWADM package
summary of subprograms, *125-10*
DBMS_MGWMSG package
constants, *126-1*
summary of subprograms, *126-17*
DBMS_MLE package
CREATE_CONTEXT function, *127-3*
DISABLE_CTX_STDERR procedure, *127-4*
DISABLE_CTX_STDOUT procedure, *127-5*
DISABLE_DEBUGGING procedure, *127-5*
DISABLE_ICS_STDERR procedure, *127-6*
DISABLE_ICS_STDOUT procedure, *127-6*
DISABLE_STDERR procedure, *127-7*
DISABLE_STDOUT procedure, *127-7*
DROP_CONTEXT procedure, *127-8*
ENABLE_DEBUGGING procedure, *127-9*
EVAL procedure, *127-9*
EXPORT_TO_MLE procedure, *127-11*
DBMS_MLE package *(continued)*
GET_AVAILABLE_LANGUAGES Function, *127-13*
GET_CTX_ERROR_STACK Function, *127-13*
GET_ERROR_STACK Function, *127-14*
IMPORT_FROM_MLE procedure, *127-14*
PARSE_DEBUG_OUTPUT function, *127-16*
SET_CTX_STDERR procedure, *127-17*
SET_CTX_STDERR_TO_DBMS_OUTPUT procedure, *127-18*
SET_CTX_STDOUT procedure, *127-18*
SET_CTX_STDOUT_TO_DBMS_OUTPUT procedure, *127-19*
SET_ICS_STDERR procedure, *127-19*
SET_ICS_STDERR_TO_DBMS_OUTPUT procedure, *127-20*
SET_ICS_STDOUT procedure, *127-21*
SET_ICS_STDOUT_TO_DBMS_OUTPUT procedure, *127-21*
SET_STDERR procedure, *127-22*
SET_STDERR_TO_DBMS_OUTPUT procedure, *127-23*
SET_STDOUT procedure, *127-23*
SET_STDOUT_TO_DBMS_OUTPUT procedure, *127-24*
DBMS_MONITOR package
statistics tracing and gathering
DBMS_MONITOR package, *128-1*
DBMS_MVIEW package
BEGIN_TABLE_REORGANIZATION procedure, *129-2*
END_TABLE_REORGANIZATION procedure, *129-3*
EXPLAIN_MVIEW procedure, *129-4*
EXPLAIN_REWRITE procedure, *129-5*
I_AM_A_REFRESH function, *129-6*
PMARKER function, *129-6*
PURGE_DIRECT_LOAD_LOG procedure, *129-7*
PURGE_LOG procedure, *129-7*
PURGE_MVIEW_FROM_LOG procedure, *129-8*
REFRESH procedure, *129-9*
REFRESH_ALL_MVIEWS procedure, *129-11*
REFRESH_DEPENDENT procedure, *129-13*
REGISTER_MVIEW procedure, *129-15*
UNREGISTER_MVIEW procedure, *129-16*
DBMS_MVIEW_STATS Overview, *130-1*
DBMS_MVIEW_STATS package, *130-1*
DBMS_MVIEW_STATS security model, *130-1*
DBMS_ODCI package, *133-1*
ESTIMATE_CPU_UNITS function, *133-1*
methods, *133-1*
DBMS_OPTIM_BUNDLE package, *134-1*
DBMS_OUTPUT package, *136-1*
DBMS_PARALLEL_EXECUTE package, *137-1*

DBMS_PDB_ALTER_SHARING package, *141-1*
DBMS_PIPE package, *143-1*
DBMS_PLSQL_CODE_COVERAGE package, *144-1*
DBMS_PREDICTIVE_ANALYTICS package, *145-1*
  EXPLAIN procedure, *145-2*
  PREDICT procedure, *145-4*
  PROFILE Procedure, *145-6*
DBMS_PREPROCESSOR package, *146-1*
DBMS_PROCESS
  START_POOL procedure, *148-3*
  STOP_POOL procedure, *148-4*
DBMS_PROCESS package, *148-1*
  overview, *148-1*
  security model, *148-1*
DBMS_PROCESS procedure
  CONFIGURE_POOL procedure, *148-2*
DBMS_PROPAGATION_ADM package, *150-1*
DBMS_RANDOM package, *152-1*
DBMS_RESCONFIG package, *157-1*
DBMS_RESOURCE_MANAGER package, *158-1*
DBMS_RESOURCE_MANAGER_PRIVS package, *159-1*
DBMS_RESULT_CACHE package, *160-1*
DBMS_RESUMABLE package, *161-1*
DBMS_RLS package
  fine-grained access control
    DBMS_RLS package, *162-1*
DBMS_ROWID package, *164-1*
DBMS_RULE package, *165-1*
DBMS_RULE_ADM package, *166-1*
DBMS_SAGA package, *168-1*
  BEGIN_SAGA function, *168-2*
  COMMIT_SAGA procedure, *168-3*
  GET_SAGA_ID function, *168-2*
  IS_INCOMPLETE function, *168-4*
  overview, *168-1*
  ROLLBACK_SAGA procedure, *168-3*
  security model, *168-1*
  SEND_REQUEST procedure, *168-5*
  SET_INCOMPLETE procedure, *168-4*
  SET_SAGA_ID procedure, *168-2*
  summary of subprograms, *168-1*
DBMS_SAGA_ADM package, *167-1*
  ADD_BROKER procedure, *167-2*
  ADD_COORDINATOR procedure, *167-3*
  ADD_PARTICIPANT procedure, *167-4*
  DROP_BROKER procedure, *167-2*
  DROP_COORDINATOR procedure, *167-4*
  DROP_PARTICIPANT procedure, *167-6*
  overview, *167-1*
  REGISTER_SAGA_CALLBACK procedure, *167-6*
  security model, *167-1*
  summary of subprograms, *167-1*
DBMS_SCHEDULER package, *169-1*
  ADD_EVENT_QUEUE_SUBSCRIBER Procedure, *169-25*
  ADD_GROUP_MEMBER Procedure, *169-26*
  ADD_JOB_EMAIL_NOTIFICATION Procedure, *169-27*
  ADD_TO_INCOMPATIBILITY Procedure, *169-29*
  ALTER_CHAIN Procedure, *169-30*
  ALTER_RUNNING_CHAIN Procedure, *169-32*
  CLOSE_WINDOW Procedure, *169-34*
  COPY_JOB Procedure, *169-34*
  CREATE_CHAIN Procedure, *169-35*
  CREATE_CREDENTIAL Procedure, *169-36*
  CREATE_DATABASE_DESTINATION Procedure, *169-37*
  CREATE_EVENT_SCHEDULE Procedure, *169-38*
  CREATE_FILE_WATCHER Procedure, *169-39*
  CREATE_GROUP Procedure, *169-40*
  CREATE_INCOMPATIBILITY Procedure, *169-42*
  CREATE_JOB Procedure, *169-43*
  CREATE_JOB_CLASS Procedure, *169-51*
  CREATE_JOBS Procedure, *169-53*
  CREATE_PROGRAM Procedure, *169-54*
  CREATE_RESOURCE Procedure, *169-57*
  CREATE_SCHEDULE Procedure, *169-58*
  CREATE_WINDOW Procedure, *169-59*
  DEFINE_ANYDATA_ARGUMENT Procedure, *169-61*
  DEFINE_CHAIN_EVENT_STEP Procedure, *169-62*
  DEFINE_CHAIN_RULE Procedure, *169-63*
  DEFINE_CHAIN_STEP Procedure, *169-66*
  DEFINE_METADATA_ARGUMENT Procedure, *169-67*
  DEFINE_PROGRAM_ARGUMENT Procedure, *169-68*
  DISABLE Procedur, *169-70*
  DROP_AGENT_DESTINATION Procedure, *169-72*
  DROP_CHAIN Procedure, *169-73*
  DROP_CHAIN_RULE Procedure, *169-74*
  DROP_CHAIN_STEP Procedure, *169-75*
  DROP_CREDENTIAL Procedure, *169-75*
  DROP_DATABASE_DESTINATION Procedure, *169-76*
  DROP_FILE_WATCHER Procedure, *169-77*
  DROP_GROUP Procedure, *169-78*
  DROP_INCOMPATIBILITY Procedure, *169-78*
  DROP_JOB Procedure, *169-79*
  DROP_JOB_CLASS Procedure, *169-80*
  DROP_PROGRAM Procedure, *169-81*

DBMS_SCHEDULER package *(continued)*
DROP_PROGRAM_ARGUMENT Procedure, *169-81*
DROP_RESOURCE Procedure, *169-82*
DROP_SCHEDULE Procedure, *169-83*
DROP_WINDOW Procedure, *169-83*
DUMP_IN_MEMORY_TRACE Procedure, *169-84*
ENABLE Procedure, *169-84*
END_DETACHED_JOB_RUN Procedure, *169-85*
EVALUATE_CALENDAR_STRING Procedure, *169-86*
EVALUATE_RUNNING_CHAIN Procedure, *169-88*
GENERATE_JOB_NAME Function, *169-89*
GET_AGENT_INFO Function, *169-89*
GET_AGENT_VERSION Function, *169-90*
GET_ATTRIBUTE Procedure, *169-91*
GET_FILE Procedure, *169-91*
GET_SCHEDULER_ATTRIBUTE Procedure, *169-93*
JOB DEFINITION constructor function, *169-13*
JOB DEFINITION object type, *169-13*
JOB_DEFINTION_ARRAY table type, *169-17*
JOBARG constructor function, *169-12*
JOBARG object type, *169-12*
JOBARG_ARRAY table type, *169-12*, *169-13*
JOBATTR constructor function, *169-17*
JOBATTR object type, *169-17*
JOBATTR_ARRAY table type, *169-17*, *169-18*
OPEN_WINDOW Procedure, *169-94*
PURGE_LOG Procedure, *169-95*
PUT_FILE Procedure, *169-96*
REMOVE_EVENT_QUEUE_SUBSCRIBER Procedure, *169-97*
REMOVE_FROM_INCOMPATIBILITY Procedure, *169-97*
REMOVE_GROUP_MEMBER Procedure, *169-98*
REMOVE_JOB_EMAIL_NOTIFICATION Procedure, *169-99*
RESET_JOB_ARGUMENT_VALUE Procedure, *169-99*
RUN_CHAIN Procedure, *169-100*
RUN_JOB Procedure, *169-102*
Scheduler Chain Condition Syntax, *169-63*
SCHEDULER_FILEWATCHER_REQUEST Object Type, *169-21*
SCHEDULER_FILEWATCHER_RESULT Object Type, *169-20*
SCHEDULER$_EVENT_INFO Object Type, *169-19*
SCHEDULER$_STEP_TYPE object type, *169-18*
DBMS_SCHEDULER package *(continued)*
SCHEDULER$_STEP_TYPE_LIST table type, *169-19*
SET_AGENT_REGISTRATION_PASS Procedure, *169-103*
SET_ATTRIBUTE Procedure, *169-104*
SET_ATTRIBUTE_NULL Procedure, *169-119*
SET_JOB_ANYDATA_VALUE Procedure, *169-120*
SET_JOB_ARGUMENT_VALUE Procedure, *169-121*
SET_JOB_ATTRIBUTES Procedure, *169-122*
SET_RESOURCE_CONSTRAINT Procedure, *169-123*
SET_SCHEDULER_ATTRIBUTE Procedure, *169-124*
STOP_JOB Procedure, *169-125*
DBMS_SERVER_ALERT package, *171-1*
DBMS_SERVICE package, *172-1*
DBMS_SESSION package, *173-1*
DBMS_SODA
subprograms, *177-1*, *177-6*
DBMS_SODA package, *177-1*
CREATE_COLLECTION Function, *177-2*
DROP_COLLECTION Function, *177-3*
GET_AS_OF_SCN Function, *177-4*
GET_AS_OF_TIMESTAMP Function, *177-4*
GET_DEFAULT_METADATA_AS_CLOB Function, *177-4*
GET_DEFAULT_METADATA_AS_VARCHAR2 Function, *177-4*
LIST_COLLECTION_NAMES function, *177-4*
OPEN_COLLECTION Function, *177-5*
Security Model, *177-1*
TO_UUID Function, *177-10*
DBMS_SPM package, *181-1*
DBMS_SQL DESC_REC4 Record Type, *182-19*
DBMS_SQL_FIREWALL package
ADD_ALLOWED_CONTEXT procedure, *183-4*
APPEND_ALLOW_LIST procedure, *183-5*
APPEND_ALLOW_LIST_SINGLE_SQL procedure, *183-6*
constants, *183-2*
CREATE_CAPTURE procedure, *183-8*
DELETE_ALLOWED_CONTEXT procedure, *183-9*
DELETE_ALLOWED_SQL procedure, *183-10*
DISABLE procedure, *183-10*
DISABLE_ALLOW_LIST procedure, *183-11*
DROP CAPTURE procedure, *183-12*
DROP_ALLOW_LIST procedure, *183-11*
ENABLE procedure, *183-12*
ENABLE_ALLOW_LIST procedure, *183-13*
EXCLUDE procedure, *183-14*
EXPORT_ALLOW_LIST procedure, *183-14*

DBMS_SQL_FIREWALL package *(continued)*
FLUSH_LOGS procedure, *183-15*
GENERATE_ALLOW_LIST procedure, *183-16*
IMPORT_ALLOW_LIST procedure, *183-17*
INCLUDE procedure, *183-18*
overview, *183-1*
PURGE_LOG procedure, *183-18*
security model, *183-1*
START_CAPTURE procedure, *183-19*
STOP_CAPTURE procedure, *183-20*
UPDATE_ALLOW_LIST_ENFORCEMENT procedure, *183-20*
DBMS_SQLDIAG package, *186-1*
DBMS_SQLDIAG security model, *186-1*
DBMS_SQLQ package, *188-1*
ALTER_QUARANTINE procedure, *188-2*
CREATE_QUARANTINE_BY_SQL_ID function, *188-3*
CREATE_QUARANTINE_BY_SQL_TEXT function, *188-4*
CREATE_STGTAB_QUARANTINE procedure, *188-5*
DROP_QUARANTINE procedure, *188-6*
GET_PARAM_VALUE_QUARANTINE function, *188-6*
PACK_STGTAB_QUARANTINE function, *188-7*
UNPACK_STGTAB_QUARANTINE function, *188-8*
DBMS_SQLSET package, *189-1*
ADD_REFERENCE function, *189-5*
CAPTURE_CURSOR_CACHE procedure, *189-6*
CREATE_SQLSET procedure, *189-8*
CREATE_STGTAB procedure, *189-9*
data structures, *189-2*
DELETE_SQLSET procedure, *189-10*
DROP_SQLSET procedure, *189-10*
LOAD_SQLSET procedure, *189-11*
overview, *189-1*
PACK_STGTAB procedure, *189-15*
REMAP_STGTAB procedure, *189-17*
REMOVE_REFERENCE procedure, *189-18*
security model, *189-2*
SELECT_CURSOR_CACHE function, *189-19*
SELECT_SQL_TRACE function, *189-24*
SELECT_SQLPA_TASK function, *189-27*
SELECT_SQLSET function, *189-28*
SELECT_WORKLOAD_REPOSITORY function, *189-30*
summary of subprograms, *189-4*
DBMS_SQLTUNE package, *34-1*, *187-1*, *190-1*
DBMS_STAT_FUNCS package, *191-1*
DBMS_STATS
constants, *192-4*
DBMS_STREAMS_AUTH package, *240-1*
DBMS_TABLE_DATA package, *195-1*
GET_BYTES_FOR_COLUMN Procedure(), *195-2*
GET_BYTES_FOR_COLUMNS Procedure(), *195-3*
GET_BYTES_FOR_ROW Procedure(), *195-3*
overview, *195-1*
Security Model, *195-1*
DBMS_TDB package, *196-1*
DBMS_TDB.CHECK_DB Function, *196-3*
DBMS_TDB.CHECK_DB procedure, *196-4*
DBMS_TF package, *197-1*
client interface, *197-3*
COLUMN_DATA_T, *197-21*
COLUMN_METADATA_T Type, *197-22*
COLUMN_T, *197-23*
COLUMN_TYPE_NAME, *197-27*
COLUMNS pseudo-operator, *197-11*
COLUMNS_NEW_T, *197-25*
COLUMNS_T, *197-25*
COLUMNS_WITH_TYPE_T, *197-25*
constants, *197-2*
CSTORE compilation state management, *197-7*
CSTORE_GET, *197-29*
data structures, *197-19*
DESCRIBE_T, *197-23*
ENV_T, *197-23*
examples, *197-15*
execution flow, *197-12*
GET_COL, *197-30*
GET_ENV, *197-32*
GET_ROW_SET, *197-32*
GET_XID, *197-36*
method name overrides, *197-9*
parallel execution, *197-12*
PARALLEL_ENV_T, *197-24*
pass-through columns, *197-6*
PUT_COL, *197-36*
PUT_ROW_SET, *197-40*
query transformations, *197-11*
read columns, *197-6*
restrictions, *197-14*
ROW_REPLICATION, *197-43*
ROW_SET_T, *197-26*
ROW_TO_CHAR, *197-45*
server side interface, *197-5*
SUPPORTED_TYPE, *197-46*
TAB_BINARY_DOUBLE_T, *197-20*
TAB_BINARY_FLOAT_T, *197-20*
TAB_BLOB_T, *197-20*
TAB_BOOLEAN_T, *197-20*
TAB_CHAR_T, *197-20*
TAB_CLOB_T, *197-20*
TAB_DATE_T, *197-20*

DBMS_TF package *(continued)*
TAB_INTERVAL_DS_T, *197-20*
TAB_INTERVAL_YM_T, *197-20*
TAB_NATURALN_T, *197-20*
TAB_NUMBER_T, *197-20*
TAB_RAW_T, *197-20*
TAB_ROWID_T, *197-20*
TAB_TIMESTAMP_LTZ_T, *197-20*
TAB_TIMESTAMP_T, *197-20*
TAB_TIMESTAMP_TZ_T, *197-20*
TAB_VARCHAR2_T, *197-20*
TABLE_COLUMNS_T, *197-26*
TABLE_T, *197-25*
TRACE, *197-46*
XID_T, *197-26*
XSTORE execution state management, *197-9*
XSTORE_CLEAR, *197-47*
XSTORE_EXISTS, *197-48*
XSTORE_GET, *197-48*
XSTORE_REMOVE, *197-51*
XSTORE_SET, *197-52*
DBMS_TF.ROW_REPLICATION procedure, *197-43*
DBMS_TF.TRACE procedure, *197-46*
DBMS_TNS package, *198-1*
overview, *198-1*
security model, *198-1*
DBMS_TNS packageRESOLVE_TNSNAME function, *198-2*
DBMS_TRACE package, *199-1*
DBMS_TRANSACTION package, *200-1*
DBMS_TRANSFORM package, *201-1*
DBMS_TTS package, *204-1*
DBMS_UMF package, *206-1*
DBMS_UNDO_ADV package, *207-1*
DBMS_USER_CERTS package, *208-1*
ADD_CERTIFICATE Procedure(), *208-2*
ADD_COPY Procedure(), *208-2*
DROP_CERTIFICATE Procedure(), *208-3*
overview, *208-1*
Security Model, *208-1*
DBMS_USERDIAG package, *209-1*
ADD_CERTIFICATE Procedure(), *209-2*
CHECK_SQL_TRACE_EVENT Procedure(), *209-2*
overview, *209-1*
DBMS_UTILITY package, *210-1*
deprecated subprograms, *210-1*
DBMS_VECTOR package
DROP_ONNX_MODEL procedure, *211-8*
JSON Metadata Parameters, *211-15*
LOAD_ONNX_MODEL procedure, *211-13*
LOAD_ONNX_MODEL_CLOUD procedure, *211-21*
DBMS_WM package, *214-1*
DBMS_WORKLOAD_RECORD package, *215-1*
DBMS_WORKLOAD_REPLAY package, *216-1*
DBMS_XA_XID Object Type, *219-4*
DBMS_XA_XID_ARRAY Table Type, *219-5*
DBMS_XBD_VERSION package, *224-1*
DBMS_XDB Constants, *218-2*
DBMS_XDB Overview, *218-1*
DBMS_XDB package, *218-1*
methods, *218-3*
DBMS_XDB_ADMIN package, *220-1*
DBMS_XDB_VERSION package
CHECKIN function, *224-2*
CHECKOUT procedure, *224-3*
GETCONTENTSBLOBBYRESID function, *224-3*
GETCONTENTSCLOBBYRESID function, *224-3*
GETCONTENTSXMLBYRESID function, *224-4*
GETPREDECESSORS function, *224-4*
GETPREDSBYRESID function, *224-5*
GETRESOURCEBYRESID function, *224-5*
GETSUCCESSORS function, *224-5*
GETSUCCSBYRESID function, *224-6*
MAKEVERSIONED function, *224-6*
UNCHECKOUT function, *224-7*
DBMS_XDBRESOURCE package, *225-1*
DBMS_XDBZ package
DISABLE_HIERARCHY procedure, *226-3*
ENABLE_HIERARCHY procedure, *226-3*
GET_ACLOID function, *226-4*
GET_USERID function, *226-4*
IS_HIERARCHY_ENABLED function, *226-5*
PURGELDAPCACHE function, *226-6*
DBMS_XMLDOM Constants, *228-3*
DBMS_XMLDOM package, *228-1*
APPENDDATA procedure, *228-21*
CREATEATTRIBUTE function, *228-22*
CREATECDATASECTION function, *228-23*
CREATECOMMENT function, *228-23*
CREATEDOCUMENT function, *228-24*
CREATEDOCUMENTFRAGMENT, *228-24*
CREATEELEMENT function, *228-25*
CREATEENTITYREFERENCE function, *228-26*
CREATEPROCESSINGINSTRUCTION function, *228-26*
CREATETEXTNODE function, *228-27*
DELETEDATA procedure, *228-27*
description, *228-1*
exceptions, *228-4*
FINDENTITY function, *228-28*
FINDNOTATION function, *228-28*
FREEDOCFRAG procedure, *228-29*
FREEDOCUMENT procedure, *228-29*
GETATTRIBUTE function, *228-31*
GETATTRIBUTENODE function, *228-32*

DBMS_XMLDOM package *(continued)*
GETCHILDRENBYTAGNAME function, *228-34*
GETDATA function, *228-35*
GETDOCTYPE function, *228-35*
GETDOCUMENTELEMENT function, *228-36*
GETELEMENTSBYTAGNAME function, *228-36*
GETENTITIES function, *228-37*
GETIMPLEMENTATION function, *228-39*
GETLENGTH function, *228-40*
GETNAME function, *228-41*, *228-42*
GETNAMEDITEM function, *228-42*
GETNAMESPACE procedure, *228-42*
GETNEXTSIBLING function, *228-43*
GETNODENAME function, *228-44*
GETNODETYPE function, *228-43*
GETNODEVALUE function, *228-44*
GETNOTATIONNAME function, *228-46*
GETNOTATIONS function, *228-47*
GETOWNERDOCUMENT function, *228-48*
GETOWNERELEMENT function, *228-48*
GETPARENTNODE function, *228-49*
GETPREFIX function, *228-49*
GETPREVIOUSSIBLING function, *228-50*
GETPUBLICID function, *228-50*
GETQUALIFIEDNAME function, *228-51*
GETSCHEMANODE function, *228-51*
GETSPECIFIED function, *228-52*
GETSTANDALONE function, *228-52*
GETSYSTEMID function, *228-53*
GETTAGNAME function, *228-53*
GETTARGET function, *228-47*
GETVALUE function, *228-54*
GETVERSION function, *228-54*
GETXMLTYPE function, *228-55*
HASATTRIBUTE function, *228-55*
HASATTRIBUTES function, *228-56*
HASCHILDNODES function, *228-56*
HASFEATURE function, *228-57*
IMPORTNODE function, *228-57*
INSERTBEFORE function, *228-58*
INSERTDATA procedure, *228-59*
ISNULL function, *228-59*
MAKEATTR function, *228-63*
MAKECDATASECTION function, *228-63*
MAKECHARACTERDATA function, *228-64*
MAKECOMMENT function, *228-64*
MAKEDOCUMENT function, *228-65*
MAKEDOCUMENTFRAGMENT function, *228-65*
MAKEDOCUMENTTYPE function, *228-66*
MAKEELEMENT function, *228-66*
MAKEENTITY function, *228-67*
MAKEENTITYREFERENCE function, *228-67*
MAKENODE function, *228-68*

DBMS_XMLDOM package *(continued)*
MAKENOTATION function, *228-70*
MAKEPROCESSINGINSTRUCTION function, *228-70*
MAKETEXT function, *228-71*
methods
APPENDCHILD function, *228-20*
APPENDDATA procedure, *228-21*
CLONENODE function, *228-21*
CREATEATTRIBUTE function, *228-22*
CREATECDATASECTION function, *228-23*
CREATECOMMENT function, *228-23*
CREATEDOCUMENT function, *228-24*
CREATEDOCUMENTFRAGMENT function, *228-24*
CREATEELEMENT function, *228-25*
CREATEENTITYREFERENCE function, *228-26*
CREATEPROCESSINGINSTRUCTION function, *228-26*
CREATETEXTNODE function, *228-27*
DELETEDATA procedure, *228-27*
DOMAttr interface, *228-6*
DOMCDataSection interface, *228-7*
DOMCharacterData interface, *228-7*
DOMComment interface, *228-8*
DOMDocument interface, *228-8*
DOMDocumentFragment interface, *228-9*, *228-20*
DOMDocumentType interface, *228-9*, *228-20*
DOMElement interface, *228-10*, *228-20*
DOMEntity interface, *228-10*, *228-20*
DOMEntityReference interface, *228-11*
DOMImplementation interface, *228-11*
DOMNamedNodeMap interface, *228-11*
DOMNodeAPPENDCHILD function, *228-20*
DOMNodeCLONENODE function, *228-21*
DOMNodeGETATTRIBUTES function, *228-33*
DOMNodeGETCHILDNODES function, *228-34*
DOMNodeGETEXPANDEDNAME procedure, *228-38*
DOMNodeGETFIRSTCHILD function, *228-38*
DOMNodeList interface, *228-12*
DOMNotation interface, *228-12*
DOMProcessingInstruction interface, *228-12*
DOMText interface, *228-5*, *228-12*
FINDENTITY function, *228-28*
FINDNOTATION function, *228-28*
FREEDOCFRAG procedure, *228-29*

DBMS_XMLDOM package *(continued)*
methods *(continued)*
FREEDOCUMENT procedure, *228-29*
FREEELEMENT procedure, *228-30*
FREENODE procedure, *228-30*
FREENODELIST Procedure, *228-31*
GETATTRIBUTE function, *228-31*
GETATTRIBUTENODE function, *228-32*
GETATTRIBUTES function, *228-33*
GETCHILDNODES function, *228-34*
GETCHILDRENBYTAGNAME function, *228-34*
GETDATA function, *228-35*
GETDOCTYPE function, *228-35*
GETDOCUMENTELEMENT function, *228-36*
GETELEMENTSBYTAGNAME function, *228-36*
GETENTITIES function, *228-37*
GETEXPANDEDNAME procedure, *228-38*
GETFIRSTCHILD function, *228-38*
GETIMPLEMENTATION function, *228-39*
GETLASTCHILD function, *228-39*
GETLENGTH function, *228-40*
GETLOCALNAME procedure, *228-40*
GETNAME function, *228-41*, *228-42*
GETNAMEDITEM function, *228-42*
GETNAMESPACE procedure, *228-42*
GETNEXTSIBLING function, *228-43*
GETNODENAME function, *228-44*
GETNODETYPE function, *228-43*
GETNODEVALUE function, *228-44*
GETNOTATIONNAME function, *228-46*
GETNOTATIONS function, *228-47*
GETOWNERDOCUMENT function, *228-48*
GETOWNERELEMENT function, *228-48*
GETPARENTNODE function, *228-49*
GETPREFIX function, *228-49*
GETPREVIOUSSIBLING function, *228-50*
GETPUBLICID function, *228-50*
GETQUALIFIEDNAME function, *228-51*
GETSCHEMANODE function, *228-51*
GETSPECIFIED function, *228-52*
GETSTANDALONE function, *228-52*
GETSYSTEMID function, *228-53*
GETTAGNAME function, *228-53*
GETTARGET function, *228-47*
GETVALUE function, *228-54*
GETVERSION function, *228-54*
GETXMLTYPE function, *228-55*
HASATTRIBUTE function, *228-55*
HASATTRIBUTES function, *228-56*
HASCHILDNODES function, *228-56*
HASFEATURE function, *228-57*
DBMS_XMLDOM package *(continued)*
methods *(continued)*
IMPORTNODE function, *228-57*
INSERTBEFORE function, *228-58*
INSERTDATA procedure, *228-59*
ISNULL function, *228-59*
MAKEATTR function, *228-63*
MAKECDATASECTION function, *228-63*
MAKECHARACTERDATA function, *228-64*
MAKECOMMENT function, *228-64*
MAKEDOCUMENT function, *228-65*
MAKEDOCUMENTFRAGMENT function, *228-65*
MAKEDOCUMENTTYPE function, *228-66*
MAKEELEMENT function, *228-66*
MAKEENTITY function, *228-67*
MAKEENTITYREFERENCE function, *228-67*
MAKENODE function, *228-68*
MAKENOTATION function, *228-70*
MAKEPROCESSINGINSTRUCTION function, *228-70*
MAKETEXT function, *228-71*
NEWDOMDOCUMENT function, *228-71*
NORMALIZE procedure, *228-72*
REMOVEATTRIBUTE procedure, *228-72*
REMOVEATTRIBUTENODE function, *228-73*
REMOVENAMEDITEM function, *228-74*
REPLACECHILD function, *228-75*
REPLACEDATA procedure, *228-76*
RESOLVENAMESPACEPREFIX function, *228-76*
SETATTRIBUTE procedure, *228-77*
SETATTRIBUTENODE function, *228-78*
SETDATA procedure, *228-79*
SETNAMEDITEM function, *228-80*
SETNODEVALUE procedure, *228-81*
SETPREFIX procedure, *228-83*
SETSTANDALONE procedure, *228-83*
SETVALUE procedure, *228-84*
SETVERSION procedure, *228-84*
SPLITTEXT function, *228-85*
SUBSTRINGDATA function, *228-85*
WRITETOBUFFER procedure, *228-86*
WRITETOCLOB procedure, *228-87*
WRITETOFILE procedure, *228-88*
NEWDOMDOCUMENT function, *228-71*
NORMALIZE procedure, *228-72*
REMOVEATTRIBUTE procedure, *228-72*
REMOVEATTRIBUTENODE function, *228-73*
REMOVENAMEDITEM function, *228-74*
REPLACECHILD function, *228-75*
REPLACEDATA procedure, *228-76*

DBMS_XMLDOM package *(continued)*
RESOLVENAMESPACEPREFIX function, *228-76*
SETATTRIBUTE procedure, *228-77*
SETATTRIBUTENODE function, *228-78*
SETDATA procedure, *228-79*
SETNAMEDITEM function, *228-80*
SETNODEVALUE procedure, *228-81*
SETPREFIX procedure, *228-83*
SETSTANDALONE procedure, *228-83*
SETVALUE procedure, *228-84*
SETVERSION procedure, *228-84*
SPLITTEXT function, *228-85*
SUBSTRINGDATA function, *228-85*
types, *228-3*
WRITETOBUFFER procedure, *228-86*
WRITETOCLOB procedure, *228-87*
WRITETOFILE procedure, *228-88*
DBMS_XMLGEN package, *229-1*
CLOSECONTEXT procedure, *229-2*
CONVERT function, *229-3*
GETNUMROWSPROCESSED Function, *229-3*
GETXML function, *229-4*
GETXMLTYPE function, *229-5*
NEWCONTEXT function, *229-6*
RESTARTQUERY procedure, *229-7*
SETCONVERTSPECIALCHARS procedure, *229-7*
SETMAXROWS procedure, *229-8*
SETROWSETTAG procedure, *229-9*
SETROWTAG procedure, *229-9*
SETSKIPROWS procedure, *229-10*
USEITEMTAGSFORCOLL procedure, *229-10*
USENULLATTRIBUTEINDICATOR procedure, *229-11*
DBMS_XMLINDEX package, *230-1*
DBMS_XMLPARSER package, *231-1*
FREEPARSER procedure, *231-2*
GETDOCTYPE function, *231-3*
GETDOCUMENT function, *231-3*
GETRELEASEVERSION function, *231-3*
GETVALIDATIONMODE function, *231-4*
NEWPARSER function, *231-4*
PARSE function, *231-4*
PARSE procedure, *231-4*
PARSEBUFFER procedure, *231-5*
PARSECLOB procedure, *231-5*
PARSEDTD procedure, *231-6*
PARSEDTDBUFFER procedure, *231-6*
PARSEDTDCLOB procedure, *231-7*
SETBASEDIR procedure, *231-7*
SETDOCTYPE procedure, *231-8*
SETERRORLOG procedure, *231-8*
SETPRESERVEWHITESPACE procedure, *231-8*
DBMS_XMLPARSER package *(continued)*
SETVALIDATIONMODE procedure, *231-9*
SHOWWARNINGS procedure, *231-9*
DBMS_XMLSCHEMA Constants, *232-1*
DBMS_XMLSCHEMA Operational Notes, *232-4*
DBMS_XMLSCHEMA package, *232-1*
DBMS_XMLSCHEMA Views, *232-3*
DBMS_XMLSTORE package, *236-1*
CLEARKEYCOLUMNLIST procedure, *236-2*
CLEARUPDATECOLUMNLIST procedure, *236-2*
CLOSECONTEXT procedure, *236-3*
DELETEXML function, *236-3*
INSERTXML function, *236-4*
NEWCONTEXT function, *236-4*
SETKEYCOLUMN procedure, *236-5*
SETROWTAG procedure, *236-5*
SETUPDATECOLUMN procedure, *236-5*
types, *236-1*
UPDATEXML function, *236-6*
DBMS_XSLPROCESSOR package
FREEPROCESSOR procedure, *238-2*
FREESTYLESHEET procedure, *238-2*
NEWPROCESSOR function, *238-3*
NEWSTYLESHEET function, *238-3*
PROCESSXSL function, *238-4*
REMOVEPARAM procedure, *238-6*
RESETPARAMS procedure, *238-6*
SELECTNODES function, *238-6*
SELECTSINGLENODE function, *238-7*
SETERRORLOG procedure, *238-7*
SETPARAM procedure, *238-8*
SHOWWARNINGS procedure, *238-8*
TRANSFORMNODE function, *238-9*
VALUEOF procedure, *238-9*
DBMS_XSLPROCESSOR Package, *238-1*
DBMS_XSTREAM_ADM package, *239-1*
security, *239-2*
DBMSOUTPUT_LINESARRAY Object Type, *136-5*
DBSFWUSER.DBMS_SFW_ACL_ADMIN package, *174-1*
DBUriType, *307-8*
DBURITYPE function, *307-9*
DBUriType subtype, *307-8*
CREATEURI function, *307-9*
DBURITYPE function, *307-9*
GETBLOB function, *307-9*
GETCLOB function, *307-10*
GETCONTENTTYPE function, *307-10*
GETEXTERNALURL function, *307-11*
GETURL function, *307-11*
GETXML function, *307-11*
methods, *307-8*
DEAUTHORIZE_DAD Procedure, *78-5*
DEBUG_EXPTOC package, *241-1*

DEBUG_ON procedure, *67-17*
DECLARE_REWRITE_EQUIVALENCE Procedures, *16-3*
DECODEFEATURES Function, *62-28*
DECRYPT Function, *52-16*
DECRYPT Procedures, *52-17*
DECRYPT_CAPTURE Procedure, *215-4*
DEFAULT_NAME function, *56-8*
DEFINE_ARRAY procedure, *182-40*
DEFINE_COLUMN procedure, *182-43*
DEFINE_COLUMN_CHAR Procedure, *182-45*
DEFINE_COLUMN_LONG procedure, *182-46*
DEFINE_COLUMN_RAW Procedure, *182-46*
DEFINE_COLUMN_ROWID Procedure, *182-47*
DEINIT_CLEANUP procedure, *31-15*
DELETE Procedure, *15-11*
DELETE_ALL_ERRORS procedure, *24-20*
DELETE_ALLOWED_CONTEXT procedure (of DBMS_SQL_FIREWALL), *183-9*
DELETE_ALLOWED_SQL procedure (of DBMS_SQL_FIREWALL), *183-10*
DELETE_BREAKPOINT function, *67-17*
DELETE_CAPTURE_INFO Procedure, *215-5*
DELETE_CATEGORY Procedure, *158-20*
DELETE_CDB_PLAN Procedure, *158-20*
DELETE_CDB_PLAN_DIRECTIVE Procedure, *158-21*
DELETE_CDB_PROFILE_DIRECTIVE Procedure, *158-21*
DELETE_COLUMN member procedure, *310-10*
DELETE_COLUMN procedure, *239-60*
DELETE_COLUMN_STATS pocedure, *192-37*
DELETE_CONSUMER_GROUP procedure, *158-22*
DELETE_DAD_ATTRIBUTE Procedure, *78-6*
DELETE_DATABASE_PREFS procedure, *192-39*
DELETE_DATABASE_STATS procedure, *192-43*
DELETE_DICTIONARY_STATS procedure, *192-44*
DELETE_ERROR procedure, *24-20*
DELETE_FILTER Procedure, *215-5*, *216-11*
DELETE_FINDING_DIRECTIVE Procedure, *15-12*
DELETE_FIXED_OBJECTS_STATS procedure, *192-46*
DELETE_INDEX_STATS procedure, *192-47*
DELETE_OBJECT procedure, *296-3*
DELETE_OER_BREAKPOINT function, *67-18*
DELETE_PARAMETER_DIRECTIVE Procedure, *15-12*
DELETE_PENDING_STATS Procedure, *192-48*
DELETE_PLAN procedure, *158-22*
DELETE_PLAN_CASCADE procedure, *158-22*
DELETE_PLAN_DIRECTIVE procedure, *158-23*
DELETE_POLICY_GROUP Procedure, *162-13*
DELETE_PRIVILEGE Procedure, *131-16*
DELETE_PROCESSING_RATE Procedure, *192-50*
DELETE_PROCREP_EXCLUSION_OBJ procedure, *87-10*
DELETE_REPLAY_INFO Procedure, *216-11*
DELETE_RUN Procedure, *147-5*
DELETE_SCHEMA_PREFS procedure, *192-50*
DELETE_SCHEMA_STATS Procedure, *192-56*
DELETE_SEGMENT_DIRECTIVE Procedure, *15-13*
DELETE_SERVICE procedure, *172-11*
DELETE_SNAPSHOT Procedure, *14-4*, *14-5*
DELETE_SQL_CONNECTION_TEST Procedure, *20-4*, *21-7*
DELETE_SQL_DIRECTIVE Procedure, *15-13*
DELETE_SQLSET procedure, *189-10*, *190-34*
DELETE_SQLWKLD Procedure, *17-17*
DELETE_SQLWKLD_REF Procedure, *17-18*
DELETE_SQLWKLD_STATEMENT Procedure, *17-19*
DELETE_SYSTEM_STATS procedure, *192-57*
DELETE_TABLE_PREFS procedure, *192-58*
DELETE_TABLE_STATS procedure, *192-64*
DELETE_TASK Procedure, *17-21*
DELETECONTENT Procedure, *62-29*, *63-8*
DELETEDATA procedure, *228-27*
DELETEDIRECTORY Procedure, *62-29*, *63-8*
DELETEFILE Procedure, *62-30*, *63-9*
DELETEHTTPEXPIREMAPPING Procedure, *221-8*
DELETEMIMEMAPPING Procedure, *221-9*
DELETEREPOSITORYRESCONFIG Procedure, *157-4*
DELETERESCONFIG Procedure, *157-5*
DELETERESOURCE Procedure, *223-9*
DELETERESOURCEMETADATA Procedures, *223-9*
DELETESCHEMA procedure, *232-7*
DELETESCHEMALOCMAPPING Procedure, *221-9*
DELETESERVLET Procedure, *221-9*
DELETESERVLETMAPPING Procedure, *221-10*
DELETESERVLETSECROLE Procedure, *221-10*
DELETEXML function, *236-3*
DELETEXMLEXTENSION Procedure, *221-10*
DEQUEUE_ARRAY Function, *25-9*
DEQUEUE_OPTIONS_T Type, *305-9*
DEQUEUE_PARALLEL_STATEMENT Procedure, *158-23*
DEREGISTER Procedure, *50-17*
DEREGISTER_ERROR_TRANSLATION Procedure, *185-7*
DEREGISTER_SQL_TRANSLATION Procedure, *185-6*
DEREGSTORECOMMAND Function, *64-9*
DERIVE_FROM_MVIEW function, *54-34*

DESC_REC2 Record Type, *182-17*
DESC_REC3 Record Type, *182-18*
DESC_RESC Record Type, *182-16*
DESC_TAB Table Type, *182-21*
DESC_TAB2 Table Type, *182-21*
DESC_TAB3 Table Type, *182-22*
DESC_TAB4 Table Type, *182-22*
DESCRIBE function, *108-2*
DESCRIBE Procedure, *140-4*
DESCRIBE_COLUMNS procedure, *182-48*
DESCRIBE_COLUMNS2 procedure, *182-49*
DESCRIBE_COLUMNS3 Procedure, *182-50*
DESCRIBE_DIMENSION procedure, *73-1*
DESCRIBE_PROCEDURE procedure, *70-5*
DESCRIBE_STACK procedure, *58-27*
DESCRIBE_T Type, *197-23*
DESTROY procedure, *155-4*
DESTROY_REQUEST_CONTEXT Procedure, *286-35*
DETACH_SESSION procedure, *67-18*
DFN function
  of HTF package, *242-22*
DFN procedure
  of HTP package, *243-22*
DIFF_PLAN Function, *237-9*
DIFF_TABLE_STATS_IN_HISTORY Function, *192-66*
DIFF_TABLE_STATS_IN_PENDING function, *192-66*
DIFF_TABLE_STATS_IN_STATTAB Function, *192-67*
dimension maintenance (cube), *54-1*
dimensions
  optimizer statistics, *41-1*
DIRLISTCLOSE function
  of HTF package, *242-23*
DIRLISTCLOSE procedure
  of HTP package, *243-22*
DIRLISTOPEN function
  of HTF package, *242-23*
DIRLISTOPEN procedure
  of HTP package, *243-23*
DISABLE procedure, *56-9*
  of DBMS_FLASHBACK package, *82-5*
  of DBMS_OUTPUT package, *136-6*
  of OWA_CACHE package, *244-2*
DISABLE procedure (of DBMS_SQL_FIREWALL), *183-10*
DISABLE Procedures, *38-2*
DISABLE_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-11*
DISABLE_APPLICATION Procedure, *83-5*
DISABLE_ASOF_VALID_TIME Procedure, *83-5*
DISABLE_AUTO_OPTIMIZE Procedure, *98-6*
DISABLE_BREAKPOINT function, *67-18*
DISABLE_CAPTURE Procedure, *147-5*
DISABLE_CONNECTION_TEST Procedure, *20-4*, *21-8*
DISABLE_CREDENTIAL Procedure, *51-4*
DISABLE_GROUPED_POLICY Procedure, *162-14*
DISABLE_HIERARCHY procedure, *226-3*
DISABLE_ILM Procedure, *98-6*
DISABLE_JOB Procedure, *125-32*
DISABLE_MIGRATION_CHECK Procedure, *29-8*
DISABLE_POLICY Procedure, *153-22*
DISABLE_PROPAGATION_SCHEDULE Procedure, *125-33*
DISABLE_PROTECTION_COLUMN Procedure, *203-6*
DISABLE_PROTECTION_SOURCE Procedure, *203-7*
DISABLE_PROTECTION_TYPE Procedure, *203-7*
DISABLEDEFAULTTABLECREATION Procedure, *233-3*
DISABLEINDEXESANDCONSTRAINTS Procedure, *235-2*
DISABLEMAINTAINDOM Procedure, *233-4*
DISASSOCIATE_FBA Procedure, *83-5*
DISCONNECT Procedure, *68-4*
DISCONNECT_SESSION procedure, *172-11*
DISPLAY function, *237-10*
DISPLAY_AWR function, *237-12*
DISPLAY_CURSOR function, *237-15*
DISPLAY_PLAN Function, *237-18*
DISPLAY_SQL_PLAN_BASELINE Function, *237-20*
DISPLAY_SQLSET Function, *237-21*
DIST_TXN_SYNC Procedure, *219-6*
DIV function
  of HTF package, *242-24*
DIV procedure
  of HTP package, *243-23*
DLISTCLOSE function
  of HTF package, *242-24*
DLISTCLOSE procedure
  of HTP package, *243-24*
DLISTDEF function
  of HTF package, *242-25*
DLISTDEF procedure
  of HTP package, *243-24*
DLISTOPEN function
  of HTF package, *242-25*
DLISTOPEN procedure
  of HTP package, *243-25*
DLISTTERM function
  of HTF package, *242-26*
DLISTTERM procedure
  of HTP package, *243-25*
DOMAIN_LEVEL Function, *132-4*
DOMAINS Function, *132-5*

DOMAttr methods, *228-6*
DOMCDataSection methods, *228-7*
DOMCharacterData methods, *228-7*
DOMComment methods, *228-8*
DOMDocument methods, *228-8*
DOMDocumentType methods, *228-9*, *228-20*
DOMEntity methods, *228-20*
DOMNamedNodeMap methods, *228-11*
DOMText methods, *228-5*
DOWNGRADE procedure, *204-2*
DOWNLOAD_FILE procedures, *301-2*
DROFILESYSTEM Procedures, *65-5*
DROP_ACL Procedure, *131-17*
DROP_ADVISOR_TASK procedure, *192-68*
DROP_ALL function, *193-2*
DROP_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-11*
DROP_ANALYSIS_TASK Procedure, *187-5*
DROP_APPLICATION Procedure, *83-6*
DROP_APPLY procedure, *24-21*
DROP_AUTO_MVS Procedure, *36-4*
DROP_BASELINE procedure, *217-31*
DROP_BY_CAT procedure, *135-3*
DROP_CAPTURE procedure
  capture process
    dropping, *44-25*
DROP_CAPTURE Procedure, *147-6*
DROP_CAPTURE procedure (of DBMS_SQL_FIREWALL), *183-12*
DROP_CHUNKS Procedure, *137-11*
DROP_COMPARISON Procedure, *47-17*
DROP_CREDENTIAL Procedure, *51-5*
DROP_DAD Procedure, *78-7*
DROP_DIAGNOSIS_TASK Procedure, *186-11*
DROP_EMPTY_SEGMENTS Procedure, *179-6*
DROP_EVALUATION_CONTEXT procedure, *166-11*
DROP_EVOLVE_TASK Procedure, *181-5*, *181-15*
DROP_EXTENDED_STATS Procedure, *192-69*
DROP_FILE function, *193-3*
DROP_GROUPED_POLICY Procedure, *162-15*
DROP_INBOUND procedure, *239-62*
DROP_LINK procedure, *206-3*
DROP_MVIEW procedure, *54-7*, *54-35*
DROP_OBJECT procedure, *120-2*
DROP_OBJECT_DEPENDENCY procedure, *24-22*
DROP_OLD_UNIFIED_AUDIT_TABLES Procedure, *31-16*
DROP_OUTBOUND procedure, *239-62*
DROP_POLICY Procedure, *153-23*, *162-15*, *203-7*
DROP_POLICY_CONTEXT Procedure, *162-16*
DROP_POLICY_EXPRESSION Procedure, *153-23*
DROP_PROFILE Procedure, *185-8*
DROP_PROPAGATION procedure, *150-6*
DROP_PURGE_JOB procedure, *31-17*
DROP_REWRITE_EQUIVALENCE Procedure, *16-5*
DROP_RULE procedure, *166-12*
DROP_RULE_SET procedure, *166-13*
DROP_SENSITIVE_COLUMN Procedure, *202-4*
DROP_SENSITIVE_TYPE Procedure, *202-5*
DROP_SENSITIVE_TYPE_SOURCE Procedure, *202-5*
DROP_SNAPSHOT_RANGE procedure, *217-32*
DROP_SQL_PATCH Function ((amp)) Procedure, *186-11*
DROP_SQL_PLAN_BASELINE, *181-15*
DROP_SQL_PLAN_DIRECTIVE Procedure, *180-4*
DROP_SQL_PROFILE procedure, *190-35*
DROP_SQLSET procedure, *189-10*, *190-35*
DROP_STAT_TABLE procedure, *192-70*
DROP_TASK Procedure, *137-11*
DROP_TOPOLOGY procedure, *206-4*
DROP_TRANSFORMATION procedure, *201-2*
DROP_TUNING_TASK procedure, *190-36*
DROP_UNUSED procedure, *135-4*
DROP_VIEW procedure
  of DBMS_HS_PARALLEL package, *94-3*
DROP_VIRTUAL_COLUMNS procedure, *107-8*
DROPPARAMETER Procedure, *230-3*
DROPSTORE Procedure, *64-9*
DUMP_ORPHAN_KEYS procedure, *156-6*
dynamic SQL
  anonymous blocks and, *182-2*
  execution flow in, *182-5*
DYNAMIC_DEPTH Function, *281-6*

E

e-mail from PL/SQL (email), *299-4*
ECDH_GENKEYPAIR function, *52-19*
ECDHDERIVE_SHAREDSECRET function, *52-18*
EDIT_DISTANCE Function, *292-2*
EDIT_DISTANCE_SIMILARITY Function, *292-2*
EHLO function and procedure, *298-13*
EM function
  of HTF package, *242-26*
EM procedure
  of HTP package, *243-26*
EMPHASIS function
  of HTF package, *242-27*
EMPHASIS procedure
  of HTP package, *243-26*
ENABLE procedure, *56-10*, *136-6*
ENABLE procedure (of DBMS_SQL_FIREWALL, *183-12*
ENABLE Procedures, *38-3*

ENABLE_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-13*
ENABLE_APPLICATION Procedure, *83-6*
ENABLE_AT_SYSTEM_CHANGE_NUMBER procedure, *82-6*
ENABLE_AT_TIME procedure, *82-6*
ENABLE_AT_VALID_TIME Procedure, *83-6*
ENABLE_AUTO_OPTIMIZE Procedure, *98-7*
ENABLE_BREAKPOINT function, *67-19*
ENABLE_CAPTURE Procedure, *147-6*
ENABLE_CONNECTION_TEST Procedure, *20-5*, *21-10*
ENABLE_CREDENTIAL Procedure, *51-6*
ENABLE_ERROR_TRANSLATION Procedure, *185-8*
ENABLE_GG_XSTREAM_FOR_STREAMS procedure, *239-63*
ENABLE_GROUPED_POLICY Procedure, *162-16*
ENABLE_HIERARCHY procedure, *226-3*
ENABLE_ILM Procedure, *98-7*
ENABLE_JOB Procedure, *125-33*
ENABLE_OPTIM_FIXES Procedures, *134-1*
ENABLE_POLICY Procedure, *153-24*, *162-17*
ENABLE_PROPAGATION_SCHEDULE Procedure, *125-34*
ENABLE_PROTECTION_COLUMN Procedure, *203-9*
ENABLE_PROTECTION_SOURCE Procedure, *203-9*
ENABLE_PROTECTION_TYPE Procedure, *203-10*
ENABLE_REG Procedure, *50-18*
ENABLE_SERVICE procedure, *206-4*
ENABLE_SQL_TRANSLATION Procedure, *185-9*
ENABLEDEFAULTTABLECREATION Procedure, *233-5*
ENABLEDIGESTAUTHENTICATION Procedure, *221-11*
ENABLEINDEXESANDCONSTRAINTS Procedure, *235-4*
ENABLEMAINTAINDOM Procedure, *233-5*
ENCODING_DEFAULT Function, *222-2*
ENCODING_ISOLATIN1 Function, *222-3*
ENCODING_UTF8 Function, *222-3*
ENCODING_WIN1252 Function, *222-3*
ENCRYPT Function, *52-19*
ENCRYPT_CAPTURE Procedure, *215-6*
END_OPERATION Procedure, *184-3*
END_PREPARE Procedure, *76-6*
END_REPLAY_SCHEDULE Procedure, *216-12*
END_REQUEST procedure, *286-35*
END_RESPONSE procedure, *286-36*
END_SQL_BLOCK Procedure, *158-24*
END_UPGRADE Procedure, *76-6*
ENDCREATE member procedure
  of ANYDATA TYPE, *302-3*
  of ANYDATASET TYPE, *303-3*
  of ANYTYPE TYPE, *304-4*
ENQUEUE_ARRAY Function, *25-11*
ENQUEUE_OPTIONS_T Type, *305-11*
ENQUOTE_LITERAL Function, *30-2*
ENQUOTE_NAME Function, *30-2*
ENV_T Type, *197-23*
EPC_TO_ORACLE_SCHEME function, *124-8*
EQUALS_HOST Function, *132-5*
ERASE Procedures, *114-27*
error queue
  deleting errors, *24-20*
  executing errors, *24-23*, *24-24*
  getting error messages, *24-26*
ERROR_DEPTH Function, *281-6*
ERROR_MSG Function, *281-7*
ERROR_NUMBER Function, *281-7*
ESCAPE function, *300-3*
ESCAPE_SC function
  of HTF package, *242-27*
ESCAPE_SC procedure
  of HTP package, *243-27*
ESCAPE_URL function
  of HTF package, *242-28*
ESCAPEURI function, *307-16*
ESTIMATE_CPU_UNITS function, *133-1*
ESTIMATE_MVIEW_SIZE Procedure, *129-3*
ETINSTANCE member function
  of ANYDATASET TYPE, *303-6*
EVALUATE procedure, *165-3*
EVOLVE_SQL_PLAN_BASELINE Function, *181-16*
EXACT_TEXT_SIGNATURES procedure, *135-4*
EXCHANGEPOSTPROC Procedure, *235-5*
EXCHANGEPREPROC Procedure, *235-5*
EXCLUDE procedure (of DBMS_SQL_FIREWALL, *183-14*
EXEC_DDL_STATEMENT Procedure, *210-14*
EXECUTE function, *182-51*
EXECUTE member procedure, *310-2*, *310-10*
EXECUTE procedure, *67-20*
EXECUTE_ADVISOR_TASK function, *192-70*
EXECUTE_ALL_ERRORS procedure, *24-23*
EXECUTE_ANALYSIS_TASK, *187-6*
EXECUTE_AND_FETCH function, *182-52*
EXECUTE_DIAGNOSIS_TASK Procedure, *186-11*
EXECUTE_ERROR procedure, *24-24*
EXECUTE_EVOLVE_TASK Function, *181-17*
EXECUTE_ILM Procedure, *97-4*
EXECUTE_ILM_TASK Procedur, *97-5*
EXECUTE_IMMEDIATE Procedure, *95-9*
EXECUTE_NON_QUERY Function, *95-10*
EXECUTE_TASK Procedure, *17-21*

EXECUTE_TUNING_TASK Function ((amp)) Procedure, *190-37*
EXECUTE_TUNING_TASK Function and Procedure, *34-2*
EXECUTE_UPDATE Procedure, *154-9*
execution flow
in dynamic SQL, *182-5*
EXISTSNODE function, *315-5*
EXISTSRESOURCE Function, *223-10*
EXPAND_MESSAGE function, *171-6*
EXPAND_SQL_TEXT Procedure, *210-14*
EXPLAIN_SQL_TESTCASE Function, *186-12*
Exponential Smoothing Settings, *57-27*
EXPONENTIAL_DIST_FIT procedure, *191-1*
EXPORT Procedure, *84-5*
EXPORT_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-14*
EXPORT_ANALYZE Procedure, *84-5*
EXPORT_AWR Procedure, *215-6*, *216-12*
EXPORT_COLUMN_STATS procedure, *192-72*
EXPORT_DATABASE_PREFS Procedure, *192-73*
EXPORT_DATABASE_STATS procedure, *192-74*
EXPORT_DICTIONARY_STATS procedure, *192-75*
EXPORT_FIXED_OBJECTS_STATS procedure, *192-76*
EXPORT_INDEX_STATS procedure, *192-76*
EXPORT_PENDING_STATS Procedure, *192-77*
EXPORT_PROFILE Procedure, *185-10*
EXPORT_SCHEMA_PREFS Procedure, *192-78*
EXPORT_SCHEMA_STATS procedure, *192-79*
EXPORT_SQL_TESTCASE Procedures, *186-12*
EXPORT_SQL_TESTCASE_DIR_BY_INC Function, *186-16*
EXPORT_SQL_TESTCASE_DIR_BY_TXT Function, *186-18*
EXPORT_SYSTEM_STATS procedure, *192-80*
EXPORT_TABLE_PREFS Procedure, *192-80*
EXPORT_TABLE_STATS procedure, *192-81*
EXPORT_XML procedure, *54-37*, *54-39*
EXPORT_XML_TO_FILE procedure, *54-12*
EXPORTRMANBACKUP Procedure, *140-4*
EXTEND_MAPPINGS Procedure, *83-7*
EXTENT_HEAT_MAP Function, *90-3*
EXTRACT function, *315-5*

F

FAILED_AUTO_TASKS_REPORT Function, *15-14*
FASTSTART_DISABLE procedure, *101-6*
FASTSTART_ENABLE procedure, *101-7*
FASTSTART_MIGRATE_STORAGE procedure, *101-9*
FCLOSE procedure, *285-7*
FCLOSE_ALL procedure, *285-7*
FCOPY procedure, *285-8*
FEATURE_OPTIONS Table Type, *203-2*
FEATURE_T Record Type, *62-15*
FEATURENAME Function, *62-30*
FEATURES_T Table Type, *62-20*
FETCH_ROW Function, *95-10*
FETCH_ROWS function, *182-52*
FFLUSH procedure, *285-8*
FGETATTR procedure, *285-9*
FGETPOS function, *285-10*
FI_HORIZONTAL function, *85-1*
FI_TRANSACTIONAL function, *85-3*
FILECLOSE Procedure, *114-29*
FILECLOSEALL Procedure, *114-30*
FILEEXISTS Function, *114-30*
FILEGETNAME Procedure, *114-31*
FILEISOPEN Function, *114-32*
FILEOPEN Procedure, *114-33*
FILETYPE Record Type, *285-5*
FIND_AFFECTED_TABLE Procedure, *76-6*
FINDENTITY function, *228-28*
FINDNOTATION function, *228-28*
FINISH_CAPTURE Procedure, *215-7*
FINISH_REDEF_COLLECTION procedure, *177-9*
FINISH_REDEF_TABLE procedure, *154-9*
FINISH_REPORT_CAPTURE Procedure, *37-2*
FIX_CORRUPT_BLOCKS procedure, *156-7*
FLUSH Function ((amp)) Procedure, *160-4*
FLUSH procedure, *56-11*, *299-7*
FLUSH_DATA function and procedure, *149-6*
FLUSH_DATABASE_MONITORING_INFO procedure, *192-83*
FLUSH_LOGS procedure (of DBMS_SQL_FIREWALL), *183-15*
FLUSH_SQL_PLAN_DIRECTIVE Procedure, *180-4*
FLUSHCACHE Procedure, *64-10*
FONTCLOSE function
of HTF package, *242-29*
FONTCLOSE procedure
of HTP package, *243-27*
FONTOPEN function
of HTF package, *242-29*
FONTOPEN procedure
of HTP package, *243-28*
FOPEN function, *285-10*
FOPEN_NCHAR function, *285-12*
FORCE parameter
and job-to-instance affinity, *105-2*
FORMAT function, *311-7*
FORMAT_CALL_STACK Function, *210-15*
FORMAT_CELL function
of HTF package, *242-30*
FORMAT_ERROR_BACKTRACE Function, *210-15*

FORMAT_ERROR_STACK Function, *210-18*
FORMCHECKBOX function
of HTF package, *242-30*
FORMCHECKBOX procedure
of HTP package, *243-28*
FORMCLOSE function
of HTF package, *242-31*
FORMCLOSE procedure
of HTP package, *243-29*
FORMFILE function
of HTF package, *242-31*
FORMFILE procedure
of HTP package, *243-30*
FORMHIDDEN function
of HTF package, *242-32*
FORMHIDDEN procedure
of HTP package, *243-30*
FORMIMAGE function
of HTF package, *242-32*
FORMIMAGE procedure
of HTP package, *243-31*
FORMOPEN function
of HTF package, *242-33*
FORMOPEN procedure
of HTP package, *243-29*
FORMPASSWORD function
of HTF package, *242-34*
FORMPASSWORD procedure
of HTP package, *243-32*
FORMRADIO function
of HTF package, *242-34*
FORMRADIO procedure
of HTP package, *243-32*
FORMRESET function
of HTF package, *242-35*
FORMRESET procedure
of HTP package, *243-33*
FORMSELECTCLOSE function
of HTF package, *242-35*
FORMSELECTCLOSE procedure
of HTP package, *243-33*
FORMSELECTOPEN function
of HTF package, *242-36*
FORMSELECTOPEN procedure
of HTP package, *243-34*
FORMSELECTOPTION function
of HTF package, *242-37*
FORMSELECTOPTION procedure
of HTP package, *243-35*
FORMSUBMIT function
of HTF package, *242-37*
FORMSUBMIT procedure
of HTP package, *243-35*
FORMTEXT function
of HTF package, *242-38*
FORMTEXT procedure
of HTP package, *243-36*
FORMTEXTAREA function
of HTF package, *242-38*
FORMTEXTAREA procedure
of HTP package, *243-36*
FORMTEXTAREA2 function
of HTF package, *242-39*
FORMTEXTAREA2 procedure
of HTP package, *243-37*
FORMTEXTAREACLOSE function
of HTF package, *242-40*
FORMTEXTAREACLOSE procedure
of HTP package, *243-38*
FORMTEXTAREAOPEN function
of HTF package, *242-40*
FORMTEXTAREAOPEN procedure
of HTP package, *243-38*
FORMTEXTAREAOPEN2 function
of HTF package, *242-41*
FORMTEXTAREAOPEN2 procedure
of HTP package, *243-39*
FRAGMENT_DELETE Procedure, *114-34*
FRAGMENT_INSERT Function, *114-35*
FRAGMENT_MOVE Procedure, *114-36*
FRAGMENT_REPLACE Procedure, *114-37*
FRAME function
of HTF package, *242-42*
FRAME procedure
of HTP package, *243-39*
FRAMESETCLOSE function
of HTF package, *242-42*
FRAMESETCLOSE procedure
of HTP package, *243-40*
FRAMESETOPEN function
of HTF package, *242-43*
FRAMESETOPEN procedure
of HTP package, *243-40*
FREE_BLOCKS procedure, *178-7*
FREE_UNUSED_USER_MEMORY Procedure, *173-5*
FREEDOCFRAG procedure, *228-29*
FREEDOCUMENT procedure, *228-29*
FREEELEMENT procedure, *228-30*
FREENODE procedure, *228-30*
FREEPARSER procedure, *231-2*
FREEPROCESSOR procedure, *238-2*
FREERESOURCE Procedure, *225-3*
FREESTYLESHEET procedure, *238-2*
FREETEMPORARY Procedures, *114-38*
FREMOVE procedure, *285-13*
FRENAME procedure, *285-13*
FSEEK procedure, *285-14*

G

GATHER_DATABASE_STATS procedures, *192-83*
GATHER_DICTIONARY_STATS procedure, *192-87*
GATHER_FIXED_OBJECTS_STATS procedure, *192-91*
GATHER_INDEX_STATS Procedure, *192-92*
GATHER_PROCESSING_RATE Procedure, *192-94*
GATHER_SCHEMA_STATS procedures, *192-94*
GATHER_SYSTEM_STATS procedure, *192-99*
GATHER_TABLE_STATS procedure, *192-101*
GENERATE_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-16*
GENERATE_CAPTURE_SUBSET Procedure, *216-12*
GENERATE_RESULT Procedure, *147-7*
GENERATE_SCHEMA Function, *106-5*
GENERATE_STATS procedure, *192-105*
GENERATE_TASK_NAME Function, *137-11*
GET function
of OWA_COOKIE package, *245-2*
GET_ AUDIT_TRAIL_PROPERTY_VALUE Function, *31-18*
GET_ACLOID function, *226-4*
GET_ADVISOR_OPR_FILTER procedure, *192-106*
GET_ADVISOR_RECS function, *192-107*
GET_ALL procedure, *245-2*
GET_ALL_DAD_ATTRIBUTES Procedure, *78-7*
GET_ALL_DAD_MAPPINGS Procedure, *78-8*
GET_ALL_GLOBAL_ATTRIBUTES Procedure, *78-9*
GET_ALL_NAMES member function, *312-4*
GET_APPLY_HWM_SEQID function, *120-3*
GET_ASH_QUERY Function, *15-14*
GET_AUDIT_COMMIT_DELAY function, *31-17*
GET_AUTHENTICATION procedure, *286-36*
GET_BASE_TABLE_NAME member function, *310-2*
GET_BASE_TABLE_OWNER member function, *310-2*
GET_BLOCK_MAP function, *144-4*
GET_BODY_CHARSET procedure, *286-37*
GET_CALL_ERROR_MSG function, *209-3*
GET_CAPTURE_INFO Function, *215-8*
GET_CATEGORY function, *213-3*
GET_CATEGORY_ID function, *124-11*
GET_CGI_ENV function, *252-7*
GET_CLIENT_ATTRIBUTES Procedure, *38-4*
GET_CLIENT_HOSTNAME function, *250-1*
GET_CLIENT_IP function, *250-2*
GET_COL procedure, *197-30*
GET_COLUMN_STATS procedures, *192-108*
GET_COMMAND_TYPE member function, *310-34*
GET_COMMIT_SCN member function, *310-34*
GET_COMMIT_SCN_FROM_POSITION static function, *310-34*
GET_COMMIT_TIME member function, *310-34*
GET_COMPATIBLE member function, *310-34*
GET_COMPONENT function, *311-8*
GET_COMPONENTS function, *124-12*
GET_COMPRESSION_RATIO Procedure, *4-1*, *48-4*
GET_COMPRESSION_TYPE Function, *48-12*
GET_COOKIE_COUNT function, *286-38*
GET_COOKIE_SUPPORT procedure, *286-38*
GET_COOKIES function, *286-39*
GET_CPU_TIME Function, *210-18*
GET_CREDENTIAL_NAME function, *143-17*
GET_CURRENT_LIFESPAN_DIGEST Function, *83-7*
GET_CURRENT_SCHEMA member function, *310-2*
GET_DAD_ATTRIBUTE Function, *78-9*
GET_DAD_LIST Procedure, *78-10*
GET_DBFS_LINK Functions, *114-39*
GET_DBFS_LINK_STATE Procedure, *114-40*
GET_DEPENDENCY Procedure, *210-18*
GET_DETAILED_EXCP_SUPPORT procedure, *286-39*
GET_DETAILED_SQLCODE function, *286-40*
GET_DETAILED_SQLERRM function, *286-40*
GET_DIVERGING_STATEMENT Function, *216-13*
GET_EDITION_NAME member function, *310-2*
GET_ENCODINGS function, *124-13*
GET_ENDIANNESS Function, *210-19*
GET_ERROR_MESSAGE function, *24-26*
GET_ETAG function, *244-2*
GET_EXPRESSION function, *58-28*
GET_EXTRA_ATTRIBUTE member function, *310-34*
GET_FASTSTART_TABLESPACE function, *101-10*
GET_FILE procedure, *81-4*
GET_FIX_CONTROL Function, *186-19*
GET_FOLLOW_REDIRECT procedure, *286-40*
GET_GLOBAL_ATTRIBUTE Function, *78-10*
GET_HASH_VALUE Function, *210-19*
GET_HEADER procedure, *286-41*
GET_HEADER_BY_NAME procedure, *286-42*
GET_HEADER_COUNT function, *286-42*
GET_HOST_ADDRESS function, *289-2*
GET_HOST_NAME function, *289-3*
GET_INDEX_DATAGUIDE Function, *107-9*
GET_INDEX_STATS procedures, *192-110*
GET_INDEXES function, *67-21*
GET_JAVA_LOGGING_LEVEL function, *124-13*

GET_JSON_INFORMATION Member Function, *310-10*
GET_LAST_ARCHIVE_TIMESTAMP function, *31-19*
GET_LEVEL function, *244-2*
GET_LINE function, *299-8*
GET_LINE procedure, *136-7*, *285-15*
GET_LINE_MAP function, *67-23*
GET_LINE_NCHAR Function, *299-8*
GET_LINE_NCHAR procedure, *285-16*
GET_LINES procedure, *136-8*
GET_LOB_DEDUPLICATION_RATIO Function, *114-41*
GET_LOB_INFORMATION member function, *310-10*
GET_LOB_OPERATION_SIZE member procedure, *310-10*
GET_LOCATION_URI function, *143-17*
GET_LOG procedure, *56-12*
GET_LOG_SPEC function, *56-13*
GET_LOGON_USER member function, *310-2*
GET_LONG_INFORMATION member function, *310-10*
GET_LTXID_OUTCOME Procedure, *20-6*
GET_MAX_STREAMS_POOL Procedure, *26-33*
GET_MESSAGE_TRACKING function, *239-64*
GET_MIN_STREAMS_POOL Procedure, *26-34*
GET_MORE_SOURCE procedure, *67-22*
GET_MV_SQL_FOR_AV_CACHE Function, *91-6*
GET_MV_SQL_FOR_STAR_CACHE Function, *91-6*
GET_NEXT_HIT function, *165-9*, *165-10*
GET_NEXT_RESULT Procedures, *182-53*
GET_NLS_PARAMETER Function, *68-5*
GET_NODE_ID_LOCAL function, *206-5*, *209-3*
GET_NODE_NAME_LOCAL function, *206-6*
GET_NUMBER_COL_CHUNK Procedure, *137-12*
GET_OBJECT_NAME member function, *310-34*
GET_OBJECT_OWNER member function, *310-34*
GET_OBJECT_TYPE member function, *310-2*
GET_OPATCH_BUGS Function, *151-3*
GET_OPATCH_COUNT Function, *151-4*
GET_OPATCH_DATA Function, *151-4*
GET_OPATCH_FILES Function, *151-4*
GET_OPATCH_INSTALL_INFO Function, *151-5*
GET_OPATCH_LIST Function, *151-5*
GET_OPATCH_LSINVENTORY, *151-5*
GET_OPATCH_OLAYS Function, *151-5*
GET_OPATCH_PREQS Function, *151-6*
GET_OPATCH_XSLT, *151-6*
GET_OWA_SERVICE_PATH function, *252-7*
GET_P1_RESOURCES Procedure, *38-4*
GET_PACKAGE_MEMORY_UTILIZATION Procedure, *173-7*
GET_PARAM function, *192-113*
GET_PARAMETER function, *56-14*
GET_PARAMETER_VALUE Function, *210-20*
GET_PASSWORD function, *250-2*
GET_PATCH_DETAILS Function, *151-6*
GET_PERSISTENT_CONN_COUNT function, *286-43*
GET_PERSISTENT_CONN_SUPPORT procedure, *286-43*
GET_PERSISTENT_CONNS procedure, *286-44*
GET_PLSQL_LOGGING_LEVEL function, *124-14*
GET_PLSQL_TRACE_LEVEL function, *199-6*
GET_POSITION member function, *310-34*
GET_POST_PROCESSED_SOURCE Functions, *146-3*
GET_PREFS function, *192-114*
GET_PREFS Function, *180-5*
GET_PROCEDURE function, *252-8*
GET_PROXY procedure, *286-45*
GET_RAW function, *285-17*, *299-9*
GET_REC_ATTRIBUTES Procedure, *17-23*
GET_REPLAY_DIRECTORY Function, *216-14*
GET_REPLAY_INFO Function, *216-14*
GET_REPLAY_TIMEOUT Procedure, *216-15*
GET_REPORT Function, *15-15*
GET_RESPONSE function, *286-45*
GET_RESPONSE_ERROR_CHECK procedure, *286-47*
GET_ROW_SET procedure, *197-32*
GET_ROW_TEXT member procedure, *310-10*
GET_ROWID function, *248-3*
GET_ROWID_CHUNK Procedure, *137-12*
GET_RUN_REPORT Function, *92-1*
GET_RUNTIME_INFO function, *67-24*
GET_SCHEME_NAMES function, *124-15*
GET_SCN member function, *310-34*
GET_SCN_FROM_POSITION static function, *310-34*
GET_SEGADV_ATTRIB Procedure, *179-7*
GET_SESSION_TIMEOUT function, *161-2*
GET_SOURCE_DATABASE_NAME member function, *310-34*
GET_SOURCE_TIME member function, *310-34*
GET_SQL Function, *186-20*
GET_SQL_HASH Function, *210-21*
GET_SQLPATCH_STATUS, *151-7*
GET_STATS_HISTORY_AVAILABILITY function, *192-124*
GET_STATS_HISTORY_RETENTION function, *192-124*
GET_SYS_CONTEXT Function, *83-8*
GET_SYSTEM_STATS procedure, *192-124*
GET_TABLE_STATS procedure, *192-126*
GET_TAG function, *239-65*
GET_TAG member function, *310-34*
GET_TARGET_ID function, *206-6*
GET_TASK_REPORT Procedure, *17-24*

GET_TASK_SCRIPT Procedure, *17-25*
GET_TDT_XML function, *124-15*
GET_TEXT function, *299-10*
GET_TEXT_NCHAR Function, *299-11*
GET_THREAD_NUMBER member function, *310-34*
GET_THRESHOLD procedure, *171-7*
GET_TIME Function, *210-22*
GET_TIMEOUT function, *161-3*
GET_TIMEOUT_BEHAVIOUR function, *67-24*
GET_TOPOLOGY_NAME_LOCAL function, *206-6*
GET_TRANSACTION_ID function, *200-4*
GET_TRANSACTION_ID member function, *310-34*
GET_TRANSACTION_TIMEOUT function, *200-4*
GET_TRANSACTION_TYPE function, *200-4*
GET_TRANSFER_TIMEOUT procedure, *286-47*
GET_TZ_TRANSITIONS Procedure, *210-22*
GET_USER_ID function, *250-2*
GET_USERID function, *226-4*
GET_VALIDATOR function, *124-17*
GET_VALUE function, *67-25*
GET_VALUE member function, *310-10*, *312-4*
GET_VALUE Procedure, *95-11*
GET_VALUE_RAW Procedure, *95-12*
GET_VALUES member function, *310-10*
GET_VERSION procedure, *149-6*
GET_VIEW_SQL function, *107-10*
GET_WARNING_SETTING_CAT function, *213-4*
GET_WARNING_SETTING_NUM function, *213-4*
GET_WARNING_SETTING_STRING function, *213-4*
GET_WATERMARK Procedure, *26-37*
GET_WHERE_CLAUSE member procedure, *310-10*
GET_WRITE_HWM_SEQID function, *120-3*
GET_X function, *247-2*
GET_XID function, *197-36*
GET_XML_INFORMATION member function, *310-10*
GET_Y function, *247-3*
GET* member functions
  of ANYDATA TYPE, *302-4*
  of ANYDATASET TYPE, *303-3*
GETACL Function, *225-4*
GETACLDOCFROMRES Function, *225-4*
GETACLDOCUMENT Function, *223-10*
GETANYTYPEFROMPERSISTENT static function
  of ANYTYPE TYPE, *304-7*
GETAPPLICATIONDATA Function, *227-8*
GETATTRELEMINFO member function
  of ANYTYPE TYPE, *304-6*
GETATTRIBUTE function, *228-31*
GETATTRIBUTENODE function, *228-32*
GETATTRIBUTES function, *228-33*
GETAUTHOR Function, *225-4*
GETBLOB function, *307-2*, *307-5*, *307-9*, *307-12*
GETBLOBVAL function, *315-6*
GETBUGSFORBUNDLE Procedures, *134-3*
GETCHARACTERSET Function, *225-5*
GETCHARSET Function, *228-33*
GETCHILDNODES function, *228-34*
GETCHILDOID Function, *227-8*
GETCHILDRENBYTAGNAME function, *228-34*
GETCHUNKSIZE Functions, *114-44*
GETCLOB function, *307-3*, *307-6*, *307-10*, *307-13*
GETCLOBVAL function, *315-6*
GETCOMMENT Function, *225-5*
GETCONTENTBLOB Function, *223-11*, *225-5*
GETCONTENTCLOB Function, *223-11*, *225-6*
GETCONTENTREF Function, *225-6*
GETCONTENTSBLOBBYRESID function, *224-3*
GETCONTENTSCLOBBYRESID function, *224-3*
GETCONTENTSXMLBYRESID function, *224-4*
GETCONTENTTYPE function, *307-3*, *307-6*, *307-10*, *307-14*
GETCONTENTTYPE Function, *225-6*
GETCONTENTTYPE Functions, *114-42*
GETCONTENTVARCHAR2 Function, *223-12*, *225-7*
GETCONTENTXML Function, *225-7*
GETCONTENTXMLREF Function, *223-12*
GETCONTENTXMLTYPE Function, *223-13*
GETCOUNT member function
  of ANYDATASET TYPE, *303-6*
GETCREATIONDATE Function, *225-7*
GETCREATOR Function, *225-8*
GETCURRENTUSER Function, *227-9*
GETCUSTOMMETADATA Function, *225-8*
GETDATA function, *228-35*
GETDEFAULTACL Procedure, *62-31*
GETDEFAULTASOF Procedure, *62-31*
GETDEFAULTCONTEXT Procedure, *62-32*
GETDEFAULTOWNER Procedure, *62-32*
GETDEFAULTPRINCIPAL Procedure, *62-33*
GETDISPLAYNAME Function, *225-9*
GETDOCTYPE function, *228-35*, *231-3*
GETDOCUMENT function, *231-3*
GETDOCUMENTELEMENT function, *228-36*
GETELEMENTSBYTAGNAME function, *228-36*
GETENTITIES function, *228-37*
GETEVENT Function, *227-9*
GETEXPANDEDNAME procedure, *228-38*
GETEXTERNALURL function, *307-3*, *307-7*, *307-11*, *307-14*
GETFEATURES Function, *63-10*
GETFEATURESBYMOUNT Function, *62-33*
GETFEATURESBYNAME Function, *62-34*
GETFEATURESBYPATH Function, *62-34*
GETFIRST Function, *227-10*
GETFIRSTCHILD function, *228-38*

GETFTPPORT Function, *221-11*
GETHANDLERLIST Function, *227-10*
GETHTTPCONFIGREALM Function, *221-11*
GETHTTPPORT Function, *221-11*
GETIMPLEMENTATION function, *228-39*
GETINFO member function
of ANYTYPE TYPE, *304-5*
GETINTERFACE Function, *227-11*
GETLANGUAGE Function, *225-9*, *227-11*, *227-13*
GETLASTCHILD function, *228-39*
GETLASTMODIFIER Function, *225-9*
GETLENGTH function, *228-40*
GETLENGTH Functions, *114-45*
GETLINK Function, *227-12*
GETLINKNAME Function, *227-12*
GETLISTENERENDPOINT Procedure, *221-12*
GETLISTENERS Function, *157-5*
GETLOCK Function, *227-13*
GETLOCKTOKEN Procedure, *223-13*
GETMODIFICATIONDATE Function, *225-10*
GETNAME function, *228-41*, *228-42*
GETNAME Function, *227-14*
GETNAMEDITEM function, *228-42*
GETNAMESPACE procedure, *228-42*
GETNEXTSIBLING function, *228-43*
GETNODENAME function, *228-44*
GETNODETYPE function, *228-43*
GETNODEVALUE function, *228-44*
GETNODEVALUEASBINARYSTREAM Function ((amp)) Procedure, *228-45*
GETNODEVALUEASCHARACTERSTREAM Function ((amp)) Procedure, *228-46*
GETNOTATIONNAME function, *228-46*
GETNOTATIONS function, *228-47*
GETNUMBERVAL function, *315-6*
GETNUMROWSPROCESSED function, *229-3*
GETOLDRESOURCE Function, *227-15*
GETOPENACCESSMODE Function, *227-15*
GETOPENDENYMODE Function, *227-16*
GETOPTIONS Functions, *114-46*
GETOUTPUTSTREAM Function, *227-16*
GETOWNER Function, *225-10*
GETOWNERDOCUMENT function, *228-48*
GETOWNERELEMENT function, *228-48*
GETPARAMETER Function, *227-17*
GETPARENT Function, *227-17*
GETPARENTNAME Function, *227-18*
GETPARENTNODE function, *228-49*
GETPARENTOID Function, *227-18*
GETPARENTPATH Function, *227-19*
GETPAT procedure, *249-6*
GETPATH Function, *227-19*
GETPATH Procedures, *62-34*, *63-10*
GETPATHBYMOUNTID Function, *62-37*
GETPATHBYREPOSID Function, *63-11*
GETPATHBYSTOREID Function, *62-37*
GETPATHNOWAIT Procedure, *63-12*
GETPATHNOWAIT Procedures, *62-38*
GETPREDECESSORS function, *224-4*
GETPREDSBYRESID function, *224-5*
GETPREFIX function, *228-49*
GETPREVIOUSSIBLING function, *228-50*
GETPRIVILEGES Function, *223-13*
GETPUBLICID function, *228-50*
GETQUALIFIEDNAME function, *228-51*
GETREFCOUNT Function, *225-10*
GETRELEASEVERSION function, *231-3*
GETREMOTEHTTPPORT Function, *221-12*
GETREMOTEHTTPSPORT Function, *221-12*
GETREPOSITORYRESCONFIG Function, *157-6*
GETREPOSITORYRESCONFIGPATHS Function, *157-6*
GETRESCONFIG Function, *157-7*
GETRESCONFIGPATHS Function, *157-7*
GETRESOID Function, *223-14*
GETRESOURCE Function, *227-20*
GETRESOURCEBYRESID function, *224-5*
GETROOTELEMENT function, *315-7*
GETRUL function, *307-7*
GETSCHEMA Function, *227-20*
GETSCHEMAANNOTATIONS Function, *233-6*
GETSCHEMANODE function, *228-51*
GETSCHEMAURL function, *315-7*
GETSIDXDEFFROMVIEW Function, *233-7*
GETSOURCE Function, *227-21*
GETSPECIFIED function, *228-52*
GETSTANDALONE function, *228-52*
GETSTATS Procedure, *62-40*
GETSTOREBYMOUNT Function, *62-39*
GETSTOREBYNAME Function, *62-39*
GETSTOREBYPATH Function, *62-40*
GETSTOREID Function, *63-13*
GETSTOREPROPERTY Function, *64-10*
GETSTRINGVAL function, *315-7*
GETSUCCESSORS function, *224-5*
GETSUCCSBYRESID function, *224-6*
GETSYSTEMID function, *228-53*
GETTAGNAME function, *228-53*
GETTARGET function, *228-47*
GETTOKENTABLEINFO Procedure ((amp)) Function, *53-2*
GETTOKENTABLEINFOBYTABLESPACE Procedure, *53-3*
GETTRACE Function, *62-40*
GETTYPE member function
of ANYDATA TYPE, *302-6*
of ANYDATASET TYPE, *303-7*
GETTYPENAME member function
of ANYDATA TYPE, *302-7*
of ANYDATASET TYPE, *303-7*
GETUPDATEBYTECOUNT Function, *227-21*
GETUPDATEBYTEOFFSET Function, *227-22*

GETURL function, *307-4*, *307-11*, *307-14*, *307-16*
GETVALIDATIONMODE function, *231-4*
GETVALUE function, *228-54*
GETVERSION function, *228-54*
GETVERSION Function, *62-41*, *63-13*
GETVERSIONID Function, *225-11*
GETXDB_TABLESPACE Function, *223-14*
GETXDBEVENT Function, *227-22*
GETXML function, *229-4*, *307-4*, *307-7*, *307-11*, *307-14*
GETXMLTYPE function, *228-55*, *229-5*
GG_PROCEDURE_REPLICATION_ON function, *87-11*
GoldenGate
  privileges, *88-1*
GoldenGate administrator, *88-1*
GRANT_ADMIN_PRIVILEGE procedure, *88-2*
GRANT_OBJECT_PRIVILEGE procedure, *166-14*
GRANT_REMOTE_ADMIN_ACCESS procedure, *240-5*
GRANT_SCHEMA_PRIVILEGE procedure, *166-15*
GRANT_SWITCH_CONSUMER_GROUP procedure, *159-1*
GRANT_SYSTEM_PRIVILEGE procedure, *159-2*, *166-16*

H

HANDLE_COLLISIONS procedure, *24-28*
HASACLCHANGED Function, *225-11*
HASATTRIBUTE function, *228-55*
HASAUTHORCHANGED Function, *225-11*
HASBLOBCONTENT Function, *223-14*
HASCHANGED Function, *225-12*
HASCHARACTERSETCHANGED Function, *225-12*
HASCHARCONTENT Function, *223-15*
HASCHILDNODES function, *228-56*
HASCOMMENTCHANGED Function, *225-12*
HASCONTENTCHANGED Function, *225-13*
HASCONTENTTYPECHANGED Function, *225-13*
HASCREATIONDATECHANGED Function, *225-14*
HASCREATORCHANGED Function, *225-14*
HASCUSTOMMETADATACHANGED Function, *225-14*
HASDISPLAYNAMECHANGED Function, *225-15*
HASFEATURE function, *228-57*
HASH function, *52-22*
HASH_LEN function, *52-23*
HASLANGUAGECHANGED Function, *225-15*
HASLASTMODIFIERCHANGED Function, *225-15*
HASMODIFICATIONDATECHANGED Function, *225-16*
HASOWNERCHANGED Function, *225-16*
HASREFCOUNTCHANGED Function, *225-17*
HASVERSIONIDCHANGED Function, *225-17*
HASXMLCONTENT Function, *223-15*
HASXMLREFERENCE Function, *223-16*
HEADCLOSE function
  of HTF package, *242-43*
HEADCLOSE procedure
  of HTP package, *243-41*
HEADER function
  of HTF package, *242-44*
HEADER procedure
  of HTP package, *243-41*
HEADOPEN function
  of HTF package, *242-44*
HEADOPEN procedure
  of HTP package, *243-42*
HELO function and procedure, *298-14*
HELP function, *298-15*
hierarchies
  validation log table, *91-3*, *91-7*
  verify validation, *91-9*
HR function
  of HTF package, *242-45*
HR procedure
  of HTP package, *243-42*
HTML tags
  applet tags
    functions, *242-2*
    procedures, *243-2*
  atags tags
    procedures, *243-2*
  character formatting tags
    functions, *242-2*
    procedures, *243-2*
  form tags
    functions, *242-2*
    procedures, *243-2*
  frame tags
    functions, *242-2*
    procedures, *243-2*
  list tags
    functions, *242-2*
    procedures, *243-2*
  paragraph formatting tags
    functions, *242-2*
    procedures, *243-2*
  table tags
    functions, *242-2*
HTMLCLOSE function
  of HTF package, *242-45*
HTMLCLOSE procedure
  of HTP package, *243-43*

HTMLOPEN function
of HTF package, *242-45*
HTMLOPEN procedure
of HTP package, *243-43*
HTP procedures, *243-1*
HTTP_HEADER_CLOSE procedure, *252-8*
HttpUriType, *307-4*
HTTPURITYPE function, *307-8*
HttpUriType subtype, *307-4*
CREATEURI function, *307-5*
GETBLOB function, *307-5*
GETCLOB function, *307-6*
GETCONTENTTYPE function, *307-6*
GETEXTERNALURL function, *307-7*
GETRUL function, *307-7*
GETXML function, *307-7*
HTTPURITYPE function, *307-8*
methods, *307-4*

I

IGNORE_AUTO_OPTIMIZE_CRITERIA Procedure, *98-7*
IGNORE_LIST Function, *160-5*
IGNORE_LIST_ADD Procedure, *160-5*
IM FastStart
DBMS_INMEMORY_ADMIN package, *101-1*
operational notes, *101-3*
security, *101-3*
IME_CAPTURE_EXPRESSIONS procedure, *101-11*
IME_CLOSE_CAPTURE_WINDOW procedure, *101-13*
IME_DROP_ALL_EXPRESSIONS procedure, *101-14*
IME_DROP_EXPRESSIONS procedure, *100-2*
IME_OPEN_CAPTURE_WINDOW procedure, *101-15*
IME_POPOULATE_EXPRESSIONS procedure, *101-16*
IMG function
of HTF package, *242-46*
IMG procedure
of HTP package, *243-43*
IMG2 procedure
of HTP package, *243-44*
IMPLEMENT_ADVISOR_TASK function, *192-128*
IMPLEMENT_EVOLVE_TASK Procedure, *181-18*
IMPLEMENT_TASK Procedure, *17-26*
IMPLEMENT_TUNING_TASK procedure, *190-38*
IMPORT Procedure, *84-6*, *106-7*
IMPORT_ALLOW_LIST procedure (of DBMS_SQL_FIREWALL), *183-17*
IMPORT_AWR Function, *216-15*
IMPORT_AWR Procedure, *215-8*
IMPORT_COLUMN_STATS procedure, *192-130*
IMPORT_DATABASE_PREFS Procedure, *192-132*
IMPORT_DATABASE_STATS procedure, *192-132*
IMPORT_DICTIONARY_STATS procedure, *192-134*
IMPORT_DISCOVERY_RESULT Procedure, *202-6*
IMPORT_FIXED_OBJECTS_STATS procedure, *192-135*
IMPORT_HISTORY Procedure, *83-9*
IMPORT_INDEX_STATS procedure, *192-136*
IMPORT_PROFILE Procedure, *185-11*
IMPORT_SCHEMA_PREFS Procedure, *192-137*
IMPORT_SCHEMA_STATS procedure, *192-138*
IMPORT_SENSITIVE_TYPES Procedure, *202-6*
IMPORT_SQL_TESTCASE Procedures, *186-20*
IMPORT_SQLWKLD_SCHEMA Procedure, *17-27*
IMPORT_SQLWKLD_SQLCACHE Procedure, *17-28*
IMPORT_SQLWKLD_STS Procedure, *17-30*
IMPORT_SQLWKLD_SUMADV Procedure, *17-31*
IMPORT_SQLWKLD_USER Procedure, *17-33*
IMPORT_SYSTEM_STATS procedure, *192-139*
IMPORT_TABLE_PREFS Procedure, *192-140*
IMPORT_TABLE_STATS procedure, *192-141*
IMPORT_XML procedure, *54-12*, *54-40*
IMPORTNODE function, *228-57*
INCIDENTID_2_SQL Procedure, *186-22*
INCLUDE procedure (of DBMS_SQL_FIREWALL), *183-18*
INCLUDE_EXTRA_ATTRIBUTES procedure, *44-26*
INDEX_TABLE_TYPE Table Type, *210-3*
INDEXXMLREFERENCES Procedure, *235-5*
INFER_AND_GENERATE_SCHEMA Function, *106-6*
INFER_SCHEMA Function, *106-3*
INIT_CLEANUP procedure, *31-20*
INIT_MIGRATION Procedure, *29-9*
INITFS Procedure, *65-6*
INITIALIZE function, *67-27*
INITIALIZE procedure, *152-2*
INITIALIZE_CONSOLIDATED_REPLAY Procedure, *216-16*
INITIALIZE_CUBE_UPGRADE procedure, *54-11*, *54-43*
INITIALIZE_REPLAY Procedure, *216-17*
INPLACEEVOLVE Procedure, *232-8*
INSERT_AUTOBIN_NUM_EQWIDTH procedure, *58-29*
INSERT_BIN_CAT_FREQ procedure, *58-33*
INSERT_BIN_NUM_EQWIDTH procedure, *58-37*
INSERT_BIN_NUM_QTILE procedure, *58-41*
INSERT_BIN_SUPER procedure, *58-43*
INSERT_CLIP_TRIM_TAIL procedure, *58-47*
INSERT_CLIP_WINSOR_TAIL procedure, *58-49*

INSERT_FINDING_DIRECTIVE Procedure, *15-15*
INSERT_MISS_CAT_MODE procedure, *58-52*
INSERT_MISS_NUM_MEAN procedure, *58-54*
INSERT_NORM_LIN_MINMAX procedure, *58-57*
INSERT_NORM_LIN_SCALE procedure, *58-59*
INSERT_NORM_LIN_ZSCORE procedure, *58-61*
INSERT_PARAMETER_DIRECTIVE Procedure, *15-16*
INSERT_PROCREP_EXCLUSION_OBJ procedure, *87-11*
INSERT_SEGMENT_DIRECTIVE Procedure, *15-17*
INSERT_SQL_DIRECTIVE Procedure, *15-18*
INSERTBEFORE function, *228-58*
INSERTDATA procedure, *228-59*
INSERTXML function, *236-4*
INSTALLDEFAULTWALLET Function, *220-2*
Installing the DBMS_CLOUD Package, *46-1*
INSTANCE procedure, *105-5*
INSTANCE_RECORD Record Type, *210-3*
INSTANCE_TABLE Table Type, *210-3*
instantiation
  aborting database preparation, *44-3*
  aborting schema preparation, *44-4*
  aborting table preparation, *44-5*
  global SCN, *24-41*
  preparing a database for, *44-27*
  preparing a schema for, *44-28*
  preparing a table for, *44-29*, *44-30*
  schema SCN, *24-70*
  table SCN, *24-73*
INSTR Functions, *114-47*
INTEGER_ARRAY Table Type, *173-1*
INTERNAL_VERSION_CHECK function, *149-7*
internet addressing
  using UTL_INADDR, *289-1*
INTERRUPT_ADVISOR_TASK procedure, *192-142*
INTERRUPT_ANALYSIS_TASK Procedure, *187-8*
INTERRUPT_DIAGNOSIS_TASK Procedure, *186-23*
INTERRUPT_EVOLVE_TASK Procedure, *181-18*
INTERRUPT_TASK Procedure, *17-34*
INTERRUPT_TUNING_TASK procedure, *190-39*
INTERVAL procedure, *105-6*
INTERVAL_DAY_TO_SECOND_TABLE, *182-22*
INTERVAL_YEAR_TO_MONTH_TABLE Table Type, *182-22*
INVALIDATE Functions ((amp)) Procedures, *160-7*
INVALIDATE Procedure, *210-23*
INVALIDATE_OBJECT Functions ((amp)) Procedures, *160-7*
IP_ADD_ACE procedure, *174-3*
IP_ADD_PDB_ACE procedure, *174-4*
IP_REMOVE_ACE procedure, *174-4*
IP_REMOVE_ACL procedure, *174-5*
IP_REMOVE_PDB_ACE procedure, *174-5*
IP_REMOVE_PDB_ACL procedure, *174-6*
IS_BIT_SET Function, *210-25*
IS_CLEANUP_INITIALIZED function, *31-22*
IS_CLUSTER_DATABASE Function, *210-26*
IS_FAST Procedure, *165-10*
IS_GG_XSTREAM_FOR_STREAMS function, *239-65*
IS_HIERARCHY_ENABLED function, *226-5*
IS_LOCATOR function, *282-1*
IS_NULL_TAG member function, *310-34*
IS_OPEN function, *182-54*, *285-18*
IS_PATCH_INSTALLED Function, *151-7*
IS_REPLAY_PAUSED Procedure, *216-19*
IS_ROLE_ENABLED Function, *173-8*
IS_SCHEMA_VALID function, *108-3*
IS_SESSION_ALIVE Function, *173-8*
IS_TRIGGER_FIRE_ONCE function, *66-6*
IS_VALID function, *108-3*
IS_VALID procedure, *108-4*
ISFOLDER Function, *223-16*, *225-17*
ISFRAGMENT function, *315-7*
ISGLOBALPORTENABLED Function, *221-12*
ISINDEX function
  of HTF package, *242-47*
ISINDEX procedure
  of HTP package, *243-45*
ISNULL function, *228-59*
ISNULL Function, *225-18*, *227-23*
ISOPEN function, *283-2*
ISOPEN Functions, *114-49*
ISPATHLOCKED Procedure, *62-41*, *63-14*
ISREMOTE Function, *114-50*
ISSCHEMABASED function, *315-7*
ISSCHEMAVALID function, *315-8*
ISSCHEMAVALIDATED function, *315-8*
ISTEMPORARY Functions, *114-51*
ITALIC function
  of HTF package, *242-48*
ITALIC procedure
  of HTP package, *243-45*
ITEM Functions, *228-62*

J

JARO_WINKLER Function, *292-3*
JARO_WINKLER_SIMILARITY Function, *292-4*
JSON data-guide formatting, *107-1*
JSON schema types, *108-1*
JSON Schema types, *107-1*
JSON_ARRAY_T object type, *309-8*
JSON_ELEMENT_T object type, *309-1*
JSON_KEY_LIST type, *309-10*
JSON_OBJECT_T object type, *309-4*
JSON_SCALAR_T object type, *309-10*

JSON_TABLE Table Type, *182-23*
JSON_TYPE_CONVERTIBLE_CHECK procedure, *107-11*

K

KBD function
of HTF package, *242-48*
KBD procedure
of HTP package, *243-46*
KEEP Procedure, *175-2*
KEEP_COLUMNS member procedure, *310-10*
KEEP_COLUMNS procedure, *239-66*
KEYBOARD function
of HTF package, *242-49*
KEYBOARD procedure
of HTP package, *243-46*

L

LAPACK Driver Routines (Linear Equations) Subprograms, *293-5*
LAPACK Driver Routines (LLS and Eigenvalue Problems), *293-6*
LAPACK_GBSV Procedures, *293-60*
LAPACK_GEES Procedures, *293-61*
LAPACK_GEEV Procedures, *293-71*
LAPACK_GELS Procedures, *293-63*
LAPACK_GESDD Procedures, *293-65*
LAPACK_GESV Procedures, *293-67*
LAPACK_GESVD Procedures, *293-69*
LAPACK_GTSV Procedures, *293-73*
LAPACK_PBSV Procedures, *293-75*
LAPACK_POSV Procedures, *293-76*
LAPACK_PPSV Procedures, *293-78*
LAPACK_PTSV Procedures, *293-80*
LAPACK_SBEV Procedures, *293-81*
LAPACK_SBEVD Procedures, *293-83*
LAPACK_SPEV Procedures, *293-85*
LAPACK_SPEVD Procedures, *293-86*
LAPACK_SPSV Procedures, *293-88*
LAPACK_STEV Procedures, *293-90*
LAPACK_STEVD Procedures, *293-91*
LAPACK_SYEV Procedures, *293-93*
LAPACK_SYEVD Procedures, *293-94*
LAPACK_SYSV Procedures, *293-95*
LAST_ERROR_POSITION function, *182-55*
LAST_ROW_COUNT function, *182-55*
LAST_ROW_ID function, *182-55*
LAST_SQL_FUNCTION_CODE function, *182-56*
LCR_TO_XML Function, *126-18*
LCR$_DDL_RECORD type, *310-2*
LCR$_ROW_LIST type, *310-44*
LCR$_ROW_RECORD type, *310-10*
LCR$_ROW_UNIT type, *310-45*
GET_LOB_INFORMATION member function, *310-10*
GET_LOB_OPERATION_SIZE member procedure, *310-10*
GET_LONG_INFORMATION member function, *310-10*
GET_POSITION member function, *310-34*
GET_ROOT_NAME member function, *310-34*
GET_ROW_TEXT member procedure, *310-10*
GET_WHERE_CLAUSE member procedure, *310-10*
SET_LOB_INFORMATION member procedure, *310-10*
SET_LOB_OPERATION_SIZE member procedure, *310-10*
LENGTH function, *294-18*
LEXICAL_DEPTH Function, *281-8*
LINE function
of HTF package, *242-49*
LINE procedure
of HTP package, *243-47*
LINK Procedures, *223-16*
LINKREL function
of HTF package, *242-50*
LINKREL procedure
of HTP package, *243-47*
LINKREV function
of HTF package, *242-50*
LINKREV procedure
of HTP package, *243-48*
LIST Function, *62-41*, *63-14*
LIST_CONTEXT Procedures, *173-9*
LISTALLCONTENT Function, *62-42*
LISTALLPROPERTIES Function, *62-42*
LISTBUNDLESWITHFCFIXES Procedures, *134-3*
LISTHEADER function
of HTF package, *242-51*
LISTHEADER procedure
of HTP package, *243-49*
LISTINGCLOSE function
of HTF package, *242-52*
LISTINGCLOSE procedure
of HTP package, *243-49*
LISTINGOPEN function
of HTF package, *242-52*
LISTINGOPEN procedure
of HTP package, *243-49*
LISTITEM function
of HTF package, *242-52*
LISTITEM procedure
of HTP package, *243-50*
LISTMOUNTS Function, *62-42*
LISTPRINT procedure, *252-8*
LISTSTORES Function, *62-43*

LNAME_ARRAY Table Type, *173-1*, *210-4*
LOAD_LONG_SQLTEXT Procedure, *216-19*
LOAD_PLANS_FROM_AWR Function, *181-19*
LOAD_PLANS_FROM_CURSOR_CACHE Functions, *181-20*
LOAD_PLANS_FROM_SQLSET Function, *181-21*
LOAD_SQLSET procedure, *189-11*, *190-40*
LOAD_SQLSET_FROM_TCB Function, *186-23*
LOAD_TABLE procedure
  of DBMS_HS_PARALLEL package, *94-4*
LOADBLOBFROMFILE Procedure, *114-52*
LOADCLOBFROMFILE Procedure, *114-53*
LOBs
  DBMS_LOB package, *114-1*
LOCAL_AWR_DBID function, *217-33*
LOCAL_TRANSACTION_ID function, *200-5*
LOCK_DOWN_APPLICATION Procedure, *83-10*
LOCK_MAP procedure, *193-4*
LOCK_OBJECT procedure, *296-4*
LOCK_PARTITION_STATS Procedure, *192-144*
LOCK_SCHEMA_STATS procedure, *192-144*
LOCK_TABLE_STATS procedure, *192-145*
LOCKPATH Procedure, *62-43*, *63-15*
LOCKRESOURCE Function, *223-17*
log apply services
  managing initialization parameters for logical standby databases, *117-1*
logical change records (LCRs)
  DDL LCRs, *310-2*
    getting current schema, *310-2*
  determining if tag is NULL, *310-34*
  executing, *310-2*, *310-10*
  extra attributes
    excluding, *44-26*
    including, *44-26*
  GET_THREAD_NUMBER member function, *310-34*
  getting command type, *310-34*
  getting commit SCN, *310-34*
  getting commit SCN from position, *310-34*
  getting commit time, *310-34*
  getting compatibility information, *310-34*
  getting extra attributes, *310-34*
  getting LCR creation time, *310-34*
  getting object name, *310-34*
  getting object owner, *310-34*
  getting SCN, *310-34*
  getting SCN from position, *310-34*
  getting source database name, *310-34*
  getting tag, *310-34*
  getting transaction identifier, *310-34*
  LCR$_DDL_RECORD type, *310-2*
  LCR$_ROW_LIST type, *310-44*
  LCR$_ROW_RECORD type, *310-10*
  LCR$_ROW_UNIT type, *310-45*
logical change records (LCRs) *(continued)*
  row LCRs, *310-10*
  setting command type, *310-34*
  setting extra attributes, *310-34*
  setting object name, *310-34*
  setting object owner, *310-34*
  setting source database name, *310-34*
  setting tag, *310-34*
  tracking, *239-64*
  types, *310-1*
logical standby databases
  managing with DBMS_LOGSTDBY package, *117-1*
logs
  Cube Build, *56-4*
  Cube Dimension Compile, *56-5*
  Cube Operations, *56-6*
  Cube Rejected Records, *56-7*
LONGEST_QUERY function, *207-3*
LZ_COMPRESS functions and procedures, *283-3*
LZ_COMPRESS_ADD procedure, *283-4*
LZ_COMPRESS_CLOSE procedure, *283-4*
LZ_COMPRESS_OPEN function, *283-5*
LZ_UNCOMPRESS functions and procedures, *283-6*
LZ_UNCOMPRESS_CLOSE procedure, *283-8*
LZ_UNCOMPRESS_EXTRACT procedure, *283-6*
LZ_UNCOMPRESS_OPEN function, *283-7*

M

machine learning
  algorithms, *57-1*
  anomaly detection, *57-2*
  association rules, *57-2*
  attribute importance, *57-2*
  automated, *145-1*
  classification, *57-2*
  clustering, *57-2*
  confusion matrix, *57-77*
  confusion matrix part, *57-83*
  cost matrix, *57-67*
  data transformation, *57-108*
  feature extraction, *57-2*
  functions, *57-3*
  lift, *57-89*
  lift part, *57-94*
  regression, *57-2*
  ROC, *57-99*, *57-103*
  scoring, *57-74*, *57-172*
  supervised, *57-1*
  time series, *57-2*
  transactional data, *57-15*
  transformations, *58-1*
  unsupervised, *57-1*
MAIL function and procedure, *298-15*

MAILTO function
of HTF package, *242-53*
MAILTO procedure
of HTP package, *243-50*
MAKE procedure, *155-4*
MAKE_DATA_BLOCK_ADDRESS Function, *210-26*
MAKEATTR function, *228-63*
MAKECDATASECTION function, *228-63*
MAKECHARACTERDATA function, *228-64*
MAKECOMMENT function, *228-64*
MAKEDOCUMENT function, *228-65*
MAKEDOCUMENT Function, *225-18*
MAKEDOCUMENTFRAGMENT function, *228-65*
MAKEDOCUMENTTYPE function, *228-66*
MAKEELEMENT function, *228-66*
MAKEENTITY function, *228-67*
MAKEENTITYREFERENCE function, *228-67*
MAKENODE function, *228-68*
MAKENOTATION function, *228-70*
MAKEPROCESSINGINSTRUCTION function, *228-70*
MAKETEXT function, *228-71*
MAKEVERSIONED function, *224-6*
MAP_ALL function, *193-4*
MAP_DAD Procedure, *78-11*
MAP_ELEMENT function, *193-4*
MAP_FILE function, *193-5*
MAP_OBJECT function, *193-6*
map_rec record type, *144-2*
MAPCLOSE function
of HTF package, *242-54*
MAPCLOSE procedure
of HTP package, *243-51*
MAPOPEN function
of HTF package, *242-54*
MAPOPEN procedure
of HTP package, *243-51*
MARK_RECOMMENDATION Procedure, *17-35*
MARKHOT Procedure, *175-4*
MATCH function, *249-6*
MATERIALIZE_DEFERRED_SEGMENTS Procedure, *179-8*
materialized view logs
master table
purging, *129-7*, *129-8*
materialized views
comparing, *47-1*
converting from table- to cube-organized, *54-29*
optimizing cube organized, *55-1*
refreshing, *129-9*, *129-11*, *129-13*
materialized views (cube organized), *54-2*
MEMORY_REPORT Procedure, *160-8*
MENULISTCLOSE function
of HTF package, *242-54*
MENULISTCLOSE procedure
of HTP package, *243-52*
MENULISTOPEN function
of HTF package, *242-55*
MENULISTOPEN procedure
of HTP package, *243-52*
merge streams, *239-68*, *239-71*
MERGE_COL_USAGE Procedure, *192-146*
MERGE_STREAMS procedure, *239-68*
MERGE_STREAMS_JOB procedure, *239-71*
MESSAGE_PROPERTIES_ARRAY_T Type, *305-19*
MESSAGE_PROPERTIES_T Type, *305-16*
META function
of HTF package, *242-55*
META procedure
of HTP package, *243-53*
MG2 function
of HTF package, *242-46*
MGD_ID constructor function, *311-4*
MGD_ID object type, *311-3*
FORMAT function, *311-7*
GET_COMPONENT function, *311-8*
MGD_ID constructor function, *311-4*
TO_STRING function, *311-9*
TRANSLATE function, *311-10*
MGD_ID object types, *311-1*
MGD_ID_COMPONENT object type, *311-2*
MGD_ID_COMPONENT_VARRAY object type, *311-2*
MIGRATE_STORED_OUTLINE Function, *181-22*
MIME_HEADER procedure, *252-9*
MIMEHEADER_DECODE function, *284-3*
MIMEHEADER_ENCODE function, *284-4*
model transparency, *58-6*
MODIFY_BASELINE_WINDOW_SIZE Procedure, *217-33*
MODIFY_PACKAGE_STATE Procedure, *173-9*
MODIFY_SERVICE Procedure, *21-12*, *172-12*
MODIFY_SNAPSHOT_SETTINGS procedure, *217-34*
MODIFY_TRANSFORMATION procedure, *201-3*
MODIFYPARAMETER Procedure, *230-4*
MOUNT_T Record Type, *62-16*
MOUNTS_T Table Type, *62-20*
MOUNTSTORE Procedure, *62-44*
MOVE_TO_DBFS_LINK Procedures, *114-56*
MOVEXDB_TABLESPACE Procedure, *220-2*
MSGID_ARRAY_T Type, *305-20*
MV_CUBE_ADVICE function, *55-1*

N

NAME_ARRAY Table Type, *210-4*
NAME_RESOLVE Procedure, *210-27*
NAME_TOKENIZE Procedure, *210-28*

NAMELIST Table Type, *181-3*
NAMESPACE function, *315-6*
NAMESPACE_ACL Function, *222-3*
NAMESPACE_METADATA Function, *222-4*
NAMESPACE_ORACLE Function, *222-4*
NAMESPACE_RESOURCE Function, *222-4*
NAMESPACE_RESOURCE_CONFIG Function, *222-5*
NAMESPACE_RESOURCE_EVENT Function, *222-5*
NAMESPACE_XDBSCHEMA Function, *222-5*
NAMESPACE_XMLDIFF Function, *222-5*
NAMESPACE_XMLINSTANCE Function, *222-6*
NAMESPACE_XMLSCHEMA Function, *222-6*
NAMESPACEIDTABLE Procedure, *53-4*
NEW_LINE procedure, *136-8*, *285-18*
NEW_ROW_LIST function and procedure, *251-3*
NEWCONTEXT function, *229-6*, *236-4*
NEWCONTEXTFROMHIERARCHY Function, *229-6*
NEWDOMDOCUMENT function, *228-71*
NEWPARSER function, *231-4*
NEWPROCESSOR function, *238-3*
NEWSTYLESHEET function, *238-3*
NEXT_DATE procedure, *105-6*
NEXT_ITEM_TYPE function, *143-17*
NL function
  of HTF package, *242-56*
NL procedure
  of HTP package, *243-53*
NOBR function
  of HTF package, *242-56*
NOBR procedure
  of HTP package, *243-54*
NOFRAMESCLOSE function
  of HTF package, *242-57*
NOFRAMESCLOSE procedure
  of HTP package, *243-54*
NOFRAMESOPEN function
  of HTF package, *242-57*
NOFRAMESOPEN procedure
  of HTP package, *243-55*
NOOP function and procedure, *298-16*
NOOP Functions, *30-4*
NORMAL function, *152-3*
NORMAL_DIST_FIT procedure, *191-2*
NORMALIZE procedure, *228-72*
NORMALIZEPATH Functions, *62-44*
NSPREFIX_ACL_ACL Function, *222-6*
NSPREFIX_RESCONFIG_RC Function, *222-6*
NSPREFIX_RESOURCE_R Function, *222-7*
NSPREFIX_XDB_XDB Function, *222-7*
NSPREFIX_XMLDIFF_XD Function, *222-7*
NSPREFIX_XMLINSTANCE_XSI Function, *222-7*
NSPREFIX_XMLSCHEMA_XSD Functio, *222-8*
NUMBER_ARRAY Table Type, *210-4*
NUMBER_TABLE Table Type, *182-23*
NVARRAY_ADD Procedure, *126-19*
NVARRAY_FIND_NAME Function, *126-19*
NVARRAY_FIND_NAME_TYPE Function, *126-20*
NVARRAY_GET Function, *126-21*
NVARRAY_GET_BOOLEAN, *126-21*
NVARRAY_GET_BYTE, *126-22*
NVARRAY_GET_DATE Function, *126-22*
NVARRAY_GET_DOUBLE Function, *126-23*
NVARRAY_GET_FLOAT Function, *126-23*
NVARRAY_GET_INTEGER, *126-24*
NVARRAY_GET_LONG Function, *126-24*
NVARRAY_GET_RAW Function, *126-25*
NVARRAY_GET_SHORT, *126-25*
NVARRAY_GET_TEXT Function, *126-26*

O

OBJECT_BLACK_LIST Function, *160-8*
OBJECT_BLACK_LIST_ADD Procedure, *160-8*
OBJECT_BLACK_LIST_CLEAR procedure, *160-9*
OBJECT_BLACK_LIST_REMOVE procedure, *160-9*
OBJECT_BLOCK_LIST Function, *160-10*
OBJECT_BLOCK_LIST_ADD Procedure, *160-10*
OBJECT_BLOCK_LIST_CLEAR procedure, *160-11*
OBJECT_BLOCK_LIST_REMOVE procedure, *160-11*
OBJECT_DEPENDENT_SEGMENTS function, *178-9*
OBJECT_GROWTH_TREND function, *178-10*
OBJECT_HEAT_MAP Procedure, *90-4*
OLAP aggregation operators, *54-19*
OLAP logs
  Cube Build, *56-4*
  Cube Dimension Compile, *56-5*
  Cube Operations, *56-6*
  Cube Rejected Records, *56-7*
OLAP metadata
  upgrading OLAP 10g to OLAP 12c, *54-8*
OLAP PL/SQL packages
  DBMS_AW_STATS, *41-1*
  DBMS_CUBE, *54-1*
  DBMS_CUBE_ADVISE, *55-1*
  DBMS_CUBE_LOG, *56-1*
OLAP templates
  validation, *54-47*
  writing to XML files, *54-39*
OLAP_DBA role, *54-2*
OLAP_USER role, *54-2*
OLD_CURRENT_SCHEMA Function, *210-29*
OLD_CURRENT_USER Function, *210-29*
OLISTCLOSE function
  of HTF package, *242-57*

OLISTCLOSE procedure
of HTP package, 243-55
OLISTOPEN function
of HTF package, 242-58
OLISTOPEN procedure
of HTP package, 243-55
ONLINE_INDEX_CLEAN Function, 156-8
OPATCH_COMPARE_CURRENT Function, 151-7
OPATCH_COMPARE_NODES Function, 151-8
OPEN Procedures, 114-57
OPEN_CONNECTION function, 299-12
OPEN_CONNECTION functions, 298-17
OPEN_CURSOR function, 182-56
OPEN_CURSOR Function, 95-13
OPEN_DATA function and procedure, 298-19
OPEN_WINDOW Procedure, 169-94
optimizer statistics, 41-1
OR REPLACE clause
for creating packages, 1-3
ORA-00600, 72-1
ORA-28104 error, 162-5
ORA-31703: cannot grant INHERIT PRIVILEGES privilege on behalf of other user, 122-15
Oracle Data Guard, 163-1
Oracle Data Mining
See Oracle Machine Learning for SQL
Oracle Machine Learning for SQL, 57-1
Oracle XStream
creating queues, 239-91
data dictionary
removing information, 239-74
Oracle-supplied types
logical change record (LCR) types, 310-1
rule types, 312-1
OVERLAY function, 294-18
OVERRIDE_PRIORITY Procedures, 38-5
OWA_COOKIE package, 245-1
OWA_PATTERN package, 249-1
OWA_SEC package, 250-1
OWNER Function, 281-8

P

PACK_MESSAGE procedures, 143-18
PACK_STGTAB procedure, 189-15
PACK_STGTAB_BASELINE Function, 181-24
PACK_STGTAB_DIRECTIVE Function, 180-5
PACK_STGTAB_SQLPATCH Procedure, 186-23
PACK_STGTAB_SQLPROF Procedure, 190-44
PACK_STGTAB_SQLSET Procedure, 190-45
package
DBMS_ODCI, 133-1
DBMS_XDB, 218-1
DBMS_XMLDOM, 228-1
DBMS_XMLGEN, 229-1
DBMS_XMLPARSER, 231-1
package *(continued)*
DBMS_XMLSCHEMA, 232-1
UriFactory, 307-15
Package - UriFactory, 307-15
package overview, 1-1
package variables
i_am_a_refresh, 129-6
packages
creating, 1-2
DBMS_MGD_ID_UTL, 124-1
referencing, 1-5
where documented, 1-5
PARA function
of HTF package, 242-58
PARA procedure
of HTP package, 243-56
PARAGRAPH function
of HTF package, 242-59
PARAGRAPH procedure
of HTP package, 243-56
PARALLEL_ENV_T Type, 197-24
PARAM function
of HTF package, 242-59
PARAM procedure
of HTP package, 243-57
PARSE procedure, 182-57, 231-4
PARSE Procedure, 95-14
PARSEBUFFER procedure, 231-5
PARSECLOB procedure, 231-5
PARSEDTD procedure, 231-6
PARSEDTDBUFFER procedure, 231-6
PARSEDTDCLOB procedure, 231-7
PATCHREPOSITORYRESCONFIGLIST Procedure, 157-8
PATH_ITEM_T Record Type, 62-16
PATH_ITEMS_T Table Type, 62-20
PATHIDTABLE Function, 53-4
PAUSE_PROFILER function and procedure, 149-7
PAUSE_REPLAY Procedure, 216-19
PIECEWISE member procedure
of ANYDATA TYPE, 302-7
of ANYDATASET TYPE, 303-8
PING procedure, 67-28
PKDECRYPT Function, 52-25
PKENCRYPT function, 52-26
PL/SQL
datatypes, 70-2
numeric codes for, 70-4
functions
DBMS_MGWADM package subprograms, 125-10
DBMS_MGWMSG package subprograms, 126-17

PL/SQL *(continued)*
procedures
DBMS_MGWADM package subprograms, *125-10*
DBMS_MGWMSG package subprograms, *126-17*
PL/SQL package
DBMS_HS_PARALLEL, *94-1*
PL/SQL packages
DBMS_ACTIVITY, *14-1*
DBMS_AUTO_MV, *36-1*
PLAINTEXT function
of HTF package, *242-60*
PLAINTEXT procedure
of HTP package, *243-57*
plan stability, *135-1*
PLSQL_TRACE_VERSION procedure, *199-6*
pointer to
CTX_ANL package, *4-1*
POISSON_DIST_FIT procedure, *191-2*
POLICY_CONDITIONS Table Type, *203-2*
POPULATE procedure, *100-3*, *120-3*
POPULATE_DIVERGENCE Procedure, *216-20*
POPULATE_WAIT function, *101-16*
PORT_STRING Function, *210-29*
PRECLOSE function
of HTF package, *242-60*
PRECLOSE procedure
of HTP package, *243-58*
predictive analytics, *145-1*
PREOPEN function
of HTF package, *242-61*
PREOPEN procedure
of HTP package, *243-58*
PREPARE Procedure, *216-22*
PREPARE_COLUMN_VALUES procedures, *192-147*
PREPARE_COLUMN_VALUES_NVARCHAR2 procedure, *192-149*
PREPARE_COLUMN_VALUES_ROWID procedure, *192-151*
PREPARE_CONSOLIDATED_REPLAY Procedure, *216-20*
PREPARE_GLOBAL_INSTANTIATION procedure, *44-27*
PREPARE_SCHEMA_INSTANTIATION procedure, *44-28*
PREPARE_SYNC_INSTANTIATION function, *44-29*
PREPARE_TABLE_INSTANTIATION procedure, *44-30*
PREVIEW_ILM Procedure, *97-6*
PRINT function
of HTF package, *242-61*
PRINT procedure
of HTP package, *243-59*
PRINT_BACKTRACE procedure, *67-28*
PRINT_CGI_ENV procedure, *252-10*
PRINT_INSTANTIATIONS procedure, *67-29*
PRINT_MULTI procedure, *251-3*
PRINT_POST_PROCESSED_SOURCE Procedure, *146-4*
PRINT_ROW_LIST procedure, *251-4*
PRINTS procedure
of HTP package, *243-59*
PRINTWARNINGS Procedur, *233-7*
privileges
GoldenGate administrator, *88-2*, *88-5*
XStream administrator, *240-2*, *240-6*
PRN function
of HTF package, *242-62*
PRN procedure
of HTP package, *243-60*
PROBE_VERSION procedure, *67-30*
PROCESS_CAPTURE Procedure, *216-26*
PROCESS_CONNECT_STRING Procedure, *68-5*
PROCESS_PENDING Procedure, *230-4*
PROCESSLINKS Procedure, *223-18*
PROCESSXSL function, *238-4*
PROGRAM_INFO Record Type, *67-12*
PROP_ITEM_T Record Type, *62-18*
PROP_ITEMS_T Table Type, *62-21*
propagations
altering, *150-2*
creating, *150-3*, *239-8*, *239-20*, *239-31*, *239-39*
DBMS_PROPAGATION_ADM package, *150-1*
dropping, *150-6*
rules
defining global, *239-8*
defining schema, *239-20*
defining subset, *239-31*
defining table, *239-39*
starting, *150-7*
stopping, *150-7*
PROPANY Functions, *62-45*
PROPERTIES_T Table Type, *62-21*
PROPERTIESH2T Function, *62-46*
PROPERTIEST2H Function, *62-46*
PROPERTY_T Record Type, *62-19*
PROPNUMBER Function, *62-47*
PROPRAW Function, *62-47*
PROPTIMESTAMP Function, *62-47*
PROPVARCHAR2 Function, *62-48*
PS procedure
of HTP package, *243-61*
PUBLISH_PENDING_STATS Procedure, *192-153*
PURGE procedure, *143-19*
PURGE Procedure, *175-5*
PURGE_COMPARISON Procedure, *47-18*
PURGE_CONTEXT Procedure, *83-10*

PURGE_LOG procedure, *183-18*
PURGE_LOST_DB_ENTRY procedure, *200-5*
PURGE_MIXED procedure, *200-7*
PURGE_PROCESSED_CHUNKS Procedure, *137-13*
PURGE_QUEUE_MESSAGES Procedure, *29-11*
PURGE_QUEUE_TABLE Procedure, *26-39*
PURGE_REFRESH_STATS procedure, *130-2*
PURGE_REFRESH_STATS Procedure, *194-10*
PURGE_ROW_CACHE procedure, *175-6*
PURGE_SOURCE_CATALOG procedure, *239-74*
PURGE_SQL_DETAILS Procedure, *217-35*
PURGE_STATS procedure, *192-154*
PURGE_TOMBSTONES procedure, *87-12*
PURGEALL Procedure, *62-48*, *63-16*
PURGELDAPCACHE function, *226-6*
PURGEPATH Procedure, *62-49*, *63-16*
PURGERESOURCEMETADATA Procedure, *223-18*
PURGESCHEMA Procedure, *232-10*
PUT procedure, *285-19*
PUT procedures, *136-9*
PUT_COL procedure, *197-36*
PUT_FILE procedure, *81-6*
PUT_LINE Procedure, *285-19*
PUT_LINE procedures, *136-10*
PUT_LINE_NCHAR procedure, *285-20*
PUT_NCHAR procedure, *285-21*
PUT_RAW function, *285-24*
PUT_ROW_SET procedure, *197-40*
PUTF procedure, *285-22*
PUTF_NCHAR procedure, *285-23*
PUTPATH Procedures, *62-49*, *63-17*

Q

QNAMEIDTABLE Function, *53-4*
QUALIFIED_SQL_NAME Function, *30-4*
QUERY_LINK_INFO procedure, *206-7*
QUERY_NODE_INFO procedure, *206-7*
QUEUE_PROPS_T Type, *305-12*
queues
  AnyData
    creating, *239-91*
    removing, *239-77*
QUICK_TUNE Procedure, *17-36*
QUIT function and procedure, *298-20*
QUOTED_PRINTABLE_DECODE function, *284-5*
QUOTED_PRINTABLE_ENCODE function, *284-5*

R

RANDOM procedure, *152-3*
RANDOMBYTES Function, *52-26*
RANDOMINTEGER Function, *52-27*
RANDOMNUMBER Function, *52-27*
RBU_MIGRATION function, *207-3*
RCPT function, *298-20*
RE$ATTRIBUTE_VALUE type, *312-3*
RE$ATTRIBUTE_VALUE_LIST type, *312-3*
RE$COLUMN_VALUE type, *312-3*, *312-7*
RE$COLUMN_VALUE_LIST type, *312-4*
RE$NAME_ARRAY type, *312-4*
RE$NV_ARRAY type, *312-4*
RE$NV_LIST type, *312-4*
  ADD_PAIR member procedure, *312-4*
  GET_ALL_NAMES member function, *312-4*
  GET_VALUE member function, *312-4*
  REMOVE_PAIR member procedure, *312-4*
RE$RULE_HIT type, *312-7*, *312-8*
RE$RULE_HIT_LIST type, *312-8*
RE$TABLE_ALIAS type, *312-8*
RE$TABLE_ALIAS_LIST type, *312-9*
RE$TABLE_VALUE type, *312-10*
RE$TABLE_VALUE_LIST type, *312-10*, *312-11*
RE$VARIABLE_TYPE type, *312-11*
RE$VARIABLE_TYPE_LIST type, *312-13*
RE$VARIABLE_VALUE type, *312-13*
RE$VARIABLE_VALUE_LIST type, *312-13*
READ Procedures, *114-58*
READ_CLIENT_INFO procedure, *23-2*
READ_LINE function, *299-14*
READ_LINE procedure
  of UTL_HTTP, *286-48*
READ_MODULE procedure, *23-2*
READ_ONLY procedure, *200-7*
READ_RAW function, *299-15*
READ_RAW procedure
  of UTL_HTTP, *286-49*
READ_TEXT function, *299-16*
READ_TEXT procedure
  of UTL_HTTP, *286-50*
READ_WRITE procedure, *200-8*
REAL_TIME_ADDM_REPORT Function, *15-19*
Real-Time SQL Monitoring subprograms, *190-9*
REASSOCIATE_FBA Procedure, *83-10*
REBUILD Procedure, *117-13*
REBUILD_FREELISTS procedure, *156-9*
REBUILDHIERARCHICALINDEX Procedure, *220-2*
RECEIVE_MESSAGE function, *143-20*
RECHECK Function, *47-18*
RECOMP_PARALLEL procedure, *295-2*
RECOMP_SERIAL procedure, *295-4*
RECONFIGCACHE Procedure, *64-11*
RECOVER Procedure, *140-5*
RECOVER_MIGRATION Procedure, *29-12*
RECOVER_OPERATION procedure, *239-75*
REDEF_TABLE Procedure, *154-10*
REDIRECT_URL procedure, *252-10*
REEXECUTE_FAILED_AUTO_TASKS Procedure, *15-20*

REFPARTITIONEXCHANGEIN Procedure, *235-6*
REFPARTITIONEXCHANGEOUT Procedure, *235-7*
refresh
  materialized views, *129-9*, *129-11*, *129-13*
REFRESH procedure, *155-6*
REFRESH Procedure, *36-7*
REFRESH_CATEGORY function, *124-20*
REFRESH_GROUPED_POLICY Procedure, *162-18*
REFRESH_MVIEW procedure, *54-7*, *54-44*
REFRESH_POLICY Procedure, *162-18*
REG_END Procedure, *50-19*
REGISTER Procedure, *19-4*, *106-9*
REGISTER_APPLICATION Procedure, *83-11*
REGISTER_DEPENDENT_OBJECT procedure, *154-12*
REGISTER_ERROR_TRANSLATION Procedure, *185-12*
REGISTER_FOREIGN_QUEUE Procedure, *125-34*
REGISTER_MVIEWS Procedure, *194-10*
REGISTER_NODE function and procedure, *206-8*
REGISTER_PARTITION_OPERATION Procedure, *194-11*
REGISTER_REMOTE_DATABASE procedure, *217-36*
REGISTER_SQL_TRANSLATION Procedure, *185-13*
REGISTERPARAMETER Procedure, *230-5*
REGISTERSCHEMA procedure, *232-10*
REGISTERSTORE Procedure, *62-51*
REGISTERSTORECOMMAND Procedur, *64-12*
REGISTERURI procedure, *232-14*
REGISTERURLHANDLER procedure, *307-17*
RELEASE function, *115-7*
REMAP_CONNECTION Procedure, *216-27*
REMAP_STAT_TABLE Procedure, *192-154*
REMAP_STGTAB procedure, *189-17*
REMAP_STGTAB_SQLPROF Procedure, *190-46*
REMAP_STGTAB_SQLSET Procedure, *190-47*
REMOVE procedure
  of DBMS_ALERT package, *19-5*
  of DBMS_JOB package, *105-7*
  of OWA_COOKIE package, *245-3*
REMOVE Procedure, *227-24*
REMOVE_AGENT Procedure, *125-35*
REMOVE_AUTO_CDR procedure, *87-12*
REMOVE_AUTO_CDR_COLUMN_GROUP procedure, *87-13*
REMOVE_AUTO_CDR_DELTA_RES procedure, *87-13*
REMOVE_CAPTURE Procedure, *216-28*
REMOVE_CATEGORY procedure, *124-21*
REMOVE_COLORED_SQL Procedure, *217-36*
REMOVE_DISCOVERY_RESULT Procedure, *202-7*
REMOVE_FROM_ILM Procedure, *97-6*
REMOVE_HOST_ACE Procedure, *131-17*
REMOVE_JOB Procedure, *125-36*
REMOVE_LINK procedure, *140-6*, *141-3*
REMOVE_MSGSYSTEM_LINK Procedure, *125-36*
REMOVE_OPTION Procedure, *125-37*
REMOVE_PAIR member procedure, *312-4*
REMOVE_PIPE function, *143-24*
REMOVE_POOL Procedure, *49-7*
REMOVE_PROXY procedure, *124-22*
REMOVE_QUEUE procedure, *239-77*
REMOVE_REFERENCE procedure, *189-18*
REMOVE_RULE procedure, *166-18*, *239-78*
REMOVE_SCHEDULE_ORDERING Procedure, *216-29*
REMOVE_SCHEME procedure, *124-22*
REMOVE_SQLSET_REFERENCE procedure, *190-49*
REMOVE_SUBSCRIBER Procedure, *125-38*
REMOVE_SUBSET_OUTBOUND_RULES procedure, *239-79*
REMOVE_TABLE_FROM_APPLICATION, *83-12*
REMOVE_WALLET_ACE Procedure, *131-18*
REMOVE_XSTREAM_CONFIGURATION procedure, *239-80*
REMOVEALL procedure, *19-5*
REMOVEANYSTORAGE Procedure, *233-8*
REMOVEATTRIBUTE procedure, *228-72*
REMOVEATTRIBUTENODE function, *228-73*
REMOVECHILD Function, *228-74*
REMOVEDEFAULTTABLE Procedure, *233-8*
REMOVEMAINTAINDOM Procedure, *233-9*
REMOVENAMEDITEM function, *228-74*
REMOVEOUTOFLINE Procedure, *233-9*
REMOVEPARAM procedure, *238-6*
REMOVESQLCOLLTYPE Procedure, *233-10*
REMOVESQLNAME Procedure, *233-11*
REMOVESQLTYPE Procedure, *233-12*
REMOVESQLTYPEMAPPING Procedure, *233-12*
REMOVETABLEPROPS Procedure, *233-13*
REMOVETIMESTAMPWITHTIMEZONE Procedure, *233-13*
RENAME_BASELINE Procedure, *217-37*
RENAME_COLUMN member procedure, *310-10*
RENAME_COLUMN procedure, *239-82*
RENAME_COLUMN Procedure, *107-12*
RENAME_QUEUE Procedure, *29-13*
RENAME_SCHEMA procedure, *239-84*
RENAME_TABLE procedure, *239-86*
RENAMECOLLECTIONTABLE Procedure, *235-8*
RENAMEPATH Procedure, *63-18*
RENAMEPATH Procedures, *62-51*
RENAMERESOURCE Procedure, *223-18*

REPLACECHILD function, *228-75*
REPLACEDATA procedure, *228-76*
REPLAY_SQL_TESTCASE Function, *186-24*
replication
split and merge, *239-68*, *239-71*, *239-93*
REPOPULATE procedure, *100-4*
REPORT Function, *215-9*, *216-29*
REPORT_ADVISOR_TASK function, *192-155*
REPORT_ANALYSIS_TASK Function, *187-8*
REPORT_AUTO_EVOLVE_TASK Function, *181-25*
REPORT_AUTO_TUNING_TASK function, *34-3*
REPORT_AUTO_TUNING_TASK Function, *190-50*
REPORT_COL_USAGE Procedure, *192-159*
REPORT_DIAGNOSIS_TASK Function, *186-27*
REPORT_EVOLVE_TASK Function, *181-26*
REPORT_GATHER_AUTO_STATS Function, *192-160*
REPORT_GATHER_DATABASE_STATS Functions, *192-162*
REPORT_GATHER_DICTIONARY_STATS Functions, *192-167*
REPORT_GATHER_FIXED_OBJ_STATS Function, *192-170*
REPORT_GATHER_SCHEMA_STATS Functions, *192-173*
REPORT_GATHER_TABLE_STATS Function, *192-178*
REPORT_PERFHUB Procedure, *142-2*
REPORT_REPOSITORY_DETAIL Function, *37-2*
REPORT_REPOSITORY_DETAIL_XML Function, *37-3*
REPORT_REPOSITORY_LIST_XML Procedure, *37-3*
REPORT_SESSION Function, *142-3*
REPORT_SINGLE_STATS_OPERATION Function, *192-183*
REPORT_SQL_DETAIL Function, *190-51*
REPORT_SQL_MONITOR Function, *184-4*, *190-56*
REPORT_SQL_MONITOR_LIST Function, *184-8*, *190-61*
REPORT_SQL_MONITOR_LIST_XML Function, *184-10*
REPORT_SQL_MONITOR_XML Function, *184-8*
REPORT_STATS_OPERATIONS Function, *192-185*
REPORT_TUNING_TASK function, *190-64*
REPORT_TUNING_TASK_XML function, *190-66*
REQ Type, *286-12*
REQUEST function, *115-8*, *286-51*
REQUEST_CONTEXT_KEY Type, *286-12*
REQUEST_PIECES function, *286-53*
REQUIRED_RETENTION function, *207-4*
REQUIRED_UNDO_SIZE function, *207-5*
RESET_ADVISOR_TASK procedure, *192-187*
RESET_ANALYSIS_TASK Procedure, *187-10*
RESET_BUFFER procedure, *143-24*
RESET_COL_USAGE Procedure, *192-188*
RESET_DIAGNOSIS_TASK Procedure, *186-28*
RESET_EVOLVE_TASK Procedure, *181-25*
RESET_JOB Procedure, *125-39*
RESET_PACKAGE Procedure, *173-13*
RESET_PARAM_DEFAULTS Procedure, *192-189*
RESET_SQLWKLD Procedure, *17-38*
RESET_SUBSCRIBER Procedure, *125-40*
RESET_TASK Procedure, *17-39*
RESET_TUNING_TASK procedure, *190-67*
RESETPARAMS procedure, *238-6*
RESOLVENAMESPACEPREFIX function, *228-76*
RESP Type, *286-13*
RESTARTQUERY procedure, *229-7*
RESTORE function, *193-6*
RESTORE_DATAFILE_PERMISSIONS Procedure, *75-2*
RESTORE_DATBASE_STATS procedure, *192-189*
RESTORE_DEFAULTS Procedure, *49-8*
RESTORE_DICTIONARY_STATS procedure, *192-190*
RESTORE_FIXED_OBJECTS_STATS procedure, *192-191*
RESTORE_SCHEMA_STATS procedure, *192-191*
RESTORE_SYSTEM_STATS procedure, *192-192*
RESTORE_TABLE_STATS procedure, *192-193*
RESTOREALL Procedure, *62-52*, *63-19*
RESTOREPATH Procedure, *62-53*, *63-19*
RESUME_ADVISOR_TASK procedure, *192-194*
RESUME_ANALYSIS_TASK Procedure, *187-10*
RESUME_DIAGNOSIS_TASK Procedure, *186-28*
RESUME_EVOLVE_TASK Procedure, *181-25*
RESUME_PROFILER function and procedure, *149-7*
RESUME_REPLAY Procedure, *216-30*
RESUME_TASK Procedure, *137-14*
RESUME_TUNING_TASK Procedure, *190-68*
RETURN_RESULT Procedures, *182-61*
REUSE_REPLAY_FILTER_SET Procedure, *216-30*
REVERSE function, *294-20*
REVOKE_ADMIN_PRIVILEGE procedure, *88-5*, *240-6*
REVOKE_OBJECT_PRIVILEGE procedure, *166-19*
REVOKE_REMOTE_ADMIN_ACCESS procedure, *240-8*
REVOKE_SCHEMA_PRIVILEGE procedure, *166-19*
REVOKE_SWITCH_CONSUMER_GROUP procedure, *159-3*

REVOKE_SYSTEM_PRIVILEGE procedure, *159-4*, *166-20*
ROLLBACK procedure, *200-8*
ROLLBACK Procedure, *154-13*
ROLLBACK_FORCE procedure, *200-8*
ROLLBACK_PLAN Procedure, *163-11*
ROLLBACK_SAVEPOINT procedure, *200-8*
rolling upgrade, *163-1*
row migration, *239-31*, *239-34*
ROW_LIST DATA TYPE, *251-1*
ROW_SET_T Type, *197-26*
ROWID datatype
  extended format, *164-10*
ROWID_BLOCK_NUMBER function, *164-4*
ROWID_CREATE function, *164-5*
ROWID_INFO procedure, *164-6*
ROWID_OBJECT function, *164-7*
ROWID_RELATIVE_FNO function, *164-8*
ROWID_ROW_NUMBER function, *164-8*
ROWID_TO_ABSOLUTE_FNO function, *164-9*
ROWID_TO_EXTENDED function, *164-10*
ROWID_TO_RESTRICTED function, *164-11*
ROWID_TYPE function, *164-12*
ROWID_VERIFY function, *164-12*
RSET Function and Procedure, *298-21*
rule sets
  adding rules to, *166-2*
  creating, *166-11*
  dropping, *166-13*
  removing rules from, *166-18*
rules
  action contexts
    adding name-value pairs, *312-4*
    getting name-value pairs, *312-4*
    getting value for name, *312-4*
    removing name-value pairs, *312-4*
  altering, *166-6*
  creating, *166-9*
  DBMS_RULE package, *165-1*
  DBMS_RULE_ADM package, *166-1*
  dropping, *166-12*
  evaluation, *165-3*
    iterators, *165-2*, *165-9*, *165-10*
  evaluation contexts
    altering, *166-4*
    creating, *166-8*
    dropping, *166-11*
  fast evaluation, *165-10*
  object privileges
    granting, *166-14*
    revoking, *166-19*
  propagations
    removing, *239-78*
  RE$ATTRIBUTE_VALUE type, *312-3*
  RE$ATTRIBUTE_VALUE_LIST type, *312-3*
  RE$COLUMN_VALUE type, *312-3*, *312-7*
rules *(continued)*
  RE$COLUMN_VALUE_LIST type, *312-4*
  RE$NAME_ARRAY type, *312-4*
  RE$NV_ARRAY type, *312-4*
  RE$NV_LIST type, *312-4*
  RE$RULE_HIT type, *312-7*, *312-8*
  RE$RULE_HIT_LIST type, *312-8*
  RE$TABLE_ALIAS type, *312-8*
  RE$TABLE_ALIAS_LIST type, *312-9*
  RE$TABLE_VALUE type, *312-10*
  RE$TABLE_VALUE_LIST type, *312-10*, *312-11*
  RE$VARIABLE_TYPE type, *312-11*
  RE$VARIABLE_TYPE_LIST type, *312-13*
  RE$VARIABLE_VALUE type, *312-13*
  RE$VARIABLE_VALUE_LIST type, *312-13*
  schema privileges
    granting, *166-15*
    revoking, *166-19*
  subset
    defining, *239-31*, *239-34*
  system privileges
    granting, *166-16*
    revoking, *166-20*
  system-created
    global apply, *239-12*
    global capture, *239-12*
    global propagation, *239-8*
    global schema, *239-24*
    removing, *239-78*
    schema capture, *239-24*
    schema propagation, *239-20*
    subset apply, *239-34*
    subset capture, *239-34*
    subset propagation, *239-31*
    table apply, *239-43*
    table capture, *239-43*
    table propagation, *239-39*
  types, *312-1*
RUN procedure, *105-7*
RUN_CHECK Procedure, *92-2*
RUN_TASK Procedure, *137-16*
RUNTIME_INFO Record Type, *67-13*

S

S function
  of HTF package, *242-63*
S procedure
  of HTP package, *243-61*
saga
  DBMS_SAGA package, *168-1*
    BEGIN_SAGA function, *168-2*
    COMMIT_SAGA procedure, *168-3*
    GET_SAGA_ID function, *168-2*
    IS_INCOMPLETE function, *168-4*

saga *(continued)*
DBMS_SAGA package *(continued)*
overview, *168-1*
ROLLBACK_SAGA procedure, *168-3*
security model, *168-1*
SEND_REQUEST procedure, *168-5*
SET_INCOMPLETE procedure, *168-4*
SET_SAGA_ID procedure, *168-2*
summary of subprograms, *168-1*
DBMS_SAGA_ADM package, *167-1*
ADD_BROKER procedure, *167-2*
ADD_COORDINATOR procedure, *167-3*
ADD_PARTICIPANT procedure, *167-4*
DROP_BROKER procedure, *167-2*
DROP_COORDINATOR procedure, *167-4*
DROP_PARTICIPANT procedure, *167-6*
overview, *167-1*
REGISTER_SAGA_CALLBACK procedure, *167-6*
security model, *167-1*
summary of subprograms, *167-1*
SAM (SQL Aggregation Management), *54-2*
SAMPLE function
of HTF package, *242-63*
SAMPLE procedure
of HTP package, *243-62*
SAVE function, *193-7*
SAVE Procedure, *225-19*
SAVEPOINT procedure, *200-9*
SCHEDULE_PROPAGATION Procedure, *26-42*, *125-40*
SCHEMA_NAME Function, *30-5*
SCHEMAELEM_RES_ACL Function, *222-8*
SCHEMAELEM_RESCONTENT_BINARY Function, *222-8*
SCHEMAELEM_RESCONTENT_TEXT Function, *222-9*
SCHEMAURL_ACL Function, *222-8*
SCHEMAURL_RESOURCE Function, *222-9*
SCHEMAURL_XDBSCHEMA Function, *222-9*
SCHEMAVALIDATE procedure, *315-8*
SCOPEXMLREFERENCES Procedure, *235-10*
SCRIPT function
of HTF package, *242-64*
SCRIPT procedure
of HTP package, *243-62*
SCRIPT_ADVISOR_TASK function, *192-195*
SCRIPT_TUNING_TASK Function, *190-71*
SDO_CS package, *253-1*
SDO_CSW_PROCESS package, *254-1*
SDO_GCDR package, *255-1*
SDO_GEOM package, *256-1*
SDO_GEOR package, *257-1*
SDO_GEOR_ADMIN package, *258-1*
SDO_GEOR_AGGR package, *259-1*
SDO_GEOR_RA package, *260-1*
SDO_GEOR_UTL package, *261-1*
SDO_LRS package documentation, *262-1*
SDO_MIGRATE package, *263-1*
SDO_NET package documentation, *264-1*
SDO_NFE package, *265-1*
SDO_OLS package, *266-1*
SDO_PC_PKG package, *267-1*
SDO_SAM package, *268-1*
SDO_TIN_PKG package, *269-1*
SDO_TOPO package, *270-1*
SDO_TOPO_MAP package, *271-1*
SDO_TUNE package, *272-1*
SDO_UTIL package, *273-1*
SDO_WFS_LOCK package, *274-1*
SDO_WFS_PROCESS package, *275-1*
SEARCH Function, *63-20*
SECURE_CONNECTION Procedure, *299-17*
security
IM FastStart, *101-3*
XStream, *239-2*
SEED procedures, *152-4*
SEED_COL_USAGE Procedures, *192-198*
SEEK_INPUT_T Type, *305-13*
SEEK_OUTPUT_T Type, *305-13*
SEGMENT_CORRUPT procedure, *179-8*
SEGMENT_DEALLOCATE_VERSIONS procedure, *100-4*
SEGMENT_DROP_CORRUPT procedure, *179-9*
SEGMENT_DUMP procedure, *179-10*
SEGMENT_FIX_STATUS procedure, *156-9*
SEGMENT_HEAT_MAP Procedure, *90-5*
SEGMENT_VERIFY procedure, *179-11*
SELECT_BASELINE_DETAILS Function, *217-37*
SELECT_BASELINE_METRIC Function, *217-38*
SELECT_CURSOR_CACHE function, *189-19*
SELECT_CURSOR_CACHE Function, *190-73*
SELECT_OBJECT procedure, *296-5*
SELECT_SQL_TRACE function, *189-24*
SELECT_SQL_TRACE Function, *190-77*
SELECT_SQLPA_TASK function, *189-27*, *190-80*
SELECT_SQLSET function, *189-28*, *190-81*
SELECT_WORKLOAD_REPOSITORY function, *189-30*
SELECT_WORKLOAD_REPOSITORY functions, *190-84*
SELECTNODES function, *238-6*
SELECTSINGLENODE function, *238-7*
SELF_CHECK procedure, *67-30*
SEM_APIS package, *276-1*
SEM_OLS package, *277-1*
SEM_PERF package, *278-1*
SEM_RDFSA package, *279-1*, *280-1*
SEND procedure, *245-3*, *291-2*
SEND_ATTACH_RAW procedure, *291-3*
SEND_ATTACH_VARCHAR2 procedure, *291-4*

SEND_MESSAGE function, *143-25*
SENDCOMMAND Procedures, *64-13*
SERV_MOD_ACT_STAT_DISABLE procedure, *128-5*
SERV_MOD_ACT_STAT_ENABLE procedure, *128-5*
SERV_MOD_ACT_TRACE_DISABLE procedure, *128-7*
SERV_MOD_ACT_TRACE_ENABLE procedure, *128-8*
SESSION _TRACE_DISABLE Procedure, *173-14*
SESSION _TRACE_ENABLE Procedure, *173-15*
SESSION_IS_ROLE_ENABLED Function, *173-14*
SESSION_TRACE_DISABLE procedure, *128-9*
SESSION_TRACE_ENABLE procedure, *128-10*
SET Procedure, *43-3*
SET_ACTION procedure, *23-3*
SET_ADVANCED_PARAMETER Procedure, *216-31*
SET_ADVISOR_TASK_PARAMETER procedure, *192-199*
SET_ANALYSIS_DEFAULT_PARAMETER Procedures, *187-15*
SET_ANALYSIS_TASK_PARAMETER Procedure, *187-11*
SET_ATTRIBUTE Procedure, *185-14*
SET_AUDIT_TRAIL_LOCATION procedure, *31-24*
SET_AUDIT_TRAIL_PROPERTY procedure, *31-26*
SET_AUTHENTICATION procedure, *286-56*
SET_AUTHENTICATION_FROM_WALLET Procedure, *286-57*
SET_AUTHORIZATION procedure, *250-3*
SET_AUTO_TUNING_TASK_PARAMETER Procedures, *34-4*
SET_BASE_TABLE_NAME member procedure, *310-2*
SET_BASE_TABLE_OWNER member procedure, *310-2*
SET_BODY_CHARSET procedures, *286-59*
SET_BREAKPOINT function, *67-31*
SET_CACHE Procedure, *244-2*
SET_CHUNK_STATUS Procedure, *137-17*
SET_CLIENT_INFO procedure, *23-4*
SET_CLIENT_SERVICE Procedure, *38-6*
SET_CNS_EXCEPTION_LOG procedure, *55-4*
SET_COLUMN_STATS procedures, *192-200*
SET_COMMAND_TYPE member procedure, *310-34*
SET_CONSUMER_GROUP_MAPPING procedure, *158-24*
SET_CONSUMER_GROUP_MAPPING_PRI procedure, *158-25*
SET_CONTEXT Procedure, *119-4*, *173-15*
SET_CONTEXT_LEVEL Procedure, *83-12*
SET_COOKIE_SUPPORT procedures, *286-60*
SET_CREDENTIAL_NAME procedure, *143-28*
SET_CURRENT_OPINST Procedure, *151-8*
SET_CURRENT_SCHEMA member procedure, *310-2*
SET_DAD_ATTRIBUTE Procedure, *78-12*
SET_DATA_LINKED procedure, *140-6*, *141-3*
SET_DATABASE_PREFS procedure, *192-202*
SET_DBFS_LINK Procedures, *114-60*
SET_DDL_TEXT member procedure, *310-2*
SET_DEFAULT_SQLWKLD_PARAMETER Procedure, *17-39*
SET_DEFAULT_TASK_PARAMETER Procedures, *17-40*
SET_DEFAULTS procedure, *19-5*
SET_DETAILED_EXCP_SUPPORT procedure, *286-61*
SET_DIAGNOSIS_TASK_PARAMETER Procedure, *186-28*
SET_DML_CONFLICT_HANDLER procedure, *24-28*
SET_DML_HANDLER procedure, *24-34*
SET_EDITION_DEFERRED Procedure, *173-17*
SET_EDITIONING_VIEWS_READ_ONLY Procedure, *77-2*
SET_ENQUEUE_DESTINATION procedure, *24-38*
SET_EVOLVE_TASK_PARAMETER Function, *181-27*
SET_EXCEPTION_MODE procedure, *209-4*
SET_EXECUTE procedure, *24-40*
SET_EXPIRES Procedure, *244-3*
SET_EXPRESSION procedure, *58-63*
SET_EXT_DATA_LINKED procedure, *140-7*, *141-4*
SET_EXTRA_ATTRIBUTE member procedure, *310-34*
SET_FIX_CONTROLS Procedures, *134-4*
SET_FOLLOW_REDIRECT procedures, *286-62*
SET_GLOBAL_ATTRIBUTE Procedure, *78-14*
SET_GLOBAL_INSTANTIATION procedure, *24-41*
SET_GLOBAL_PREFS procedure, *192-209*
SET_HEADER procedure, *286-63*
SET_HEAT_MAP_ALL Procedure, *98-7*
SET_HEAT_MAP_START Procedure, *98-8*
SET_HEAT_MAP_TABLE Procedure, *98-8*
SET_HOST_ACL Procedure, *131-18*
SET_IDENTIFIER, *173-17*
SET_INDEX_STATS procedures, *192-220*
SET_INITIAL_CONSUMER_GROUP procedure, *158-26*
SET_JAVA_LOGGING_LEVEL procedure, *124-23*
SET_KEY_COLUMNS procedure, *24-45*
SET_LAST_ARCHIVE_TIMESTAMP procedure, *31-28*

SET_LOB_INFORMATION member procedure, *310-10*
SET_LOB_OFFSET member procedure, *310-10*
SET_LOB_OPERATION_SIZE member procedure, *310-10*
SET_LOCATION_URI procedure, *143-28*
SET_LOG_LEVEL Procedure, *125-42*
SET_LOG_SPEC procedure, *56-17*
SET_LOGON_USER member procedure, *310-2*
SET_MAILHOST Procedure, *27-1*
SET_MAILPORT Procedure, *27-2*
SET_MAX_STREAMS_POOL Procedure, *26-47*
SET_MESSAGE_TRACKING Procedure, *239-88*
SET_METADATA_LINKED procedure, *140-8*, *141-5*
SET_MIN_STREAMS_POOL Procedure, *26-47*
SET_MODULE procedure, *23-5*
SET_MVREF_STATS_PARAMS procedure, *130-3*
SET_NLS Procedure, *173-18*
SET_NLS_PARAMETER Procedure, *68-6*
SET_NOT_MODIFIED Procedure, *244-3*
SET_NULL_COLUMN_VALUES_TO_EXPR Procedure, *77-3*
SET_OBJECT_NAME member procedure, *310-34*
SET_OBJECT_OWNER member procedure, *310-34*
SET_OBJECT_TYPE member procedure, *310-2*
SET_OER_BREAKPOINT function, *67-32*
SET_OPTION Procedure, *125-42*
SET_P1_RESOURCES Procedure, *38-6*
SET_PARAM procedure, *192-223*
SET_PARAM Procedure, *154-13*
SET_PARAMETER procedure, *44-31*, *56-17*, *239-89*
  apply process, *24-48*
SET_PARAMETER Procedure, *163-5*
SET_PERSISTENT_CONN_SUPPORT procedure, *286-64*
SET_PLSQL_LOGGING_LEVEL procedure, *124-23*
SET_PLSQL_TRACE procedure, *199-6*
SET_PREFS Procedure, *180-7*
SET_PROCESSING_RATE Procedure, *192-227*
SET_PROFILE_EXPLICIT procedure, *140-9*, *141-6*
SET_PROTECTION_REALM procedure, *250-4*
SET_PROXY procedure, *124-24*, *286-67*
SET_PURGE_JOB_INTERVAL procedure, *31-30*
SET_PURGE_JOB_STATUS procedure, *31-31*
SET_REPERROR_HANDLER procedure, *24-69*
SET_REPLAY_DIRECTORY Procedure, *216-32*
SET_REPLAY_TIMEOUT Procedure, *216-32*
SET_RESPONSE_ERROR_CHECK procedure, *286-68*
SET_ROLE Procedure, *173-18*
SET_ROLE_EXPLICIT procedure, *140-9*, *141-6*
SET_ROWID_THRESHOLD Procedure, *50-19*
SET_SCHEMA_INSTANTIATION procedure, *24-70*
SET_SCHEMA_PREFS procedure, *192-229*
SET_SEGADV_ATTRIB Procedure, *179-11*
SET_SENDFROM Procedure, *27-2*
SET_SESSION_LONGOPS procedure, *23-5*
SET_SESSION_TIMEOUT procedure, *161-3*
SET_SOURCE_DATABASE_NAME member procedure, *310-34*
SET_SQL_MAPPING Procedure, *216-33*
SET_SQL_TRACE Procedure, *173-18*
SET_SQLWKLD_PARAMETER Procedure, *17-41*
SET_SURROGATE_CONTROL Procedure, *244-4*
SET_SYSTEM_DEFAULT Procedure, *130-4*
SET_SYSTEM_STATS procedure, *192-236*
SET_TABLE_INSTANTIATION procedure, *24-73*
SET_TABLE_PREFS procedure, *192-237*
SET_TABLE_STATS procedure, *192-246*
SET_TAG member procedure, *310-34*
SET_TAG procedure, *239-90*
SET_TASK_PARAMETER Procedure, *17-43*
SET_THRESHOLD procedure, *171-8*
SET_TIMEOUT function, *67-32*
SET_TIMEOUT procedure, *161-4*
SET_TIMEOUT_BEHAVIOUR procedure, *67-33*
SET_TRACEFILE_IDENTIFIER procedure, *209-4*
SET_TRANSFER_TIMEOUT procedure, *286-68*
SET_TRANSFORM procedure, *58-65*
SET_TRIGGER_FIRING_PROPERTY procedure, *66-7*
SET_TUNING_TASK_PARAMETER Function, *190-86*
SET_UP_QUEUE procedure, *239-91*
SET_UPDATE_CONFLICT_HANDLER procedure, *24-75*
SET_USER_EXPLICIT procedure, *140-10*, *141-7*
SET_USER_MAPPING Procedure, *216-34*
SET_VALUE function, *67-33*
SET_VALUE member procedure, *310-10*
SET_VALUE_DEPENDENCY procedure, *24-79*
SET_VALUES member procedure, *310-10*
SET_WALLET procedure, *286-69*
SET_WALLET_ACL Procedure, *131-19*
SET_WARNING_SETTING_STRING procedure, *213-5*
SET_WATERMARK Procedure, *26-47*
SET_XML_INFORMATION member procedure, *310-10*
SET* member procedures
  of ANYDATA TYPE, *302-8*
  of ANYDATASET TYPE, *303-9*
SETACL Procedure, *223-19*, *225-19*
SETANYSTORAGE Procedure, *233-14*
SETATTRIBUTE procedure, *228-77*

SETATTRIBUTENODE function, *228-78*
SETAUTHOR Procedure, *225-19*
SETBASEDIR procedure, *231-7*
SETCHARACTERSET Procedure, *225-20*
SETCHARSET Procedure, *228-78*
SETCOMMENT Procedure, *225-20*
SETCONTENT Procedures, *225-20*
SETCONTENTTYPE Procedure, *114-61*, *225-21*
SETCONVERTSPECIALCHARS procedure, *229-7*
SETCUSTOMMETADATA Procedure, *225-21*
SETDATA procedure, *228-79*
SETDEFAULTACL Procedure, *62-53*
SETDEFAULTASOF Procedur, *62-54*
SETDEFAULTCONTEXT Procedure, *62-54*
SETDEFAULTOWNER Procedure, *62-55*
SETDEFAULTPRINCIPAL Procedure, *62-55*
SETDEFAULTTABLE Procedure, *233-15*
SETDISPLAYNAME Procedure, *225-22*
SETDOCTYPE procedure, *231-8*
SETDOCTYPE Procedure, *228-79*
SETDVERSION procedure, *228-84*
SETERRORLOG procedure, *231-8*, *238-7*
SETFTPPORT Procedure, *221-13*
SETHTTPCONFIGREALM Procedure, *221-13*
SETHTTPPORT Procedure, *221-13*
SETHTTPSPORT Procedure, *221-14*
SETINFO member procedure
  of ANYTYPE TYPE, *304-2*
SETKEYCOLUMN procedure, *236-5*
SETLANGUAGE Procedure, *225-22*
SETLISTENERENDPOINT Procedure, *221-14*
SETLISTENERLOCALACCESS Procedure, *221-15*
SETMAXROWS procedure, *229-8*
SETNAMEDITEM function, *228-80*
SETNODEVALUE procedure, *228-81*
SETNODEVALUEASBINARYSTREAM Function ((amp)) Procedure, *228-81*
SETNODEVALUEASCHARACTERSTREAM Function ((amp)) Procedure, *228-82*
SETNULLHANDLING Procedure, *229-8*
SETOPTIONS Procedure, *114-62*
SETOUTOFLINE Procedure, *233-15*
SETOWNER Procedure, *225-23*
SETPARAM procedure, *238-8*
SETPATH Procedure, *63-20*
SETPATH Procedures, *62-56*
SETPREFIX procedure, *228-83*
SETPRESERVEWHITESPACE procedure, *231-8*
SETREMOTEHTTPPORT Procedure, *221-15*
SETREMOTEHTTPSPORT Procedure, *221-15*
SETRENDERPATH Procedure, *227-24*
SETRENDERSTREAM Procedure, *227-25*
SETROWSETTAG procedure, *229-9*
SETROWTAG procedure, *236-5*
SETSCHEMAANNOTATATIONS Procedure, *233-17*
SETSCHEMAVALIDATED procedure, *315-9*
SETSKIPROWS procedure, *229-10*
SETSQLCOLLTYPE Procedure, *233-18*
SETSQLNAME Procedure, *233-19*
SETSQLTYPE Procedure, *233-20*
SETSQLTYPEMAPPING Procedure, *233-21*
SETSTANDALONE procedure, *228-83*
SETSTATS Procedure, *62-56*
SETSTOREPROPERTY Procedure, *64-13*
SETTABLEPROPS Procedure, *233-21*
SETTIMESTAMPWITHTIMEZONE Procedure, *233-22*
SETTRACE Procedure, *62-57*
SETUPDATECOLUMN procedure, *236-5*
SETVALIDATIONMODE procedure, *231-9*
SETVALUE procedure, *228-84*
SHOW_BREAKPOINTS procedures, *67-35*
SHOW_EXTENDED_STATS_NAME Function, *192-249*
SHOW_FRAME_SOURCE procedure, *67-35*
SHOW_SOURCE procedures, *67-36*
SHOW_STATS procedure, *297-14*
SHOW_STATS_HTML procedure, *297-16*
SHOWPAGE procedure, *252-11*
SHOWSOURCE procedure, *252-12*
SHOWWARNINGS procedure, *231-9*, *238-8*
SHRINK_TABLESPACE procedure, *178-11*
SHUTDOWN Procedure, *125-44*
SIGN function, *52-27*
SIGNAL procedure, *19-6*
SIGNATURE procedure, *252-12*
SIMPLE_SQL_NAME Function, *30-5*
SIZES procedure, *175-7*
SKIP_CORRUPT_BLOCKS procedure, *156-10*
SLEEP procedure, *173-19*
SMALL function
  of HTF package, *242-64*
SMALL procedure
  of HTP package, *243-63*
SODA Type
  REMOVE_ONE Function, *313-10*
  REPLACE_ONE Function, *313-10*
  REPLACE_ONE_AND_GET Function, *313-11*
SODA Types, *313-1*
  ACQUIRE_LOCK Function, *313-20*
  AS_OF_SCN Function, *313-21*
  AS_OF_TIMESTAMP Function, *313-21*
  CLOSE Function, *313-30*
  COUNT Function, *313-22*
  CREATE_INDEX Function, *313-3*
  CREATE_VIEW_FROM_DG Function, *313-4*
  DROP_INDEX Function, *313-5*
  FILTER Function, *313-22*
  FIND Function, *313-6*

SODA Types *(continued)*
FIND_ONE Function, *313-6*
GET_BLOB Function, *313-14*
GET_CLOB Function, *313-14*
GET_CREATED_ON Function, *313-15*
GET_CURSOR Function, *313-23*
GET_DATA_GUIDE Function, *313-7*, *313-23*
GET_DATA_TYPE Function, *313-15*
GET_INDEX Function, *313-7*
GET_JSON Function, *313-15*
GET_KEY Function, *313-16*
GET_LAST_MODIFIED Function, *313-16*
GET_MEDIA_TYPE Function, *313-16*
GET_METADATA Function, *313-8*
GET_NAME Function, *313-8*
GET_ONE Function, *313-24*
GET_VARCHAR2 Function, *313-17*
GET_VERSION Function, *313-17*
HAS_NEXT Function, *313-30*
HINT Function, *313-24*
INSERT_ONE Function, *313-8*
INSERT_ONE_AND_GET Function, *313-9*
KEY Function, *313-25*
KEYS Function, *313-25*
LIMIT Function, *313-26*
LIST_INDEXES Function, *313-9*
NEXT Function, *313-30*
Overview, *313-1*
REMOVE Function, *313-26*
REPLACE_ONE Function, *313-26*
REPLACE_ONE_AND_GET Function, *313-27*
SAMPLE Function, *313-28*
SAVE Function, *313-12*
SAVE_AND_GET Function, *313-12*
Security Model, *313-1*
SKIP Function, *313-28*
SODA_Collection_T Type, *313-2*
SODA_CURSOR_T Type Type, *313-29*
SODA_Document_T Function, *313-17*
SODA_Document_T Type, *313-13*
SODA_OPERATION_T Type Type, *313-18*
Summary, *313-2*
TRUNCATE Function, *313-13*
VERSION Function, *313-29*
SOURCE_LINES_T Table Type, *146-2*
SPACE_ERROR_INFO function, *161-4*
SPACE_USAGE procedure, *178-13*
SPACEUSAGE Procedure, *62-57*, *63-21*
split streams, *239-93*
SPLIT_STREAMS procedure, *239-93*
SPLITPATH Procedure, *223-19*
SPLITTEXT function, *228-85*
SQL Apply
managing logical standby databases, *117-1*
managing with DBMS_LOGSTDBY package, *117-1*
SQL generation, *310-10*
SQL Performance Reporting Subprograms, *190-10*
SQL_DIAGNOSE_AND_REPAIR Function, *186-29*
SQL_HASH Function, *185-15*
SQL_ID Function, *185-16*
SQL_OBJECT_NAME Function, *30-6*
SQL*Plus
creating a sequence, *1-4*
SQLID_TO_SQLHASH Function, *210-29*
SQLSET_ROW Object Type, *189-2*, *190-5*
SQLSET_ROW type, *189-2*
SQLTEXT_TO_SIGNATURE function, *190-88*
STACK_BIN_CAT procedure, *58-66*
STACK_BIN_NUM procedure, *58-68*
STACK_CLIP procedure, *58-70*
STACK_COL_REM procedure, *58-72*
STACK_MISS_CAT procedure, *58-74*
STACK_MISS_NUM procedure, *58-76*
STACK_NORM_LIN procedure, *58-78*
staging
queues
creating, *239-91*
removing, *239-77*
START__CONSOLIDATED_REPLAY Procedure, *216-34*
START_APPLY procedure, *24-80*
START_CAPTURE procedure, *44-46*
START_CAPTURE Procedure, *215-10*
START_CAPTURE procedure (of DBMS_SQL_FIREWALL), *183-19*
START_COVERAGE Function, *144-4*
START_OUTBOUND procedure, *239-97*
START_POOL Procedure, *49-8*
START_PROFILER functions and procedures, *149-7*
START_PROFILING Procedure, *93-4*
START_PROPAGATION procedure, *150-7*
START_REDEF_COLLECTION procedure, *177-9*
START_REDEF_TABLE procedure, *154-14*
START_REPLAY Procedure, *216-35*
START_REPORT_CAPTURE Procedure, *37-5*
START_SERVICE procedure, *172-16*
START_TRANSACTION Function, *200-9*
STARTTLS Function and Procedure, *298-22*
STARTUP Procedure, *125-44*
STARTUP_EXTPROC_AGENT procedure, *241-2*
STATUS Function, *160-9*
STATUS_LINE procedure, *252-13*
STEP_ID function, *200-10*
STOP_APPLY procedure, *24-81*
STOP_CAPTURE procedure, *44-46*
STOP_CAPTURE procedure (of DBMS_SQL_FIREWALL), *183-20*
STOP_COVERAGE Procedure, *144-5*

STOP_ILM Procedure
DBMS_ILM, *97-7*
STOP_OUTBOUND procedure, *239-97*
STOP_POOL Procedure, *49-9*
STOP_PROFILER function and procedure, *149-8*
STOP_PROFILING Procedure, *93-5*
STOP_PROPAGATION procedure, *150-7*
STOP_SERVICE procedure, *172-17*
STOP_TASK Procedure, *137-18*
STORE_T Record Type, *62-19*
STORE_VALUES procedure, *248-3*
stored outlines
DBMS_OUTLN, *135-1*
OUTLN_PKG package, *135-1*
STOREPUSH Procedure, *64-17*
STREAM2MULTI procedure, *251-4*
STRIKE function
of HTF package, *242-65*
STRIKE procedure
of HTP package, *243-64*
STRING function, *152-4*
STRONG function
of HTF package, *242-65*
STRONG procedure
of HTP package, *243-64*
STYLE function
of HTF package, *242-66*
STYLE procedure
of HTP package, *243-65*
SUB procedure
of HTP package, *243-65*
SUBMIT procedure, *105-8*
SUBMIT_PENDING_AREA procedure, *158-27*
SUBPROGRAM Function, *281-9*
SUBSTR function, *294-20*
SUBSTR Functions, *114-63*
SUBSTRINGDATA function, *228-85*
SUBTRACT procedure, *155-7*
Summary of DBMS_AQELM Subprograms, *27-1*
Summary of DBMS_AQMIGTOOL Subprograms, *29-2*
Summary of DBMS_DIMENSION Subprograms, *73-1*
Summary of DBMS_ERRLOG Subprograms, *79-1*
Summary of DBMS_MVIEW Subprograms, *129-2*
Summary of DBMS_MVIEW_STATS Subprograms, *130-1*
summary of DBMS_TF subprograms, *197-26*
Summary of DBMS_XDBZ Subprograms, *226-2*
Summary of DBMS_XMLDOM Subprograms, *228-13*
Summary of DBMS_XMLSCHEMA Subprograms, *232-4*
Summary of UTL_LMS Subprograms, *290-1*
SUMMARY procedure, *191-3*
SUP function
of HTF package, *242-67*
SUP procedure
of HTP package, *243-66*
SUPPORTED_TYPE function, *197-46*
SUSPEND_TRANSACTION Procedure, *200-10*
SWITCH_CONSUMER_GROUP_FOR_SESS procedure, *158-27*
SWITCH_CONSUMER_GROUP_FOR_USER procedure, *158-28*
SWITCH_CURRENT_CONSUMER_GROUP Procedure, *173-19*
SWITCH_DESTINATION procedure, *206-9*
SWITCH_PLAN procedure, *158-28*
SYNC_INTERIM_COLLETION procedure, *177-10*
SYNC_INTERIM_TABLE procedure, *154-16*
SYNCHRONIZE function, *67-37*
SYNCHRONIZE_PARTITIONS_FOR_HIVE Procedure, *89-3*
synchronous capture
altering, *44-11*
instantiation
preparing a table for, *44-29*
SYNCINDEX Procedure, *230-6*
synonyms
comparing, *47-1*
SYS.MGW_MQSERIES_PROPERTIES Object Type, *125-5*
SYS.MGW_PROPERTIES Object Type, *125-7*
SYS.MGW_PROPERTY Object Type, *125-8*
SYS.MGW_TIBRV_PROPERTIES Object Type, *125-9*

T

t_map_rec table type, *144-3*
TABLE_COLUMNS_T Type, *197-26*
TABLE_CREATE procedure, *56-19*
TABLE_T Type, *197-25*
TABLE_TO_COMMA Procedures, *210-30*
TABLECAPTION function
of HTF package, *242-67*
TABLECAPTION procedure
of HTP package, *243-66*
TABLECLOSE function
of HTF package, *242-68*
TABLECLOSE procedure
of HTP package, *243-67*
TABLEDATA function
of HTF package, *242-68*
TABLEDATA procedure
of HTP package, *243-67*
TABLEHEADER function
of HTF package, *242-69*
TABLEHEADER procedure
of HTP package, *243-68*

TABLEOPEN function
of HTF package, *242-70*
TABLEOPEN procedure
of HTP package, *243-68*
TABLEPRINT function, *252-13*
TABLEROWCLOSE function
of HTF package, *242-70*
TABLEROWCLOSE procedure
of HTP package, *243-69*
TABLEROWOPEN function
of HTF package, *242-71*
TABLEROWOPEN procedure
of HTP package, *243-69*
tables
comparing, *47-1*
table items as arrays, *182-27*
TABLESPACE_FIX_BITMAPS procedure, *179-12*
TABLESPACE_FIX_SEGMENT_STATES procedure, *179-13*
TABLESPACE_HEAT_MAP Procedure, *90-6*
TABLESPACE_MIGRATE_FROM_LOCAL procedure, *179-13*
TABLESPACE_MIGRATE_TO_LOCAL procedure, *179-14*
TABLESPACE_REBUILD_BITMAPS procedure, *179-15*
TABLESPACE_REBUILD_QUOTAS procedure, *179-15*
TABLESPACE_RELOCATE_BITMAPS procedure, *179-16*
TABLESPACE_VERIFY procedure, *179-17*
tags
GET_TAG function, *239-65*
SET_TAG procedure, *239-90*
TARGET_FILE function, *56-20*
TARGET_LOB function, *56-20*
TARGET_PROGRAM_RUNNING procedure, *67-38*
TARGET_TABLE function, *56-20*
TARGET_TRACE function, *56-21*
TELETYPE function
of HTF package, *242-71*
TELETYPE procedure
of HTP package, *243-70*
TERMINATE procedure, *152-5*
TEXT_DECODE function, *284-6*
TEXT_ENCODE function, *284-7*
The DBMS_XDB_REPOS package, *223-1*
TIME_TABLE Table Type, *182-23*
TIME_WITH_TIME_ZONE_TABLE Table Type, *182-23*
TIMESTAMP_TABLE Table Type, *182-23*
TIMESTAMP_WITH_LTZ_TABLE Table Type, *182-24*
TIMESTAMP_WITH_TIME_ZONE_TABLE Table Type, *182-24*
TITLE function
of HTF package, *242-72*
TITLE procedure
of HTP package, *243-70*
TO_CURSOR_NUMBER Function, *182-62*
TO_REFCURSOR Function, *182-63*
TO_STRING function, *311-9*
TODATE function, *252-16*
TOOBJECT procedure, *315-9*
TOUCHRESOURCE Procedure, *223-20*
TRACE procedure, *55-5*, *209-4*
TRACE Procedure, *62-58*
TRACEENABLED Function, *62-59*
tracking LCRs, *239-64*
TRANSACTION_BACKOUT Procedures, *82-7*
TRANSFER_STATS Procedure, *192-249*
TRANSFER_UNIFIED_AUDIT_RECORDS procedure, *31-32*
TRANSFORM function, *315-9*
transformations
rule-based
adding a column, *239-5*
deleting a column, *239-60*
keeping columns, *239-66*
renaming a column, *239-82*
renaming a schema, *239-84*
renaming a table, *239-86*
TRANSFORMNODE function, *238-9*
TRANSLATE function, *294-21*, *311-10*
TRANSLATE_ERROR Procedure, *185-17*
TRANSLATE_SQL Procedure, *185-18*
TRANSLITERATE Function, *294-23*
TRANSPORT_SET_CHECK procedure, *204-2*
TRIM Procedures, *114-65*
TUNE_MVIEW Procedure, *17-57*
TYPE_BUILD function, *56-21*
TYPE_DIMENSION_COMPILE function, *56-22*
TYPE_OPERATIONS function, *56-22*
TYPE_REJECTED_RECORDS function, *56-23*
types
COLUMN_T, *197-23*
DBMS_HIERARCHY package, *91-1*
MGD_ID, *311-1*, *311-3*
MGD_ID_COMPONENT, *311-2*
MGD_ID_COMPONENT_VARRAY, *311-2*
TABLE_T, *197-25*

U

ULISTCLOSE function
of HTF package, *242-72*
ULISTCLOSE procedure
of HTP package, *243-71*
ULISTOPEN function
of HTF package, *242-73*

ULISTOPEN procedure
of HTP package, *243-71*
UNASSIGN_ACL Procedure, *131-20*
UNASSIGN_WALLET_ACL Procedure, *131-21*
UNCHECKOUT function, *224-7*
UNCL_ARRAY Table Type, *210-4*
UNCONFIGURE_NODE procedure, *206-10*
UNDERLINE function
of HTF package, *242-73*
UNDERLINE procedure
of HTP package, *243-72*
UNDO_ADVISOR function, *207-6*
UNDO_AUTOTUNE function, *207-7*
UNDO_HEALTH function, *207-7*
UNDO_INFO function, *207-9*
UNESCAPE function, *300-4*
UNESCAPEURI function, *307-16*
UNIFORM_DIST_FIT procedure, *191-4*
UNIQUE_SESSION_ID Function, *173-21*
UNIQUE_SESSION_NAME function, *143-29*
UNIT_LINE Function, *281-8*
UNKEEP procedure, *175-7*
UNLOCK_MAP procedure, *193-7*
UNLOCK_PARTITION_STATS Procedure, *192-250*
UNLOCK_SCHEMA_STATS procedure, *192-251*
UNLOCK_TABLE_STATS procedure, *192-251*
UNLOCKPATH Procedure, *62-60*, *63-22*
UNLOCKRESOURCE Function, *223-20*
UNMAP_DAD Procedure, *78-15*
UNMARKHOT Procedure, *175-8*
UNMOUNTSTORE Procedure, *62-60*
UNMOUNTVOLUME Procedure, *75-2*
UNPACK_MESSAGE procedures, *143-29*
UNPACK_STGTAB procedure, *189-32*
UNPACK_STGTAB_BASELINE Function, *181-29*
UNPACK_STGTAB_DIRECTIVE Function, *180-8*
UNPACK_STGTAB_SQLPATCH Procedure, *186-30*
UNPACK_STGTAB_SQLPROF procedure, *190-89*
UNPACK_STGTAB_SQLSET procedure, *190-90*
UNREGISTER_DEPENDENT_OBJECT procedure, *154-16*
UNREGISTER_FOREIGN_QUEUE Procedure, *125-45*
UNREGISTER_NODE procedure, *206-10*
UNREGISTER_REMOTE_DATABASE procedure, *217-38*
UNREGISTERSTORE Procedure, *62-61*
UNREGISTERURLHANDLER procedure, *307-17*
UNSCHEDULE_PROPAGATION Procedure, *125-46*
UNUSED_SPACE procedure, *178-16*
UPDATE_ALLOW_LIST_ENFORCEMENT procedure (of DBMS_SQL_FIREWALL), *183-20*
UPDATE_BY_CAT procedure, *135-4*
UPDATE_CATEGORY Procedure, *158-29*
UPDATE_CDB_AUTOTASK_DIRECTIVE Procedure, *158-29*
UPDATE_CDB_DEFAULT_DIRECTIVE Procedure, *158-30*
UPDATE_CDB_PLAN Procedure, *158-31*
UPDATE_CDB_PLAN_DIRECTIVE Procedure, *158-32*
UPDATE_CDB_PROFILE_DIRECTIVE Procedure, *158-33*
UPDATE_CONSUMER_GROUP procedure, *158-34*
UPDATE_CREDENTIAL Procedure, *51-6*
UPDATE_DATAFILE_INFO procedure, *217-39*
UPDATE_FULL_REDACTION_VALUES Procedure, *153-25*
UPDATE_OBJECT procedure, *296-6*
UPDATE_OBJECT Procedure, *17-59*
UPDATE_OBJECT_INFO Function, *217-39*
UPDATE_PLAN procedure, *158-35*
UPDATE_PLAN_DIRECTIVE procedure, *158-36*
UPDATE_POLICY_EXPRESSION Procedure, *153-26*
UPDATE_REC_ATTRIBUTES Procedure, *17-61*
UPDATE_SIGNATURES procedure, *135-4*
UPDATE_SQLSET procedures, *189-33*, *190-91*
UPDATE_SQLWKLD_ATTRIBUTES Procedure, *17-62*
UPDATE_SQLWKLD_STATEMENT Procedure, *17-63*
UPDATE_TASK_ATTRIBUTES Procedure, *17-65*
UPDATERESOURCEMETADATA Procedures, *223-20*
UPDATEXML function, *236-6*
UPGRADE_AW Procedure, *54-46*
UPGRADE_DATABASE Procedure, *76-7*
UPGRADE_SCHEMA Procedure, *76-8*
UPGRADE_STAT_TABLE procedure, *192-252*
UPGRADE_TABLE Procedure, *76-10*
UPGRADE_VALIDATE_LOG_TABLE procedure
of DBMS_HIERARCHY package, *91-7*
URI Types
description, *307-1*
UriFactory package, *307-15*
ESCAPEURI function, *307-16*
GETURL function, *307-16*
methods, *307-15*
REGISTERURLHANDLER procedure, *307-17*
UNESCAPEURI function, *307-16*
UNREGISTERURLHANDLER procedure, *307-17*

UriType supertype, *307-1*
GETBLOB function, *307-2*
GETCLOB function, *307-3*
GETCONTENTTYPE function, *307-3*
GETEXTERNALURL function, *307-3*
GETURL function, *307-4*
GETXML function, *307-4*
methods, *307-1*
UROWID_TABLE Table Type, *182-24*
USE_FILTER_SET Procedure, *216-35*
USE_ROLLBACK_SEGMENT procedure, *200-11*
USEBINARYSTREAM Function, *228-86*
USEDPORT Procedure, *221-16*
USEITEMTAGSFORCOLL procedure, *229-10*
USENULLATTRIBUTEINDICATOR procedure, *229-11*
USER_EXPORT procedures, *105-10*
UTL Streams Types, *314-1*
UTL_BINARYINPUTSTREAM Type, *314-1*
UTL_BINARYOUTPUTSTREAM Type, *314-2*
UTL_CHARACTERINPUTSTREAM Type, *314-3*
UTL_CHARACTEROUTPUTSTREAM Type, *314-4*
UTL_I18N package
DETECT_CHARSET procedure, *287-7*
DETECT_LANGUAGE procedure, *287-9*
DETECT_LANGUAGE_CHARSET procedure, *287-11*
ESCAPE_REFERENCE function, *287-14*
GET_COMMON_TIME_ZONES function, *287-15*
GET_DEFAULT_CHARSET function, *287-15*
GET_DEFAULT_ISO_CURRENCY function, *287-16*
GET_DEFAULT_LINGUISTIC_SORT function, *287-17*
GET_LOCAL_LANGUAGES function, *287-17*
GET_LOCAL_LINGUISTIC_SORTS function, *287-18*
GET_LOCAL_TERRITORIES function, *287-19*
GET_LOCAL_TIME_ZONES function, *287-20*
GET_MAX_CHARACTER_SIZE function, *287-21*
GET_TRANSLATION function, *287-22*
MAP_CHARSET function, *287-23*
MAP_FROM_SHORT_LANGUAGE function, *287-24*
MAP_LANGUAGE_FROM_ISO function, *287-25*
MAP_LOCALE_TO_ISO function, *287-26*
MAP_TERRITORY_FROM_ISO function, *287-27*
MAP_TO_SHORT_LANGUAGE function, *287-27*
RAW_TO_CHAR function, *287-28*
UTL_I18N package *(continued)*
RAW_TO_NCHAR function, *287-30*
STRING_TO_RAW function, *287-31*
TRANSLITERATE function, *287-32*
UNESCAPE_REFERENCE function, *287-34*
VALIDATE_CHARACTER_ENCODING functions, *287-35*
UTL_LMS package
FORMAT_MESSAGE function, *290-1*
GET_MESSAGE function, *290-2*
UTL_MATCH package, *292-1*
UTL_RAW package, *294-1*
UTL_RECOMP package, *295-1*
UTL_RPADV package, *297-1*
UUDECODE function, *284-8*
UUENCODE function, *284-9*

V

v$vpd_policies, *162-3*
VALIDATE Procedure, *210-30*, *234-2*
VALIDATE_ANALYTIC_VIEW function
of DBMS_HIERARCHY package, *91-7*
VALIDATE_CHECK_SUCCESS function
of DBMS_HIERARCHY package, *91-9*
VALIDATE_DIMENSION procedure, *73-2*
VALIDATE_HIERARCHY function
of DBMS_HIERARCHY package, *91-10*
VALIDATE_PENDING_AREA procedure, *158-40*
VALIDATE_REPORT function, *108-5*
VALIDATE_REWRITE_EQUIVALENCE Procedure, *16-6*
VALIDATE_SCHEME function, *124-25*
VALIDATE_XML procedure, *54-47*
VALUE functions, *152-5*
VALUEOF procedure, *238-9*
VARCHAR2_TABLE Table Type, *182-24*
VARCHAR2A Table Type, *182-24*
VARCHAR2S Table Type, *182-25*
VARIABLE function
of HTF package, *242-74*
VARIABLE procedure
of HTP package, *243-72*
VARIABLE_VALUE procedures, *182-65*
VARIABLE_VALUE_PKG Procedure, *182-66*
VECTOR_TABLE Table Type, *182-25*
VERBOSE_ACTION function, *56-16*
VERBOSE_DEBUG function, *56-15*
VERBOSE_INFO function, *56-16*
VERBOSE_NOTICE function, *56-15*
VERBOSE_STATS function, *56-14*
VERIFY function, *52-29*
VERIFY_BLOCKCHAIN_LIFESPAN Procedure, *83-12*
VERIFY_VALUES function, *248-4*
VERSION function, *56-23*

views
comparing, *47-1*
virtual dependency definitions
object dependencies
creating, *24-19*
dropping, *24-22*
value dependencies, *24-79*
Virtual Private Database. See VPD, *162-1*
VPD
viewing current cursors and policy predicates, *162-3*
VPD use of DBMS_RLS, *162-1*
VRFY function, *298-23*

W

WAIT_ON_PENDING_DML Procedure, *210-31*
WAITANY procedure, *19-6*
WAITONE procedure, *19-7*
WBR function
of HTF package, *242-74*
WBR procedure
of HTP package, *243-73*
WEIBULL_DIST_FIT procedure, *191-5*
WHAT procedure, *105-10*
WHO_CALLED_ME procedure, *252-17*
WRAP Functions, *66-8*
WRITE Procedures, *114-66*
WRITE_DATA procedure, *298-24*
WRITE_END procedure, *120-4*
WRITE_LINE function, *299-17*
WRITE_LINE procedure, *286-70*
WRITE_RAW function, *299-18*
WRITE_RAW procedure, *286-71*
WRITE_RAW_DATA procedure, *298-25*
WRITE_TEXT function, *299-18*
WRITE_TEXT procedure, *286-72*
WRITEAPPEND Procedures, *114-68*
WRITES_FLUSH procedure, *121-2*
WRITETOBUFFER procedure, *228-86*
WRITETOCLOB procedure, *228-87*
WRITETOFILE procedure, *228-88*

X

XA_COMMIT Function, *219-6*
XA_END Function, *219-7*
XA_FORGET Function, *219-8*
XA_GETLASTOER Function, *219-8*
XA_PREPARE Function, *219-8*
XA_RECOVER Function, *219-9*
XA_ROLLBACK Function, *219-9*
XA_SETTIMEOUT Function, *219-10*
XA_START Function, *219-10*
XDBEvent Type Subprograms, *227-2*
XDBHandler Type Subprograms, *227-4*
XDBHandlerList Type Subprograms, *227-4*
XDBLink Type Subprograms, *227-5*
XDBPath Type Subprograms, *227-5*
XDBRepositoryEvent Type Subprograms, *227-3*
XDBSCHEMA_PREFIXES Function, *222-9*
XDBUriType, *307-11*
XDBURITYPE function, *307-15*
XDBUriType subtype, *307-11*
CREATEURI function, *307-12*
GETBLOB function, *307-12*
GETCLOB function, *307-13*
GETCONTENTTYPE function, *307-14*
GETEXTERNALURL function, *307-14*
GETURL function, *307-14*
GETXML function, *307-14*
methods, *307-11*
XDBURITYPE function, *307-15*
XDBZ Constants, *226-1*
XFORM_BIN_CAT procedure, *58-80*
XFORM_BIN_NUM procedure, *58-82*
XFORM_CLIP procedure, *58-84*
XFORM_COL_REM procedure, *58-86*
XFORM_EXPR_NUM procedure, *58-87*
XFORM_EXPR_STR procedure, *58-89*
XFORM_MISS_CAT procedure, *58-92*
XFORM_MISS_NUM procedure, *58-94*
XFORM_NORM_LIN procedure, *58-95*
XFORM_STACK procedure, *58-97*
XGBoost settings, *57-47*
XML_TO_LCR Function, *126-27*
XMLType
CREATENONSCHEMABASEDXML function, *315-2*
CREATESCHEMABASEDXML function, *315-2*
CREATEXML function, *315-3*
EXISTSNODE function, *315-5*
EXTRACT function, *315-5*
GETBLOBVAL function, *315-6*
GETCLOBVAL function, *315-6*
GETNUMBERVAL function, *315-6*
GETROOTELEMENT function, *315-7*
GETSCHEMAURL function, *315-7*
GETSTRINGVAL function, *315-7*
ISFRAGMENT function, *315-7*
ISSCHEMABASED function, *315-7*
ISSCHEMAVALID function, *315-8*
ISSCHEMAVALIDATED function, *315-8*
NAMESPACE function, *315-6*
SCHEMAVALIDATE procedure, *315-8*
SETSCHEMAVALIDATED procedure, *315-9*
TOOBJECT procedure, *315-9*
TRANSFORM function, *315-9*
XMLTYPE function, *315-10*
XMLTYPE function, *315-10*
XPATH2TABCOLMAPPING Function, *235-10*

XRANGE function, *294-25*
XSD_ATTRIBUTE Function, *222-10*
XSD_COMPLEX_TYPE Function, *222-10*
XSD_ELEMENT Function, *222-10*
XSD_GROUP Function, *222-10*
XSTORE_CLEAR, *197-47*
XSTORE_EXISTS function, *197-48*
XSTORE_GET procedure, *197-48*
XSTORE_REMOVE procedure, *197-51*
XSTORE_SET procedure, *197-52*
XStream
  DBMS_XSTREAM_ADM package, *239-1*
  privileges, *240-1*
  removing configuration, *239-80*
  security, *239-2*
XStream administrator, *240-1*
XSTREAM_APPLY procedure, *240-2*